

Precision Weak Gravitational Lensing Using Velocity Fields: Fisher Matrix Analysis

DAVID WITTMAN¹ AND MATTHEW SELF¹

¹*Physics Department, University of California, Davis, CA 95616*

ABSTRACT

Weak gravitational lensing measurements based on photometry are limited by shape noise, the variance in the unknown unlensed orientations of the source galaxies. If the source is a disk galaxy with a well-ordered velocity field, however, velocity field data can support simultaneous inference of the shear, inclination, and position angle, virtually eliminating shape noise. We use the Fisher Information Matrix formalism to forecast the precision of this method in the idealized case of a perfectly ordered velocity field defined on an infinitesimally thin disk. For nearly face-on targets one shear component, γ_x , can be constrained to $0.003 \frac{I_0}{n_{\text{pix}}} \frac{25}{n_{\text{pix}}}$ where I_0 is the S/N of the central intensity pixel and n_{pix} is the number of pixels across a diameter enclosing 80% of the light. **This precision degrades with inclination angle, by a factor of three by $i=50^\circ$. Uncertainty on the other shear component, γ_+ , is about 1.5 (7) times larger than the γ_x uncertainty for targets at $i = 10^\circ$ (50°).** For arbitrary galaxy position angle on the sky, these forecasts apply not to γ_+ and γ_x as defined on the sky, but to two eigenvectors in $(\gamma_+, \gamma_x, \mu)$ space where μ is the magnification. We also forecast the potential of less expensive partial observations of the velocity field such as slit spectroscopy. We conclude by outlining some ways in which real galaxies depart from our idealized model and thus create random or systematic uncertainties not captured here. **In particular, our forecast γ_x precision is currently limited only by the data quality rather than scatter in galaxy properties because the relevant type of scatter has yet to be measured.**

Keywords: gravitational lensing: weak

1. INTRODUCTION

Weak gravitational lensing is a key technique in modern cosmology, in which the gravitational field of a celestial object is reconstructed from the distortion it imprints on background sources of light; see [Bartelmann & Maturi \(2017\)](#) for a recent review. The distortion is described in terms of shear, defined as stretching the image in one direction and compressing it in the perpendicular direction, and convergence, defined as an isotropic stretching. Shear can be depicted as a headless vector with a dimensionless magnitude and a position angle (PA) on the sky modulo 180° , or in terms of two components separated by 45° in PA. Shear is inferred from the observed shapes of source galaxies, under the assumption that galaxies have no preferred orientation in the absence of lensing. The fundamental source of noise in this approach is the large intrinsic *scatter* in galaxy orientations, called shape noise. This scatter is such that the shear on a single galaxy is uncertain by at least 0.2 in each component, while the relevant signal is usually much smaller. Averaging over many source galaxies in a given patch of sky builds the signal-to-noise ratio (S/N),

but correspondingly decreases the angular resolution of the reconstruction.

Techniques to measure convergence also face substantial amounts of noise. Convergence leads to magnification, which increases the flux of sources while decreasing the effective area of sky probed. This can shift the counts of sources as a function of apparent magnitude (eg, [Morrison et al. 2012](#); [Garcia-Fernandez et al. 2016](#)). This is again a technique that relies on aggregation of many sources due to the low information content of each individual source.

To increase the information content of an individual source, we must know more about its unlensed state. A recent idea in this regard is that a source with a well-ordered velocity field, such as a rotating disk galaxy, can potentially provide that information. The velocity in each pixel provides a tag that helps place that pixel in the source plane—a more specific tag than is possible with the intensity field. Although velocity measurements are more expensive than intensity measurements, the gain in per-galaxy precision is potentially quite large. This paper aims to quantify that gain with a Fisher information matrix analysis.

First, we briefly outline the history of the velocity field idea. Blain (2002) first recognized that shear perturbs the symmetry of the velocity field. He used a rotating ring toy model to show how velocity measurements could constrain the shear component at 45° to the source galaxy’s unlensed photometric axes, which we call γ_x . Morales (2006) extended the velocity-field idea to full disk galaxies, and provided a clear picture of how γ_x causes the major and minor velocity axes to deviate from perpendicularity. A version of this method has been implemented by de Burgh-Day et al. (2015), who infer the shear by determining the transformation required to restore symmetry to the velocity map. They find that shears as small as 0.01 are measurable in simulations, and they find shears consistent with zero, with uncertainties ~ 0.01 , on unlensed nearby disk galaxies. However, their approach is still insensitive to the component of shear along the unlensed photometric axes because that component, which we call γ_+ , preserves the symmetry of the velocity field.

γ_+ does change the observed axis ratio, so Huff et al. (2013) proposed constraining this component as follows. They propose predicting the total rotation speed of the galaxy using the Tully-Fisher relation (Tully & Fisher 1977), then comparing this prediction with the measured line-of-sight rotation speed to find the inclination of the disk. Assuming the disk to be circular when viewed face-on, the inclination uniquely predicts the unlensed axis ratio, which effectively removes the problem of shape noise. The Huff et al. (2013) goal of designing an efficient large cosmic shear survey led them to propose minimal velocity-field measurements per galaxy (slit spectra along the apparent photometric axes) and to assume approximations, such as the low-shear limit and negligible magnification, that may fail in more general lensing situations. Considering that de Burgh-Day et al. (2015) needed the full velocity field of a very well-resolved nearby galaxy to infer γ_x , it is not clear that shear could be measured precisely using only crossed slits along the photometric axes. Nevertheless, the insight of Huff et al. (2013)—that symmetry is not the only source of information in the velocity field—is potentially powerful and deserves further investigation.

This paper uses the Fisher Information Matrix formalism to forecast the best achievable performance in the case of perfectly ordered rotation and an infinitesimally thin disk. This is highly idealized, but the point is to determine whether the method is promising enough to justify further development. Therefore, we forecast the best possible performance across a wide range of scenarios: from zero-shear lines of sight on up to higher-shear lines of sight, from nearly face-on targets to nearly

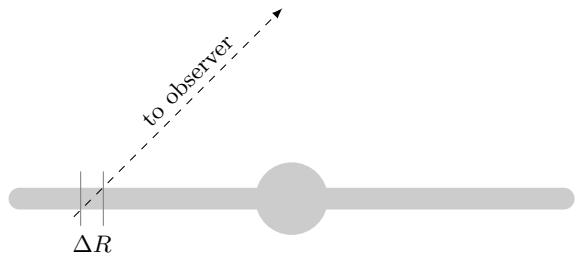


Figure 1. Effect of finite disk thickness. A line of sight probes particles at a range of cylindrical galactocentric distances R depending on their height above or below the midplane. Where the rotation curve is approximately linear in R across the range ΔR , the above- and below-plane contributions are approximately equal and opposite, which preserves the mean velocity but increases the linewidth. Hence to first order the disk can be modeled as an infinitesimally thin disk but with greater linewidth.

edge-on targets, from full velocity-field observations to crossed slits and so on.

The remainder of this paper is organized as follows. In §2 we describe and illustrate the method; in §3 we present the resulting forecasts; and in §4 we discuss the implications.

2. METHOD

We assume an infinitesimally thin disk galaxy with a polar (R, ϕ) coordinate system specifying particle locations. Viewed at inclination i (where $i = 0$ is face-on) but before lensing, we define an (x, y) coordinate system, in which

$$x = R \cos(\phi - \phi_0) \cos i \quad (1)$$

$$y = R \sin(\phi - \phi_0) \sin i \quad (2)$$

where ϕ_0 is the unlensed PA of the apparent major axis. The velocity field is assumed to be a function only of R , with measured line-of-sight velocity $v_{los} = v(R) \sin(\phi - \phi_0) \sin i$.

Note that, to first order, a finite-thickness disk can be modeled as an infinitesimally thin disk but with greater linewidth. Figure 1 illustrates the argument: stars along the line of sight above and below the disk depart from the midplane value of R in equal and opposite ways. Therefore the mean velocity for this line of sight is unchanged if the rotation curve is linear in R across the range of R probed by the line of sight. The line of sight does, however, encounter a wider range of velocities than would be the case for an infinitesimally thin disk, leading to a greater linewidth unless the rotation curve is approximately flat across the range of R probed by a given line of sight. Real galaxies will present additional complications, such as bulges and warps. We stress that

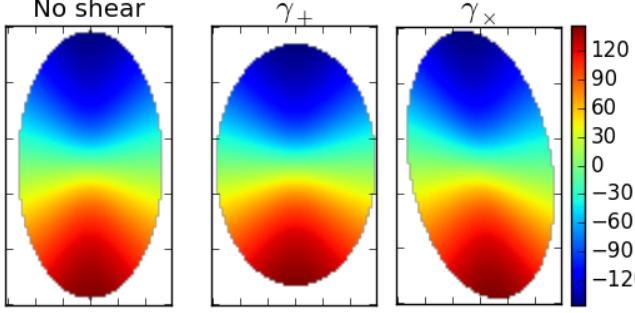


Figure 2. Velocity fields before lensing (left), after applying $\gamma_+ = 0.1$ (middle), and after applying $\gamma_x = 0.1$ (right). The galaxy has maximum rotation speed of 220 km/s and is inclined at 1 radian to the line of sight. The colorbar shows units of km/s.

our approach here is to explore the optimal case of a bulgeless, dynamically cold thin disk in order to establish the limits of this method, reveal parameter degeneracies and requirements for priors, and identify key assumptions that will need to be explored further.

Lensing transforms the coordinates described above to observed coordinates, which we denote with primes:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = A^{-1} \begin{bmatrix} x \\ y \end{bmatrix} \quad (3)$$

where

$$A^{-1} = \mu \begin{pmatrix} 1 - \kappa + \gamma_+ & -\gamma_x \\ -\gamma_x & 1 - \kappa - \gamma_+ \end{pmatrix} \quad (4)$$

Here κ is the convergence, which is proportional to the surface mass density; $\mu = \frac{1}{(1-\kappa)^2 - \gamma^2}$ is the magnification, and $\gamma = \sqrt{\gamma_+^2 + \gamma_x^2}$ is the magnitude of the shear. We choose to parametrize the shear in terms of γ_+ and γ_x , which are dimensionless quantities with identical ranges, rather than a magnitude and a PA. Then, the lensing matrix can be completed by specifying either κ or μ . We choose μ because prior information on μ is more likely to be available through other methods.

With this in mind, the left panel of Figure 2 shows an unlensed model velocity field for $i = 60^\circ$ and fortuitously aligned with the coordinate axes. The middle and right panels show the same field after lensing by γ_+ and γ_x respectively. (All fields in this figure are cropped at a consistent physical radius; this guides the eye but may overstate the power of the method, because such cuts and comparisons will not be available to the data analyst.) The right panel displays the asymmetry discussed in the introduction, which we will associate with γ_x throughout the paper. Our formalism defines the shear components with respect to sky coordinate axes rather than the galaxy axes, so in practice

the asymmetry-causing component need not be γ_x as defined on the sky. Although the physical distinction is between shear components aligned and not aligned with the apparent unlensed galaxy axes, we choose not to *define* the components this way because in practice the unlensed axes are unknown. By defining shear components on the sky, we adopt the basis in which shear will actually be used. That said, to highlight physical behaviors we will typically align the galaxy as in Figure 2 and refer to γ_x as causing the asymmetry.

A key assumption is that the unlensed velocity field has the symmetry shown. Under this assumption, the data analyst can determine γ_x because the relevant unlensed condition is known. The effect of γ_+ is to change the apparent axis ratio, so measuring γ_+ requires knowledge of the unlensed axis ratio. That axis ratio is set by the inclination, an effect distinct from that of γ_+ in that inclination *also* changes the line-of-sight velocity. It is conceptually useful to consider the extreme case of a uniform observed velocity field, from which we can deduce that the galaxy must be viewed face-on. This implies a unlensed axis ratio of unity, so we can deduce γ_+ from the observed axis ratio, with no shape noise.¹ The key is the ability to deduce a unlensed axis ratio from the velocity field amplitude; this is a way of restating the idea of Huff et al. (2013).

To go beyond this conceptual understanding we must choose quantitative models for the intensity and velocity fields. First, we define the parameter r_{80} , which is the radius that encircles 80% of the galaxy light. For an exponential disk, this is 2.99 times the exponential scale length. The intensity field is specified by $I = I_0 \exp(-\frac{2.99R}{r_{80}})$, where the parameter I_0 represents the central intensity. We set the intensity uncertainty in each pixel to unity, so I_0 represents the S/N of the intensity measurement in the central pixel. The intensity uncertainty field is uniform because sky noise, rather than photon noise from the galaxy itself, is the dominant uncertainty in broadband imaging of most galaxies. We set the fiducial value of I_0 to 90, which is a high S/N reflecting the fact that bright galaxies are the likeliest targets for integral field spectroscopy. The velocity uncertainty is set by $\sigma_{v,0}$, the uncertainty in the central pixel (with a fiducial value of 10 km/s) and grows exponentially with R because source photon noise is likely to be the limiting factor.

We adopt a simple arctan rotation curve: $v = v_{\max} \frac{2}{\pi} \arctan \frac{R}{r_0}$, where the factor $\frac{2}{\pi}$ ensures that $v \rightarrow v_{\max}$ as $R \rightarrow \infty$ given an arctan function that

¹ In practice, there will still be some uncertainty due to uncertainty in the intrinsic circularity of face-on disks.

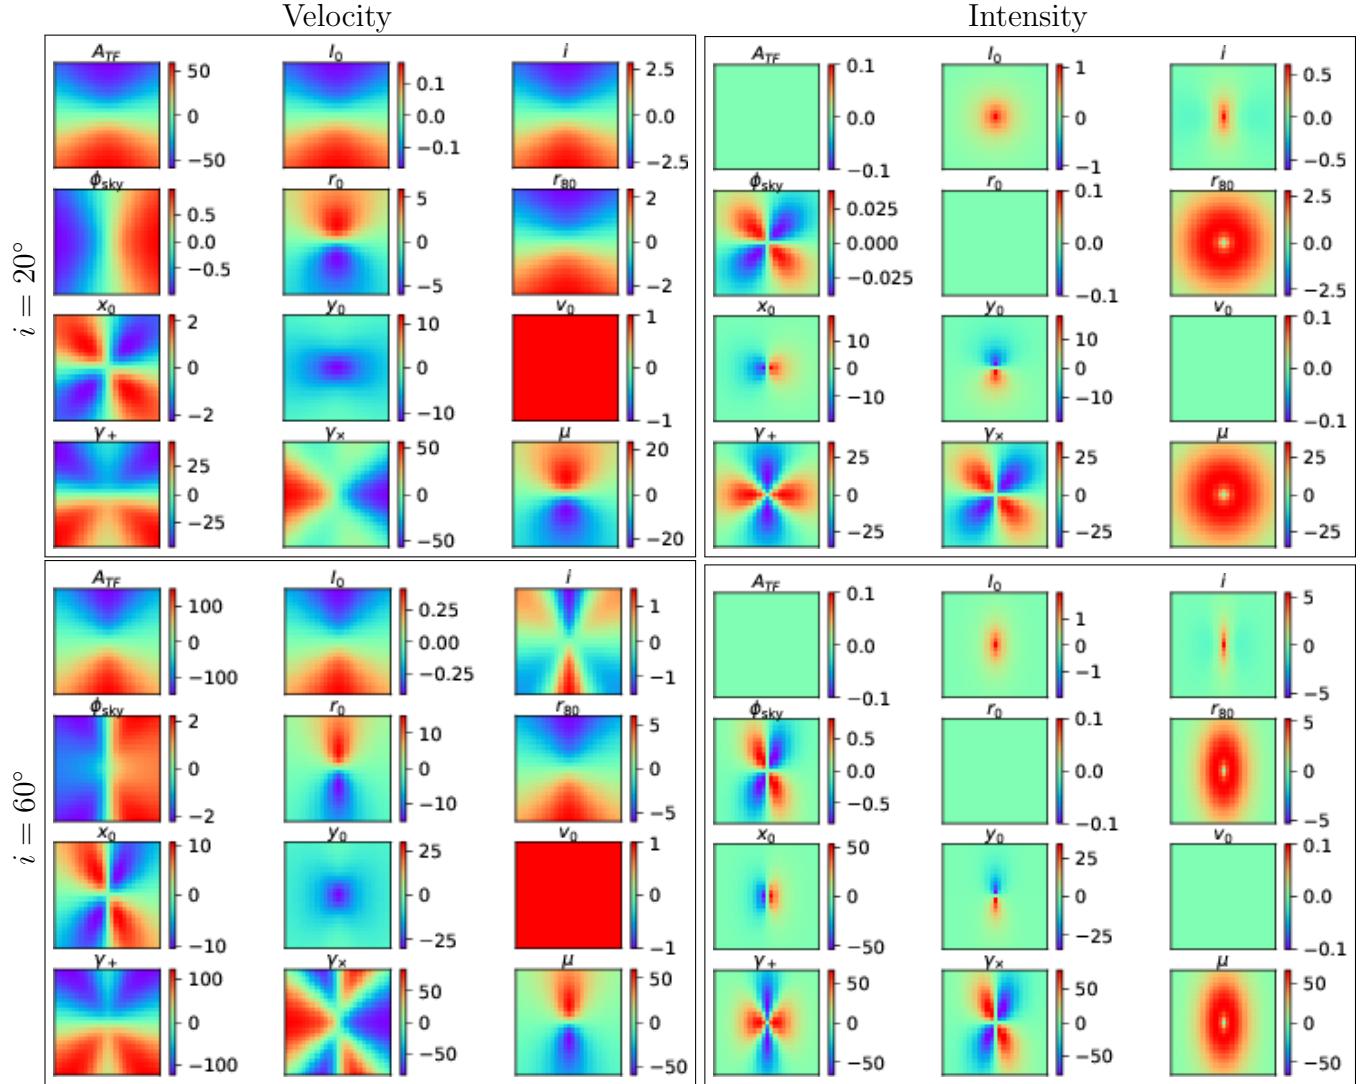


Figure 3. Partial derivatives of the velocity (left) and intensity (right) fields with respect to each of the parameters, for an inclination of 20° (top row) and 60° (bottom row). The colorbar units are km/s on the left, and arbitrary intensity units on the right. To keep the scales roughly the same across panels, we show the change in velocity per 0.01 change in shear and convergence. The form of the velocity field itself can be seen in the A_{TF} panels on the left because that field is linear in A_{TF} . Similarly, the form of the intensity field can be seen in the I_0 panels on the right.

returns radians. We also investigated the more complicated Universal Rotation Curve (URC; Persic et al. 1996; Salucci et al. 2007) and found the results to be nearly identical; a few minor differences will be discussed in §3.8. With either form, the rotation curve has a scale length independent of the scale length describing the intensity field. If these two scales were the same, the model would be more constrained and yield higher precision, but the scales do appear to differ in observed galaxies.

v_{\max} is related to the intensity field via the Tully-Fisher relation (TFR) as follows. The TFR empirically states that $L \propto v_{\max}^n$ where $n \approx 4$, with a scatter in luminosity or stellar mass of about 16% (Miller et al. 2011). This implies that at fixed L the scatter in v_{\max} is

about 4%. For an exponential disk, the total luminosity is $L \propto I_0 r_{80}^2$, so the TFR predicts $v_{\max} \propto (I_0 r_{80}^2)^{0.25}$. With our fiducial values of I_0 and r_{80} (12.5 pixels), we need $v_{\max} = 20(I_0 r_{80}^2)^{0.25}$ to produce a typical rotation speed around 200 km/s.² Hence we define a Tully-Fisher amplitude, A_{TF} , with a fiducial value of unity, such that $v_{\max} = 20A_{TF}(I_0 r_{80}^2)^{0.25}$. We then place a prior of $\pm 4\%$ on A_{TF} .

Table 1 summarizes the parameters for this model, including the nuisance parameters x_0, y_0, v_0 describing the galaxy position on the sky and systemic radial velocity.

² More precisely, $v_{\max} = 218$ in this case, but note that in the arctan model v_{\max} is reached only as $R \rightarrow \infty$.

The units listed in this table are relevant to the forecast precision plots presented below; units are omitted for dimensionless quantities. The results can be quite sensitive to the inclination angle i , so i will be varied in many plots rather than remaining fixed at a fiducial value.

We construct velocity and intensity fields extending to a radius of r_{80} , thus encompassing 25×25 pixels each. We compute partial derivatives numerically to a relative precision of order 10^{-11} using the algorithm in Section 5.7 of Press et al. (1992), which we re-implement in Python. Figure 3 shows the partial derivatives of the velocity and intensity fields with respect to each parameter at two different inclinations. These figures will help readers understand which parameters are highly correlated. Note that A_{TF} , I_0 , i , and r_{80} have nearly identical effects on the velocity field. For A_{TF} this is broken by its lack of effect on the intensity field, but I_0 and r_{80} also have nearly identical effects on that field—with opposite sign, but the sign is not relevant for determining degeneracy and correlation. The effect of i on the intensity field is not identical to that of I_0 , but there is a good deal of overlap, indicating that the three parameters I_0 , i , and r_{80} will be highly correlated. Magnification (μ) joins this family because its effect on both velocity and intensity fields is much like $-I_0$, and its effect on the intensity field is identical to changing the intensity scale length r_{80} . Finally, r_0 is linked with all these parameters because, as a rotation curve scale length, its effect on the velocity field is identical to that of magnification μ . The strength of these correlations will vary with the specific values of inclination, shear, and so on: Figure 3, for example, shows that by $i = 60^\circ$ perturbations in i affect the velocity field differently than perturbations in I_0 and r_{80} .

For any given value of i , we concatenate the velocity and intensity fields into a Python data structure representing a generalized data field we denote \vec{D} . Denoting the set of parameters as P , the Fisher matrix elements are then

$$\mathcal{F}_{ij} = \sum_{\text{pixels}} \vec{\sigma}^{-2} \left(\frac{\partial \vec{D}}{\partial P_i} \right) \left(\frac{\partial \vec{D}}{\partial P_j} \right) \quad (5)$$

where i and j index the parameters, and $\vec{\sigma}$ is the uncertainty field associated with the data field. We then invert the Fisher matrix to obtain the covariance matrix C . We also compute the correlation matrix $\rho \equiv D^{-1} C D^{-1}$ where $D \equiv \sqrt{\text{diag}(C)}$.

3. RESULTS

3.1. Degeneracies

We find that $\frac{\partial \vec{D}}{\partial \gamma_+}$ is a linear combination of the other partial derivative fields. The coefficients

depend on the parameter values themselves, but for concreteness we display the coefficients for our fiducial scenario at $i = 30^\circ$:

$$\frac{\partial \vec{D}}{\partial \gamma_+} + 199 \frac{\partial \vec{D}}{\partial i} - 180 \frac{\partial \vec{D}}{\partial I_0} + 4 \frac{\partial \vec{D}}{\partial r_0} + 12.5 \frac{\partial \vec{D}}{\partial r_{80}} - 6 \frac{\partial \vec{D}}{\partial A_{TF}} = 0 \quad (6)$$

Hence, a model can be transformed into another model with a different γ_+ value that predicts the same data, providing that we:

- increment i to preserve the apparent axis ratio despite the change caused by γ_+ . (In our setup, the unlensed apparent major axis is in the “ y ” direction while positive γ_+ acts to stretch the “ x ” direction, hence one must make the galaxy more edge-on to counteract positive γ_+ .)
- decrement I_0 to preserve the apparent surface brightness. (In our model the galaxy is transparent, so making it more edge-on had the side effect of increasing the apparent surface brightness.)
- increment r_{80} to preserve the observed angular size of the major axis of the intensity field. In concert with the change in inclination angle, this also preserves the apparent minor axis.
- increment r_0 to preserve the observed angular scale of the rotation curve’s rise. The coefficients on r_{80} and r_0 here are equal to their fiducial values, confirming that $d\gamma_+$ equals the fractional change in each apparent size, as it should when $\kappa = 0$.
- finally, we have to preserve the Tully-Fisher relation. To preserve the amplitude of the observed velocity-field pattern despite being more edge-on, our model must suppose a lower rotation speed, thus decreasing A_{TF} . Alternately, the same effect can be achieved by adjusting μ , which allows a lower-luminosity model to fit the intensity field.

The specific linear combination depends on the scenario, but always involves A_{TF} , μ , I_0 , i , r_0 , and r_{80} . It also involves ϕ_{sky} if the γ_x is nonzero, and γ_x if ϕ_{sky} is nonzero. We tested a parametrization in terms of reduced shear ($g_{+(*x)} \equiv \frac{\gamma_{+(*x)}}{1-\kappa}$) rather than shear. This did not change the set of interdependencies. We find similar dependencies parametrizing in terms of κ rather than μ .

Table 1. Model parameters

Symbol	Fiducial value	Unit	Description
A_{TF}	1	-	v_{\max} as a fraction of the Tully-Fisher prediction
I_0	90	-	intensity S/N at center
i	varies	deg	inclination angle
ϕ_{sky}	0	deg	sky position angle of unlensed major axis
r_0	4	pixel	rotation curve scale length
r_{80}	12.5	pixel	radius of 80% encircled light
x_0	0	pixel	center of galaxy in x coordinate
y_0	0	pixel	center of galaxy in y coordinate
v_0	0	km/s	galaxy systemic radial velocity
γ_+	0	-	shear parallel to sky coordinates
γ_x	0	-	shear at 45° to sky coordinates
μ	1	-	magnification
Data parameters			
n_{pix}	25	pixel	field diameter
$\sigma_{v,0}$	10	km/s	uncertainty in v , central pixel

This degeneracy prevents the Fisher matrix from being inverted.

The TFR should be effective at breaking this degeneracy, in the sense that a fractional change in γ_+ requires a large fractional change in A_{TF} , in tension with the TFR. But the fact that a tweak in μ can substitute for a tweak in A_{TF} leaves the Fisher matrix still noninvertible.

However, when *finite* steps are taken along the degeneracy direction, the data change quadratically with the step size, suggesting that data can indeed constrain the model. We have confirmed this with Markov Chain Monte Carlo (MCMC) explorations of the likelihood surface: our fiducial data constrain μ to slightly better than ± 0.1 at all inclinations. Hence we support the Fisher forecast by placing a prior of ± 0.1 on μ . The fact that data constrain μ as well as all the other parameters is potentially important and will be further explored in a subsequent paper using higher order expansions of the likelihood surface (Heavens 2016) and/or MCMC techniques.

The physical context is that $\mu = 1$ in the absence of lensing; only the densest lines of sight have μ approaching 2 or more; and for those lines of sight the fact that μ is high will generally be known in advance. We also note that the weak lensing formalism used here breaks down at high magnification. Specifically, we assume that the matrix A (hence the parameters γ_+ , γ_x , and μ) is constant over the extent of the target galaxy, and this is not generally the case along strongly lensed lines of sight. In those cases, more traditional strong-lensing techniques will be preferred, although it is possible that the veloc-

Table 2. Priors

Parameter	Width (Gaussian σ)
A_{TF}	0.04
μ	0.1

ity field can complement the intensity field in constraining the strong-lensing reconstruction (Rizzo et al. 2018). For all these reasons, the μ uncertainty on a typical weak lensing line of sight would approach 0.1 in any case, so our prior is a good match to the physical situation. Table 2 lists the priors applied as part of our standard forecast.

With this prior in place, we invert the Fisher matrix. Numerical instability in matrix inversion generally becomes important if the inverse of the condition number, the ratio between largest and smallest modulus eigenvalue, is not much larger than the inaccuracies in our knowledge of the matrix elements (Vallisneri 2008). With the μ prior we find condition numbers ranging from $\approx 10^{6.5}$ at $i = 10^\circ$ to $\approx 10^{11.5}$ at $i = 80^\circ$. The inverse of the latter overlaps the $\sim 10^{-11}$ uncertainty in our numerical differentiation cited above.³ Hence, we cannot make numerically stable forecasts for galaxies close to edge-on. We limit our forecasts to those with condition number $< 10^8$ ($i \leq 50^\circ$). Appendix A shows that forecasts with these condition numbers match very well with MCMC explorations of the likelihood

³ For standard double precision arithmetic, the relative rounding error is $\approx 10^{-16}$, so machine precision is a subdominant uncertainty here.

surface. In the verified results below, the shear constraints consistently degrade as the inclination increases from 10° to 50° , so there is little reason to push the forecast to higher inclination.

Figure 4 shows the resulting correlation matrices for low- and high-inclination cases. The parameters in the family discussed above (A_{TF} , I_0 , i , r_0 , r_{80} , γ_+ , and μ) are indeed correlated, with some increase in correlation at higher inclination. Separately, there is an anticorrelation between ϕ_{sky} and γ_x , which is moderately strong at $i = 20^\circ$ but quite strong at $i = 50^\circ$. This suggests that higher inclinations will yield looser constraints for both shear components. These correlations set the stage for understanding our primary products, forecasts of precision on each parameter.

3.2. Fiducial forecast

We repeat the process of building and inverting the Fisher matrix in order to present these forecasts as a function of i , as shown in Figure 5. The main features are:

- The i -dependence is dramatic: face-on targets yield much more information. This is perhaps counterintuitive because such a target will have a featureless velocity field, but in our idealized model such a featureless field carries the information that the unlensed image is exactly circular, which is most sensitive to shear.
- The γ_x precision is tighter than the γ_+ precision by a factor of 1.5 (at $i=10^\circ$) to 7 (at $i=50^\circ$). This is because γ_+ inference depends crucially on prior knowledge of v_{max} and μ while γ_x inference depends on a more fundamental symmetry argument. The precision of that symmetry argument depends, of course, on the assumption that real galaxy velocity fields have negligible shearlike modes, so this assumption is one that should be tested in further work.
- In this high-S/N and well-resolved scenario, both shear components can be inferred to a precision of 0.01 or better if the target is nearly face-on. At 50° (close to a typical value for randomly selected targets) γ_x can still be inferred to this precision but the constraint on γ_+ is less useful. (§3.4 will show that a linear combination of γ_+ and μ can still be constrained at this inclination.)

3.3. Dependence on Tully-Fisher prior

Tightening the Tully-Fisher (TF) prior has no effect on the γ_+ forecast because the dominant source of uncertainty for γ_+ is uncertainty in μ , at least in our fiducial setup. This raises the question of how loose a TF prior is tolerable. We found a $\lesssim 10\%$ relative effect on γ_+ uncertainty when the prior is loosened from 0.04 to 0.08, and $\approx 30\%$ relative effect when further loosened to 0.16 (i.e., 16% scatter in rotation speed at fixed luminosity, or almost a factor of two scatter in luminosity at fixed speed). In summary, we see a substantial effect on γ_+ when the TF prior becomes looser than the prior on μ . Conversely, the current level of TFR scatter is low enough that uncertainty in μ will remain the factor driving the γ_+ uncertainty for the foreseeable future.

Note that neither μ nor TF priors affect the γ_x forecast. In our idealized model, the limiting factor on γ_x precision is merely the precision and resolution of the velocity field measurements. This is unlikely to be the case in nature, where velocity fields are not perfectly orderly. An important task beyond the scope of this paper is to quantify the leading sources of γ_x uncertainty and systematic error due to natural variations from this idealized model.

3.4. Eigenvector decomposition

A striking feature of our results so far is the dramatic growth of γ_+ uncertainty with inclination, from about **1.5 times the γ_x uncertainty at $i = 10^\circ$ to about 7 times the γ_x uncertainty at $i = 50^\circ$** . In this subsection we show that this is largely due to greater mixing of γ_+ and μ as i increases. **To better illustrate what happens at high inclination, we go slightly beyond our standard range of $10^\circ - 50^\circ$ and use a very loose μ prior of ± 1 .**

Figure 6 illustrates the constraints in the (γ_+, μ) plane at three representative inclinations. At low inclination the constraints on γ_+ and μ are nearly orthogonal. This makes sense because μ should be irrelevant in the face-on case: given a uniform velocity field, the unlensed galaxy is circular so both components of shear can be determined precisely. At higher inclination, however, the constraint ellipse rotates in the (γ_+, μ) plane. With the μ uncertainty remaining ± 1 , this rotation greatly expands the uncertainty on γ_+ .

Some of the precision could be recaptured by parametrizing the lensing in terms of eigenvectors of the (γ_+, μ) submatrix of the covariance matrix. These are represented graphically by the major and minor axes of the ellipses in Figure 6. Although the minor axis does increase with i , it increases only about one-fifth as much as the γ_+ uncertainty; the increase in the latter is mostly due to the eigenvector rotation.

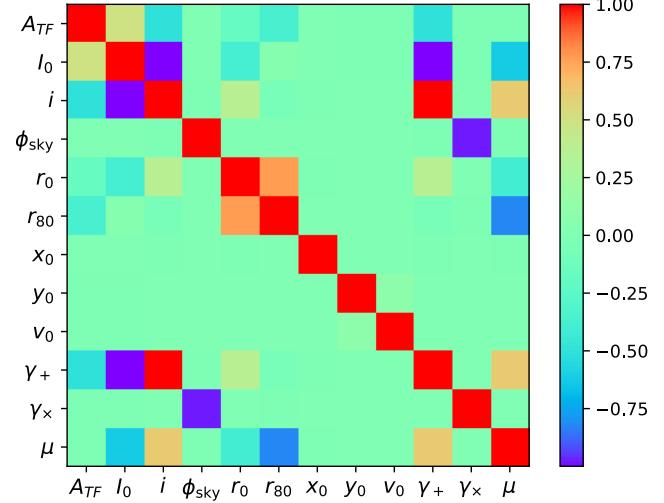
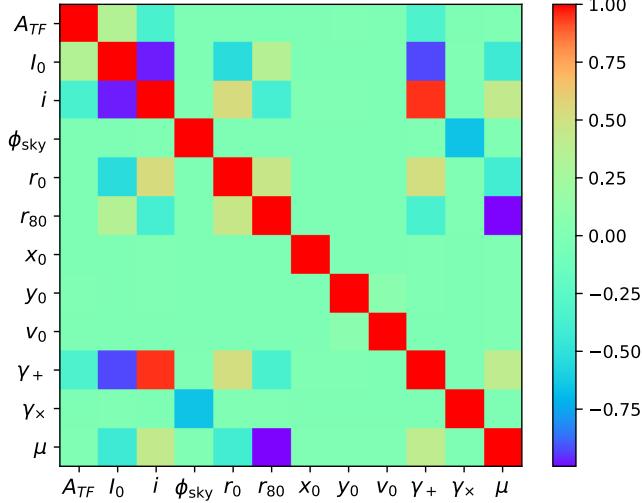


Figure 4. Correlation matrices for the $i = 20^\circ$ (left) and $i = 50^\circ$ (right) cases.

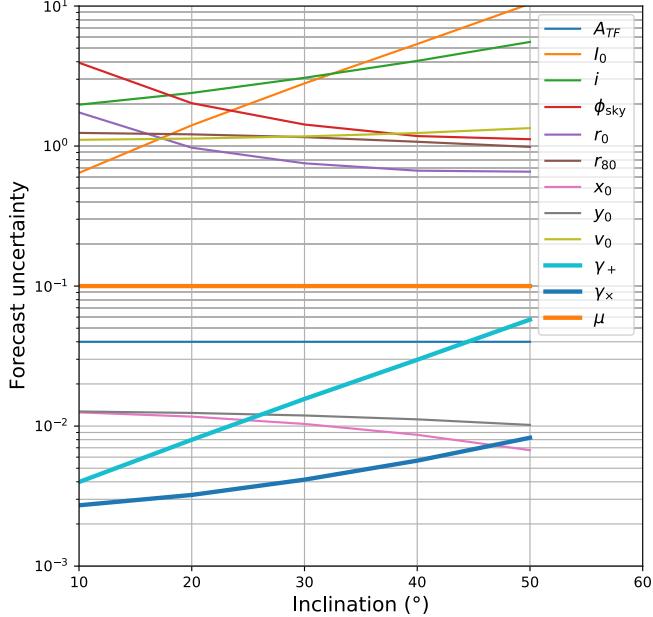


Figure 5. Forecast constraints as a function of inclination angle i .

The eigenvector decomposition could potentially be used to improve precision even at low inclination. Even the γ_+ width of the $i = 10^\circ$ ellipse in Figure 6 is due largely to its μ dependence. The eigenvector decomposition defines a γ_+ -like component with an uncertainty around 0.004, nearly as good as for γ_x .

The practical impact of this reparametrization may depend on the application. It may not be useful in a cosmic shear survey. When fitting mass profiles to lenses, however, each profile predicts both γ_+ and μ along a given line of sight. In other words, it will predict a point in the (γ_+, μ) plane depicted in Figure 6. Hence, the $i = 60^\circ$ ellipse may have substantial power

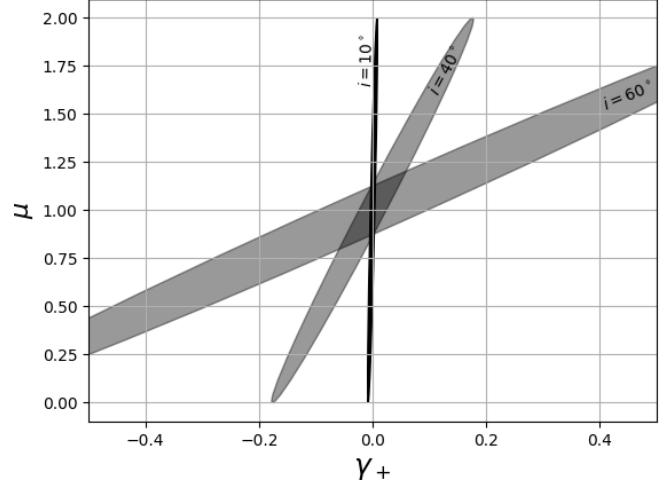


Figure 6. Constraints on γ_+ and μ as a function of inclination angle.

to discriminate between models despite the fact that it is compatible with a range of γ_+ as well as a range of μ .

Even so, it is evident that high inclinations are much less constraining than low inclinations. Taking the inverse of the area of this ellipse as a figure of merit, we find that the merit degrades by a factor of six from $i = 10^\circ$ to $i = 40^\circ$, and by another factor of four from there to $i = 60^\circ$.

We find similar behavior when parametrizing the lensing matrix in terms of κ rather than μ .

3.5. Dependence on target position angle

In our fiducial setup the unlensed major axis position angle is aligned with the sky coordinates ($\phi_{\text{sky}} = 0$), so there is no distinction between a coordinate system fixed to the sky and one fixed to the source galaxy. In principle, a coordinate system fixed to the galaxy

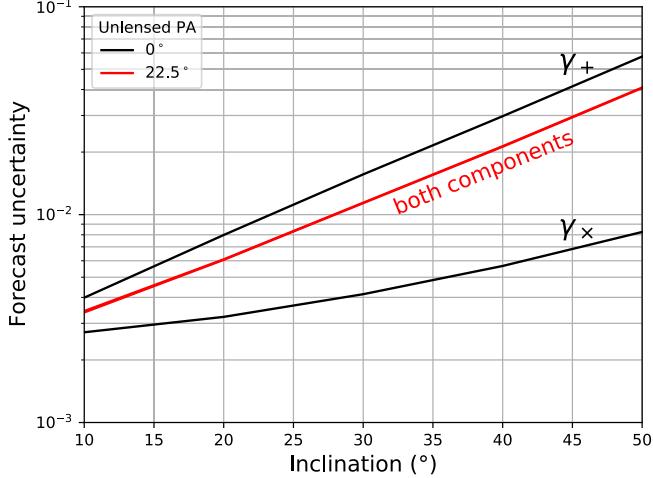


Figure 7. The effect of ϕ_{sky} on the uncertainty of shear components defined on the sky. Unless ϕ_{sky} is near a multiple of 45° or the target is nearly face-on, neither component of shear can be measured precisely.

cleanly separates the shear components into a precisely constrained one (related to the broken symmetry of the velocity field) and a less well constrained one (correlated with μ). In practice, sky-based shear components must ultimately be used to interpret the shear—to relate it to a lens, for example. Hence we have defined γ_+ and γ_\times based on sky coordinates. Our forecast uncertainties have included marginalizing over the parameter ϕ_{sky} , but our fiducial setup is still close enough to the “pure” galaxy basis that γ_+ and γ_\times are clearly distinct. In this subsection we show that for general values of ϕ_{sky} , γ_\times is no longer an eigenvector of the $(\gamma_+, \gamma_\times, \mu)$ space.

Figure 7 shows, along with the fiducial results, a Fisher matrix forecast for $\phi_{\text{sky}} = 22.5^\circ$, where one might expect equal precision for γ_+ and γ_\times . Indeed, the two red curves representing γ_+ and γ_\times are identical and a factor of $\sqrt{2}$ below the fiducial γ_+ forecast, indicating that the uncertainty is maximally mixed between the two components. If desired, an eigenvector decomposition could be used to define two linear combinations of γ_+ , γ_\times , and μ that are well measured and one that is constrained only by the prior on μ . This is just one snapshot of the ϕ_{sky} -dependent mixing: at $\phi_{\text{sky}} = 45^\circ$ (not shown) γ_+ becomes a well-constrained eigenvector, and there are intermediate degrees of mixing for intermediate values of ϕ_{sky} .

A reasonable approach to forecasting precision for randomly oriented sources would be to use the “both components” forecast in Figure 7. If, for example, we are concerned with the tangential shear of sources scattered around an axisymmetric lens, the per-source precision will vary between the γ_+ and γ_\times curves in Figure 7, with a mean given by the “both components” curves.

In this case, targets do need to be close to face-on to reach 0.01 precision; this could be limiting, as only 13% of randomly oriented disks will be within 30° of face-on. In other applications it may be possible to extract more information using the eigenvector decomposition.

We remind readers that the low uncertainty for two eigenvectors stems from the idealized assumption that disk galaxies intrinsically have no shearlike modes (their intrinsic major and minor axes are equal, and their velocity fields are symmetric and locked to their intensity fields). To the extent that real galaxies depart from these assumptions, the noise floor for the eigenvectors will be higher, rendering the eigenvector decomposition (and the velocity-field method overall) less advantageous.

3.6. Dependence on shear

Because the observed velocity field is not a linear function of shear, we expect the forecast precision to depend on the shear itself. Figure 8 presents the shear constraint forecast for increasing levels of γ_\times , with γ_+ held fixed at zero. As the true γ_\times increases, the γ_\times precision degrades while the γ_+ precision is nearly unaffected. The degradation occurs particularly at low inclination where the γ_\times precision had been excellent, so the γ_\times forecast becomes more nearly independent of inclination. In fact, as the true γ_\times becomes substantial the γ_\times parameter becomes correlated with μ and its family of correlated parameters. As a result, the forecast γ_\times uncertainty scales with the μ prior as well as the true γ_\times . For low true γ_\times , though, the γ_\times uncertainty hits an inclination-dependent floor. This is consistent with the [de Burgh-Day et al. \(2015\)](#) result that $\gamma_\times \approx 0.01 \pm 0.01$ for two nearby galaxies with presumably negligible shear, while suggesting that such precision is unobtainable along lines of sight with $\gamma_\times \gtrsim 0.1$.

Figure 9 shows the effect of varying the true γ_+ : the γ_+ precision can change in either direction by factors of a few to several depending on the inclination, but there is no dramatic overall trend. There is no effect on the γ_\times precision.

We also tested scenarios with mixed shear. The precision of each component appears to depend only on the true amount of that component, regardless of the true amount of the other component.

These patterns can also be understood in terms of the eigenvector decomposition. The presence of shear alters the mixing between γ_+ , γ_\times , and μ , and μ contamination is particularly noticeable in cases where the precision had been excellent (for γ_\times generally, and for γ_+ at low i). The presence of γ_+ actually *reduces* the correlation with μ at higher inclinations and thus improves γ_+ inference

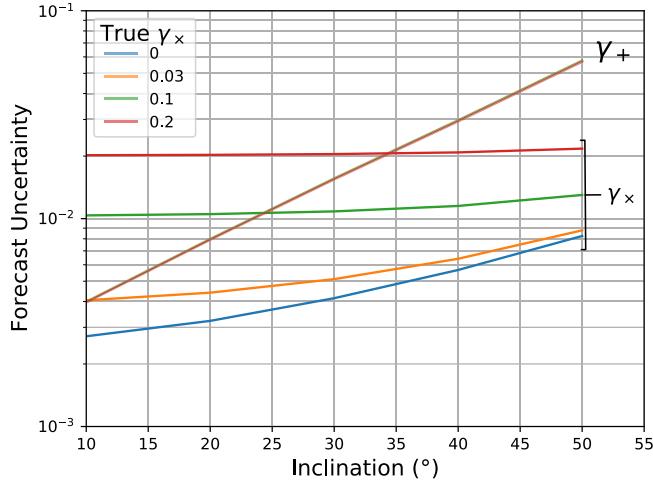


Figure 8. The uncertainty on γ_x increases when the true γ_x is substantial, due to increasing correlation with μ and related parameters. In the presence of substantial γ_x the γ_x uncertainty scales with both γ_x and the size of the μ prior, which is ± 0.1 here. The γ_+ uncertainty is, in contrast, nearly unaffected by the true level of γ_x . The true γ_+ is fixed at zero in this figure.

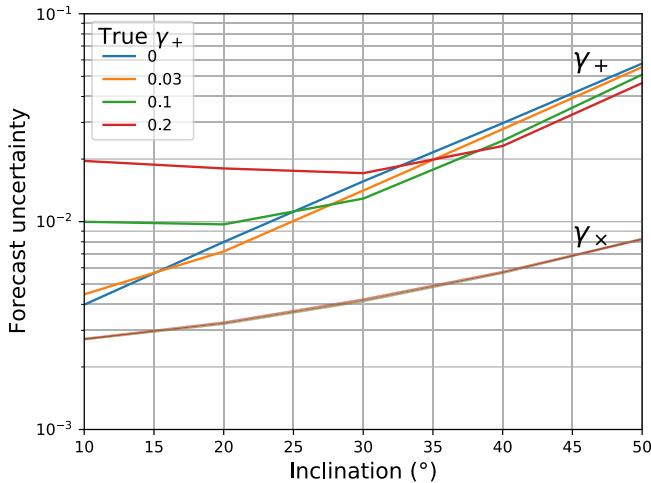


Figure 9. The γ_+ uncertainty depends on the true γ_+ , but much less dramatically than the γ_x uncertainty depends on the true γ_x (Figure 8). The γ_x uncertainty is nearly unaffected by the true level of γ_+ . The true γ_x is fixed at zero in this figure.

there, but a factor of ≈ 2 reduction from a fairly high baseline looks less dramatic on a logarithmic plot.

The eigenvector decomposition may enable useful lensing constraints in high-shear regions. Not evident in Figures 8 and 9 is the fact that the eigenvalues are nearly identical regardless of shear. An eigenvector composed mostly of γ_x but with an admixture of γ_+ and μ can still be constrained to high precision even at

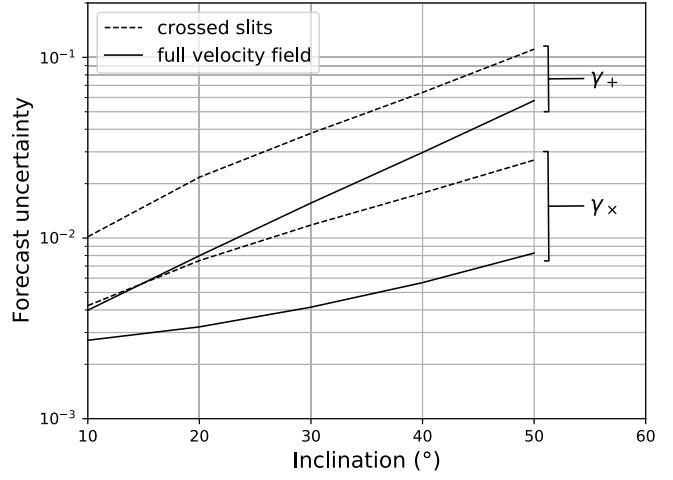


Figure 10. Shear constraints with full velocity-field observations (solid) versus crossed slits along the major and minor axes (dashed).

$\gamma_x = 0.1$. For fitting mass models, this could still be highly constraining as explained in §3.4.

3.7. Crossed slits

Obtaining full velocity-field data can be expensive, so we investigate the suggestion of Huff et al. (2013) that slit spectra be taken across the apparent major and minor photometric axes. (In our fiducial case with zero shear, these are the same as the velocity-field axes.) We implement this by masking out most pixels in the velocity-field partial derivative fields. In Figure 10 we plot the crossed-slit forecast as dashed curves, along with the standard full-field forecast as solid curves. The crossed slits increase the uncertainty by a factor of 2–3, depending only slightly on inclination.

Although crossed slits appear to perform well for nearly face-on targets, further work is required to assess their robustness. For example, consider a galaxy whose velocity field departs from our idealized model in a few places due to substructures, one of which falls in a slit. To compare the robustness of the full velocity field and the crossed slits across a variety of realistic galaxies, simulations will be required.

3.8. Rotation Curve Model

We also made forecasts with the Universal Rotation Curve (URC; Persic et al. 1996; Salucci et al. 2007) model, which links r_{80} to the radius at which the rotation curve becomes flat; in this case r_0 is still an independent parameter describing the steepness of the rise, which can lead to an overshoot. We found that the URC model leads to a small improvement when nearly edge-on (ie when the rotation curve is most apparent

in observations) but otherwise yields remarkably similar shear constraints. We attribute this rotation-curve insensitivity to the basic mechanisms underlying the inference of each shear component. Inference of γ_x relies on a symmetry argument that should be insensitive to the specific form of the rotation curve. Inference of γ_+ is limited by lack of knowledge of μ , a factor which is in no way ameliorated by adopting the URC model.

3.9. Dependence on resolution and signal-to-noise

Our fiducial setup uses velocity and intensity fields with 625 independent pixels (25 square) which, along the apparent major axis, just encloses $r_{80} = 12.5$ pixels. We deliberately made our forecast agnostic as to the target redshift and instrument details, but for context, a typical disk scale length is about 4 kpc (Fathi et al. 2010). This yields $r_{80} \approx 12$ kpc, so one can think of each fiducial pixel as representing about 1 kpc. Sources behind lenses are likely to have redshifts ≈ 0.4 and up, hence their angular diameter distances set a scale of about 5–8 kpc per arcsecond (this applies to arbitrarily high redshift sources, due to the broad maximum in angular diameter distances as a function of redshift). Therefore, a 1 kpc pixel will subtend 0.1–0.2 arcsec.

Seeing. To this point we have assumed that each pixel is completely independent, but most instruments are designed with pixel sizes smaller than the point-spread function (PSF). Therefore we tested the effect of blurring the intensity and velocity fields with a Gaussian of $\sigma = 2$ pixels. This degrades the Fisher matrix forecast by only about 10% in relative terms. Hence, useful observations of the lowest-redshift targets may be possible from the ground with excellent seeing or with low-order adaptive optics, while high-redshift targets are better pursued from space or from the ground with good adaptive optics systems.

Resolution. We tested the effect of halving or doubling the angular resolution, with the field size still just enclosing r_{80} . For simplicity, we will describe results only for the favorable inclination $i = 10^\circ$. **Doubling (halving) the resolution halved (doubled) the forecast uncertainty on each shear component.** Uncertainty decreases as pixels are added, roughly following the trend n_{pix}^{-1} where n_{pix} is the number of pixels encompassing $\pm r_{80}$ across the source major axis (and equaling the square root of the total number of pixels).

Velocity precision. Figure 11 shows the effect of varying $\sigma_{v,0}$, the uncertainty in the velocity measurement of the central pixel, at $i = 10^\circ$. Both curves can be well fit by a simple model in which a term depending linearly on $\sigma_{v,0}$ is added in quadrature to a constant

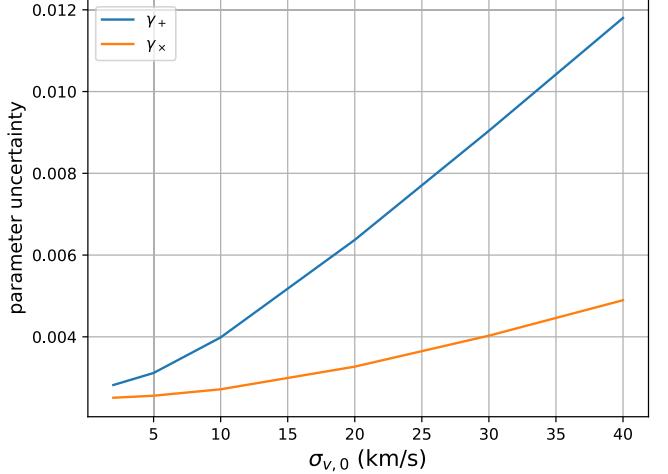


Figure 11. Shear constraints as a function of the velocity measurement uncertainty in the central pixel, at $i = 10^\circ$.

noise floor; however, the constants depend on the component and the parametrization. The γ_x component is relatively insensitive to $\sigma_{v,0}$, while the γ_+ component degrades more quickly. Nevertheless, Figure 11 shows that good constraints can still be obtained, at least at favorable inclinations, with lower-precision velocity data than assumed in our fiducial model. A reasonable target is 10–20 km/s: efforts to go below that meet with diminishing returns. Sacrificing additional velocity precision (i.e., exposure time) to allow targeting of more galaxies may also be a reasonable strategy.

Intensity S/N. We reduced I_0 , the S/N of the central intensity pixel, from its fiducial value of 90. We found that the precision scales nearly inversely with this S/N.

Summary. The forecast precision at $i = 10^\circ$ scales roughly as $\frac{90}{I_0} \frac{25}{n_{\text{pix}}}$ where I_0 is the S/N of the central intensity pixel and n_{pix} is the number of pixels encompassing $\pm r_{80}$ across the source major axis. The dependence on velocity measurement uncertainty is not so simply captured, but 10–20 km/s is a reasonable goal for observations, with little motivation to push below that. As a caveat, we have not investigated the extent, if any, to which constraints on μ degrade with data sets less precise than our fiducial one, hence we have not quantified the effect this may in turn have on shear constraints. Simulations will be required to address this issue.

4. SUMMARY AND DISCUSSION

Our approach has been to assume idealized and well-measured ($\sigma_{v,0} = 10$ km/s) velocity fields in order to explore the potential of velocity field lensing. Our fiducial result is that at the favorable inclination $i = 10^\circ$ the γ_x constraint can reach $0.003 \frac{90}{I_0} \frac{25}{n_{\text{pix}}}$ where I_0 is the S/N of the central intensity pixel and n_{pix} is the number of

pixels encompassing $\pm r_{80}$ across the source major axis. The γ_+ constraint is 1.5–7 times looser, depending on inclination. Both constraints degrade substantially at higher inclinations.

In more detail, we find:

- **The model is degenerate under infinitesimal displacements in specific directions in parameter space. However, the data change quadratically with step size along these directions, so the data can constrain all parameters. The quadratic effect can be emulated in the Fisher matrix formalism by putting a ± 0.1 prior on μ .**
- For our fiducial zero-shear scenario, constraints on γ_x are precise to better than 0.01 for targets inclined by less than $\approx 55^\circ$ —nearly half of all randomly inclined disks. This precision is a useful benchmark because it is roughly 20 times better than the per-galaxy precision for standard weak lensing, and also matches that found by [de Burgh-Day et al. \(2015\)](#). This precision, if true for both shear components, would make one velocity-field target worth roughly $20^2 = 400$ galaxy images, thus providing strong motivation to obtain the more expensive velocity-field observation.
- This precision is more difficult to reach for γ_+ , the shear component parallel/perpendicular to the unlensed apparent major axis. With our default prior on μ (± 0.1) only targets with $i < 25^\circ$ reach 0.01 precision on γ_+ . This is a small minority of randomly inclined disks. Furthermore, for this select group of targets the assumption of face-on circularity is likely to be crucial, and bears further investigation.
- For either component, constraints degrade with increasing i . For γ_+ the trend is somewhat steeper so targets with substantial inclination become uninteresting. The precision can be improved somewhat if a tighter prior on μ can be justified.
- the Tully-Fisher relation is not a limiting factor. A fractional velocity scatter smaller than the prior on μ is sufficient.
- The notion of a well-measured γ_x and a less well-measured γ_+ is useful for conceptual understanding, but for general source PA the result is more complicated. Of the three parameters $(\gamma_+, \gamma_x, \mu)$ two eigenvectors can be well measured and the third is constrained only by the prior on μ . In the fiducial case, γ_x is an eigenvector but the γ_+ -like

eigenvector includes a μ component, hence constraints on γ_+ look worse. As i increases that eigenvector rotates to include more μ , so the pure γ_+ constraints degrade more rapidly than the γ_x constraints. If one chooses to measure the γ_+ -like eigenvector rather than γ_+ , the constraints degrade somewhat less rapidly with i .

- In the presence of shear, the nominal γ_+ and γ_x constraints degrade, but this is due to eigenvector rotation in the $(\gamma_+, \gamma_x, \mu)$ space. The eigenvalues are equally well constrained in the presence or absence of shear.
- A per-pixel velocity uncertainty of 10–20 km/s is adequate, with smaller uncertainties yielding only marginal improvements.
- Observing a subset of the velocity field via crossed slits may be a viable strategy for reducing observing expense. In the fiducial case ($\phi_{\text{sky}} = 0$) this causes **a factor of 2–3 degradation in the precision of each component**. A more realistic assessment of crossed slits versus full velocity fields will require exploration of how real disk galaxies depart from our idealized assumptions **as well as slit placement uncertainty**.

Our model is highly idealized. It assumes:

- The galaxy is circular when viewed face-on.
- The velocity field is well ordered and completely described by a simple analytical function. The choice of rotation curve does not appear to matter, but the azimuthal symmetry surely matters.
- The velocity and intensity fields share a single inclination angle and PA. With the arctan rotation curve, there is no other link between the two fields (apart from the Tully-Fisher relation). With the URC, there is a link via r_{80} but this does not lead to tighter constraints because the limiting factors lie elsewhere.
- The disk is infinitesimally thin. The finite thickness of real disks will likely loosen the constraints at higher inclinations, because our forecast does not account for the increased velocity width in each pixel nor for extinction.
- No additional structure such as bulges, bars, or warps. Bulges may add noise, but bars and warps seem more concerning in terms of biases. Nevertheless, [de Burgh-Day et al. \(2015\)](#) did succeed in inferring a plausible γ_x ($\approx 0.01 \pm 0.01$) for radio

observations of an unlensed nearby galaxy with a prominent gas warp, so it is possible that warps do not disturb the velocity-field symmetry in the same way that shear does. **On the other hand, if the precision cited by de Burgh-Day et al. (2015) is due to typical galaxy features, future forecasts will need to account for this, with ~ 0.01 becoming the γ_x noise floor.** More work is needed to address this question.

The salience of warps may hinge on the velocity-field tracer: gas or stars. Gas is a convenient tracer for both radio and optical spectroscopy, but is also susceptible to inflows and outflows as well as warps. If this leads to the velocity equivalent of shape noise, the velocity-field method could become much less attractive. Stellar velocity fields are more orderly, but obtaining velocity fields from stellar absorption lines will require much more observational effort.

These observational choices are also tied to the question of whether the velocity and intensity fields must come from the same tracer. In our model the two fields are linked by a common center, inclination, and PA. The fact that our forecast is sensitive to the intensity field S/N suggests that reaching the 0.01 level does require constraints on the disk center, inclination, and PA beyond those derivable from the velocity field itself. Therefore, misalignments between intensity and velocity fields are a potential source of concern.

Recent observations indicate the potential for such misalignments. Figure 9 of Contini et al. (2016) compares the difference between the kinematic PA, as extracted from observations with the MUSE integral field spectrograph at the VLT, with the morphological PA as extracted from HST/F814W broadband images. They find one galaxy (of 27) with a large PA difference that cannot be related to poor resolution or by being nearly face-on (where PA is less well defined): the source of this difference is a bar. Even among the nearly face-on cases, they attribute some of the PA differences to structures such as spiral arms, bars, or clumps. Similarly, Wisnioski et al. (2015) find some significant offsets between the PA of broadband light and of the velocity field as traced by H α emission with the KMOS integral field spectrograph at the VLT. It is possible that such offsets would be reduced (albeit at additional observational expense) if stars were used to trace both velocity and intensity fields. Other potential steps to mitigate this source of error could be to model bars and spiral arms out of the intensity field, and/or to introduce a nuisance parameter representing the intensity-velocity PA offset and marginalize over it.

This concludes a long list of sources of uncertainty, yet to be quantified, that could prevent this method from being of practical use. Yet there are substantial strengths to this method as well:

- At favorable inclinations, tight constraints are achievable even with uninformative priors on μ .
- The method may work well with fitting mass models to lenses. Each background source will yield a constraint that may span a range of γ_+ , γ_x , and μ but is a long, narrow ellipsoid in $(\gamma_+, \gamma_x, \mu)$ space. Because a mass model predicts, for a given line of sight, a unique point in that space, the ellipsoid is likely to be highly constraining regardless of how it is oriented in that space. That said, the most highly constrained principal axis of this ellipsoid corresponds to our fiducial γ_x forecast, so this argument does not allow parameter inference better than our fiducial forecast. Rather, inferences that cannot take advantage of the eigenvectors may be limited to the precisions presented in Figures 8 and 9.
- This is a method of obtaining a high-precision shear measurement along a *single* line of sight, whereas traditional weak lensing enables this precision only after averaging over a large area of sky. These are different and potentially complementary types of information. The velocity field method, for example, may yield more information about localized substructures, which are effective probes of certain aspects of dark matter (see, e.g., Drlica-Wagner et al. 2019 for an overview).

Morales (2006) also argued that this method avoids some of the major systematic errors of traditional weak lensing. For example, he argues that the PSF is no longer a first-order contributor to systematics. However, our assumption that the source is well-resolved implies that the PSF would be largely irrelevant for these sources regardless of the method. He also argues that this method is less susceptible to contamination by intrinsic alignments. It is indeed robust against scenarios in which source galaxies are aligned in the absence of lensing, because the shear is measured *independently* on each target. But there are more subtle intrinsic alignment scenarios (Hirata & Seljak 2004). Imagine that Galaxy A sits in a gravitational tidal field that directly affects its velocity field by perturbing the orbits of its stars, while Galaxy B is a background source lensed by that gravitational tidal field. To the extent that the velocity field perturbation in Galaxy A mimics lensing modes, it will have an “intrinsic shear” that is corre-

lated with the lensing shear on Galaxy B. In fact, this is perhaps the most important open question here: can an external tidal field, perhaps due to a neighbor or satellite, induce shearlike perturbations in a disk's velocity field? If so, marginalizing over a range of such velocity field models could introduce significant uncertainty. Whatever their origin, natural sources of uncertainty will degrade γ_x more than γ_+ because the γ_x forecast currently is limited only by the precision of the velocity measurements.

A possible extension to this method is to analyze the velocity *dispersion* field as well (which requires no additional observations). The dispersion field is nonuniform because the disk's radial, tangential, and vertical dis-

persions contribute differently to the line-of-sight dispersion, depending on azimuth. This yields unlensed symmetry that differs from that of the velocity field: it is symmetric about both major and minor axes. However, it is unlikely that this would contribute substantially to the Fisher information, because the azimuthal variations are small.

This work was supported in part by NSF grants 1518246 and 1911138. We thank Bryant Benson, Gary Bernstein, Brian Lemaux, Hunter Martin, and Kevin Bundy for useful discussions. We also thank the anonymous referee for significant constructive feedback.

APPENDIX

A. VERIFYING THE FISHER FORECAST WITH AN EXPLORATION OF THE LIKELIHOOD SURFACE

We employ the Markov Chain Monte Carlo (MCMC) code `emcee` (Foreman-Mackey et al. 2013) to sample the likelihood surface. `emcee` implements an affine-invariant sampling algorithm, and hence performs well even on highly degenerate systems, provided they are not strongly multimodal (Goodman & Weare 2010).

For each case we generate mock data and add the fiducial amount of noise to the velocity and intensity fields. In addition to the Tully-Fisher prior used in the main text (± 0.04 on A_{TF}) we place flat step function priors on other parameters to keep the model physically well defined: $A_{\text{TF}}, I_0, r_0, r_{80} \geq 0$; i in $[0, 90^\circ]$; ϕ_{sky} in $[0, 45^\circ]$; and $\mu > 0$. **Note that we do not apply the ± 0.1 prior on μ used for the Fisher forecasts; the data already constrain μ to this level due to their quadratic dependence on steps in the degeneracy direction.** We initialize one thousand walkers in a small ball around the correct values. We run the Markov chain for ≈ 1000 autocorrelation times as a burn in, followed by an additional ≈ 1000 autocorrelation times to record the positions of the walkers.

We first examine our fiducial case at the extremes of $i = 10^\circ$ (Figure 12) and $i = 50^\circ$ (Figure 13). In each figure, the colorscale represents the density of MCMC samples, the cyan contour represents the MCMC 68% confidence region, and the black ellipse represents the Fisher forecast for the 68% confidence region. The three least interesting parameters (x_0 , y_0 , and v_0) are omitted for clarity. A good match is evident throughout all panels at $i = 10^\circ$. At $i = 50^\circ$ there is a hint that the forecast is becoming more pessimistic than the MCMC samples. This effect is more noticeable at higher inclinations. We attribute this to increasing numerical errors as one goes to higher inclinations: the condition number at $i = 50^\circ$ (60°) is 7×10^7 (3×10^8). Hence, for the fiducial setup we provide forecasts only for $i \leq 50^\circ$, and more generally we provide forecasts only where the condition number is $< 10^8$.

Finally, we present a case far from our fiducial scenario: with $\phi_{\text{sky}} = 10^\circ$ and $\gamma_+ = \gamma_x = 0.07$, these three parameters are highly mixed. We have also changed r_0 (r_{80}) to 6 (10) kpc, I_0 to 75, and μ to 1.2, with an inclination angle of 35° . Figure 14 shows that the forecast still accurately predicts the MCMC precision.

REFERENCES

Bartelmann, M., & Maturi, M. 2017, Scholarpedia, 12, 32440

Blain, A. W. 2002, ApJL, 570, L51

Blakeslee, J. P. 2001, arXiv Astrophysics e-prints, astro-ph/0108253

Contini, T., Epinat, B., Bouché, N., et al. 2016, A&A, 591, A49

de Burgh-Day, C. O., Taylor, E. N., Webster, R. L., & Hopkins, A. M. 2015, MNRAS, 451, 2161

Drlica-Wagner, A., Mao, Y.-Y., Adhikari, S., et al. 2019, arXiv e-prints, arXiv:1902.01055

Fathi, K., Allen, M., Boch, T., Hatziminaoglou, E., & Peletier, R. F. 2010, MNRAS, 406, 1595

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

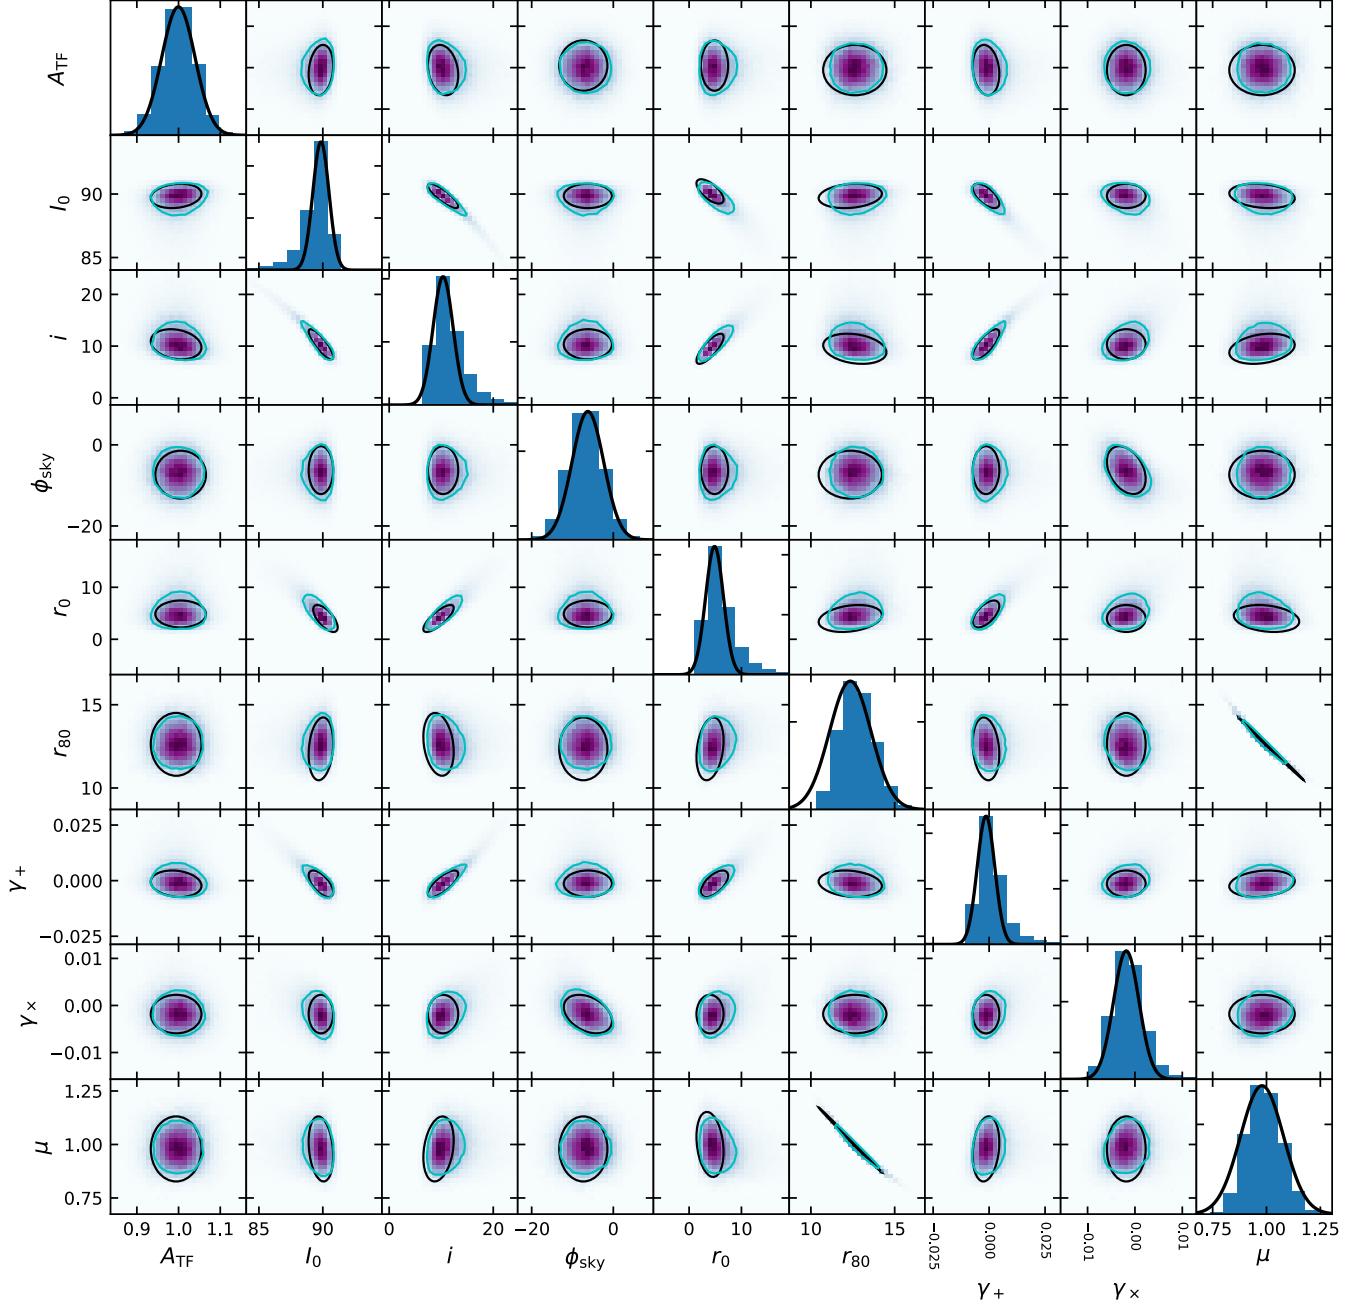


Figure 12. Corner plot for the fiducial case at $i = 10^\circ$, showing agreement between MCMC confidence intervals and the Fisher forecast. The colorscale represents the density of MCMC samples, the cyan contour represents the MCMC 68% confidence region, and the black ellipse represents the Fisher forecast for the 68% confidence region. The forecast ellipses are centered on the densest MCMC pixel.

Garcia-Fernandez, M., Sánchez, E., Sevilla-Noarbe, I., et al.

2016, ArXiv e-prints, arXiv:1611.10326

Goodman, J., & Weare, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65

Heavens, A. 2016, Entropy, 18, 236

Hirata, C. M., & Seljak, U. 2004, PhRvD, 70, 063526

Huff, E. M., & Graves, G. J. 2014, ApJL, 780, L16

Huff, E. M., Krause, E., Eifler, T., George, M. R., &

Schlegel, D. 2013, ArXiv e-prints, arXiv:1311.1489

Miller, S. H., Bundy, K., Sullivan, M., Ellis, R. S., & Treu, T. 2011, ApJ, 741, 115

Morales, M. F. 2006, ApJL, 650, L21

Morrison, C. B., Scranton, R., Ménard, B., et al. 2012, MNRAS, 426, 2489

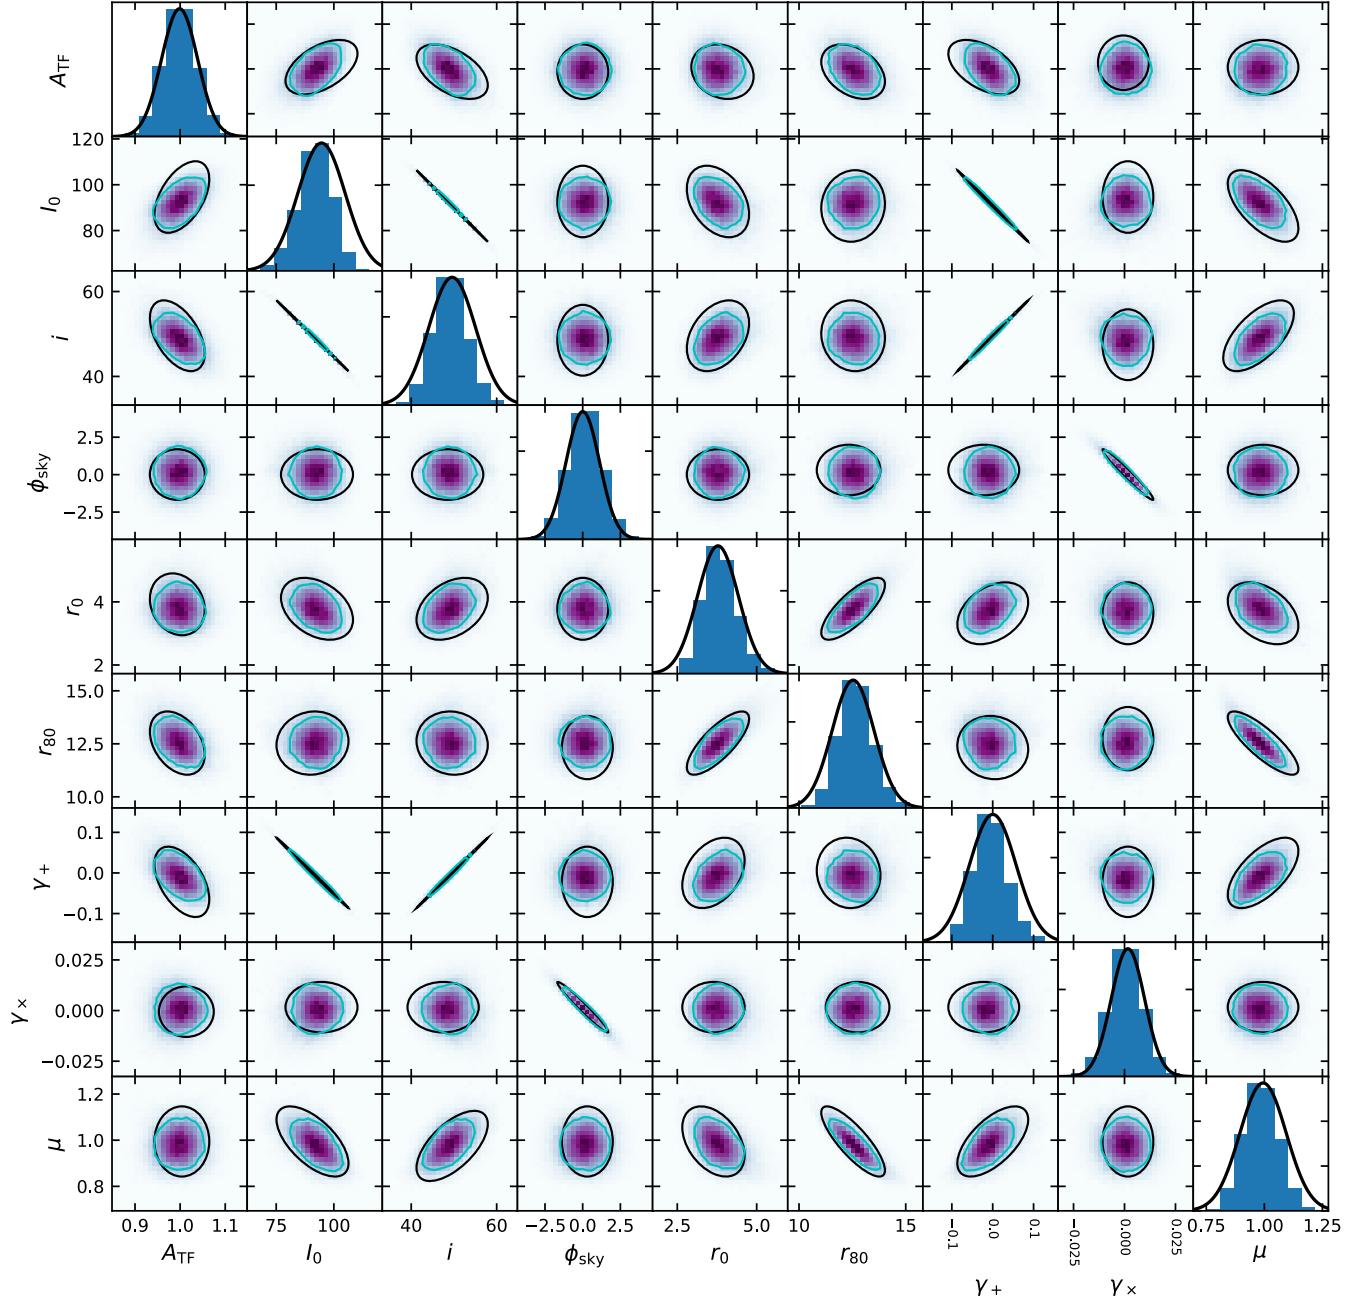


Figure 13. As for Figure 12, but at $i = 50^\circ$.

Persic, M., Salucci, P., & Stel, F. 1996, MNRAS, 281, 27
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., &
 Flannery, B. P. 1992, Numerical recipes in C. The art of
 scientific computing
 Rizzo, F., Vegetti, S., Fraternali, F., & Di Teodoro, E.
 2018, MNRAS, 481, 5606

Salucci, P., Lapi, A., Tonini, C., et al. 2007, MNRAS, 378,
 41
 Tully, R. B., & Fisher, J. R. 1977, A&A, 54, 661
 Vallisneri, M. 2008, PhRvD, 77, 042001
 Wisnioski, E., Förster Schreiber, N. M., Wuyts, S., et al.
 2015, ApJ, 799, 209

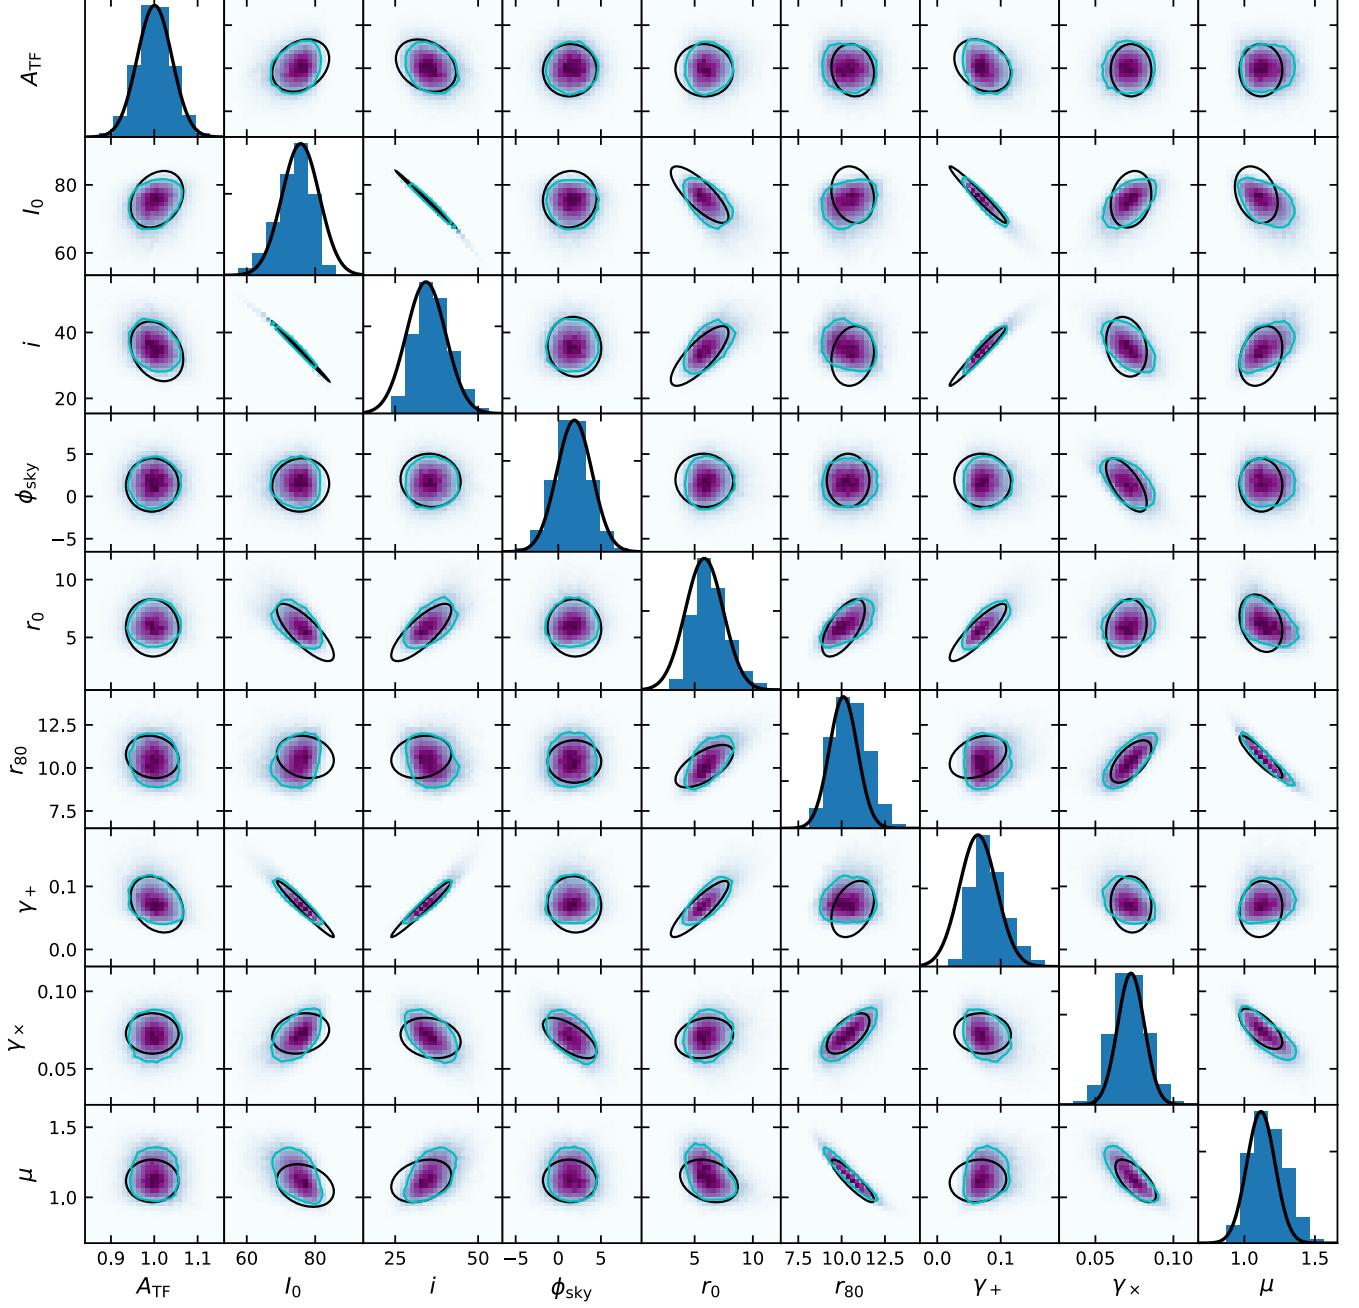


Figure 14. As for Figure 12, but for the far-from-fiducial setup described in the text. Note that the forecast contours are centered on the densest MCMC pixel, rather than the MCMC centroid.