Modern Observations: the Accelerating Universe

So far we're caught up to the early or possibly mid 1990s in basic cosmological observations. In this lecture, however, we are going to address the greatest surprise in cosmology in the last decade: the universal expansion is accelerating. Combined with other observations (notably of the cosmic microwave background, which we will address in a separate lecture), this implies that the universe is dominated by a truly mysterious component that has been dubbed dark energy. Most of the evidence for acceleration has come from observations of a particular type of supernova. We will therefore discuss these Type Ia supernovae and why they are especially useful for this purpose.

Type Ia Supernovae

To put this category of events into perspective, recall that in principle the measurements of luminosity distance *and* redshift from a set of objects can distinguish between different cosmological models. However, as a look at Figure 3 from Hogg's article (astro-ph/9905116) indicates, models differ significantly only at reasonably large redshifts, as in $z > 0.2$ at least, and indeed $z \sim 1$ is preferable. However, the standard candles we have discussed thus far have no chance of being observed at that distance. For example, consider $z = 0.1$. This is still in the nonrelativistic regime, so it means an apparent recession speed of roughly $cz = 30,000$ km s^{-1}. This in turn implies a distance of about 400 Mpc, from $d = v/H_0$ and $H_0 = 72$ km s^{-1} Mpc^{-1}. At that distance, though, even a very bright star with an absolute magnitude of $M_v = -5$ (such as a bright Cepheid) would have an apparent magnitude of $m_v = 33$, which is well beyond the capabilities of even Hubble to detect. Indeed, even for space-based telescopes there is a foreground glow (due to zodiacal dust) that would make observations at that level very difficult.

As a result, measuring large luminosity distances requires much brighter standard candles. One possibility is to try something with galaxies instead of stars, because we can certainly see galaxies to large redshifts. The key, as it is for Cepheids or other sources, is to identify some other property of the source that correlates strongly with the luminosity. Measuring that property and the flux from the source then tells you the luminosity distance.

For spiral galaxies, for example, one finds that the rotational velocity (as measured by Doppler shifts) tends to increase from the center of the galaxy, but asymptote to a roughly constant value $v_{\rm flat}$ at large enough distances from the center. This behavior is actually not what one would predict from the visible mass, and was one of the early indicators that dark matter might exist. However, for our current purposes it is mainly important that there exists a correlation (the Tully-Fisher relation) between the asymptotic orbital speed and the luminosity of the galaxy. Therefore, if one can get a good enough spectrum to measure $v_{\rm flat}$, one knows the luminosity and can therefore get the luminosity distance. The problem is that the relation is not quite tight enough at large enough distances to distinguish between competing cosmological models. The same goes for a similar relation for elliptical galaxies (the "fundamental plane").

One of my favorite tries of this type is slightly different. Consider two identical galaxies at different distances. Suppose that your angular resolution is one arcsecond. That means that for each galaxy you measure flux from many separate square arcseconds. Each of those areas will have different numbers of bright stars, which contribute most of the light. As a result, the flux will vary from pixel to pixel. However, the square arcsecond of the more distant galaxy will have more of those, hence the variation will be less from the more distant galaxy. This method of surface brightness fluctuations is a pretty good way to tell the distance to galaxies of known types. Again, however, it isn't quite good enough.

Where, then, shall we turn? In the last 10-15 years, the primary answer has been to look at supernovae.

Supernovae come in two basic types. To confuse matters, these are *not* Type I and Type II! It's no one's fault, but the original naming was observationally based and, reasonably enough, divided supernovae into categories that had (Type II) or did not have (Type I) hydrogen lines in their spectra. Now, however, it is understood that there are actually two rather distinct types of event: those that produce Type Ia supernovae, and those that produce everything else.

The "everything else" category involves classic supernovae: a massive star evolves, and eventually its core fuses to iron, which doesn't release energy when fused and therefore builds up passively and supported by degenerate electrons. When the core mass reaches roughly $1.4\,M_\odot$ (the Chandrasekhar mass), degeneracy support fails and the core implodes. The resulting energy of $\sim 10^{46}$ J is almost all ($\sim 99\%$) released in neutrinos, but the remaining $\sim 10^{44}$ J in photons and kinetic energy is enough to blow the stars to bits and leave behind a neutron star or black hole. We'll call these core collapse supernovae.

In contrast, the Type Ia category are believed to be binary systems with a white dwarf accreting matter from a companion. The model that seems to have emerged victorious is that if the white dwarf is very close to the Chandrasekhar mass, additional accretion can push it over. The white dwarf then starts to collapse. However, unlike the case with core collapse supernovae, the white dwarf is *not* pure iron, containing instead lighter nuclei such as carbon, oxygen, neon, and magnesium. Such nuclei do release energy when they fuse, and the collapse triggers a spectacular blast of fusion that blows the star apart. We will call these white dwarf supernovae.

In both cases, one sees the glimmerings of hope that these sources could be used as standard candles. The reason is that for each source, a characteristic mass (the Chandrasekhar mass) is involved, so the the total energy release should not be overwhelmingly different from event to event. Note, though, that core collapse supernovae and white dwarf supernovae give significantly different photon luminosities (which is what is measured), so one can't lump all supernovae in one bin.

In the 1990s, people focused in particular on white dwarf supernovae, because observations of them in galaxies with independent distance estimates suggested that their luminosities are clustered more than those of core collapse supernovae. There was still, however, enough variation from event to event (>20%) that they were not quite suitable for precision distinction between cosmological models.

Then, in 1993, a breakthrough occurred. In that year Mark Phillips published a paper showing that the peak luminosity of a white dwarf supernova is positively correlated with the time that it takes to decay from the maximum to some fraction of the maximum (or equivalently, how much it decays in the first 15 days after maximum). This is critical, because an independently measurable quantity gives a hint about the already pretty standard luminosity of these events. With this, or similar, corrections, white dwarf supernovae appear to be standard candles to better than 10%. This is good enough to do serious work.

The result is in Figure 1, which is taken from the Supernova Cosmology Project. The plot is essentially of luminosity distance on the vertical axis versus redshift on the horizontal, although the actual observable quantity is apparent magnitude instead of luminosity distance. That is, this is flux versus redshift. The curves show what is expected for different cosmological models. Here Ω_M is for matter, and Ω_Λ is for a cosmological constant.

The ultimate conclusion is that the universe is accelerating in its expansion, hence some bizarre "dark energy" is in play. We do have to be careful, though. First, from the plot we see that the scatter in the observations is dramatically larger than the differences between the curves. It is therefore necessary to rely on statistics: with many measurements, the centroid curve can be determined pretty well. More importantly, we have to worry about systematic effects. The fundamental observation is actually that white dwarf supernovae at a given redshift are slightly dimmer than expected if $\Omega = 1$ and there is no cosmological constant [compare the (1,0) curve with the (0.28,0.72) curve, and remember that higher magnitude means *lower* flux]. As Carl Sagan liked to say, extraordinary claims require extraordinary proof. What other effects might cause slight dimming at large distances and/or earlier times, and can we rule them out?

Mundane Explanations?

One possibility is that light is getting absorbed or scattered along the way. This would

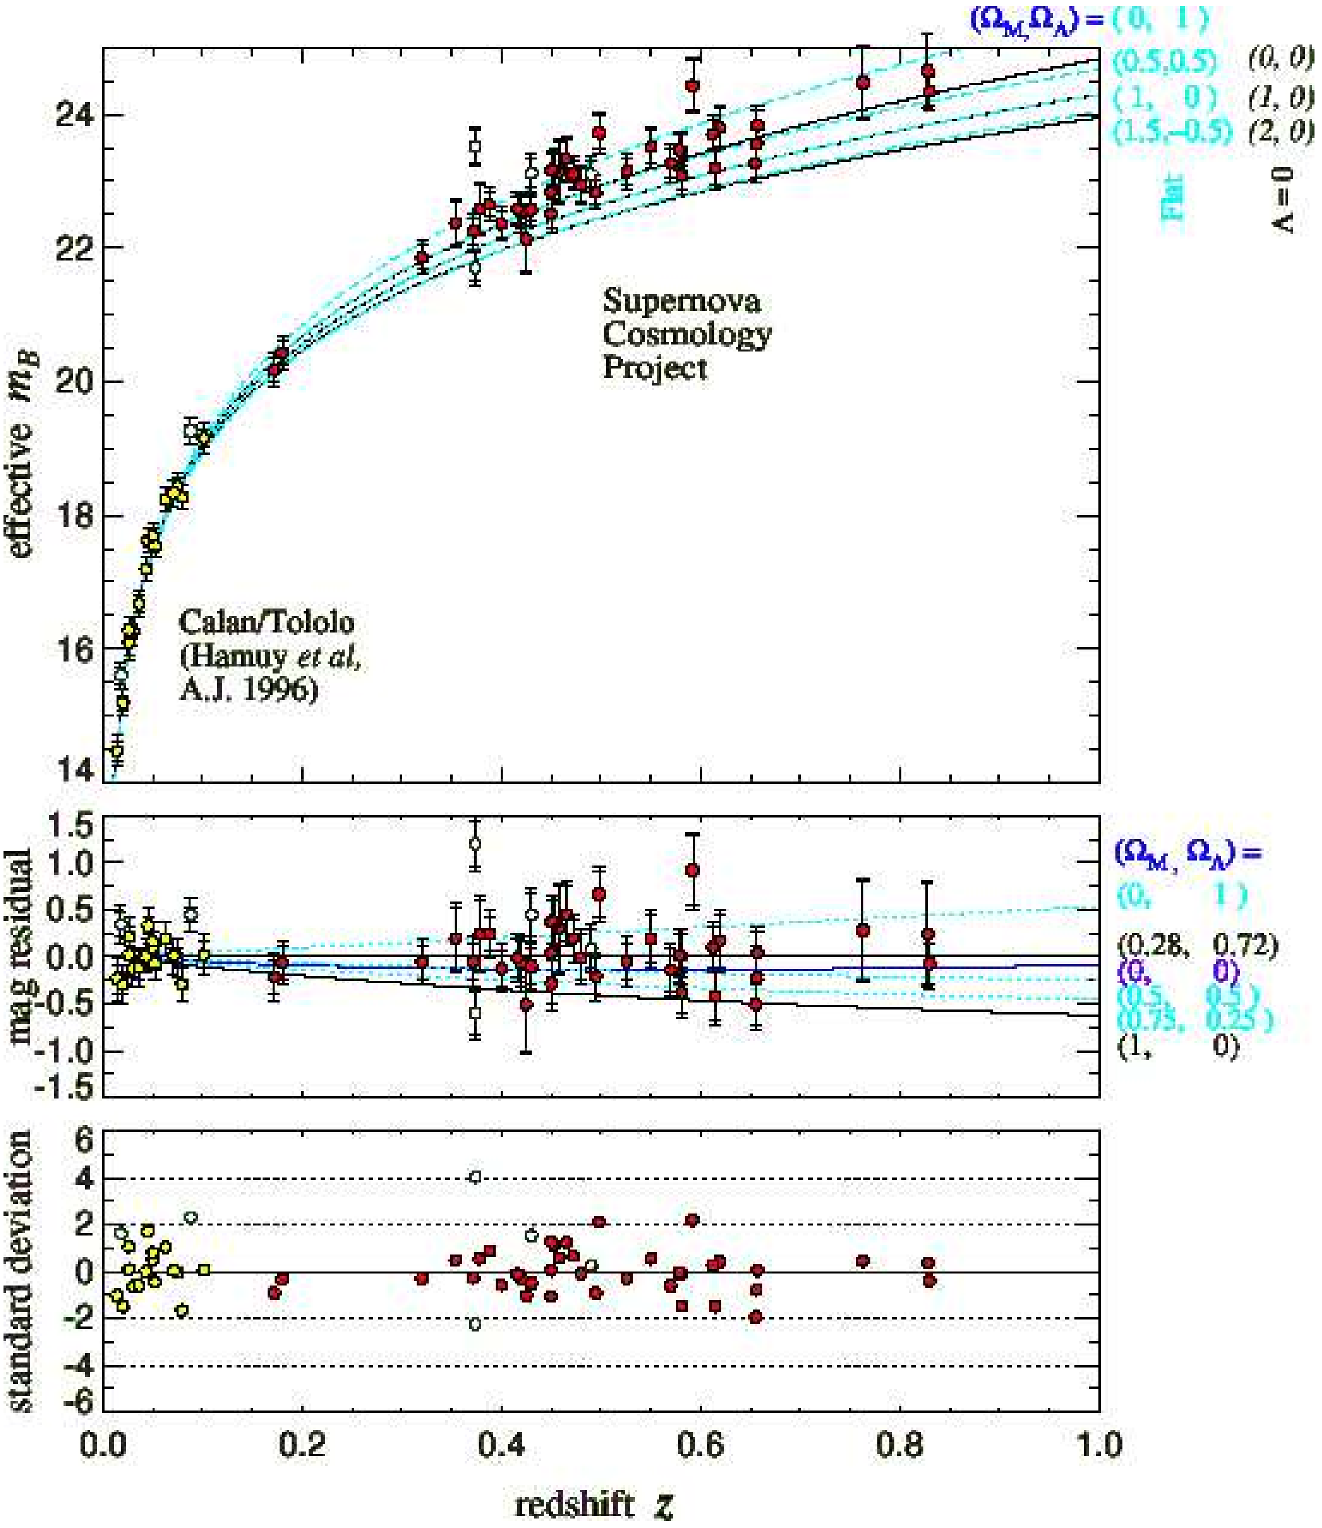

Fig. 1.— Magnitude vs. redshift of Type Ia supernovae. From the Supernova Cosmology Project.

work in the right direction, because more distant supernovae have to traverse a greater distance, and hence a larger amount of stuff, than nearer supernovae. **Ask class:** how would we test this?

The key is to realize that in the optical, where the observations are made, scattering is strongly wavelength dependent. The result is that if a lot of flux is lost to scattering, one also expects that the source should look redder than it does normally (for the same reason that sunsets are red; blue light is scattered out, but red light is not scattered as much). Therefore, an easy way to check for this effect is to see whether distant, dim supernovae look redder than close, bright white dwarf supernovae. The answer: no. In the spirit of checking all possibilities, we should note that if the "dust" is sufficiently large then the scattering is wavelength-independent for optical light. Such "grey dust", however, has not been seen, so it is rather ad hoc and not believed.

Another possibility is that white dwarf supernovae at redshifts of $z > 0.5$ genuinely were dimmer than they are now. This is something that must be taken seriously, for the simple reason that our theoretical understanding of white dwarf supernovae is relatively primitive (in contrast to, for example, Cepheid variables). **Ask class:** what might we do about this?

In general, we cannot test all hypotheses of this type, because maybe there is some unknown effect. What we can do, however, is make a specific guess at why supernovae at high redshifts are different from those now, and test that guess. One that seems physically reasonable is that the luminosity vs. decay rate relation somehow depends on the fraction of elements heavier than helium (aka the "metallicity"; astronomers are weird in that they call all elements heavier than helium "metals"!). For example, it could be that this affects the nuclear energy release or the accretion properties just enough to make a difference. One would then expect systematics with redshift, because the farther back in time we go, the less time the universe has had to synthesize metals in stars.

For this specific hypothesis, though, we can apply local tests. The reason is that even among local galaxies the metallicity varies all over the place, and in particular there are plenty of galaxies with metallicities as low as those at $z > 1$. Systematics would then be expected locally as well as with redshift, but they aren't seen. Whatever the details of white dwarf supernovae are, they appear robust against metallicity.

Another future test would be to look at white dwarf supernovae in a wider range of redshifts, say up to $z = 2$ or so. The point is that the flux versus redshift is predicted to have a specific form if it is cosmologically driven, but if there is some intrinsic effect it would almost certainly have a different form. Therefore, the better the statistics the greater our confidence.

Yet another possibility has to do with gravitational lensing. Light is deflected by mass, so the light from a white dwarf supernova takes some (minor!) turns as it wends its way to us. If we are in some favored location, then the flux we see can be amplified compared to what it would be with a smooth distribution of matter. However, lensing doesn't create light (it merely moves it around), so for every location with enhanced flux there are several with slightly diminished flux. At larger distances there are more chances to pass by lumps of matter, so this effect might be enhanced with increased redshift. **Ask class:** how might we test this out?

This is another effect that gets settled by statistics. Each new white dwarf supernova is an independent measurement, and we can therefore test whether the distribution of fluxes is as predicted from (1) a standard universe with gravitational lensing, or if it is better predicted by (2) a universe with accelerating expansion. As more observations have been made, it has become clear that (2) is the better option.

The net result is that although the supernova observers have been admirably open-minded to other effects, accelerating expansion has come through as the best interpretation. Cosmologists have therefore been forced to accept this bizarre effect as reality, and to try to explain it. In the next lecture we will also accept this, and talk about one form of dark energy: the cosmological constant.

Intuition Builder

Black holes absorb light independently of wavelength, so there would be dimming but no reddening involved. Can you rule out the possibility that black holes are the cause of the dimming? Suppose you are restricted to current data rather than being able to appeal to future high-redshift data sets.