Weak Gravitational Lensing

Tessa Baker

March 2, 2017

0 References

- *Modern Cosmology*, Scott Dodelson. Primary source for this lecture, and a really good all-round textbook for beginning grad student-level cosmology.
- *Weak Gravitational Lensing of the CMB*, Anthony Lewis & Anthony Challinor, arXiv: astro-ph/0601594. Comprehensive and pedagogical review paper covering both theory and applications of CMB lensing. §3.1 contains all that is needed for this lecture.
- *Gravitation, Foundations and Frontiers*, Thanu Padmanabhan. An excellent, if sometimes idiosyncratic, GR textbook – one of my favourites. Though it doesn't contain much on gravitational lensing, I reference it here for a proof of the conservation of surface brightness (see §3.2).

More heavy-duty:

- *Gravitational Lensing: Strong, Weak and Micro*, P. Schneider, C. Kochanek & J. Wambsganss. Lecture notes from the Saas-Fee summer school. All four parts are available online; section 3 (weak lensing) is at `https://arxiv.org/abs/astro-ph/0509252`.
- *Gravitational Lenses*, P. Schneider, J. Ehlers & E. Falco. A classic, and with good reason. Rigorous and full of gems; far more detail than is needed for this course.

1 Introduction

Review of some basic facts we already know:

- massless particles like photons move along the null geodesics of a spacetime;
- the presence of mass induces spacetime curvature, i.e. the metric $g_{\mu\nu}$ is no longer the (flat) Minkowski metric of Special Relativity;
- the null geodesics of a curved spacetime are not necessarily straight lines. The geodesic equation confirms this qualitative notion:

$$\frac{d^2x^\alpha}{d\lambda^2} + \Gamma^\alpha_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0 \tag{1}$$

  If ${\Gamma^\alpha}_{\mu\nu} \neq 0$ then this is no longer the equation of a straight line.

Taken together, these statements imply that light rays will be deflected – bent off their original path – by massive objects. Note that when we look at the sky with our eyes or a telescope, we infer the position of objects by tracing back photons from them *in a straight line.* This means where we perceive an object to be ('image position') may be different to its true location ('source position'), see Fig. 1. In some cases a bundle of light rays emitted from a source can be deflected onto multiple different paths. This results in us seeing multiple images of the same source in different locations, see Fig. 2.

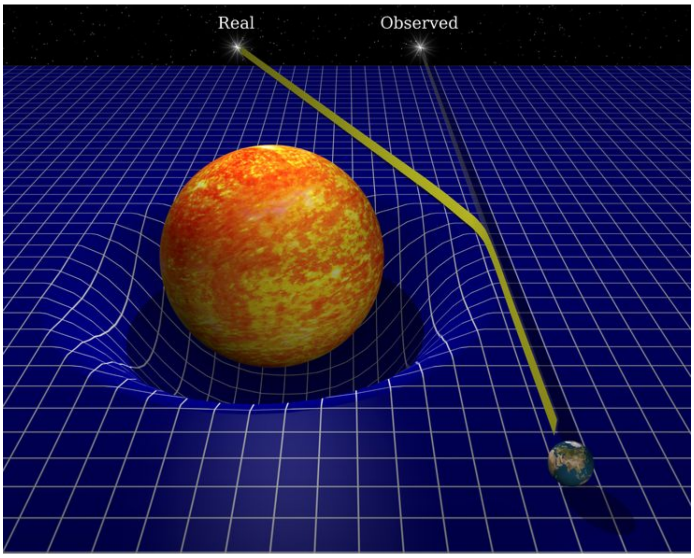

Figure 1: Deflection of light in a curved spacetime. In this case, the light of a distant star is deflected by the gravitational field of the Sun. Note that the location of the image corresponds to tracing the received ray back in a straight line.

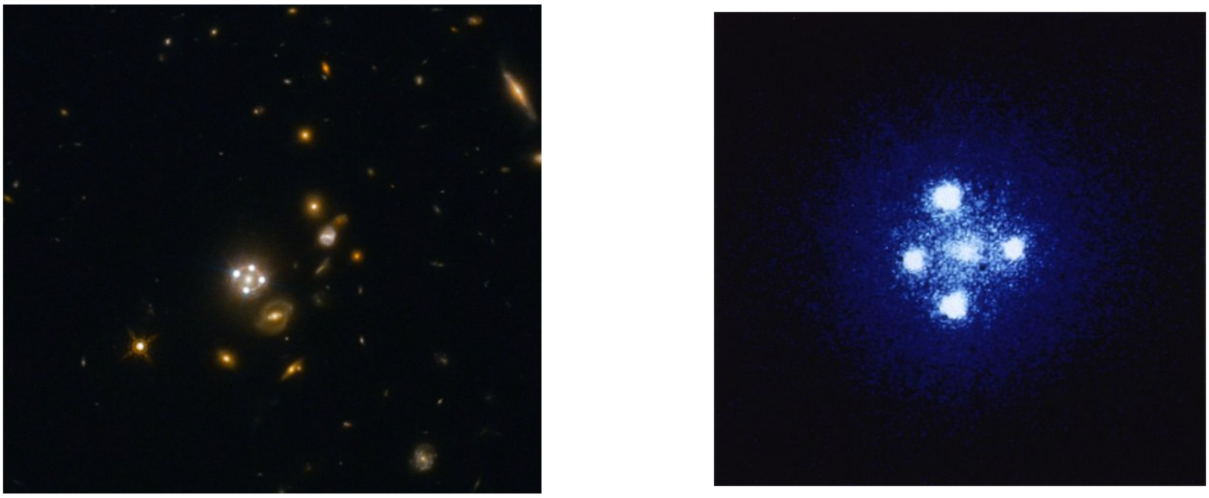

Figure 2: The Einstein Cross is a multiply imaged quasar, lensed by an intervening galaxy that sits almost exactly in front of it. The quasar is located about 8 billion light years away, whilst the lensing galaxy is at 400 million light years. The angular size of the cross in the sky is roughly1.6 x 1.6 arcseconds; the right-hand panel shows a zoom-in. This is an example of **strong** gravitational lensing. *Image credits: NASA/ESA/Hubble/STSci.*

Two kinds of gravitational lensing appear in a cosmological context[1]: strong and weak. In strong lensing we consider virtually all the deflection to be caused by a single massive object along the line of sight to the source – there is one, extreme lensing event. For example, we observe that images of very distant galaxies can be lensed (and hence distorted) by other (clusters of) galaxies closer to us, as shown in Fig. 2.

In contrast, weak lensing involves multiple, less extreme lensing events. These multiple lensing events occur as a photon from a distant galaxy[2] propagates towards us through the large-scale 'cosmic web' of dark matter that fills the universe, see Fig. 3. These dark matter structures bend the path of the photon gently, with the result that our galaxy image gets lightly squished – see Fig. 4.

[1]There are also non-cosmological examples of gravitational lensing: lensing by the Sun during the solar eclipse of 1919 was famously the first attempted experimental test of General Relativity. Also, there are all sorts of interesting optical distortion effects that can happen close to black hole event horizons, as a result of lensing by **very** strong gravitational fields.

[2]At the end of these notes we'll consider a case where the source is not a galaxy, but the CMB itself.

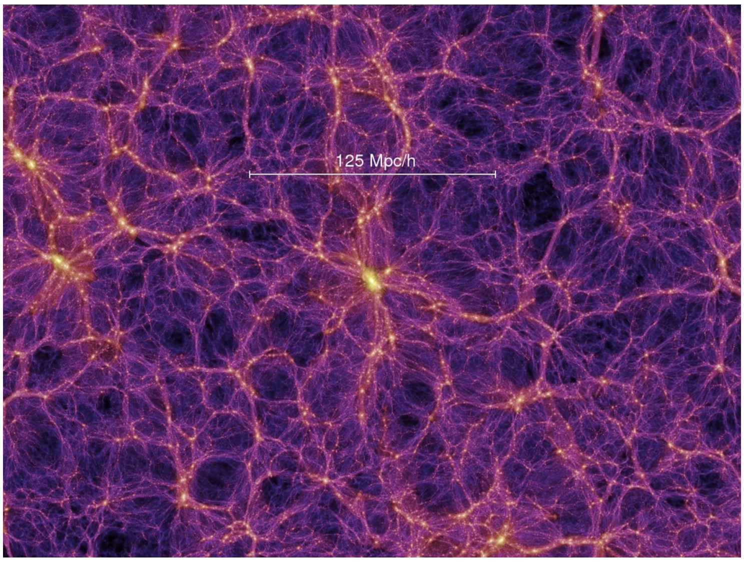

Figure 3: Snapshot from the Millennium Simulation, showing the large-scale cosmic web. The purple filaments indicate dark matter; the baryons/galaxies we observe are coloured gold/yellow. *Image credit: the Virgo Collaboration.*

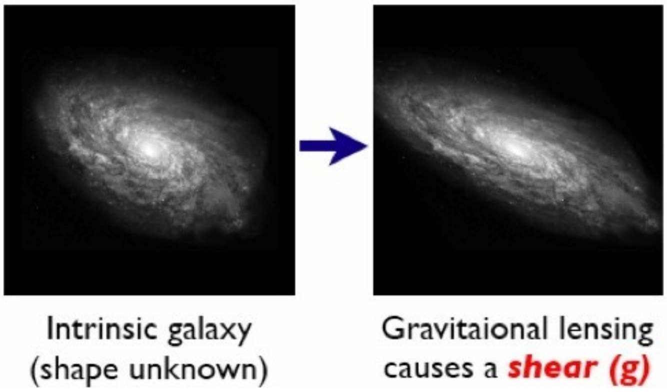

Figure 4: Schematic indicating the subtle changes to a galaxy image induced by weak gravitational lensing. The image on the RHS has an increased ellipticity along the axis running 4 o'clock to 10 o'clock. Remember that the left panel is unobservable – we would only see the image on the right. *Image credit: Sarah Bridle, Uni. Manchester.*

As we'll see shortly, we can use the effects of weak lensing – when measured for tens of thousands of galaxies – to learn about the contents of the universe, the statistics of the cosmic web structure, and to test General Relativity and ideas about dark energy.

Although strong lensing *does* have the potential to measure the Hubble constant H_0 (ask me at the end of the lecture if interested), at the moment it's not as powerful a tool for cosmology as weak lensing. This is because to fully understand strong lensing measurements you need a detailed mass model of the low-redshift cluster that is acting as the lens. This is difficult to obtain. Even if a good mass model can be constructed for a particular system, the process needs to be repeated for thousands of systems in order to make statistically significant measurements of cosmological parameters.

The strategy employed in weak lensing is very different. Instead of requiring detailed knowledge about the mass and structure of the lens, one looks for correlations between lensed galaxies in the same patch of sky (where 'correlations' means that they have a tendency to be elongated in the same direction). There are both current and upcoming experiments optimised to measure these correlations over large patches of sky quickly, e.g. the Dark Energy Survey (DES), the Large Synoptic Survey Telescope (LSST), the ESA Euclid satellite.

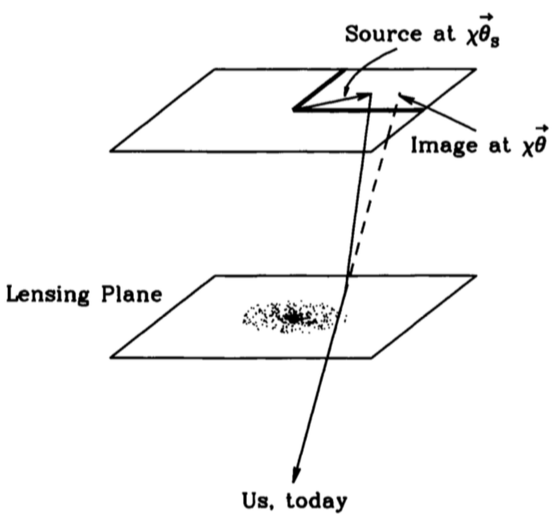

Figure 5: Schematic showing the set-up for our calculation, and the locations $\chi\vec{\theta}$ and $\chi\vec{\theta}^S$. *Credit: S. Dodelson.*

We will build up a description of weak lensing in the following stages:

1. We'll start by using the GR and cosmology you already know to study the propagation of photons in a perturbed cosmological spacetime.

2. We'll show how the paths of these geodesics can be used to construct a rank-2 tensor called the *distortion tensor*. We'll also show how the distortion tensor describes the *ellipticity* of a galaxy, i.e. its deviation from a sphere (or really, a circle once projected on the sky).

3. We then set about building a power spectrum of this distortion tensor that encodes the correlations between galaxy images as described above. In fact, there will be several different power spectra we can talk about – because the distortion tensor is a rank-2 object, we can choose to correlate its components in various ways. We'll tie up the results of this calculation to some real-world data from recent galaxy surveys.

4. If time permits, we'll also discuss the related topic of CMB lensing.

2 Deflected Photon Geodesics

Fig. 5 shows the set-up we need. A photon is emitted from a point of the source (remember, it's an extended object like a galaxy) and deflected by the gravitational field of another object in the lens plane[3] To start with, this looks a lot like a strong lensing scenario. This is just an artefact of drawing a convenient diagram – in reality there is a whole sequence of lensing masses along the line of sight between the observer and the source.

We set up a system of coordinates as follows. There is a radial coordinate, χ, that describes radial comoving distances from the observer according to:

$$\chi(a \text{ or } t) = \int_t^{t_0} \frac{d\tilde{t}}{\tilde{a}(\tilde{t})} = \int_a^1 \frac{1}{\tilde{a}^2 H(\tilde{a})} d\tilde{a} \tag{2}$$

where I've indicated that the scale factor a can be used as an alternative to physical time.

We'll use two further coordinates in the plane perpendicular to the line of sight: since we expect galaxies to be spherical on average, we can use polar coordinates. Making use of the small-angle approximation, we can then denote a general point by the vector $\chi\vec{\theta} = \chi(\theta^1, \theta^2, 1)$. We want to find the mapping between a point in the source, $\chi\vec{\theta}_S$ and it's location in the image plane.

[3]Since the size of the object doing the lensing is much smaller than the total distance travelled by the photon, we can ignore it's extent along the line of sight. *This is called the thin-lens approximation.*

Recall the following definition of the four-momentum of a photon, and the magnitude of its spatial component p:

$$P^\mu = \frac{dx^\mu}{d\lambda} \qquad\qquad p^2 = g_{ij}P^iP^j \tag{3}$$

where λ is an affine parameter. Of course, since a photon is null we must have $P^\mu P_\mu = 0$; using the line element for a perturbed FRW spacetime:

$$ds^2 = -(1+2\Psi)dt^2 + a(t)^2(1-2\Phi)\delta_{ij}dx^idx^j \tag{4}$$

we can expand $P^\mu P_\mu = 0$ to find:

$$-(1+2\Psi)(P^0)^2 + p^2 = 0 \qquad \Rightarrow \qquad P^0 = p(1+2\Psi)^{-\frac{1}{2}} \approx p(1-\Psi) \tag{5}$$

We'll need this relation in a moment. We'll also need the perturbed Christoffel symbols, which you know how to calculate from your GR course. You should be able to show that for the above line element these are:

$$\Gamma^0_{00} = \dot\Psi \qquad\qquad \Gamma^i_{00} = \partial^i\Psi \tag{6}$$

$$\Gamma^i_{j0} = \delta^i_j\left(H - \dot\Phi\right) \qquad\qquad \Gamma^i_{jk} = \delta^{im}\delta_{ij}\partial_m\Phi - \delta^i_k\partial_j\Phi - \delta^i_j\partial_k\Phi \tag{7}$$

Note that since Φ and Ψ are small quantities, we've dropped all terms of order Φ^2, Ψ^2 and higher when working these out. Now let's tackle the spatial component of the geodesic equation, eq.(1). We'll bravely develop both LHS and RHS in parallel; our first step is to convert all derivatives to be w.r.t. χ by using the chain rule:

$$\begin{aligned}
\frac{d^2x^i}{d\lambda^2} &= -\Gamma^i_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} \\
\frac{dt}{d\lambda}\frac{d\chi}{dt}\frac{d}{d\chi}\left(\frac{d(\chi\theta^i)}{d\chi}\frac{d\chi}{dt}\frac{dt}{d\lambda}\right) &= -\Gamma^i_{\mu\nu}\frac{dx^\mu}{d\chi}\frac{dx^\nu}{d\chi}\left(\frac{dt}{d\lambda}\right)^2\left(\frac{d\chi}{dt}\right)^2 \\
\frac{d}{d\chi}\left(\frac{d(\chi\theta^i)}{d\chi}\frac{d\chi}{dt}\frac{dt}{d\lambda}\right) &= -\left[\Gamma^i_{00}\left(\frac{dx^0}{d\chi}\right)^2 + \Gamma^i_{0j}\frac{dx^0}{d\chi}\frac{dx^j}{d\chi} + \Gamma^i_{jk}\frac{dx^j}{d\chi}\frac{dx^k}{d\chi}\right]\left(\frac{dt}{d\lambda}\right)\left(\frac{d\chi}{dt}\right) \\
\frac{d}{d\chi}\left(-\frac{p}{a}\frac{d(\chi\theta^i)}{d\chi}\right) &\approx \frac{p}{a}\left[a^2\partial^i\Psi - 2H\frac{d(\chi\theta^i)}{d\chi} + \delta^{im}\partial_m\Phi\right]
\end{aligned} \tag{8}$$

where in the second line we've written $\vec{x} = \chi\vec{\theta}$. In the third line we've expanded out the index summation and cancelled some factors. In the fourth line we've recognised that $dt/d\lambda = P^0$, and used eqs. 2, 5 & 7. We've also used the fact that θ^i is a small quantity, so terms like $\theta^i \times \Phi$ can be neglected to first order. This is the reasoning behind the last term on the RHS; the only non-negligible contribution is when $j = k = 3$, i.e. the coordinate along the line of sight.

You know already (Baumann notes eq.1.2.50) that to first order the spatial 3-momentum redshifts as $p \propto 1/a$. Using this in eq.8 (we don't need to specify the constant of proportionality since it cancels from both sides):

$$\frac{d}{d\chi}\left(\frac{1}{a^2}\frac{d(\chi\theta^i)}{d\chi}\right) \approx \left[-\partial^i(\Psi+\Phi) + 2\frac{H}{a}\frac{d(\chi\theta^i)}{d\chi}\right] \tag{9}$$

$$\frac{1}{a^2}\frac{d^2(\chi\theta^i)}{d\chi^2} - \frac{2}{a^3}\frac{da}{dt}\left(\frac{dt}{d\chi}\right)\frac{d(\chi\theta^i)}{d\chi} \approx \left[-\partial^i(\Psi+\Phi) + 2\frac{H}{a}\frac{d(\chi\theta^i)}{d\chi}\right] \tag{10}$$

$$\frac{d^2(\chi\theta^i)}{d\chi^2} \approx -a^2\partial^i(\Psi+\Phi) = -2\delta^{ij}\partial_j\Phi \tag{11}$$

where the second equality in the line directly above uses that at late times in the universe (so neutrinos and radiation are negligible) $\Phi = \Psi$ to a very good approximation[4].

[4]This comes is comes from the Einstein equations, when evaluated to first order in linear perturbations. The equality of Φ and Ψ is a feature (almost) unique to General Relativity. In lots of theories of modified gravity – put forwards as alternatives to the cosmological constant explanation of accelerated expansion – we have $\Phi \neq \Psi$. By combining observations in the right way we can test for the equality of the metric potentials, and hence test ideas about dark energy.

Integrating once, twice (and going to pains to make the arguments of the metric potentials clear):

$$\frac{d(\chi\theta^i)}{d\chi} = -\delta^{ij}\int_0^{\chi} d\tilde{\chi}\;\partial_j[\Psi(\vec{x}(\tilde{\chi}) + \Phi(\vec{x}(\tilde{\chi})] \;+\; \text{constant} \tag{12}$$

$$\theta^i_S = -\frac{\delta^{ij}}{\chi}\int_0^{\chi} d\chi' \int_0^{\chi'} d\tilde{\chi}\;\partial_j[\Psi(\vec{x}(\tilde{\chi}) + \Phi(\vec{x}(\tilde{\chi})] \;+\; \text{constant} \tag{13}$$

The LHS has picked up a subscript S in the last line because $\theta^i(\chi) = \theta^i_S$, the original location of the point in the source plane. We now realise that the constant must be equal to θ^i, the point that θ^i_S gets mapped to in the image: because if there is no lens present, the integral vanishes and we must be left with the trivial relation $\theta^i_S = \theta^i$.

Reversing the order of integration in eq.(13):

$$\theta^i_S = \theta^i - \frac{\delta^{ij}}{\chi}\int_0^{\chi} d\tilde{\chi}\int_{\tilde{\chi}}^{\chi} d\chi'\;\partial_j[\Psi(\vec{x}(\tilde{\chi}) + \Phi(\vec{x}(\tilde{\chi})] \tag{14}$$

The innermost integral is now trivial, since the integrand only depends on $\tilde{\chi}$.

$$\theta^i_S = \theta^i - \delta^{ij}\int_0^{\chi} d\tilde{\chi}\;\left\{\partial_j[\Psi + \Phi]\left(1 - \frac{\tilde{\chi}}{\chi}\right)\right\} \tag{15}$$

where I've suppressed the arguments of the potentials again, for clarity. This is the key expression we're seeking. It tells us that the difference between the the the original position of a point in our source and it's observed position in the image plane is controlled by the gradient of the gravitational potential due to the lens, modulo a weighting factor (the piece in round brackets) that encodes information about the relative conformal distances to the lens and the source.

The additional complication is that, in the weak lensing case, this fairly intuitive[5] quantity is integrated over a whole distribution of roughly equal-size lensing masses along the line of sight. In the strong lensing case we would expect this integral to be dominated by the contribution at the distance of the main, extreme lensing event.

3 The Distortion Tensor

3.1 Definition

Eq.15 only describes one component of our position vector $\chi\vec{\theta}$ at a time. What we really want is an object that gives us the mapping between both transverse components at once (the longitudinal component along the χ axis is less interesting, since we can't observe it directly[6]).

A common notation is to define a 2x2 matrix $\mathcal{A}$ using the derivative of eq.15:

$$\mathcal{A}^i_j = \frac{\partial\theta^i_S}{\partial\theta^j} \tag{16}$$

Note that the first term on the RHS of eq.15 will give an identity matrix contribution to $\mathcal{A}$. Once again, this corresponds to the trivial limit where no lens is present and the true and apparent positions are identical. So we introduce a second quantity, *the distortion tensor*, that describes any non-trivial difference between (the transverse components of) $\chi\vec{\theta}_S$ and $\chi\vec{\theta}$:

$$\psi^i_j = \mathcal{A}^i_j - \mathbf{I}_2 = \begin{pmatrix} -\kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & -\kappa + \gamma_1 \end{pmatrix} \tag{17}$$

$$= -\delta^{ik}\int_0^{\chi} d\tilde{\chi}\;\left\{\partial_j\partial_k[\Psi + \Phi]\,\tilde{\chi}\left(1 - \frac{\tilde{\chi}}{\chi}\right)\right\} \tag{18}$$

where the expansion into elements on the first line serves to define the quantities κ and γ_i that we'll explore in a moment. Beware a devious $\tilde{\chi}$ that has appeared in the integrand above. This is because, for the

[5]Admittedly, perhaps only intuitive in hindsight.

[6]We *can* still measure it, if we're prepared to whip out a spectrograph and measure the redshift of spectral features in our source galaxy.

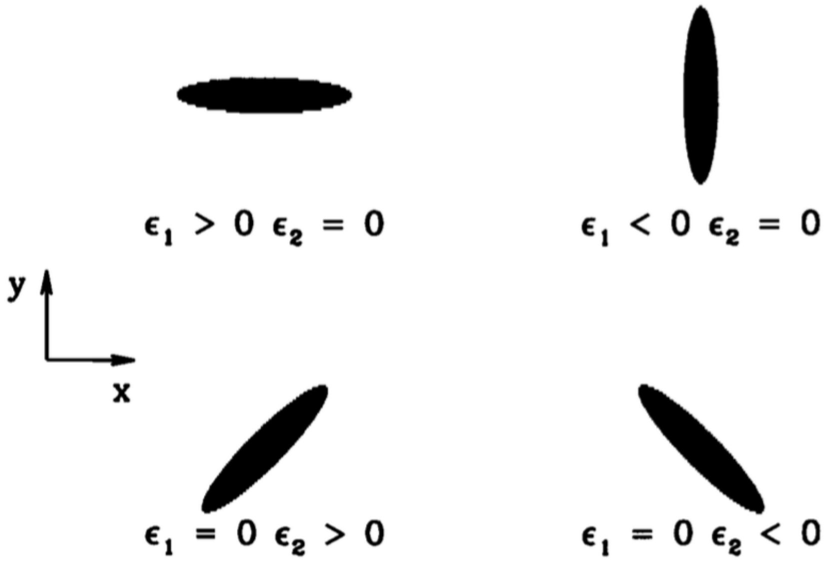

Figure 6: Image ellipticity parameters. Note that the values of ϵ_i assigned to each image depend on the choice of axes orientation. *Credit: S. Dodelson.*

distortion tensor, we really wanted to differentiate under the integrand w.r.t. θ. However, the derivative already appearing in eq.15 is w.r.t. x^j, so we have used the chain rule again to match it:

$$\frac{\partial \theta_S^i}{\partial \theta^j} = \frac{\partial \theta_S^i}{\partial x^k}\frac{\partial x^k}{\partial \theta^j} = \frac{\partial \theta_S^i}{\partial x^k}\frac{\partial (\chi \theta^k)}{\partial \theta^j} = \frac{\partial \theta_S^i}{\partial x^k}\chi \delta_j^k = \frac{\partial \theta_S^i}{\partial x^j}\chi \tag{19}$$

In eq.(17) we have introduced some new variables to describe parts of the distortion tensor. These are the **convergence**, κ, and the two components of the **shear**, γ_1 and γ_2. Broadly speaking, κ describes the magnification of an image with respect to the original source, and the γ_i describes how the shape of image has been distorted with respect to the original source by lensing.

3.2 Ellipticities and Shear

To explain the shear parameters more quantitatively, we first need to think about how to quantify the ellipticity of an image. Imagine we have a function $I_{\mathrm{obs}}(\theta)$ that describes how the brightness of an image varies with position. Integrating this function over the image would just give us the total flux of the source. So to get some measure of the shape of the source, we instead think about taking moments of the brightness distribution.

For images like those shown in the top line of in Fig. 6, the dipole moments vanish because there are equal portions in all four quadrants. The same is true of the images in the bottom line. The first non-vanishing moments are instead the quadrupoles:

$$q_{ij} = \int d\theta \, I_{\mathrm{obs}}(\theta)\, \theta_i \theta_j \tag{20}$$

These don't vanish because the factor $\theta_i\theta_j$ has the same sign in opposite quadrants, so they no longer cancel each other. By rotational symmetry, a perfectly circular image has $q_{xx} = q_{yy}$ and $q_{xy} = 0$. So, we can assess the non-circularity of an image via the following quantities:

$$\epsilon_1 = \frac{q_{xx} - q_{yy}}{q_{xx} + q_{yy}} \qquad\qquad \epsilon_2 = \frac{2q_{xy}}{q_{xx} + q_{yy}} \tag{21}$$

Fig. 6 indicates how the signs of ϵ_1 and ϵ_2 describe the orientation of an elliptical image. Plugging eq.(20) into the first of eq.(21):

$$\epsilon_1 = \frac{\int d^2\theta I_{\mathrm{obs}}(\theta)\,[\theta_x\theta_x - \theta_y\theta_y]}{\int d^2\theta I_{\mathrm{obs}}(\theta)\,[\theta_x\theta_x + \theta_y\theta_y]} \tag{22}$$

Now, we use the fact that surface brightness is conserved between the source and the image[7], such that $I_{\rm obs}(\theta) = I_{\rm true}(\theta^S)$. Using eq.(16), we can write that for small deflections the position of points in the source and image planes are related by $\theta_i = \left(\mathcal{A}^{-1}\right)_i^j \theta_j^S$. Hence:

$$\epsilon_1 = \frac{\int d^2\theta^S \,|\det \mathcal{A}^{-1}|\, \left[\left(\mathcal{A}^{-1}\right)_x^i \left(\mathcal{A}^{-1}\right)_x^j - \left(\mathcal{A}^{-1}\right)_y^i \left(\mathcal{A}^{-1}\right)_y^j\right] I_{\rm true}(\theta^S)\,\theta_i^S\theta_j^S}{\int d^2\theta^S \,|\det \mathcal{A}^{-1}|\, \left[\left(\mathcal{A}^{-1}\right)_x^i \left(\mathcal{A}^{-1}\right)_x^j + \left(\mathcal{A}^{-1}\right)_y^i \left(\mathcal{A}^{-1}\right)_y^j\right] I_{\rm true}(\theta^S)\,\theta_i^S\theta_j^S} \tag{23}$$

where $\det \mathcal{A}^{-1}$ has appeared when we changed the integration variable to θ^S. Note that the $\mathcal{A}$-matrices can now be pullled outside of the integral, since they don't depend on θ^S (indeed, the contents of the $\mathcal{A}$ matrices describe what happens to the photon *after* leaving the source plane).

We will take the average underlying, true source image of a galaxy to be circular. Of course this is not true for any given object, but when we average over thousands of randomly-oriented galaxies in a survey, it will be. Then the quadrupole moments of the true image vanish unless we have $i = j$; so the integral appearing in both the numerator and the denominator must be proportional to δ_{ij}. Since the $\mathcal{A}$ matrices have been taken outside, we now find we have exactly the same integrated quantity in both the numerator and denominator. Cancelling them (and using the fact that $\mathcal{A}$ is symmetric), we arrive at:

$$\begin{aligned}
\epsilon_1 &= \frac{\left[\left(\mathcal{A}^{-1}\right)_x^i \left(\mathcal{A}^{-1}\right)_x^j - \left(\mathcal{A}^{-1}\right)_y^i \left(\mathcal{A}^{-1}\right)_y^j\right]\delta_{ij}}{\left[\left(\mathcal{A}^{-1}\right)_x^i \left(\mathcal{A}^{-1}\right)_x^j + \left(\mathcal{A}^{-1}\right)_y^i \left(\mathcal{A}^{-1}\right)_y^j\right]\delta_{ij}} && (24)\\
&= \frac{\left[\left(\mathcal{A}^{-1}{}_x^x\right)^2 - \left(\mathcal{A}^{-1}{}_y^y\right)^2\right]}{\left[\left(\mathcal{A}^{-1}{}_x^x\right)^2 + 2\left(\mathcal{A}^{-1}{}_y^x\right)^2 + \left(\mathcal{A}^{-1}{}_x^x\right)^2\right]} && (25)
\end{aligned}$$

We can straightforwardly find the inverse matrix $\mathcal{A}^{-1}$ (but remember to add the identity matrix back onto eq.(17)!)

$$\mathcal{A}^{-1} = \frac{1}{(1-\kappa)^2 - \gamma_1^2 - \gamma_2^2}\begin{pmatrix} 1-\kappa+\gamma_1 & \gamma_2 \\ \gamma_2 & 1-\kappa-\gamma_1 \end{pmatrix} \tag{26}$$

Plugging the components of this into eq.(25):

$$\begin{aligned}
\epsilon_1 &= \frac{(1-\kappa+\gamma_1)^2 - (1-\kappa-\gamma_1)^2}{(1-\kappa+\gamma_1)^2 + 2\gamma_2^2 + (1-\kappa-\gamma_1)^2} && (27)\\
&= \frac{4\gamma_1(1-\kappa)}{2(1-\kappa)^2 + 2\gamma_1^2 + 2\gamma_2^2} && (28)
\end{aligned}$$

If the distortions and magnifications induced by lensing are small (which is the case for *weak* gravitational lensing), then we can drop quantities that are second-order in κ and γ_i. This leads to:

$$\epsilon_1 \simeq \frac{4\gamma_1}{2(1-2\kappa)} \simeq 2\gamma_1 \tag{29}$$

You can show, via a totally analogous calculation for ϵ_2, that $\epsilon_2 \simeq 2\gamma_2$. Hence measuring the shapes of many galaxy images – and hence getting a statistical measurement of ϵ_1 and ϵ_2 in a particular direction on the sky – we can get estimates of γ_1 and γ_2. And these shear parameters, we know from eq.(17), are related to (derivatives) of the gravitational potential field.

[7] Gravitational lensing does not create or destroy photons – it merely alters their paths. True, a lens can deflect photons such that they reach an observer they otherwise would have missed. This means that the total integrated flux the observer receives from that source is higher. However, the price paid for this is that the area of that source appears larger on the sky. Therefore the surface brightness – the flux per unit source area (per unit time per frequency interval etc.) is conserved. See the Padmanabhan reference given at the start of these notes, p461, for a nice proof of this result using phase space densities.

3.3 Source Distribution

We nearly have all the pieces we need to start calculating observable quantities. However, so far we have discussed only the distortion of light rays from a single source object (galaxy). As mentioned at the start, weak gravitational lensing involves correlating the distortions of a whole population of galaxies in order to probe the large-scale gravitational field of cosmological structure. To do this, we need to integrate eq.(18) over a source population to find the total distortion tensor (which, with an abuse of notation, we will also denote as ψ^i_j).

We introduce a function $W(\chi)$ which describes the distribution of the redshifts of our source galaxies. The simplest example here would be a Gaussian peaked at (say) $z \sim 2$. (Of course, this is not a realistic example, as we'd expect our source distribution to die away more rapidly at high redshift where galaxies become fainter and thus harder to detect.) We'll take the function $W(\chi)$ to be appropriately normalised such that $\int W(\chi)\, d\chi = 1$. Our total distortion tensor is then:

$$\psi^i_j = -\delta^{ik} \int_0^{\chi_\infty} d\chi\, W(\chi) \int_0^{\chi} d\chi' \,\left\{ \partial_j \partial_k [\Psi(\tilde{\vec{x}}) + \Phi(\tilde{\vec{x}})]\, \chi' \left(1 - \frac{\chi'}{\chi}\right) \right\} \tag{30}$$

where χ_∞ is the furthest limit of our source population, and we've reminded ourselves that the gravitational potentials are functions of the position vector $\vec{x} = \chi(\theta_1, \theta_2, 1)$. Reversing the order of integration like we did before, we can rewrite this as:

$$\psi^i_j = -\delta^{ik} \int_0^{\chi_\infty} d\chi\, \partial_j \partial_k [\Phi(\vec{x})]\, g(\chi) \qquad \text{where} \qquad g(\chi) = 2\chi \int_\chi^{\chi_\infty} \left(1 - \frac{\chi}{\chi'}\right) W(\chi')\, d\chi' \tag{31}$$

To keep life simple, I've also adopted the GR case $\Psi = \Phi$ in the line above. We'll stick with this limit from now on.

4 Power Spectra

We can finally bite the bullet and calculate power spectra of the components of the distortion tensor. Remember, a power spectrum is just the Fourier-space version of a correlation function. And a correlation function – informally speaking – is simply something which tells you how closely related two points are[8] as a function of the distance between them. If you select any two points at random, there's a chance that by fluke they might have very similar field values. So instead, we have to average over many pairs of points. The general kind of relation we will need multiple times in what follows is then (shown here for Φ):

$$\langle \tilde{\Phi}(\vec{k}) \tilde{\Phi}^*(\vec{k}') \rangle = (2\pi)^3\, \delta^{(3)}(\vec{k} - \vec{k}')\, P_\Phi(k) \tag{32}$$

Inverting this relation, one has:

$$P_\Phi(k) = \int \frac{d^3k'}{(2\pi)^3} \langle \tilde{\Phi}(\vec{k}) \tilde{\Phi}^*(\vec{k}') \rangle \tag{33}$$

Note that: a) the definition of the power spectrum involves a complex conjugate. Because although $\Phi(\vec{x})$ is real, its Fourier transform will pick up complex exponentials; and b) the power of 2π and kind of delta function on the RHS depend on the dimensionality of the Fourier-space variable. In the line above appears a 3D wavevector $\vec{k}$, which is conjugate to real-space 3D position vector $\vec{x}$. In what follows, however, we will also need the 2D Fourier variable $\vec{\ell}$, which is conjugate to our 2D angular position vector $\vec{\theta}$.

4.1 Power Spectrum of the Distortion Tensor

We'll first form the power spectrum that correlates two general components of the distortion tensor (note that this will be a four-index object). We'll then see how to tease this apart into magnification and shear components. Here goes:

$$P_{\psi_{ijpq}}(\ell) = \int \frac{d^2\ell'}{(2\pi)^2} \langle \tilde{\psi}_{ij}(\vec{\ell}) \tilde{\psi}^*_{pq}(\vec{\ell}') \rangle \tag{34}$$

[8] More precisely, the field value at those two points.

where

$$\tilde{\psi}_{ij}(\vec{\ell}) = \int d^2\theta \, \psi_{ij}(\vec{\theta}) \, e^{-i\vec{\ell}\cdot\vec{\theta}} \tag{35}$$

$$= -\int d^2\theta \int_0^{\chi_\infty} d\chi \, \partial_i\partial_j[\Phi(\vec{x})] \, g(\chi) \, e^{-i\vec{\ell}\cdot\vec{\theta}} \tag{36}$$

$$= \int d^2\theta \int_0^{\chi_\infty} d\chi \int \frac{d^3k}{(2\pi)^3} \, k_i k_j \tilde{\Phi}(\vec{k}) \, g(\chi) \, e^{-i\vec{\ell}\cdot\vec{\theta}} e^{i\vec{k}\cdot\vec{x}} \tag{37}$$

where the first equality shows my conventions for a Fourier transform, the second equality uses eq.(31) and the third line converts the real-space potential into its Fourier transform also. Note the derivatives have become factors of k_i.

We squidge two copies of eq.(37) together (one C.C.) and use this in eq.(34). Note that this gives us a horrendous seven-dimensional integral!

$$\begin{aligned} P_{\psi_{ijpq}}(\vec{\ell}) = \int \frac{d^2\ell'}{(2\pi)^2} \int_0^{\chi_\infty} d\chi \int_0^{\chi_\infty} d\chi' \int d^2\theta \int d^2\theta' \int \frac{d^3k}{(2\pi)^3} \int \frac{d^3k'}{(2\pi)^3} \\ \times \, k_i k_j k'_i k'_j \, \langle \tilde{\Phi}(\vec{k}) \tilde{\Phi}^*(\vec{k}') \rangle \, g(\chi) g(\chi') \, e^{-i\left(\vec{\ell}\cdot\vec{\theta} - \vec{\ell}'\cdot\vec{\theta}'\right)} e^{i\left(\vec{k}\cdot\vec{x} - \vec{k}'\cdot\vec{x}'\right)} \end{aligned} \tag{38}$$

Don't panic; we're about to kill off a lot of these integrals. We start by replacing $\langle\Phi\Phi^*\rangle$ using eq.(32); the delta-function then kills the k' integral.

$$P_{\psi_{ijpq}}(\vec{\ell}) = \int \frac{d^2\ell'}{(2\pi)^2} \int_0^{\chi_\infty} d\chi \int_0^{\chi_\infty} d\chi' \int d^2\theta \int d^2\theta' \int \frac{d^3k}{(2\pi)^3} \, k_i k_j k_i k_j \, P_\Phi(k) \, g(\chi) g(\chi') \, e^{-i\left(\vec{\ell}\cdot\vec{\theta} - \vec{\ell}'\cdot\vec{\theta}'\right)} e^{i\vec{k}\cdot(\vec{x}-\vec{x}')} \tag{39}$$

Next up will be the θ and θ' integrals. Note that the two exponentials, upon expansion of their arguments, are:

$$\exp\left[-i\left(\ell_1\theta_1 + \ell_2\theta_2 - k_1\chi\theta_1 - k_2\chi\theta_2\right)\right] \times \exp\left[i\left(\ell'_1\theta'_1 + \ell'_2\theta'_2 - k_1\chi'\theta'_1 - k_2\chi'\theta'_2\right)\right] \times \exp\left[ik_3\left(\chi - \chi'\right)\right]$$

When we integrate over θ and θ' the first two factors above will yield the delta functions $\delta^{(2)}\left(\vec{\ell} - \chi\vec{k}_{2D}\right)$ and $\delta^{(2)}\left(\vec{\ell}' - \chi'\vec{k}_{2D}\right)$, alongside factors of 2π. $\vec{k}_{2D}$ here explicitly refers to the first two components of $\vec{k}$, which reside in the image plane. The third component, corresponding to the direction along the line of sight, is separated out in the last exponential above.

$$\begin{aligned} P_{\psi_{ijpq}}(\vec{\ell}) = \int d^2\ell' \int_0^{\chi_\infty} d\chi \int_0^{\chi_\infty} d\chi' \int \frac{d^3k}{2\pi} \, k_i k_j k_i k_j \, P_\Phi(k) \, g(\chi) g(\chi') \\ \times \, \delta^{(2)}\left(\vec{\ell} - \chi\vec{k}_{2D}\right) \, \delta^{(2)}\left(\vec{\ell}' - \chi'\vec{k}_{2D}\right) \, e^{[ik_3(\chi-\chi')]} \end{aligned} \tag{40}$$

Now the k_3 part of the d^3k integral gives $2\pi\,\delta(\chi - \chi')$, which can be used to kill the χ' integral.

$$P_{\psi_{ijpq}}(\vec{\ell}) = \int d^2\ell' \int_0^{\chi_\infty} d\chi \int d^2k_{2D} \, k_i k_j k_i k_j \, P_\Phi(k) \, g(\chi)^2 \, \delta^{(2)}\left(\vec{\ell} - \chi\vec{k}_{2D}\right) \, \delta^{(2)}\left(\vec{\ell}' - \chi\vec{k}_{2D}\right) \tag{41}$$

Our LHS is just a function of ℓ, but currently we have both ℓ and k on the RHS. So we'll use one delta function to replace k by ℓ'/χ, and rewrite the second delta function. Note that we can then remove the outer integral in ℓ', since it's now taken care of by the innermost one:

$$P_{\psi_{ijpq}}(\vec{\ell}) = \int_0^{\chi_\infty} d\chi \int \frac{d^2\ell'}{\chi^2} \, \frac{\ell'_i \ell'_j \ell'_p \ell'_q}{\chi^4} \, P_\Phi(\ell/\chi) \, g(\chi)^2 \, \delta^{(2)}\left(\vec{\ell}' - \vec{\ell}\right) \tag{42}$$

Finally, the last delta function kills the ℓ' integral, leaving us with our result:

$$P_{\psi_{ijpq}}(\vec{\ell}) = \int_0^{\chi_\infty} d\chi \, \frac{g(\chi)^2}{\chi^2} \, \frac{\ell_i \ell_j \ell_p \ell_q}{\chi^4} \, P_\Phi(\ell/\chi) \tag{43}$$

4.2 Power Spectra for Shear and Convergence, E & B Modes

$P_{\psi_{ijpq}}(\vec{\ell})$ above is the power spectrum showing the correlation between any two components ij and pq of the distortion tensor; hence it's a four-index object. We are particularly interested in the correlations for the components of ψ_{ij} that correspond to shear and convergence. Using eq.(17), we see that:

$$\kappa = -\frac{1}{2}\left(\psi_{11} + \psi_{22}\right) \tag{44}$$

$$\gamma_1 = \frac{1}{2}\left(\psi_{22} - \psi_{11}\right) \tag{45}$$

$$\gamma_2 = -\psi_{12} \tag{46}$$

So we can see that the power spectrum of the convergence is:

$$P_\kappa(\vec{\ell}) = \langle \kappa\kappa* \rangle = \frac{1}{4}\langle(\psi_{11} + \psi_{22})\,(\psi^*_{11} + \psi^*_{22})\rangle \tag{47}$$

$$= \frac{1}{4}\left[\langle \psi_{11}\psi^*_{11})\rangle + \langle \psi_{22}\psi^*_{22})\rangle + \langle \psi_{11}\psi^*_{22})\rangle + \langle \psi_{22}\psi^*_{11})\rangle\right] \tag{48}$$

$$= \frac{1}{4}\left[P_{\psi\,1111}(\vec{\ell}) + P_{\psi\,2222}(\vec{\ell}) + 2P_{\psi\,1122}(\vec{\ell})\right] \tag{49}$$

Recall that $\vec{\ell}$ is a 2D vector, the Fourier conjugate to $\vec{\theta}$. We're going to switch from describing it via two components $\{\ell_1, \ell_2\}$ to using a magnitude and an angle, i.e. $\ell_1 = \ell\cos\phi$ and $\ell_2 = \ell\sin\phi$. This will enable us to make use of trig identities to simplify things.

Using this, together with eqs.(43) and (49), we get:

$$P_\kappa(\ell) = \left[\sin^4\phi + \cos^4\phi + 2\sin^2\phi\cos^2\phi\right]\,\frac{\ell^4}{4}\int_0^{\chi_\infty} d\chi\,\frac{g(\chi)^2}{\chi^6}\,P_\Phi(\ell/\chi) \tag{50}$$

$$\Rightarrow \qquad P_\kappa(\ell) = \frac{\ell^4}{4}\int_0^{\chi_\infty} d\chi\,\frac{g(\chi)^2}{\chi^6}\,P_\Phi(\ell/\chi) \tag{51}$$

You can see that the leading bracket in the first line above is equal to unity. In a similar vein, the power spectra of γ_1 and γ_2 are:

$$P_{\gamma_1}(\ell,\,\phi) = \cos^2(2\phi)\,\frac{\ell^4}{4}\int_0^{\chi_\infty} d\chi\,\frac{g(\chi)^2}{\chi^6}\,P_\Phi(\ell/\chi) = \cos^2(2\phi)\,P_\kappa(\ell) \tag{52}$$

$$P_{\gamma_2}(\ell,\,\phi) = \sin^2(2\phi)\,\frac{\ell^4}{4}\int_0^{\chi_\infty} d\chi\,\frac{g(\chi)^2}{\chi^6}\,P_\Phi(\ell/\chi) = \sin^2(2\phi)\,P_\kappa(\ell) \tag{53}$$

However, the two lines above say something slightly odd. They tell us that the power spectra for the shear components depend on ϕ, which is the angle made with an arbitrarily chosen axis in the (Fourier-space) image plane. Clearly our choice of axis can't have an effect on the underlying physics. This suggests that the P_{γ_i} are not the most sensible variables to work with. It turns out that a particular linear combination of the shear components produces a power spectrum that is independent of ϕ.

Consider the following combinations:

$$E(\ell,\phi) = \cos(2\phi)\gamma_1(\ell,\phi) + \sin(2\phi)\gamma_2(\ell,\phi) \tag{54}$$

$$B(\ell,\phi) = -\sin(2\phi)\gamma_1(\ell,\phi) + \cos(2\phi)\gamma_2(\ell,\phi) \tag{55}$$

Now look what happens when we take the power spectrum of E (suppressing arguments for ease of notation):

$$P_E(\ell) = \cos^2(2\phi)\,P_{\gamma_1}(\ell,\phi) + \sin^2(2\phi)\,P_{\gamma_2}(\ell,\phi) + 2\sin(2\phi)\cos(2\phi)\,P_{\gamma_1\,\gamma_2}(\ell,\phi) \tag{56}$$

We need

$$P_{\gamma_1\,\gamma_2}(\ell,\phi) = -\frac{1}{2}\langle \psi_{12}\,(\psi_{22} - \psi_{11})\rangle \tag{57}$$

$$= -\frac{1}{2}\left[P_{\psi\,1222} - P_{\psi\,1211}\right] \tag{58}$$

$$= -\frac{\ell^4}{2}\left[\sin^3\phi\cos\phi - \sin\phi\cos^3\phi\right]\int_0^{\chi_\infty} d\chi\,\frac{g(\chi)^2}{\chi^6}\,P_\Phi(\ell/\chi) \tag{59}$$

$$= 2\sin\phi\cos\phi\left[\cos^2\phi - \sin^2\phi\right]P_\kappa(\ell) \tag{60}$$

$$= \sin(2\phi)\cos(2\phi)P_\kappa(\ell) \tag{61}$$

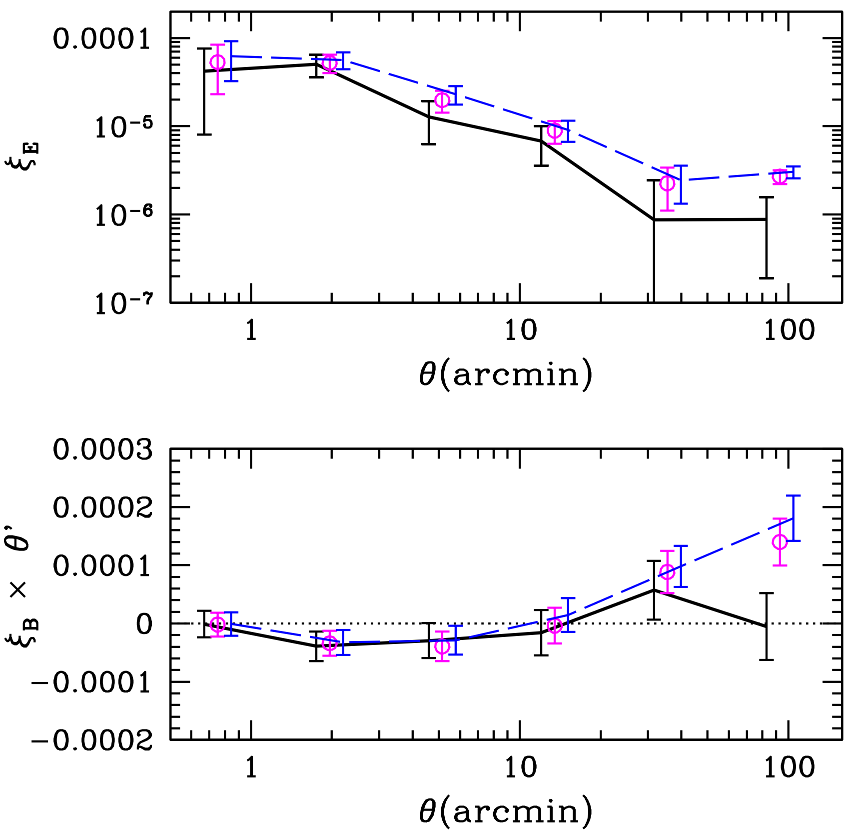

Figure 7: E- and B-mode correlation functions from the KiloDegree Survey (KiDS), Kuijken et al. (2014). See text for description.

Stick this into eq.(56) and use eqs.(52) and (53):

$$P_E(\ell) = \cos^2(2\phi)\, P_{\gamma_1}(\ell,\phi) + \sin^2(2\phi)\, P_{\gamma_2}(\ell,\phi) + 2\sin^2(2\phi)\cos^2(2\phi)\, P_\kappa(\ell) \tag{62}$$

$$\Rightarrow \qquad P_E(\ell) = P_\kappa(\ell) \tag{63}$$

Whereas for the power spectrum of B we find:

$$P_B(\ell) = \sin^2(2\phi)\, P_{\gamma_1}(\ell,\phi) + \cos^2(2\phi)\, P_{\gamma_2}(\ell,\phi) - 2\sin(2\phi)\cos(2\phi)\, P_{\gamma_1\,\gamma_2}(\ell,\phi) \tag{64}$$

$$= \left[\sin^2(2\phi)\cos^2(2\phi) + \cos^2(2\phi)\sin^2(2\phi) - 2\sin^2(2\phi)\cos^2(2\phi)\right]\, P_\kappa(\ell) \tag{65}$$

$$\Rightarrow \qquad P_B(\ell) = 0 \tag{66}$$

These results are massively useful. Firstly, eq.(63) tells us how to extract real, physical information that is independent of any observer-imposed coordinate choice. What's more, if we use our shear measurements to calculate the power spectrum of the E-mode, we get the convergence (magnification) power spectrum for free. Eq.(66) is also extremely important because it allows us to check for systematics (i.e. unmodelled sources of error) in our measurements. If we've done our job properly[9], then the B-mode power spectrum should vanish.

Fig. 4.2 shows measurements of the E and B-mode correlation functions from the KiloDegree Survey, the largest dedicated weak lensing survey to date. Note that these plots show the **real-space** correlation function, which is the Fourier transform of the power spectra calculated above. Nevertheless, the basic information

[9]There are a lot of subtleties to be accounted for when measuring lensing shear, such as errors in our galaxy redshifts (a necessary evil of trying to survey large numbers of galaxies very quickly) and intrinsic alignments (the fact that galaxies near each other are likely be aligned anyway due to local gravitational fields, and not just because intervening dark matter structures lens their photons in the same direction).

content is the same. Note also that in the B-mode plot, the correlation function has been multiplied by θ to emphasise deviations from zero at large angular separations.

The blue, dashed line shows a naive interpretation of the raw data. The pink (open) points show the data after removing from the sample some patches of sky where the observations were particularly poor in quality (due to bad weather, obscuring dust in our own galaxy etc.) The black, solid line shows the data after applying further corrections for errors introduced by instrumental effects, i.e. miscalibrations or other errors introduced by the telescope itself. You can see that after sufficient error budgeting, the data are consistent with a zero B-mode spectrum[10]. However, these curves also show just how crucial error management is for correct interpretation of weak lensing data!

5 Lensing of the Cosmic Microwave Background

In previous lectures you've studied how the cosmic microwave background (CMB) is produced during the early thermal history of the universe. You've also seen that features in the CMB – anisotropies – can tell us something about inflation and cosmological parameters, and so are an intense object of study.

Problem: in this lecture we've calculated how images of distant galaxies get distorted as their photons pass through large-scale potential wells in the universe. Shouldn't exactly the same thing be happening to the CMB photons[11]? Doesn't this mean, then, that the CMB anisotropies we're so keen to study actually get 'moved around', changed by lensing en route to us?

This turns out to be exactly correct. Fortunately, as we will estimate below, the effect of lensing on the CMB is not so large as to eradicate all the useful information from it. However, it is large enough to have a measurable effect that must be carefully modelled, and is of interest in its own right.

5.1 Order of Magnitude Estimates

Recall from your GR course that the angular deflection of a particle by a point mass is:

$$\alpha = \frac{4GM}{bc^2} \sim \frac{2\Phi}{c^2} \tag{67}$$

where b is the impact parameter between the particle and the point mass. You've learnt in this course that the large-scale potential wells in the universe have a depth of about $\sim 10^{-5}$. They also have an average comoving size of around 300 Mpc[12]. The CMB itself is at a comoving distance of $\sim$ 14 Gpc from us, so we expect a CMB photon to have passed through roughly $14 \times 10^3/300 \sim 50$ such potential wells. Therefore we could estimate its total deflection from its original direction of motion (at the time of last scattering) to be roughly $50 \times 10^{-5} \sim 10^{-4}$ radians, which is of order an arcminute in degrees.

At the same time, the angular scale subtended by our 'average' potential well on the sky – taking it to be roughly halfway between us and the CMB – is about $300/1400 \sim 2^\circ$. So, although our CMB photons are only being diverted by an arcminute or so, we expect their deflections to be *coherent over angular scales of degrees.*

Recall that the first peak in the power spectrum of the CMB is at scales of around 1°. Remember also that the power spectrum shows information effectively averaged over all directions in the sky. So although individual anisotropies may get deflected in a particular direction, the net effect of lensing on the CMB power spectrum is that it 'blurs out' the scale of the CMB peaks in general. The size of the first acoustic peak anisotropies gets increased by $2'/1^\circ \sim 3\%$. As one goes to smaller scales in the CMB, the relative effect of this blurring becomes larger.

[10] The error bars in Fig. 4.2 show the 1σ errors; remember you need a 3σ effect (at least!) to claim a statistically significant deviation from zero.

[11] Of course the CMB photons, being in the microwave part of the EM spectrum, are of much lower frequency than optical galaxy images, but this makes no difference. Gravitational lensing is **achromatic**, i.e. it does not depend on photon frequency.

[12] This is rather sloppy. There are potential wells of all physical sizes in the universe, and defining an average is a bit meaningless. I really mean that the peak of the matter power spectrum is at scales corresponding to $\sim$ 300 Mpc.

5.2 Power Spectrum of the CMB Lensing Potential

Consider that CMB photons originally travelling in the direction $\hat{\vec{n}}$ get deflected through an angle $\vec{\alpha}$ such that their direction is $\hat{\vec{n}}' = \hat{\vec{n}} + \vec{\alpha}$. We write the CMB temperature measured in that new direction as:

$$T\left(\hat{\vec{n}}'\right) = T\left(\hat{\vec{n}} + \vec{\alpha}\right) \tag{68}$$

The expression for the deflection angle $\vec{\alpha}$ is exactly what we worked out in eq.(15). The only change we need make is changing the upper limit of the integral to a fixed distance χ_*, the conformal distance to the CMB. Unlike the galaxy weak lensing case, where we needed to integrate over a source distribution $W(\chi)$, for CMB lensing all our source photons come from the same redshift[13].

To lowest order, we can write $\vec{\alpha}$ as the gradient of a scalar potential β, i.e.

$$\vec{\alpha} = \nabla\beta = \frac{\nabla_{\text{ang}}\beta}{\chi} \tag{69}$$

where ∇_{ang} is the angular derivative on the sphere at conformal distance χ. Using our eq.(15), we can then write:

$$\beta(\hat{\vec{n}}) = -2\int_0^{\chi_*} d\chi\, \left(\frac{\chi_* - \chi}{\chi\,\chi_*}\right)\, \Phi[\vec{x}(\chi,\, \hat{\vec{n}})] \tag{70}$$

where β is the **CMB lensing potential** and we have shown explicitly the direction of observation as an argument. The CMB lensing potential is often denoted by ψ; I've avoided this here to prevent confusion with the distortion tensor of weak galaxy lensing. Once again we have specialised to the GR case of $\Phi = \Psi$ above.

We now want to find the power spectrum of β. The calculation is quite similar to that of §4, but somewhat easier because of the lack of indices on β. We start by Fourier transforming the potential as usual (this time sticking to k and x instead of ℓ and θ):

$$\Phi(\vec{x}) = \int \frac{d^3k}{(2\pi)^3}\, \tilde{\Phi}(\vec{k})\, e^{i\vec{k}\cdot\vec{x}} \tag{71}$$

Taking the two-point correlation function of β (and writing $f(\chi) = (\chi_* - \chi)/\chi\,\chi_*$):

$$\langle\beta(\vec{n})\beta^*(\vec{n}')\rangle = 4\int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi'\; f(\chi)\, f(\chi') \int \frac{d^3k}{(2\pi)^3}\int \frac{d^3k'}{(2\pi)^3}\, \langle\tilde{\Phi}(\vec{k})\tilde{\Phi}(\vec{k}')\rangle\, e^{i(\vec{k}\cdot\vec{x} - \vec{k}'\cdot\vec{x}')} \tag{72}$$

$$= 4\int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi'\; f(\chi)\, f(\chi') \int \frac{d^3k}{(2\pi)^3}\, P_\Phi(\vec{k},\chi)\, e^{i\vec{k}\cdot(\vec{x} - \cdot\vec{x}')} \tag{73}$$

This time we'll make use of the following identity for the expansion of Fourier basis functions:

$$e^{i\vec{k}\cdot\vec{x}} = 4\pi\sum_{\ell m} i^\ell\, j_\ell(k\chi)\, Y^*_{\ell m}(\hat{\vec{n}})\, Y_{\ell m}(\hat{\vec{k}}) \tag{74}$$

where j_ℓ are the spherical Bessel functions. Using this and breaking the d^3k integral into angular and radial parts:

$$\langle\beta(\vec{n})\beta^*(\vec{n}')\rangle = 64\pi^2\int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi'\; f(\chi)\, f(\chi') \int\int_{\theta=0}^{\pi}\int_{\phi=0}^{2\pi} \frac{k^2\,dk}{(2\pi)^3}\, d\theta\, d\phi$$

$$\times\, P_\Phi(\vec{k},\chi) \sum_{\ell\ell' mm'} i^{\ell-\ell'}\, j_\ell(k\chi)\, j_\ell(k\chi')\, Y^*_{\ell m}(\vec{n})\, Y_{\ell m}(\vec{k})\, Y_{\ell' m'}(\vec{n}')\, Y^*_{\ell' m'}(\vec{k}') \tag{75}$$

$$= 16\int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi'\; f(\chi)\, f(\chi') \int \frac{k^2\,dk}{2\pi}\, P_\Phi(\vec{k},\chi) \sum_{\ell\ell' mm'} j_\ell(k\chi)\, j_\ell(k\chi')\, Y^*_{\ell m}(\vec{n})\, Y_{\ell' m'}(\vec{n})\, \delta_{\ell\ell'}\delta_{mm'} \tag{76}$$

where we have used the orthogonality condition of the spherical harmonics to reach the second equality.

[13] We are approximating last scattering as being an instantaneous event. Of course this is not strictly true.

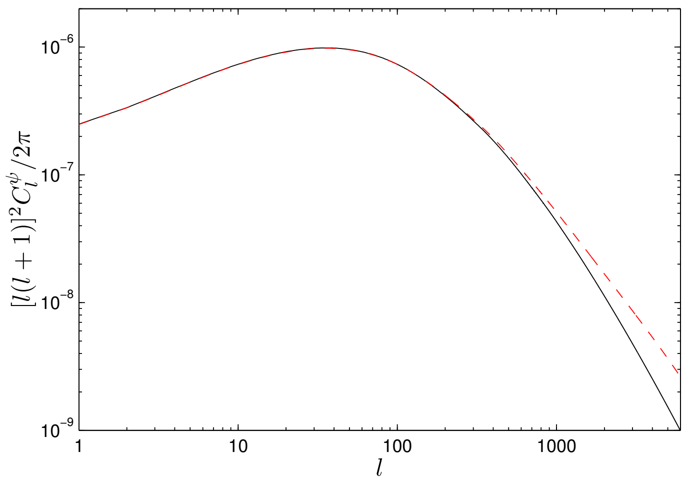

Figure 8: Power spectrum of the CMB lensing potential, as per eq.(81). The dashed line shows a model that accounts for corrections beyond linear perturbation theory; these become increasingly important at small angular scales (high ℓ).

Now we work on the LHS side a little. For CMB-related quantities, it's usual to express power spectra in terms of the 'C_ℓ'. These are the the two-point correlators of the Fourier coefficients when a quantity is expanded in terms of spherical harmonics. That is,

$$\beta(\vec{\hat{n}}) = \sum_{\ell m} \beta_{\ell m} Y_{\ell m}(\vec{\hat{n}}) \tag{77}$$

$$\Rightarrow \qquad \langle \beta_{\ell m} \beta^*_{\ell' m'} \rangle = \delta_{\ell \ell'}\, \delta_{m m'} C_\ell^\beta \tag{78}$$

The superscript β on C_ℓ here is just to make it clear these are the C_ℓ describing the CMB lensing potential. You'll also find in the literature/textbooks C_ℓ describing the CMB temperature power spectrum, E and B polarisation modes, etc. Using the two lines above in the LHS of eq.(76), we can read off (i.e. stripping off the delta functions and $Y_{\ell m}$s):

$$C_\ell^\beta = 16 \int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi' \, f(\chi)\, f(\chi') \int \frac{k^2\, dk}{2\pi} P_\Phi(\vec{k}, \chi)\, j_\ell(k\chi)\, j_\ell(k\chi') \tag{79}$$

One final simplification: in linear theory, we can relate the potential at a given instant of time to a primordial perturbation (at the same value of k) via a *transfer function*, i.e.

$$\Phi(\vec{k}, \chi) = T_\Phi(k, \chi)\, \mathcal{R}(\vec{k}) \tag{80}$$

where we're implicitly using conformal distance χ as a time variable here. $\mathcal{R}$ is a primordial perturbation laid down during inflation, and T_Φ describes how that initial perturbation has grown. In general T_Φ is a complicated function (and sometimes can't be written down analytically), but it can be calculated numerically. Our final result can then be written as:

$$C_\ell^\beta = \frac{8}{\pi} \int k^2\, dk\, P_{\mathcal{R}}(\vec{k}, \chi) \left[\int_0^{\chi_*} d\chi\, f(\chi)\, T_\Phi(k, \eta)\, j_\ell(k\chi) \right]^2 \tag{81}$$

where $P_{\mathcal{R}}$ is the primordial power spectrum. It is described by a small number of cosmological parameters, which are relatively well-measured. Fig. 8 shows the results of this calculation for the power spectrum of the CMB lensing potential, and Fig. 9 shows the latest data from the Planck CMB satellite and the ground-based ACT and SPT telescopes. Note the x-axis is only a partial log in the second figure, hence the slightly different shapes of Figs. 8 & 9.

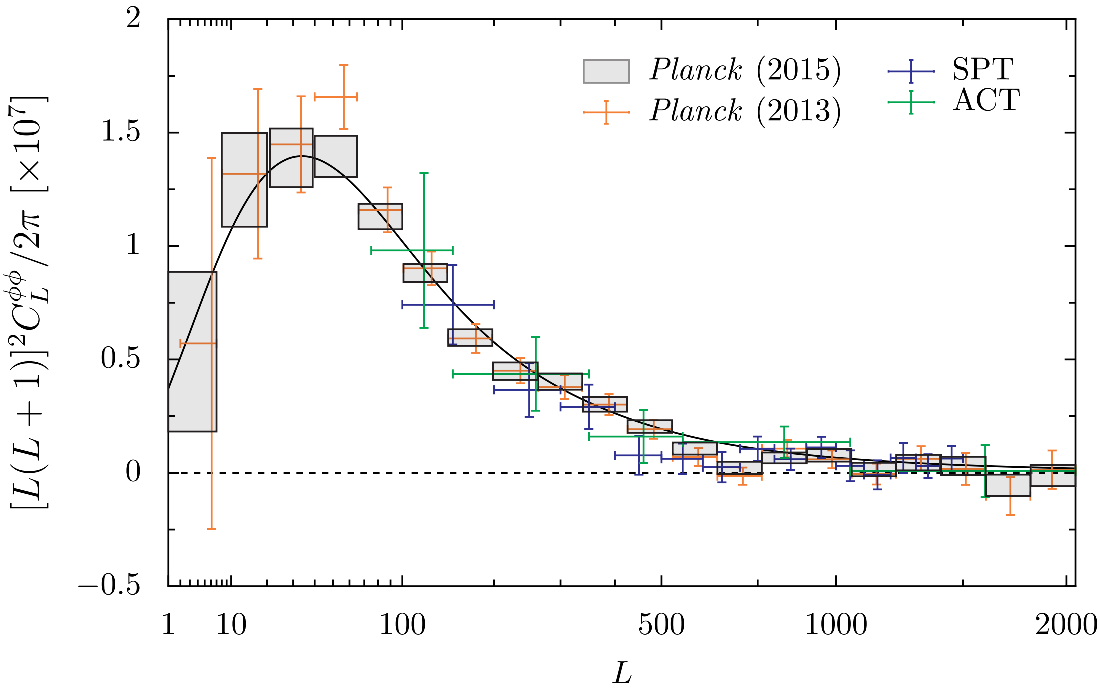

Figure 9: Measurements of the CMB lensing potential power spectrum from Planck, the South Pole Telescope (SPT) and the Atacama Cosmology Telescope (ACT). The solid line shows the prediction from the standard cosmological model, ΛCDM.