Gravitational Lensing

Gravitational lensing, which is the deflection of light by gravitational fields and the resulting effect on images, is widely useful in cosmology and, at the same time, a source of irreducible uncertainty in certain measurements. Here we discuss the basics, followed by some real-world complications and applications. Much of the development of this lecture follows the Living Reviews in Relativity article by Wambsganss.

Basics

It is sometimes stated that the deflection of light by gravity was first predicted with general relativity. This is something of an oversimplification. In Newtonian theory, you could imagine light as made of particles that traveled at the speed c, and such particles would feel gravity just like anything else. If they passed at an impact parameter b from a point mass M, they would experience a total deflection angle in radians of

$$\tilde{\alpha} = \frac{2GM}{c^2 b} \qquad \text{(wrong)} . \tag{1}$$

This Newtonian expression was first derived by Johann Soldner in 1804! It is also the expression derived by Einstein using a preliminary formulation of general relativity in 1911. He asked observers to try to measure the deflection past the limb of the Sun in the total solar eclipse of 1914. Fortunately for Einstein (and unfortunately for millions of other people), World War 1 intervened. He then found that in the final version of general relativity the actual expression is a factor of two larger:

$$\tilde{\alpha} = \frac{4GM}{c^2 b} \qquad \text{(correct)} . \tag{2}$$

This was confirmed in an expedition led by Sir Arthur Stanley Eddington in 1919 (although the error bars were larger than they reported). This English confirmation of a German scientist's correction of an Englishman's theory (i.e., Newton), so soon after a war where England and Germany were on opposite sides, made Einstein a household name.

Fundamentally, therefore, gravitational lensing just acts like classical geometric optics. Curved spacetime causes light bundles to deflect, and also to shear and expand. The result is that a background light source can have its apparent position, shape, and flux changed by a foreground gravitational lens. Keep in mind, though, that the surface brightness of a lensed object is *not* changed, because the light goes from flat spacetime to flat spacetime after going through the lens, and as we discussed earlier the surface brightness is altered only by redshifts. Also note that light is neither created nor destroyed by lensing, just redistributed.

That means that if you see source brightening due to lensing, some other observer with a different line of sight would have to see diminished flux to compensate.

The Lens Equation

Consider a spherically symmetric lens, with a mass $M(\xi)$ within an impact parameter ξ from the center (see Figure 1). The basic deflection equation then gives

$$\tilde{\alpha} = \frac{4GM(\xi)}{c^2\xi} \ . \tag{3}$$

Suppose that the lens is at an angular diameter distance D_L from us, that the source is at angular diameter distance D_S from us, and that the angular diameter distance between the two is D_{LS}. Note that for cosmologically significant distances we have to be careful: $D_{LS} \neq D_S - D_L$ in general! Let β be the angle between the lens and the actual position of the source, θ be the angle between the lens and the image of the source, and $\alpha(\theta) \equiv \theta - \beta$. One can then show that

$$\alpha = (D_{LS}/D_S)\tilde{\alpha} \ . \tag{4}$$

If the lens is not spherically symmetric then one has vectorial angles, and the two-dimensional lens equation is

$$\vec{\beta} = \vec{\theta} - \vec{\alpha}(\vec{\theta}) \ . \tag{5}$$

We can then note that since $\xi = \theta D_L$, we get (for spherical symmetry)

$$\beta(\theta) = \theta - \frac{D_{LS}}{D_L D_S}\frac{4GM}{c^2\theta} \ . \tag{6}$$

If the source is exactly behind the lens, $\beta = 0$, then we get the angular *Einstein radius*

$$\theta_E = \sqrt{\frac{4GM}{c^2}\frac{D_{LS}}{D_L D_S}} \ . \tag{7}$$

This gives an approximate scale for the angular deflection you would expect from a given lens. If the lens is roughly halfway to the source, then the Einstein radius (which, remember, is in radians and thus dimensionless) is approximately the square root of the ratio of the Schwarzschild radius ($R_s = 2GM/c^2$; this is the radius of a nonspinning black hole with mass M) to the distance to the lens.

Note also that this is approximately the angular distance from the lens needed such that a background source will have multiple images (well, for a point source it will always have multiple images; for things other than black holes, though, this need not be the case). Since the solid angle is proportional to $\theta_E^2 \propto M$, this means, for example, that the probability of multiple lensing from a star cluster or galaxy is simply proportional to the total mass,

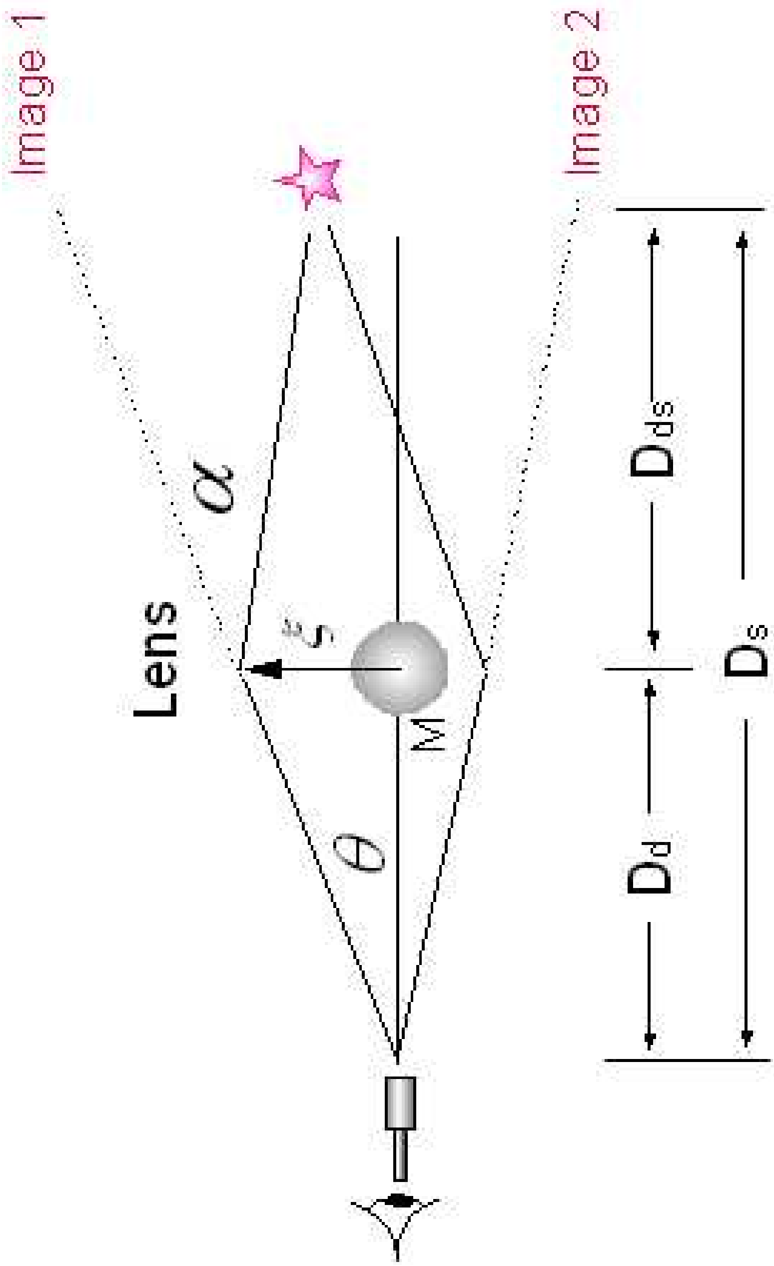

Fig. 1.— Geometry of gravitational lensing. Light from a source at angular diameter distance D_s is deflected through an angle α by a lens of mass M at angular diameter distance D_d, causing the image to shift by an apparent angle θ. From http://www.extinctionshift.com/lens_diagram.gif

independent of how that mass is distributed. Among other things, this means that although black holes are unquestionably way cool, they account for only a small fraction of multiple imaging because they account for only a small fraction of mass.

For typical cosmological distances the Einstein radius is

$$\theta_E \approx \text{few}\sqrt{M/10^{12}\,M_\odot}\ \text{arcsec} \tag{8}$$

and for a lens star in the Milky Way and a source near the Galactic bulge we have

$$\theta_E \approx 5\times 10^{-4}\sqrt{M/1\,M_\odot}\ \text{arcsec}\ . \tag{9}$$

This is why individual images have been seen cosmologically, but for microlensing (i.e., by individual stars) one only sees the increase in flux due to multiple images.

The lensing equation for a spherically symmetric source can then be rewritten as

$$\beta = \theta - \theta_E^2/\theta \tag{10}$$

and hence the image positions are

$$\theta_{1,2} = \frac{1}{2}\left(\beta \pm \sqrt{\beta^2 + 4\theta_E^2}\right)\ . \tag{11}$$

Note, as a check, that this is θ_E when $\beta = 0$. We can then use this to find the magnification of each image, that is, the ratio of the flux from the image to the flux of the unlensed source. This stems from the conservation of surface brightness, and is

$$\mu = \frac{\theta}{\beta}\frac{d\theta}{d\beta}\ . \tag{12}$$

For a symmetric source, and for a scaled angular impact parameter $u \equiv \beta/\theta_E$, the two magnifications are

$$\mu_{1,2} = \frac{u^2+2}{2u\sqrt{u^2+4}} \pm \frac{1}{2} \tag{13}$$

for a total magnification of

$$\mu_{\rm tot} = \mu_1 + \mu_2 = \frac{u^2+2}{u\sqrt{u^2+4}}\ . \tag{14}$$

This value is *always* larger than unity, even when $u \rightarrow \infty$.

Wait a second! Doesn't that contradict my previous statement that lensing doesn't create or destroy light? No, there is a subtlety in how one does the comparison. If you compare the brightness of a source *in an empty universe* with one that has matter, you find that overall convergence of the rays is larger when matter is present. If, however, you compare a universe consisting of smoothly distributed matter with one that has the same amount of matter spread lumpily, the total light is conserved.

Another important aspect of basic lensing concerns time delay. Consider the paths traveled by rays that give us two images. In general they will have (1) traveled different distances, and (2) traversed different gravitational potentials (this is known as the Shapiro delay, and is also seen from binary pulsars). If the (negative!) potential traveled is ψ, then the time delay function is

$$\tau(\vec{\theta},\vec{\beta}) = \tau_{\rm path} + \tau_{\rm grav} = \frac{1+z_L}{c}\frac{D_L D_S}{D_{LS}}\left[\frac{1}{2}(\vec{\theta}-\vec{\beta})^2 - \psi(\vec{\theta})\right] \; . \tag{15}$$

Here z_L is the redshift of the lens. If you play around with this expression you will find that the time delay between two different images is typically of the order of the light crossing time of the Schwarzschild radius of the lens, times the redshift factor [i.e., something like $(2GM/c^3)(1+z_L)$]. This can be months or years for cosmological sources lensed by galaxies, but is much too small to measure for microlensing events.

Our last comment on basic principles concerns weak lensing (where only a single image is seen). If the source is extended, e.g., a galaxy, then shear tends to distort the image. If, for example, the source is an elliptical galaxy, then the image looks like a ring arc that is concave towards the lens. Now, galaxies have lots of intrinsic shapes, so observation of just one doesn't do a lot of good. However, images of many galaxies behind a galaxy cluster can be analyzed to look for patterns. In fact, this has become a major tool to estimate the masses of clusters in a way that is independent of galaxy motions or gas temperatures. As such, it has played an important role in the evidence for dark matter, which we will discuss after the midterm.

Applications and Caveats

Lensing has a host of potential applications. Let's consider a few of them, along with reasons to be cautious at this point.

One of my favorites has always been that the probability of multiple imaging of quasars at given redshifts is strongly dependent on Ω_Λ. Therefore, in principle, all one has to do is a large survey of quasars (e.g., from the Sloan Digital Sky Survey), figure out what fraction are multiply imaged, and voila!

I still think this might yield important results down the line. At this stage, however, there are two major difficulties. The first is an issue that bedevils all sorts of cosmological observations, and is uninformatively called Malmquist bias. It is the simple statement that intrinsically bright things are easier to see (and can be seen farther) than intrinsically dim things. The issue for lensing is that strong lensing is guaranteed to give an image that is brighter than the normal source. Therefore, one would infer a much higher fraction of multiple lensing events than exists in reality. In fact, otherwise invisible quasars can be seen this way (potentially allowing extra-deep probes). Correcting for this effect would require detailed knowledge of the luminosity function of quasars below the detectable limit, which doesn't yet exist at the required level.

This effect, incidentally, is what accounts for the observation that quasars are statistically more likely to be found angularly near galaxies. Halton Arp interpreted this as evidence that quasars are ejected from galaxies, but has managed for decades to neglect the point that we should in that case see lots of blueshifted quasars ejected towards us. The full effect turns out to be somewhat complicated. Magnification means we can see dimmer quasars in the direction of a galaxy, but at the same time the background images are spread out (think of looking at something with a magnifying glass: the letters look bigger, but the distance between them also gets larger). This leads to very specific predictions about the angular correlation of galaxies with quasars as a function of quasar magnitude, and these are borne out beautifully with the Sloan data.

The second difficulty is identifying multiple lensing cases. You might think it was simple: look at two quasars separated by at most a few arcseconds, which have the same spectra and which have correlated time variability separated by a time that is consistent with the image separation. This sequence actually has worked in a few cases, notably in the first case, Q0957+561. However, it is actually pretty difficult in many cases. The problem is that spectra aren't perfect, so spectra can be similar but from genuinely different quasars. In addition, quasars are highly variable, again measurements aren't perfect, and local effects (e.g., different absorption through the different ray paths, or microlensing within the main lens) can complicate issues further. Not trivial.

Another application of lensing takes advantage of the effect that when the lens itself is complicated, many images can be produced. The locations of all the images, plus time delays between them, can in principle be used as an independent measure of H_0. In addition, as pointed out first by Chris Kochanek, the distribution of lens redshifts as a function of source redshifts can be used as a measure of Ω_Λ. In practice, though, there have been numerous difficulties in applying these methods. These are largely due to the complexities of modeling the lenses.

To me, this does not say that lensing cannot be used as a cosmological probe. Instead, I feel that it illustrates an important general point: modern day cosmology is part of astrophysics, and therefore one has to understand the real life complications of sources to reap the benefits.

Intuition Builder

Explain why lensing serves as a limiting source of uncertainty for luminosity distances to sources of all types.