Supernovae, Dark Energy, and the Accelerating Universe: The Status of the Cosmological Parameters

Saul Perlmutter
Supernova Cosmology Project
Lawrence Berkeley National Laboratory
Berkeley, CA 94720

1 Introduction

I was asked to present the status of the cosmological parameters, and in particular the status of the recent results concerning the accelerating universe—and the possible cosmological constant or dark energy that is responsible for the universe's acceleration. This result comes most directly from the recent type Ia supernova work, so although I will mention a few of the approaches to the cosmological parameters, I will emphasize the work with the type Ia supernovae. I will try to give you a sense of exactly how we reached the current conclusions and what the current level of confidence is in that conclusion.

Because this presentation will be emphasizing the supernova work, I would like to highlight the strong team of scientists in our Supernova Cosmology Project [1, 2]; you will probably recognize some of the particle physics heritage of this particular team. I also would like to list the names of the members of the other supernova team, led by Brian Schmidt of the Australian National University, that has been working in this field [3]. These two groups together comprise a good fraction of the entire community of scientists working on supernovae.

2 Supernovae as a simple, direct cosmological measurement tool

The particular advantage of the type Ia supernova approach to measuring the cosmological parameters is the fact that you can explain essentially the entire method in one graph. The basic idea is that you want to find an object of known brightness, a "standard candle," and then plot it on the astronomer's Hubble diagram (Fig. 1), which is a plot of brightness (magnitude) against redshift. We should interpret this graph as follows: for an object of known brightness, the fainter the object the farther away it is and the further back in time you are looking, so you can treat the y-axis as the time axis. The x-axis, the redshift, is a very direct measurement of the relative expansion of the universe, because as the universe expands the wavelengths of the photons travelling to us stretch exactly proportionately—and that *is* the redshift. Thus the Hubble diagram is showing you the "stretching" of the universe as a function of time. As you look farther and farther away, and further back in time, you can find the deviations in the expansion rate that are caused by the cosmological parameters.

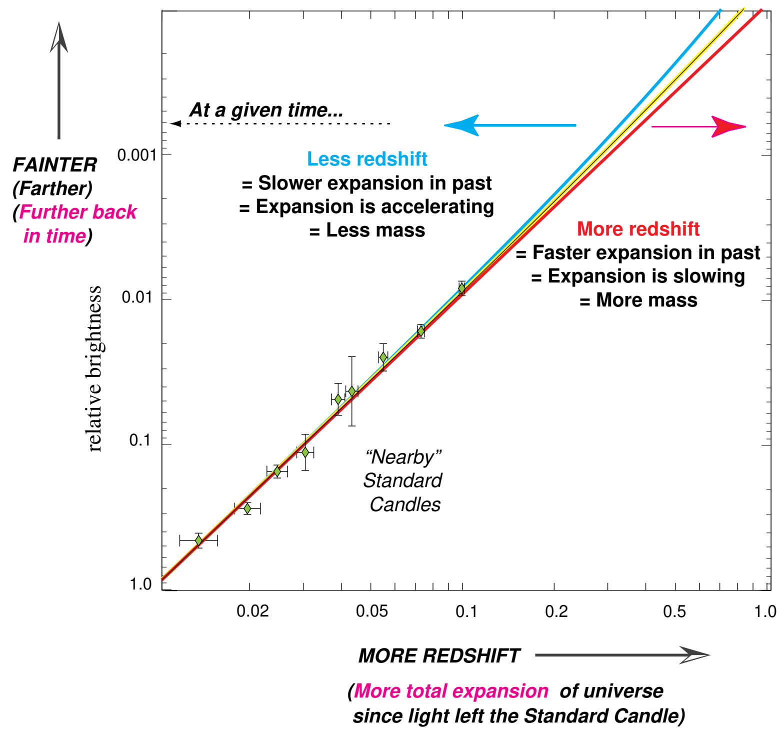

Figure 1: The Hubble plot: A history of the "size" of the Universe.

In particular, you can think of making a measurement of a supernova explosion at one given time in history. If, for example, you found more redshift at that time than expected from the current expansion rate, that would imply that the expansion was faster in the past and has been slowing down. This would lead you to conclude that there was a higher mass density in the universe.

3 Type Ia supernovae as 'difficult' standard candles: A strategy to make them manageable

For a standard candle, we use the type Ia supernova, the brightest of the supernova types. They brighten in just a few weeks and fade away within a few months, so it is necessary to explain what I mean here when I refer to this standard candle's brightness. Generally, we will use the magnitude at peak, which turns out to be a very consistent brightness (after a bit of calibration, as discussed below). These are the brightest of the supernova explosion events by about a factor of six, so at high redshift most of the supernovae that we find are type Ia.

Although such a bright, standard candle should make an excellent cosmological measurement tool, the problem with using the supernovae is that they turn out to be a real "pain in the neck" for any kind of research work: one can never predict a supernova explosion and supernovae only explode a couple of times per millennium in any given galaxy. This makes it very difficult to apply for the world's largest telescopes, which are necessary to observe the most distant supernovae that exploded far back in time. Proposals for telescope time must be written six months in advance, and of course it is necessary to guarantee that there will be something to observe. The first steps of this project, then, involved developing the strategies to make it possible to guarantee the discovery of the supernovae.

We came up with what is essentially a "batch" approach (see Fig. 2) [4]. We observe a number of wide fields of apparently empty sky, out of the plane of our Galaxy. If you open the shutter long enough, any patch of sky will have hundreds of distant galaxies in it. We observe tens of thousands of galaxies, in a few patches of sky, and then come back in three weeks to observe the same galaxies over again. In these tens of thousands of galaxies there will be a dozen to two dozen supernova explosions that were not there three weeks earlier. This is a new batch of supernovae that are ready to be followed up, and the photometry and spectroscopy observations can now be *scheduled in advance* to follow the supernovae as they brighten to peak and fade away. With this three-week time baseline, the supernovae generally do not have time to reach peak brightness, so almost all of the discoveries are pre-maximum. This guarantees that you can track the supernovae over maximum with scheduled follow-up observations. This strategy thus turns a rare, random event into something that can be studied in a systematic way.

To make this a little more concrete, you can see what the actual data looks like in Fig. 3. All of the faintest specs and smudges in this picture are the distant galaxies that we are studying. We have to find a new spot of light among these tens of thousands of smudges that was not there three weeks before. As shown in the blow-up figure, these are very minor-looking events, even though they are among the most dramatic, energetic events in the universe. With computer image analysis, however,

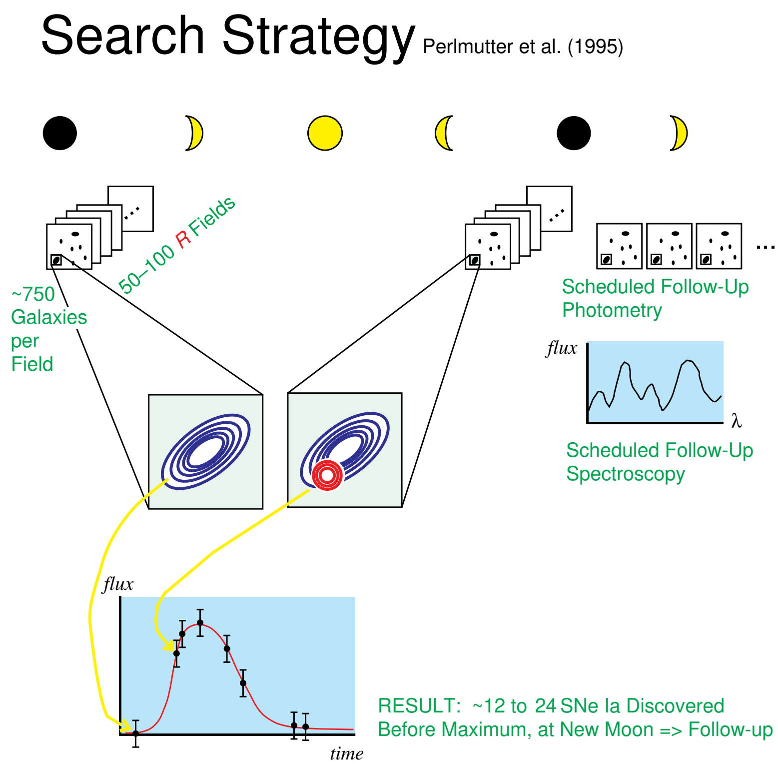

Figure 2: Search strategy to discover "batches" of supernovae in a scheduled, systematic procedure [4].

we are able to find these events and trigger the sequence of follow up observations at telescopes around the world—and also at the Hubble Space Telescope. (To win time on the Hubble Space Telescope, not only do you have to predict the date of a new supernova discovery, but you also have to tell in advance which square degree of the sky will have the supernova in it.)

This strategy turns out a batch of supernova every time we go out to the telescopes. We do this typically once per semester, and study one or two dozen events. By now, we have built up a sample at these very high redshifts of more than 80 supernovae. In the redshift distribution histogram (Fig. 4) the color coding indicates the search semester. You can see that we have found more and more each semester and worked our way out to higher and higher redshifts. The most recent supernovae have all been followed with the Hubble Space Telescope.

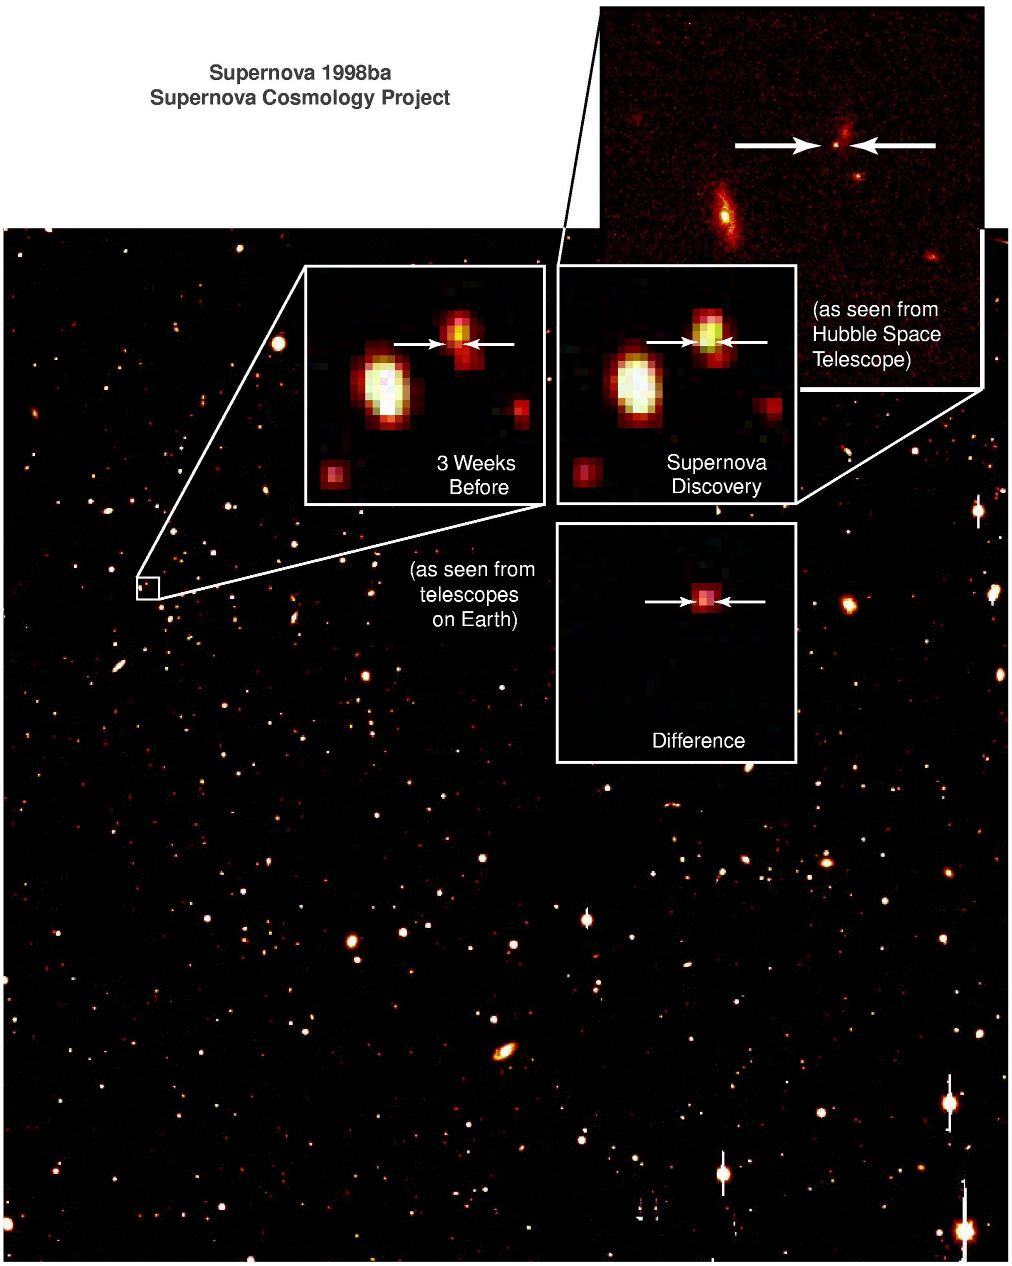

Figure 3: Supernova 1998ba, an example of a supernova discovery using the "batch approach."

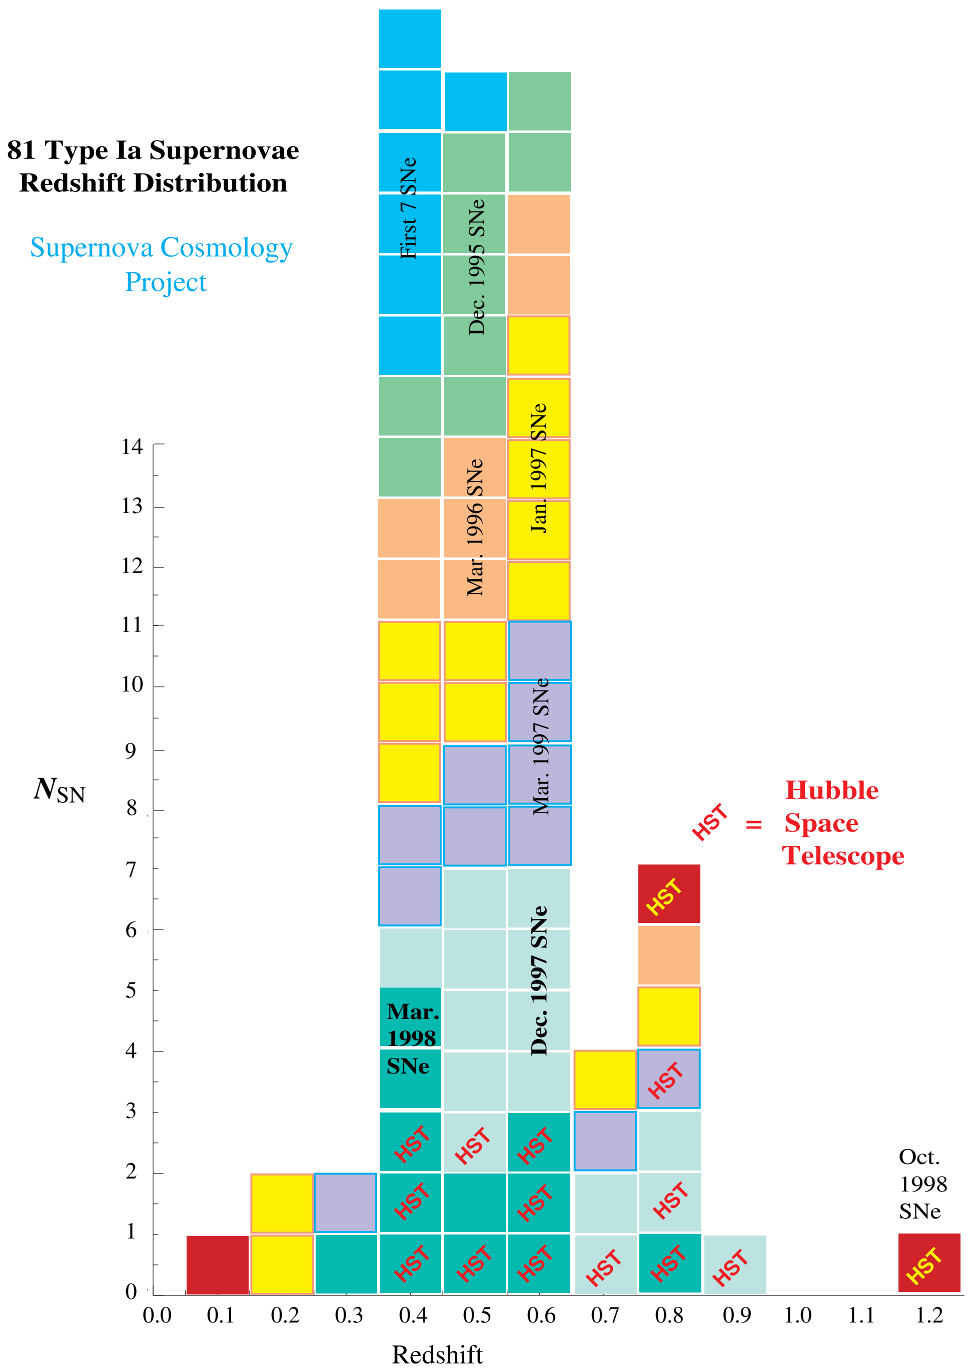

Figure 4: The redshift distribution of the first 81 high-redshift supernovae discovered by the Supernova Cosmology Project. The label “HST” indicates supernovae that were followed with Hubble Space Telescope photometry.

4 The (only) two other analysis steps needed

Before this high-redshift supernova data can be plotted on the Hubble diagram and the cosmological parameters fitted, there are two small additional analysis steps necessary in order to compare the distant supernovae to the nearby supernovae on the same Hubble plot. First of all, although most type Ia supernovae follow a very similar light curve, there are a few outliers that are a little bit brighter or a little bit fainter. In the early 1990s, it was pointed out by Mark Phillips at Cerro Tololo Interamerican Observatory in Chile that there is an easy way to distinguish these supernovae, and recognize the slightly brighter ones and slightly fainter ones, using the timescale of the events. Phillips [5] noted that the decline rate in the first 15 days after maximum provides a good parameterization of the timescale, and that this is a good predictor of how bright the supernova will be. Later, Riess, Press, and Kirshner [6] showed another elegant statistical method which effectively added and subtracted shoulders on the light curve to achieve the same sort of timescale characterization. Finally, our group developed a third method [4, 7], which we call the timescale stretch factor method, in which we simply stretch or contract the timescale of the event by a linear stretch factor, s. This also predicts very nicely the brightness of the supernova: The $s > 1$ supernovae are the brighter ones and the $s < 1$ supernovae are the fainter ones.

We can now put together the whole range of type Ia supernovae light curves on a single plot. The upper panel of Fig. 5 (from Kim, et al., in preparation) shows a sample of relatively nearby supernovae from the Calan Tololo survey [8], for which we can use the redshifts to give us the relative distances. Their relative brightnesses can then be compared, after adjusting for the different distances. Most of the supernovae follow the typical $s = 1$ light curve on this graph, but there are some brighter ones and some fainter ones. We fit the stretch of the light curve time scale, use this to predict the supernova luminosity, and then normalize the supernova's light curve to the standard $s = 1$ luminosity. After also accounting for small differences in supernova reddening (as discussed below), this calibration procedure results in the remarkably tight distribution of light curves shown in the lower panel of Fig. 5. The dispersion at peak is approximately 10 to 12 percent, which makes this one of the most impressive standard candles available in astronomy.

There is one other piece of information that you need to know in order to make a comparison between the low redshift supernovae and the high redshift supernovae. We observe the low redshift supernovae with a blue "B band" filter that captures the peak of their spectrum (see Fig. 6). At high redshifts, we wish to observe the same part of the spectrum, so we have to use the red "R band" filter. However you'll notice that slightly different parts of the spectrum come through the B filter at low redshift and the R filter at high redshift, and you have to make a small correction to account for that difference (see [9]).

With this so-called "K-correction", together with the stretch correction on the

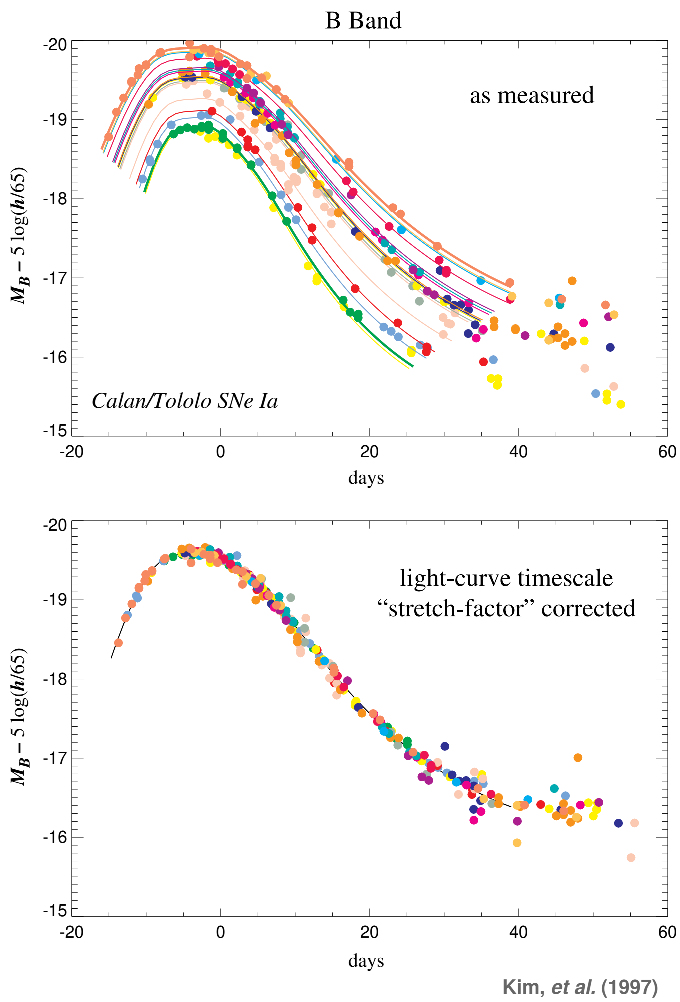

Figure 5: *Upper panel:* The range of lightcurve for low-redshift supernovae discovered by the Calan/Tololo Supernova Survey. At these redshifts, the relative distances can be determined (from redshift), so their relative brightnesses are known. *Lower panel:* The same lightcurves after calibrating the supernova brightness using the "stretch" of the timescale of the lightcurve as an indicator of brightness (and the color at peak as an indicator of dust absorption) [11].

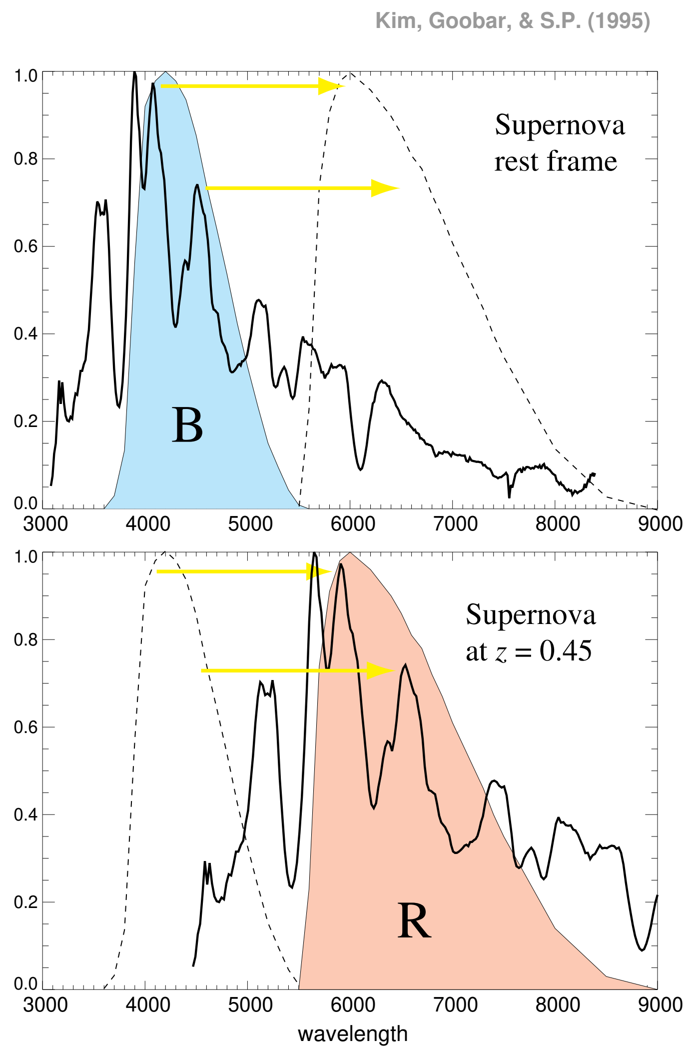

Figure 6: Slightly different parts of the supernova spectrum are observed through the "B filter" transmission function at low redshift (*upper panel*) and through the "R filter" transmission function at high redshift (*lower panel*). This small difference is accounted for by the "cross-filter K-correction" [9].

supernova luminosity, the low and high redshift supernova light curves can now be compared with each other on the same diagram. One striking cosmological effect is immediately apparent: events at a high redshift, $z \sim 0.5$, last 1.5 times longer than events at low redshift [10]. This is one of the most dramatic examples of a macroscopic time dilation that you will get to see. If you take out that $(1+z)$ time dilation, and also remove the small variations in the stretch factor, the low redshift and high redshift composite light curves now lay right on top of each other. This shows that the supernovae are very similar across redshifts and that the K-correction does an excellent job in bringing them in line with each other.

5 Cosmological results from the Hubble plot

We can now plot the low and high redshift supernovae together on the Hubble diagram and look to see which curves—representing different values of the cosmological parameters—fit best. If we find more redshift in the past this would imply that the expansion has been slowing down and hence there is more mass in the universe to slow it down. However, we now know that there are other possible cosmological parameters that can work in the other direction; for example, a vacuum energy density can make the universe expand faster. Thus there is a degeneracy here; it is hard to tell apart a situation with more mass or less vacuum energy density.

In 1995, we presented an approach to this problem [12]. We pointed out that if you have a hypothetical supernova at a redshift of 0.5 (at that time, a supernova discovery at this redshift was still hypothetical), then a measurement of its magnitude, with its associated error bar, selects out a strip on the Ω_M–Ω_Λ plane. This shows the degeneracy, because the strip can include high Ω_M and high Ω_Λ or low Ω_M and low Ω_Λ. However, because the mass density and the cosmological constant enter into the equation for apparent magnitude with different powers of redshift, this strip will rotate as you use supernovae at higher redshifts. If we can find a supernova, for example, at a redshift of 1, we will get a rotated strip on the Ω_M–Ω_Λ plane and we can look for the overlap region with the strip from the z = 0.5 supernova (see Fig. 7). This allows us to read off the values of the two parameters separately.

We first demonstrated how this could work in a January 1998 Nature article [13]. We compared the confidence-region strip from five supernovae at a redshift of about 0.4 with one from a supernova at a redshift of 0.83 and showed such an overlap region. Within a week after this article appeared, we were ready to go ahead and put an additional batch of almost 40 new supernovae onto the Hubble diagram [14]. With this much more dense coverage on the Hubble diagram, we found a much more tightly constrained confidence region on the Ω_M–Ω_Λ plane (see Fig. 8) [15].

Because this confidence region plot is key to much of the rest of the discussion, a few comments on its interpretation might be helpful here. First of all, almost anyone

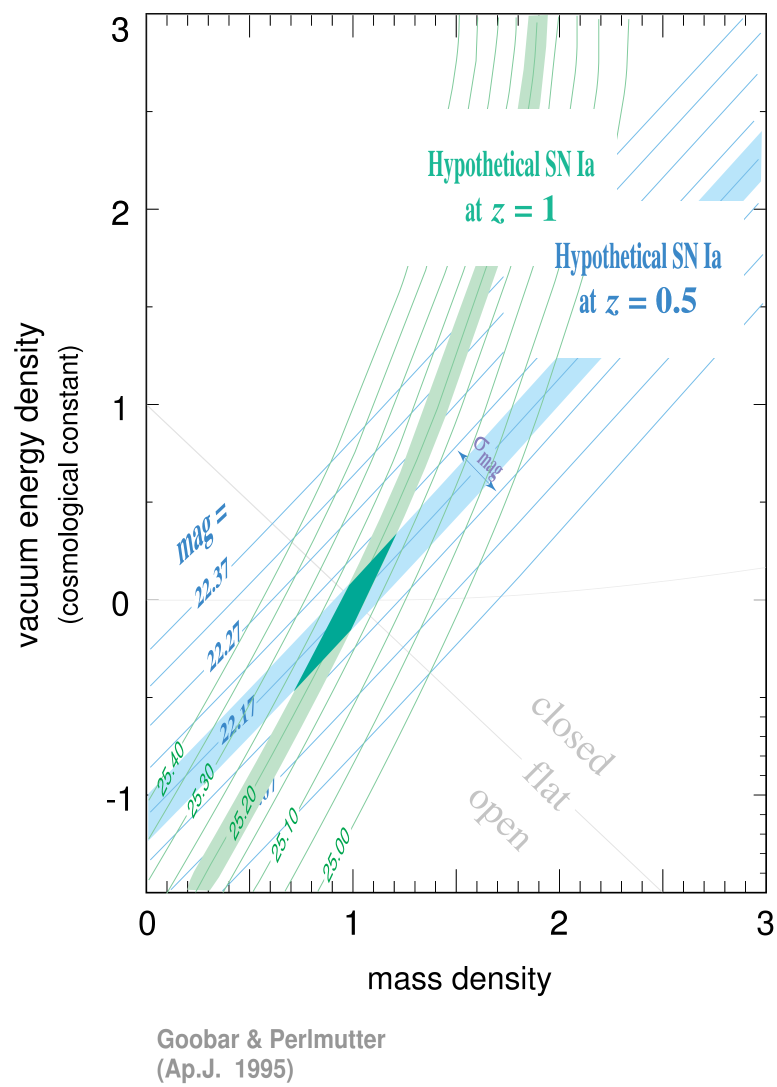

Figure 7: The approach to separating the apparently degenerate measurements of mass density and vacuum energy density by studying supernovae at different redshifts [12]. This plot is for the hypothetical (and, we now believe, counterfactual) case of a flat universe with no vacuum energy.

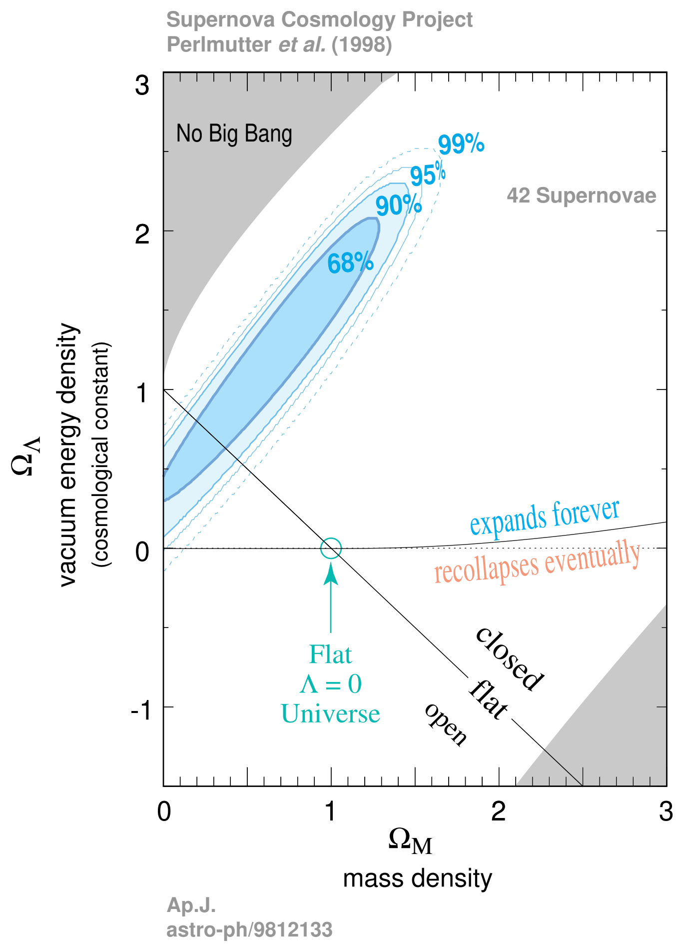

Figure 8: Confidence region on the Ω_M–Ω_Λ plane [15] based on data from 42 supernovae at redshift $z \approx 0.35 - 0.85$ (Supernova Cosmology Project) compared with 18 supernovae at low redshift (Calan/Tololo Supernova Survey).

who has taken a course in cosmology in the previous 25 years will have gotten used to the idea that the curvature of the universe determines its destiny. Thus, a universe that is spatially closed, *i.e.*, one that curves in on itself, will eventually slow to a halt in its expansion and then collapse again; so it will come to an end. On the other hand, a universe that is either flat or curved open will expand forever. This tight relationship between curvature and destiny, however, is only true if we ignore the cosmological constant.

With the cosmological constant in the story, all four scenarios are possible: for example, we could have a closed universe that expands forever or an open universe that recollapses eventually. The confidence region for the 42 supernovae already addresses the question of the universe's fate. It appears from this dataset that the universe will expand forever and, moreover, the universe appears to be accelerating in its expansion. These data do not, however, tell us anything about the curvature of the universe. The long thin confidence region extends on either side of the flat universe line, which divides closed from open curvatures. The other important conclusion that you can draw from the 42-supernova confidence region is that the data are very far away from being consistent with the simplest cosmology, one which is flat and has zero cosmological constant (indicated by a circle in Fig. 8).

This result can also be interpreted as a measurement of the age of the universe, if you know the current expansion rate (*i.e.*, the current Hubble constant). This is shown on the plot with isochrons of age for a given mass density and cosmological constant (see Fig. 9). The supernova confidence region picks off a value of about 14.5 billion years, or 15 billion years along the flat universe line. These values are for a Hubble constant of 65 km/s/Mpc; if we had chosen a Hubble constant that was 10 percent higher we would have found an age that was 10 percent lower. In either case, there no longer appears to be an "age crisis" in which the oldest stars seemed to be older than the age of universe.

6 Comparison with other results in the field

We were soon able to compare this Ω_M–Ω_Λ confidence region with results presented by the group led by Brian Schmidt [16, 17, 18]. Their results from 16 high redshift supernovae matched beautifully with our results. Thus, two independent analyses based on mostly independent supernovae reached the same conclusion.

These results can also be compared with those from other methods for measuring the cosmological parameters. In particular, we can ask to what extent do we know that we live along the flat universe line, because our measurement does not constrain that very well. The cosmic microwave background, the leftover glow from the very dense period at the beginning of the big bang, is a very good indicator of how curved the universe is. We are beginning to see CMB data coming in that is starting to

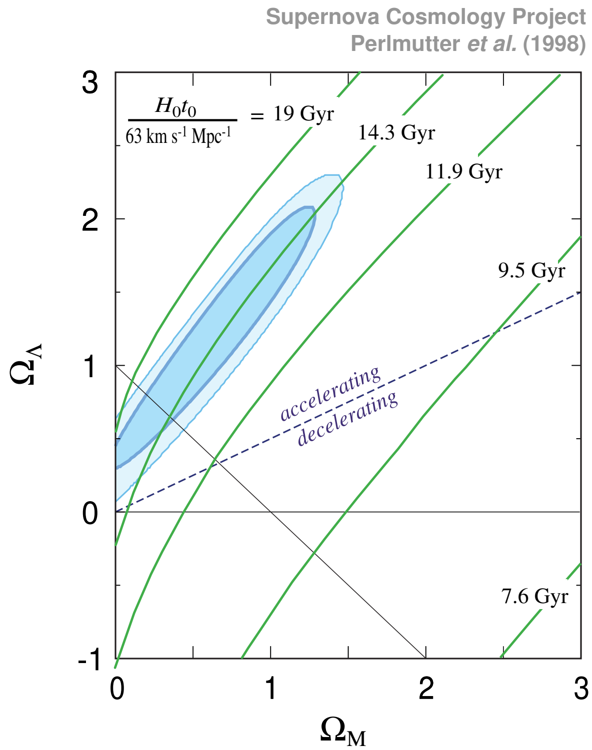

Figure 9: Isochrons of the universe's age (from [15]) overplotted on the confidence region of Figure 8.

constrain the curvature. Although much better data should be available within the next few years, we can already begin to rule out the upper right ("over-closed") and lower left (very "open") regions of the Ω_M–Ω_Λ plane. This has been taken by some as suggestive that we will find the answer to be a flat universe. If you put the CMB data together with the supernova data, you find a result that centers quite close to the flat universe line, with mass density approximately 0.3 and vacuum energy density approximately 0.7.

Finally, we can compare these results with the one other main source of information on the cosmological parameters: the dynamics of clusters of galaxies and the evolution of clusters, which are both sensitive indicators of the mass density of universe. Currently, the cluster data also indicates a mass density of around 1/3, consistent with the supernova and CMB results. We thus find ourselves at a historical moment in which the three main approaches to the cosmological parameters are all in good agreement with a cosmological constant of 2/3 and a mass density of 1/3. (For a full review of these current results, see [19].) Perhaps when other measurement techniques are developed, the situation will become more complicated, but so far we appear to have a rather simple-looking scenario—with the only complication being the existence of the cosmological constant.

7 What level of confidence can be placed on this result?

Given the surprising nature of this result, it is important to ask how strongly it can be believed. First of all, if you believe that we have evidence for a flat universe from the cosmic microwave background data—or if you like the inflationary universe model which predicts a flat universe—then the supernova results are very strong. They are very many standard deviations away from a flat, $\Lambda = 0$ universe. There would have to be a systematic error as large as almost 50 percent of the brightness of the supernovae for there to be no cosmological constant, if the universe is flat. This is highly unlikely.

However, if you prefer to take a more agnostic view and consider the flatness of the universe to be still unproven, then we must ask if it is possible that we live off of the flat universe line in a universe with low mass density and no cosmological constant. This case is still ruled out at the 99 percent level statistically, but now it would only take a 15 to 20 percent systematic error to make the data consistent with this open universe. We have spent most of our effort in the last year or so looking for such systematics. Possible candidates include absorption of light by dust, supernova evolution, selection bias in the supernova detection, and gravitational lensing effects. We will focus here on the first two of these.

Because the supernova light can be dimmed by dust, the concern is that the high redshift supernova look fainter not because of a positive cosmological constant, but rather because their light has traveled through dust. Fortunately, there is a straightforward way to measure this. All dust that has so far been observed in the universe absorbs blue light more than red light, so sources that are seen through dust appear redder. We can compare the colors of the nearby supernovae and the distant supernovae to see if the latter are in fact redder. We find no statistical difference, and a full statistical analysis indicates that ordinary dust that reddens cannot account for our supernova data without a cosmological constant.

Is it possible that supernovae have evolved between the high redshift epoch and the low redshift epoch? In other words, perhaps the supernovae that we see today are older than the high redshift supernovae, and perhaps older supernovae are intrinsically brighter. There is a good way to look for such physics changes in a supernova event, which is to look at the time series of supernova spectra. As a supernova brightens and fades away, different features in the spectrum appear or disappear. A spectrum thus provides a rather distinctive signature indicating the day along the light curve. This spectral signature tells us about the physical state (elements' abundances, velocities, and temperatures) of a part of the supernova atmosphere from which the light was emitted. As the supernova explosion proceeds, the atmosphere expands and the outer layers become more transparent, allowing us to see further and further in to the core. Over time, then, the spectrum samples the physical state of the supernova from the outer shells in towards the inner core.

We can thus compare the physical states of nearby and distant supernovae by comparing their spectra at matching days along the light curve. One excellent example is a supernova that we discovered at a redshift of 0.36. We followed this supernova with photometry over its light curve and obtained a spectrum just before maximum, while the other supernova group collaborated with us by obtaining spectra on two different dates after maximum. All three spectra are remarkable matches to spectra on the corresponding day (in the supernova rest frame) for a low redshift type Ia supernova. In other words, as we scan into the core of the supernovae, the physics appears to be the same at low and high redshift. We have carried this sort of comparison out to redshifts as high as 0.83 [13] and found that the supernova spectrum falls right where it should in the time series of appropriately redshifted nearby supernova spectra (see Fig. 10).

There are other approaches to comparing the low redshift and high redshift supernovae. For example, the excellent match between the low and high redshift light curves' shape also suggests that the physics has not changed over time. A recent suggestion that there may be a difference between low and high redshift "rise time" between explosion and peak [20] was not born out by more extensive examination of the data [21].

Our overall "scorecard" of uncertainties (see [15]) now looks something like this: the statistical uncertainties are small enough that the presence of a cosmological constant is highly significant. Most of the systematic errors can be constrained to stay well below these statistical uncertainties. We have also performed several additional crosschecks of effects that could affect the results. For example, we have checked the sensitivity of the results to the exact form of the width-luminosity relation that we use to calibrate the brightness of the supernovae. We find that this relation makes very little difference in the results, because almost all of the supernovae have approximately the same width light curve at both low and high redshifts.

7.1 Plugging the remaining loopholes

However, there are two "loopholes" among the proposed systematic uncertainties, which are not yet as tightly constrained. These could conceivably allow the supernova data to agree with a $\Lambda = 0$, low mass universe (although probably not with a flat $\Lambda = 0$ universe). The first of these is the possibility that there is a new kind of dust, a dust that is gray, which does not absorb blue light more than red. Such dust would have escaped detection in our hunts for reddening of the supernova light. The other loophole is the possibility that the beautiful match of low and high redshift light

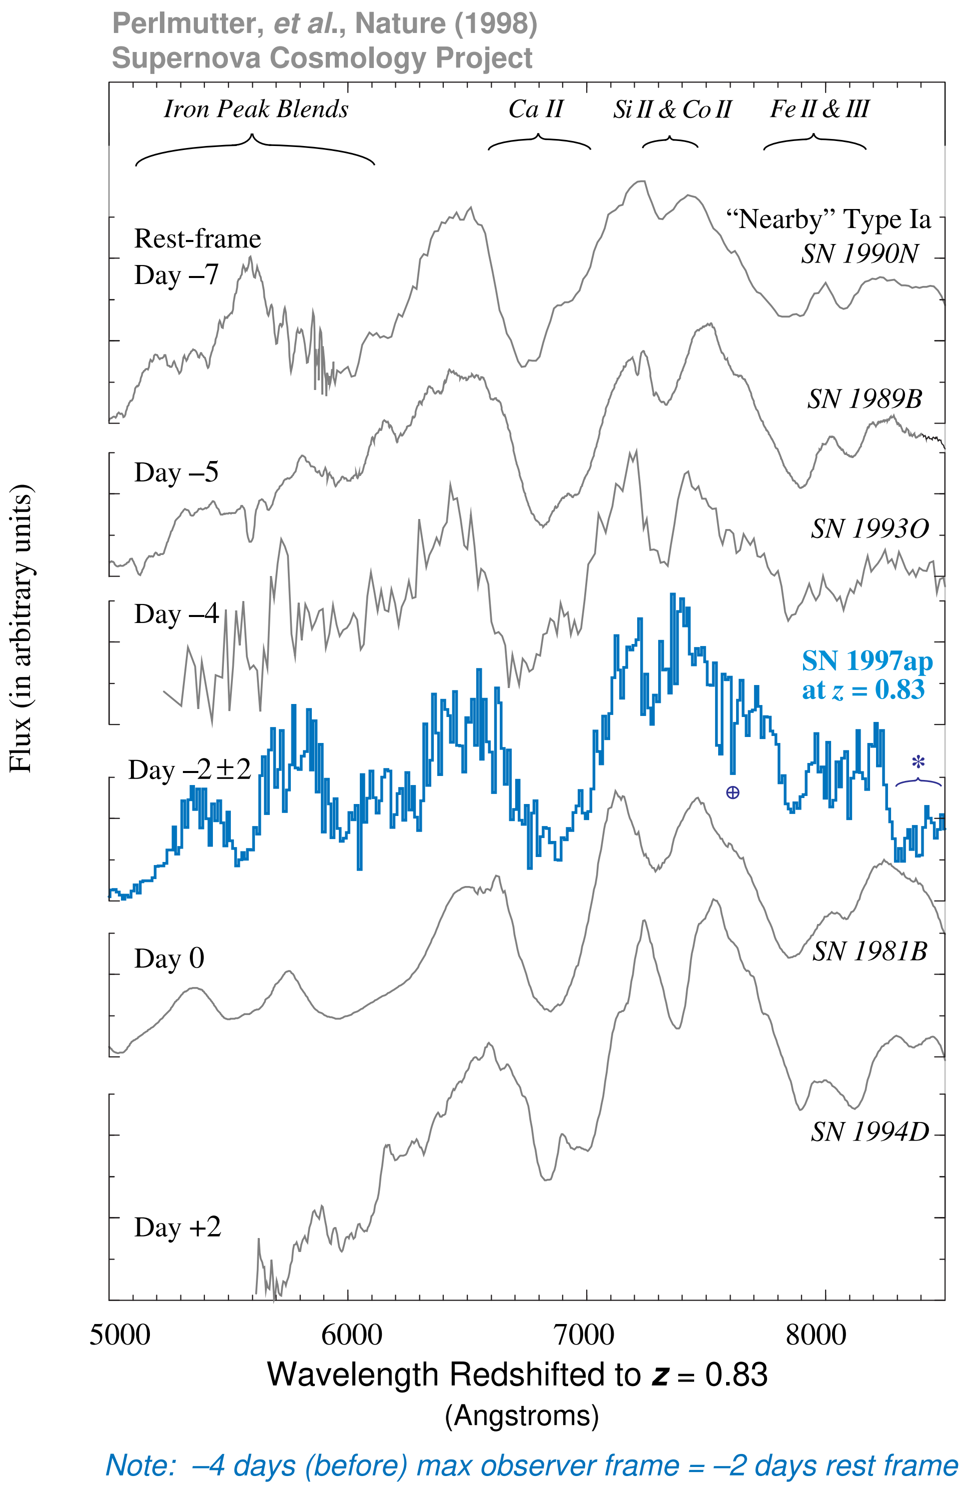

Figure 10: Supernova 1997ap at $z = 0.83$, set in the time sequence of nearby supernova spectra (from [13]).

curves and spectra that we have seen so far is a statistical fluke due to low number statistics. Perhaps more complete samples will show physical differences. Fortunately, we do have empirical approaches to allow us to study both of these loopholes; these studies are now in progress.

To understand how the studies deal with the questions of supernova evolution, it is important to be clear that evolution is not assumed to be a monotonic function of the age of the universe; in other words, we do not expect that supernovae are uniformly fainter as you study them at higher redshifts. Rather, the main concern is that the typical environment in which a supernova explodes may on average be a little bit different at high redshift from that at low redshift. For example, a host galaxy that has undergone many generations of star formation will have built up a higher density of the heavy elements (the astronomers call this "metallicity"), and one might imagine that this might lead to a supernova explosion of differing brightness. The key point, however, is that different galaxies have begun their life at different times in history, so at any given redshift there will be a wide distribution of galaxy ages, and hence metallicities. The demographics may shift as we go back in redshift, such that the peak of the host-galaxy age distribution becomes a little bit younger, but there are still examples of both young and old host-galaxy environments even in the nearby supernova population.

We therefore can study supernova evolution effects simply by looking at nearby supernovae across a wide range of host galaxy ages. For the relatively small samples that have already been studied, the light curve width-luminosity relation appears to account for any evolutionary differences quite well, as we have seen. However, we would like to be able to examine hundreds of nearby supernovae to find even small departures from, and refinements of, this calibration relation.

To this end, we have been working on new methods to find and study nearby supernovae with the same systematic predictability that our search approach made possible at high redshifts. The problem here is that one has to be able to observe and analyze much larger areas of sky at low redshift to scan the same number of galaxies as in our high redshift searches. Working with several other groups, we ran a pilot study last spring in which we analyzed tens of gigabytes of data per night. We succeeded in finding 40 supernovae in one month, 29 of which were type Ia [22]. (This is comparable to the number of supernovae discovered in three years of searching with previous photographic plate searches.) We are now scaling up this project and intend to find several hundred low redshift type Ia supernovae. The quality of the follow-up is as important as the quantity; each supernova will be studied with a time series of flux-calibrated spectra over the first few months of its light curves.

At the other extreme, we have also been developing techniques to find and study supernovae at significantly higher redshifts. In a pilot study for this project, using the Keck 10 meter telescope, we were able to find a supernova at $z = 1.2$, twice the typical redshifts of the earlier searches [23]. The spectrum of this supernova, nicknamed "Albinoni", matched a type Ia supernova a few days before its light curve peak. We followed Albinoni with photometry in the optical and infrared at the Hubble Space Telescope. Now that Albinoni has mostly faded away, we will soon be able to complete its analysis and place a new data point on the Hubble diagram.

Supernovae at such higher redshifts are useful in addressing both loopholes, evolution and gray dust. This is because the curve on the Hubble diagram that is predicted for a cosmology with a positive cosmological constant is fainter at $z = 0.5$ than the curve for a universe with no cosmological constant, but at much higher redshifts the two curves come back together and cross. This behavior is very difficult to mimic with an evolutionary effect. Similarly, any hypothetical gray dust that would be sufficient to dim the supernovae at $z = 0.5$, making it falsely appear that there is a positive cosmological constant, should also dim supernovae at Albinoni's redshift by almost 30 percent relative to a cosmology with a real cosmological constant but no gray dust. Of course, it will require more than one supernova at Albinoni's redshift to answer these loophole questions with statistical and systematic significance, and so the very-high-redshift searching is ongoing.

8 New CCD detector technology for high redshifts

To work at these very high redshifts, we have been developing new detector technology. In the course of developing detectors for the SSC, Stephen Holland at Lawrence Berkeley National Laboratory came up with a high-resistivity silicon process that we have used to fabricate astronomical CCDs [24]. These new CCDs have much higher quantum efficiency in the wavelengths between 0.8 and 1 micron, where the very-high-redshift supernovae are brightest. These new devices do not require the highly thinned silicon that make conventional CCDs fragile, expensive to process, and subject to interference-pattern "fringing"—which degrades spectra and images in the far red wavelengths. Our new CCDs will thus greatly improve our ability to study supernovae like Albinoni.

9 "Dark energy": What's wrong with a Λ?

Why do we find a positive cosmological constant disturbing? What is wrong with a vacuum energy density? First of all, what's surprising is that the value of the cosmological constant we are finding is so small. One might have expected from the zero-point energies of all the particles and fields that a vacuum energy density would show up around the Planck scale, or at least the supersymmetry breaking scale. These guesses would have been off by more than 120 orders of magnitude or more than 50 orders of magnitude respectively. It had therefore been assumed for many years that there must be some perfect cancellation that makes the vacuum energy density exactly zero. Now, presumably, the cancellations must work to better than a part in 10^{50} but then still leave a tiny vacuum energy density that we see accelerating the universe's expansion. This seems an awkward explanation.

There is also a problem with the cosmological constant that might be termed the "Why now?" problem. The mass density of the universe has been dropping by many orders of magnitude over the life of the universe as the universe has expanded; meanwhile the vacuum energy density has stayed constant. Why should the two densities be within a factor of two or three of each other today? This, too, is an uncomfortable coincidence.

Several new fundamental physics theories have been proposed to try to avoid these two problems with the cosmological constant. These alternative "dark energy" theories have been characterized by their different equations of state, and in particular the ratio between pressure and density (in appropriately dimensionless units) in this equation. We can measure this ratio, w, as a function of the mass density of the universe using the supernova data. We find a range of possible values from $w = -0.6$ to -1 [15, 25]. If we add some constraints from other astrophysical data, suggesting that the mass density is around 0.3, then the confidence region favors values nearer to $w = -1$, which would be the equation-of-state ratio for the cosmological constant [26]. Several dynamical-scalar-field "quintessence" models (see, *e.g.*, [27]) have somewhat less negative equation-of-state ratios, and so could be tested if the measurement uncertainties were smaller. Dark energies can also have equations of state that vary with time, and this would require much higher precision data finely spaced over a wide range of redshifts.

10 A definitive supernova cosmology measurement from a satellite

It is clear that if we are to study the physics of this dark energy that accelerates the universe's expansion, we must make a major step forward in the quality, quantity, and redshift extent of the supernova data. We have recently proposed a new satellite experiment to accomplish this. (Its current name is "SNAP," for "Supernova/Acceleration Probe".) The concept includes a 2-m class telescope, a wide-field one-square-degree imager with almost one billion pixels, a near-infrared imager, and a spectrograph that covers the wavelength range from 0.3 to 1.7 microns in three subsystems [28].

With these instruments, a satellite experiment would be able to find and study 2000 SNe Ia per year, almost two orders of magnitude more supernovae per year than the current ground-based work. The Hubble diagram could now be filled with thousands of supernovae extending to redshifts as far as $z = 1.7$. Moreover, every one of these 2000 supernovae would be followed with a much more extensive photometric and spectroscopic study than almost any supernova to date. This detailed follow-up is crucial in reducing all of the sources of systematic uncertainty. (Such a large statistical sample would serve no purpose unless the systematic uncertainties are reduced in this way, because the current statistical uncertainties are already within a factor of two of the systematics.)

With this kind of comprehensive data set, the confidence region in the Ω_M – Ω_Λ plane could be narrowed down (as shown in Fig. 11) to a ± 0.02 measurement of mass density and a ± 0.05 measurement of the cosmological constant, or a ± 0.01 measurement of both assuming a flat universe. This supernova measurement by itself would determine the curvature of the universe to ± 0.06—a curvature measurement that is independent of the cosmic microwave background data. This provides an important, independent test of the inflationary universe theory, which predicts a flat universe.

With the proposed satellite data set, it will be possible to determine the equation-of-state ratio of the dark energy to much higher precision, ± 0.05 in a flat universe. Various quintessence models and other dark energy models can then be differentiated (see Fig. 12). We can study time-varying equations of state by breaking the dataset into smaller redshift bins. One may then be able to reconstruct the potential for a dynamical scalar field model.

11 Conclusions

To summarize, the current type Ia supernova data suggests that we live in an accelerating universe, a universe with either a positive cosmological constant or some other dark energy with strongly negative pressure. This statement can be made with great confidence if you have reason to believe that we live in a flat universe, either based on the cosmological microwave background data or based on the predictions of the inflationary universe theory. The best fit in a flat universe is approximately $\Omega_M = 0.7$ and $\Omega_\Lambda = 0.3$. Even if you are not yet convinced by the current data (or inflationary theory) that we live in a flat universe, and consider the possibility that the universe is open and low-mass, the positive-Λ supernova results are still statistically quite strong and we have not yet been able to identify any systematic uncertainty that could reconcile our data with zero Λ. The loopholes that remain in this last statement will be addressed with upcoming data from the new low-redshift supernova campaign and the new very-high-redshift supernova work.

Finally, we have begun to address the question of what is this new, mysterious energy that is causing the universe to accelerate. Perhaps it will turn out to be Einstein's cosmological constant after all, but if so there are important physics problems that will need to be addressed. The proposed satellite experiment is designed to study this dark energy question, and provide a definitive supernova measurement of

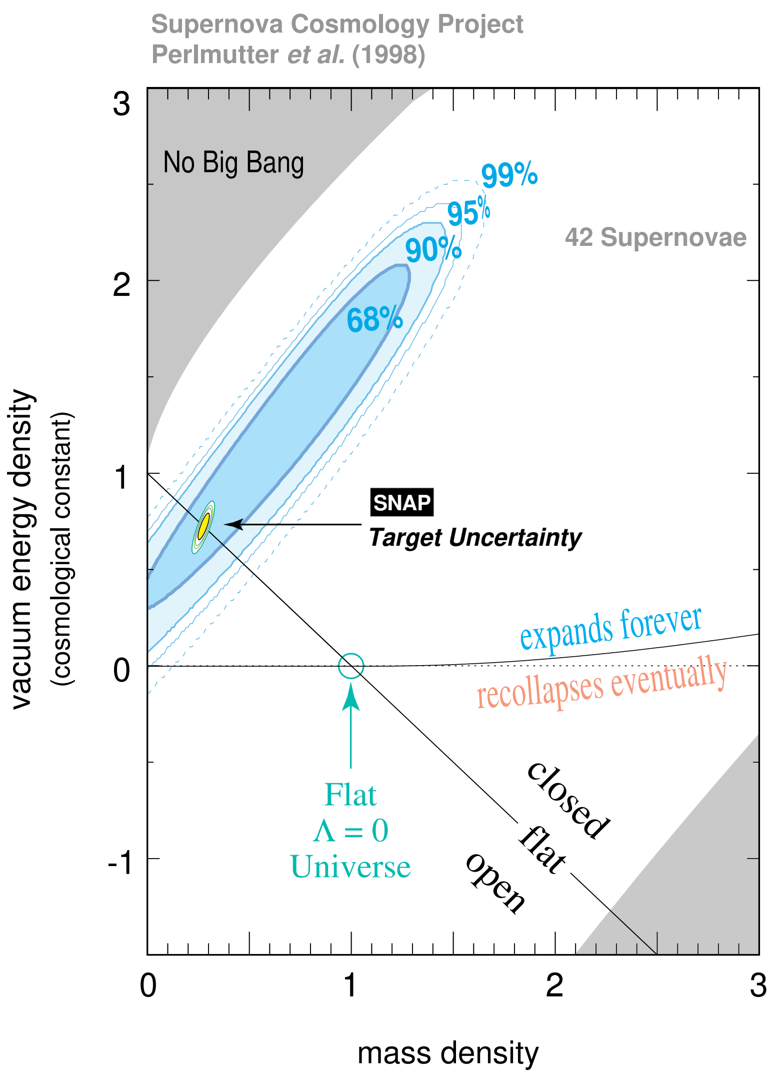

Figure 11: Target baseline confidence region in the $\Omega_M - \Omega_\Lambda$ plane for 2000 hypothetical supernovae at $z = 0 - 1.7$ discovered and measured with the proposed SNAP satellite, compared with the current confidence region from 42 supernovae at $z \approx 0.35 - 0.85$.

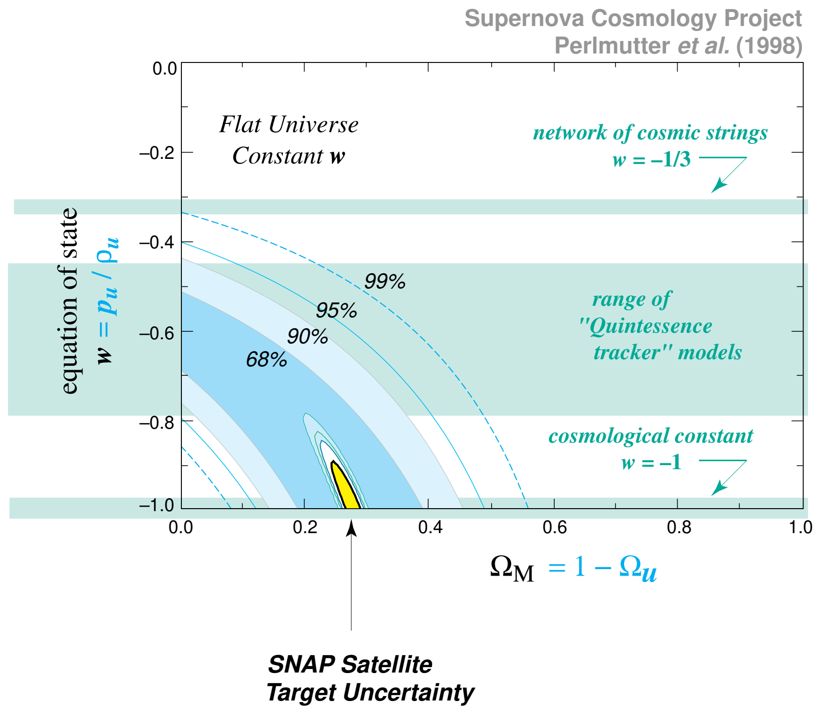

Figure 12: Target baseline confidence region in the $w - \Omega_M$ plane for 2000 hypothetical supernovae at $z = 0 - 1.7$ discovered and measured with the proposed SNAP satellite, compared with the current confidence region from 42 supernovae at $z \approx 0.35 - 0.85$.

the history of the universe's expansion.

Our current state of knowledge leaves us with a rather surprising picture of what this expansion history might look like. We had thought that we lived in a universe that would either expand forever, asymptotically slowing, or someday come to a halt and recollapse. Instead, we seem to be living in a universe that was slowing in its expansion for the first half of history when the mass density dominated, but has been accelerating ever since. It now appears that we live in a universe that will continue to accelerate forever in its expansion—but the jury is still out; and we have only begun to take our first steps in cosmology as an empirical science.

References

[1] See `http://supernova.lbl.gov/`.

[2] The Supernova Cosmology Project: G. Aldering, B. J. Boyle, P. G. Castro, W. J. Couch, S. Deustua, S. Fabbro, R. S. Ellis, A. V. Filippenko, A. Fruchter, G.

Goldhaber, A. Goobar, D. E. Groom, I. M. Hook, M. Irwin, A. G. Kim, M. Y. Kim, R. A. Knop, J. C. Lee, T. Matheson, R. G. McMahon, H. J. M. Newberg, C. Lidman, P. Nugent, N. J. Nunes, R. Pain, N. Panagia, C. R. Pennypacker, S. Perlmutter, R. Quimby, P. Ruiz-Lapuente, B. Schaefer, and N. Walton.

[3] The High-Z Supernova Search: P. Challis, A. Clocchiatti, A. Dierks, A. V. Filippenko, P. Garnavich, R. Gilliand, C. Hogan, S. Jha, R. Kirshner, B. Leibundgut, M. Phillips, D. Reiss, A. Riess, B. Schmidt, R. Schommer, C. Smith, J. Spyromilio, C. Stubbs, N. Suntzeff, and J. Tonry.

[4] S. Perlmutter, *et al.*, 1995. in Presentations at the NATO ASI in Aiguablava, Spain, LBL-38400, page I.1, and in Thermonuclear Supernova, P. Ruiz-Lapuente, R. Canal, and J. Isern, eds. (Dordrecht: Kluwer, 1997).

[5] M. M. Phillips, Ap. J. Lett. **413**, L105 (1993).

[6] A. G. Riess, W. H. Press, and R. P. Kirshner, Ap. J. **438**, L17 (1995), **473**, 88 (1996).

[7] S. Perlmutter *et al.*, Ap. J. **483**, 565 (1997).

[8] M. Hamuy, M., M. M. Phillips, J. Maza, N. B. Suntzeff, R. A. Schommer, and R. Aviles, AJ **112**, 2391 (1996).

[9] A. Kim, A. Goobar, and S. Perlmutter, PASP **108**, 190 (1996).

[10] G. Goldhaber, in preparation.

[11] A. Kim *et al.,* in preparation.

[12] A. Goobar and S. Perlmutter, Ap. J. **450**, 14 (1995).

[13] S. Perlmutter *et al.*, Nature **391**, 51 (1998), E **392**, 311 (1998).

[14] S. Perlmutter *et al.*, Presentation at Jan '98 AAS Meeting, Lawrence Berkeley National Laboratory Report No. LBNL-2230, astro-ph/9 812473 (1998).

[15] S. Perlmutter, *et al.*, Ap. J. **483**, 565 (1999), astro-ph/9812133.

[16] A. V. Filippenko and A. G. Riess, in Proceedings of the 3rd International Symposium on Sources and Detection of Dark Matter in the Universe, Feb. 1998, ed. D. Cline, Phys. Rept. **307**, 31 (1998).

[17] A. G. Riess *et al.*, AJ **116**, 1009 (1998).

[18] B. P. Schmidt *et al.*, Ap. J. **507**, 46 (1998).

[19] N. A. Bahcall, J. P. Ostriker, S. Perlmutter, and P. J. Steinhardt, Science **284**, 1481 (1999).

[20] A. G. Riess, *et al.*, AJ, (in press).

[21] G. Aldering, R. Knop, and P. Nugent, AJ, (in press).

[22] G. Aldering *et al.*, IAU Circular 7130 (1999); G. Aldering *et al.*, in preparation.

[23] G. Aldering *et al.*, IAU Circular 7046 (1998).

[24] See references at `http://ccd.lbl.gov/`.

[25] P. Garnavich, *et al.*, Ap. J. **509**, 74 (1998).

[26] S. Perlmutter, M. S. Turner, M. White, Phys. Rev. Lett. **83**, 670 (1999).

[27] R. R. Caldwell, R. Dave, R., and P. J. Steinhardt, Phys. Rev. Lett, **80**, 1582 (1998).

[28] For details, see `http://snap.lbl.gov/`.