

MEASURING THE ACCELERATION OF THE COSMIC EXPANSION USING SUPERNOVAE

Nobel Lecture, December 8, 2011

by

SAUL PERLMUTTER

University of California at Berkeley and Lawrence Berkeley National Laboratory Institute, Baltimore, MD, USA.

INTRODUCTION

The “discovery of the acceleration of the expansion of the Universe” does not flow trippingly off the tongue – which is fitting, since the work that led to it was comparably long and tortuous. In this Lecture, I would like to give you a feel for some of the science issues we were facing over the ten years leading up to the discovery. I will primarily use the graphics from the original overhead projector transparencies we were using in those years. Although they are not as beautiful as our modern-day graphics, I hope they will help give some of the texture of what was going on during that period.

AN ANCIENT QUESTION

The question that motivated all this work is something you can imagine that the very first humans might have asked when they walked out of their caves at night and found themselves looking up at the starry sky. Do we live in a Universe that goes on forever in space and will it last forever in time? I think it almost defines what it means to be the very first humans – that they could ask such questions.

For most of human history, this sort of question was a truly philosophical question. It wasn't until the 20th century that we began to have a scientific version of this question. This is partly because Einstein's theory of general relativity gave us some new conceptual tools that made it possible to think about this topic in a more rigorous way. But it's also because Edwin Hubble (1929) measured an expansion of the Universe, which meant that we started to see in more concrete terms what we could mean by the fate of the Universe.

Hubble's observations indicated that we don't live in a Universe that is standing still, but rather one in which all of the distances between galaxies are growing with time. You can then immediately start asking yourself whether it will continue to grow with the same speed over time, or might it slow down because gravity would attract all stuff to all other stuff in the Universe. In fact, you might wonder if it could be slowing enough so that someday it could come to a halt, and then collapse into a Big Crunch that could be an end of the Universe.

This is a question about the future of the Universe that you can address by looking into the past of the Universe, by looking to see what was happening billions of years ago and how much the Universe was slowing down back then. If it was slowing enough, you could predict that it is slowing enough to collapse in the future. This was understood, even back in the 1930s, in the decade following Hubble's discovery, by astronomers like Walter Baade and Fritz Zwicky who were studying supernovae. They saw that in principle, you could use a very bright exploding star, a supernova, to perhaps answer this question.

SUPERNOVAE AS PROBES OF THE HISTORY OF THE UNIVERSE

You could take the brightness of a supernova as an indicator of how far away it is: The fainter it is, the further away it is from us – and hence its light has taken more time to reach us. So with the fainter supernovae, you are looking further and further back in time. You can also use the colors of the spectral features of a supernova: a supernova would look blue if it were seen nearby, but when you see it very, very far away it looks red. How red it gets tells you how much the Universe has stretched since the supernova exploded, because while the light is traveling to us, its wavelength stretches by the exact same proportion as the Universe stretches. So, this is a very direct way of plotting how much the Universe has stretched as a function of time. In principle, if you observe enough supernovae and plot them on such a graph, you can see just how much the Universe has decelerated in the past and make a prediction about how much it will slow down in the future.

Baade (1938) wrote about the possible supernova measurement back in the 1930s. The problem was that the supernovae that they knew about at that time were not quite good enough “standard candles” – they weren't all quite the same brightness. They varied by more than a factor of two or three in brightness, so while it looked like a good idea, it was not really practical to do at that time. That's essentially where the problem stood for about the next 50 years (but see Kowal 1968 for a benchmark along the way).

HOW OUR WORK BEGAN

Two things happened in the mid-1980s that got me and others interested in the problem: First, there was the realization that the supernovae could be subdivided into subclasses. It was in the mid-1980s that the “Type Ia” subclass was identified (Panagia 1985, Uomoto and Kirshner 1985, Wheeler and Levreault 1985). It began to be clear that this subclass really was a better standard candle than the others. For example, a histogram of supernova brightnesses in a paper by Nino Panagia (1985) showed that supernovae found in spiral galaxies varied greatly, while those found in elliptical galaxies had only a small dispersion (Fig. 1). The implication was that the ellipticals might host a subclass that was a better standard candle, and also that the spiral galaxies are known to have dust that would add further dispersion. The

possibility that we now had a good standard candle suggested that we now might be ready to go back to the original idea of Baade and Zwicky.

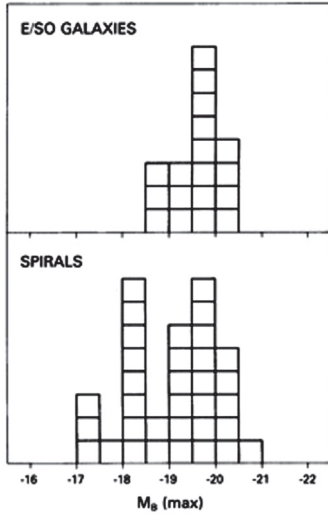


Figure 1. Histogram of supernova absolute magnitudes from Panagia (1985).

The other development of the mid-1980s was the introduction into astronomy of the new sensors, the CCD detectors, which are like the detectors in the back of the digital cameras that most people have today. These were just becoming available in the beginning of the 1980s, and I worked on one of the first astronomy projects to use CCDs. It also included the new computers that were just then becoming fast enough to analyze the large amounts of data that came out of these detectors. This was a project led by my inspiring thesis advisor, Richard Muller. It was doing a search for nearby supernovae using a small robotic telescope and automatic detection of the supernovae by the computers (Fig. 2). I worked on the software that made it possible to subtract the image of a galaxy from the image of a supernova plus the galaxy, and thus find these supernovae (Fig. 3). With the automatic




Figure 2. Saul Perlmutter and Richard Muller with a new telescope the group was automating to replace the Leuschner Observatory telescope (though it was later used for a different purpose at another site).

searching technique we had found 20 supernovae by the time the project stopped (Muller et al. 1992; Fig. 4).



Figure 3. Example of digital image subtraction. From the CCD image of a supernova and its host galaxy, we subtract an image of the galaxy before the supernova appeared (or after it disappeared), leaving an image of just the supernova.




Figure 4. Before-and-after images of one of the supernovae discovered in 1986 by our Berkeley Automated Supernova Search, led by Professor Richard Muller.

These two developments led Carl Pennypacker and myself, both of us researchers in Rich Muller's group at that time in 1987 (I had stayed on as a postdoc after just completing my PhD), to decide to try out the original idea of Baade and Zwicky to measure the deceleration of the universe's expansion using supernovae. We began working on the project with Rich's support, and it looked very promising. However, we were aware that it was not going to be easy. For a number of reasons we knew there was a lot of work to be done to make this possible. So, when we proposed the project we did not get an immediately enthusiastic response from the referees and reviewers. But it's Carl's nature to be absolutely undaunted and to be optimistic that we can do anything, and so we carried on. There were the practical problems of trying to find the more distant supernovae. There were some specific technical problems concerning how you would analyze those very faint, distant supernovae and compare them to nearby examples. There were issues about the

Why is the supernova measurement *not* easy?

1. Can they be found **far enough** -- and **enough of them** -- for cosmology?
Can they be found **early enough** to measure brightness over peak?
2. Can they be identified as Type Ia with spectra, despite how faint they will be?
Can their brightness be compared with nearby ones, despite greatly "redshifted" spectra?
3. Are the supernovae **standard enough**?
And how can one eliminate possible dust that might diminish their brightness?
4. Couldn't the supernovae **evolve over five billion years**?

Figure 5. Some of the issues that we recognized as hurdles to be crossed in order to use supernovae to measure the deceleration of the Universe's expansion.

standardness of the supernovae themselves and their consistency over time that needed to be addressed. These were the specific details that we saw as the hurdles that had to be jumped over in order to do the project (Fig. 5).

HOW CAN WE FIND SUPERNOVAE ON DEMAND?

I will try to describe how we addressed the problems, because what is interesting about this particular measurement is that it is so simple to describe that it is possible to explain the difficulties in some detail, and most people can then understand what it would take to do this particular project.

First of all, there are the questions of:

- Can you find these supernovae at all?
- Can you find them far enough away?
- Can you find enough of them?
- Can they be found early enough so that they can brighten, fade away and you can measure the peak brightness?

The peak brightness is what we knew was standard, so we couldn't use them if we found them weeks after they had already peaked. These were very difficult problems, because if you had to choose a research tool, you would never choose supernovae – they are a real pain in the neck to do research on (Fig. 6). They are rare – the Type Ia supernovae explode only a couple of times per millennium in a given galaxy. They are random – they don't tell you when they're going to explode. They are short lived – they brighten and fade away in time scales of weeks, so they don't stick around so that you can study them. So, they are just very inconvenient to do research with.

The difficulties were illustrated when we later saw the results from another supernova project, which we learned had started shortly before ours. Hans Nørgaard-Nielsen and a Danish team searched several years for very distant Type Ia supernovae but found just one, several weeks after it had already

Problems

with Type Ia Supernovae as a tool for cosmology

Rare ~1 / 500 years / galaxy

Random can't schedule telescope time
or plan discoveries at new moon

Rapid difficult to catch on the rise

Figure 6. The characteristics of Type Ia supernovae that make them difficult to find and study over their peak brightness.

passed its peak brightness (Nørgaard-Nielsen *et al.* 1989). So, while it was encouraging that such distant supernovae existed, the critics of our project said that this did not look like a very viable program. This was the first concern.

The approach that we took to this problem was to develop the capability to look at more than one galaxy at a time (since looking at one galaxy would have meant waiting the 500 years for a Type Ia supernova to explode), and even more than a small cluster of galaxies at a time (as the Danish program had done). We decided to build a wide-field camera that would allow us, with each exposure, to look at 10 to 20 times as many galaxies as you would find even in a cluster of galaxies. We had to develop an unusual optical system (Fig. 7) that would bring light from a very big field of view onto a small CCD detector. This novel instrument went onto the Anglo-Australian 4-m telescope which meant that we were able to work with a large enough telescope with a large enough field of view to be able to search for supernovae at great distances in thousands of galaxies at a time.




Figure 7. The novel F/1 wide-field CCD camera we developed for the Anglo-Australian 4-m telescope (AAT) to collect a wide enough field to search for $z > 0.3$ Type Ia supernovae in hundreds of galaxies in each image.

With our wide-field camera on the Anglo-Australian Telescope you obtain images like the one shown in Fig. 8a, in which all the small specks of light are the distant galaxies in which we were searching for supernovae. You then take an image another time in the year (Fig. 8b) and subtract it from the first one. You are left with an image that shows just the spot that got bright – and that's your supernova. Of course we did this with software, so we had to go back to our image analysis software and develop that much further. Fig. 8c shows what the computer subtraction looked like, and Fig. 9 shows some of the members of the group at that time.




Figure 8a. An image obtained November 1989 with our wide-field camera on the Anglo-Australian 4-m telescope. The small specks of light are the distant galaxies in which we were searching for supernovae.




Figure 8b. The same field as in Fig. 8a, but observed January 1990. It is reversed in grayscale to indicate that it will be subtracted from the first.



Figure 8c. Computer subtraction of the Fig. 8b image from the Fig. 8a image. The spot remaining is what a supernova would look like.



Figure 9. Members of the group in 1992 discuss images with many distant galaxies obtained with a wide-field camera for the distant supernova search. Left to right: Carl Pennypacker, Saul Perlmutter, Heidi Marvin Newberg, Gerson Goldhaber, Rich Muller.

This method of finding supernovae seemed to work. However, the remaining problems were still haunting us, because even if you could find the supernova you couldn't prove that you had a supernova, since you would need to schedule the largest telescopes in the world to obtain the identifying spectrum. You would also need to schedule the largest telescopes in the world to follow it as it brightens and fades away, to measure its brightness at peak. Of course, no telescope time assignment committee would give you the time to schedule a telescope six months in advance for a supernova that may or may not appear on the proposed observing date, say March 3rd. As Mario Hamuy *et al.* (1993a) put it, in their discussion of the Calan/Tololo Supernova Search at much lower redshifts, "Unfortunately, the appearance of a SN is not predictable. As a consequence of this we cannot schedule the follow-up observations a priori, and we generally have to rely on someone else's telescope time. This makes the execution of this project somewhat difficult."

So, we had to figure out a way to make the whole operation more systematic, and what we came up with was a new search strategy that I developed to make this possible. Fig. 10 shows the strategy on a timeline of new moon to full moon to new moon. I realized that if you collect all of your first images just after new moon, wait 2 1/2 weeks or so, and collect all of the second set of images just before the next new moon (and then subtract all the second set from the first set), you now have enough galaxies collected with a wide field imager to guarantee not just one supernova discovery, but more than half a dozen supernova discoveries. Once you reach this statistically

significant sample size of supernovae, then you can be sure that you always have some new supernovae to observe by the second new moon. The other advantage is that with this short timescale between the two sets of images, you can guarantee that the supernovae won't have enough time to reach maximum and then start fading away; since they rise in a couple of weeks and fade in a few months, this timescale ensures that you always catch them while they're still brightening. You can then guarantee that right before that second new moon (which is the time you need to do the spectroscopy) you will have new supernovae, they will be on the rise, and there will be more than one. Now you can schedule the follow-up spectroscopy and photometry on the following nights.

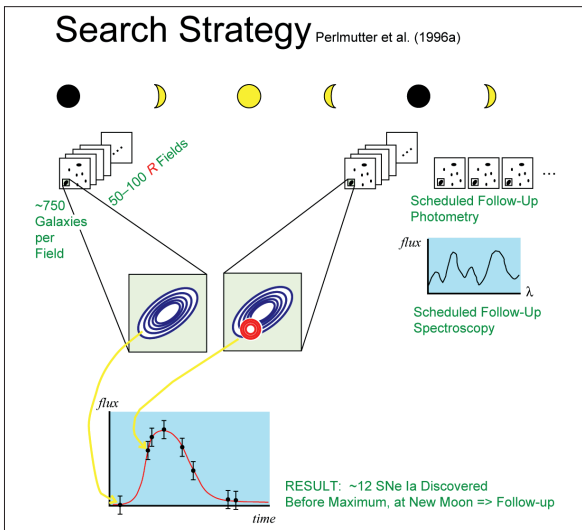


Figure 10. The “batch” observational strategy that made it possible to guarantee multiple new supernova discoveries at high redshift. They would all be on a pre-specified date (in particular just before a new moon) and all while still brightening (Perlmutter et al. 1995a). This in turn made it possible to propose many months in advance for scheduled telescope time at the largest telescopes to study the supernova over the peak of their light curves with both photometry and spectroscopy.

Once we demonstrated this, we were able to start applying for telescope time at the best telescopes in the world – the telescopes in Chile, where the weather is good enough that you can usually follow this whole time series without getting hit by clouds and rain. The first time that we tried this new “batch” observational strategy we called up the International Astronomical Union telegram service, which notifies astronomers around the world when a new supernova is discovered so it can be studied. They’re very careful about what they allow in the telegram and I wanted to be sure they would be able – and willing – to accept our results. I warned them that we would be sending them a half a dozen new supernova discoveries two weeks from now. They laughed, because nobody had ever predicted a supernova discovery before, and certainly nobody ever found more than one at a time before, so it sounded a little unusual. In fact, we then sent out such telegrams semester after semester for the following years, as this new observing technique worked, producing batches of supernovae. So, this did work and we had surmounted the problem of making the supernova a systematic tool.

MORE DIFFICULTIES TO OVERCOME

The next problem was that even if you found the supernovae, there was the question of whether you could identify the subtype, i.e. the Type Ia's. It wasn't clear that they would be bright enough to obtain a reasonable spectrum. In fact the first spectra that we obtained looked mostly like a lot of noise. We realized, however, that we could take advantage of the fact that the spectral features of supernovae are very broad on these wavelength scales, unlike the narrow spectral features of the underlying host galaxy and the noise. (This is because the supernova explosion spreads the spectral features out over a wide range of wavelengths, due to the Doppler shifts.) This meant that we could cut out all the sharp lines, and then smooth the whole spectrum, in order to bring out the broad features that then did look like a supernova (Fig. 11).




Figure 11a. Spectrum of a high redshift Type Ia supernova (Lidman et al. 2005).




Figure 11b. Same spectrum as Fig. 11a, after removing very narrow spectral features, and smoothing to bring out the broad supernova features. The red curve shows the excellent match with a spectrum of a low-redshift Type Ia supernova, as it would appear redshifted to $z = 0.55$.

Once we were able to show this, it became believable that we could identify the supernovae at these great distances. It also helped that the larger telescopes came online just at this time, e.g. the Keck telescope that was being developed just upstairs from me at Lawrence Berkeley Laboratory while I was a graduate student. At that time, I had no idea that I would be using it, but just a few years later – as it was being commissioned – it was exactly the telescope that we needed to do this kind of work.

Another problem arises because the spectrum of a low-redshift supernova gets shifted to the redder wavelengths when you look at supernovae at greater distances – that’s what we mean by “redshift”. Seen nearby, most of the light of the supernova is emitted in the blue wavelengths, so typically at low redshifts we study the supernova using the B-band filter (“B” for “blue”). But, at high redshift the same blue filter would be looking at the very faint UV tail of the spectrum (Fig.12). The question is, how are you going to compare a supernova seen in this faint tail of the spectrum with a supernova seen at the peak of the spectrum. This kind of comparison is called a K correction, and whether it could be done well enough to compare low- and high-redshift supernovae was far from clear. A paper by Bruno Leibundgut (1990) tried to show what the correction would have to be to account for the difference between the different parts of the two spectra. He did a very careful job, but the uncertainties in the correction looked like they would present a significant problem.

In 1992 we found the first high-redshift supernova that we were able to follow throughout its entire light curve. It was at a redshift of $z = 0.45$, (z stands for redshift). When we started trying to analyze it, we saw that the K correction everybody thought we’d have to use was the wrong way to go about it. Instead of trying to find what part of the spectrum we were looking at based on what the blue filter saw, what we really had to do was look at the high redshift supernova with a red filter, the R-band filter. The same part of the spectrum that comes through the R-band filter at high redshift is what comes through the B-band filter at low redshift (Fig. 13). Alex Kim, Ariel Goobar, and I wrote a paper describing this “cross-filter K correction” approach, which made us confident that we could control uncertainties that would otherwise be introduced by the single-filter K correction (Kim, Goobar and Perlmutter 1996; see also Nugent, Kim and Perlmutter 2002).

The next set of issues had to do with how standard these supernovae were. Could you rely on them as a marker of distances? The papers I mentioned earlier, e.g. Panagia (1985), were certainly suggestive of the possibility that the Type Ia supernovae could make a good standard. Bruno Leibundgut (1988) and Gustav Tammann (Tammann and Leibundgut 1990) were really responsible for making the community aware of how consistent the Type Ia supernovae were with each other, during the late 1980s and into the 1990s, so it’s important to mention their names in this context. It’s also important to note David Branch’s group (Branch and Tammann 1992, Miller and Branch 1990, Branch and Miller 1993) in Oklahoma, who were compiling sets of published supernovae, and finding that they could come up with ways




Figure 12. The standard “K correction” was intended to capture the difference between the amount of light in a given filter (here the “B” filter) seen at zero redshift and at high redshift, due to the different parts of the spectrum that would be observed in that filter. This was calculated for high-redshift supernovae by Leibundgut (1990) and by Hamuy et al. (1993b).

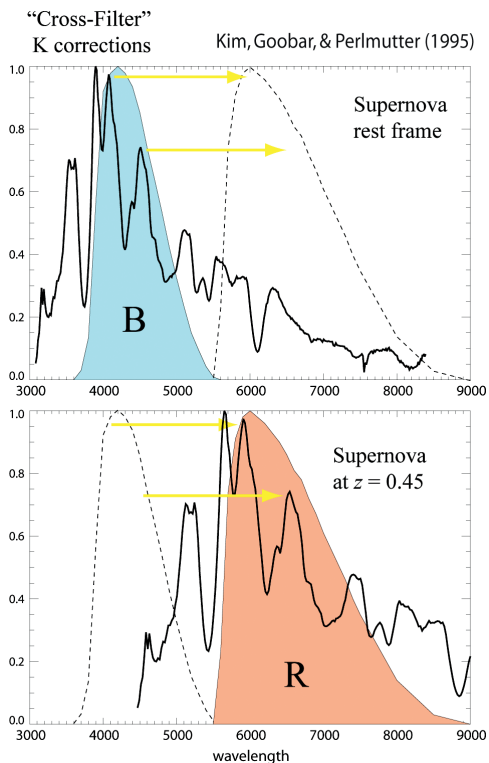


Figure 13. The new “Cross-filter K correction” developed in Kim, Goobar, and Perlmutter (1996) worked by observing the high redshift supernova with a red filter; in this case the R-band filter, so that the same part of the supernova spectrum is coming through this filter at high redshift as comes through the B-band filter at low redshift. This approach makes possible a much smaller uncertainty for the K correction.

to make them even more standard than the earlier papers indicated. One good example of this was the paper by Vaughn *et al.* (1995), which I worked on with the Oklahoma group, that showed the improvement in the dispersion of Type Ia supernovae that was possible by simply rejecting the redder supernovae (Fig. 14). This would remove supernovae that were dimmed by dust (making them appear redder), and also peculiar supernovae that were intrinsically redder. This meant that you could take the dispersion of brightness that you observed from the 40 to 50% range that was seen in the earlier dataset down to something like a 30% dispersion by removing the redder supernovae. The Vaughn *et al.* paper pointed out that this dispersion included measurement errors, and since these were not very well measured data sets the measurement errors could, in fact, already account for most of this 30% dispersion. So, it was possible that the Type Ia supernovae were perfectly good standard candles but, given the quality of data available, we just couldn't tell.

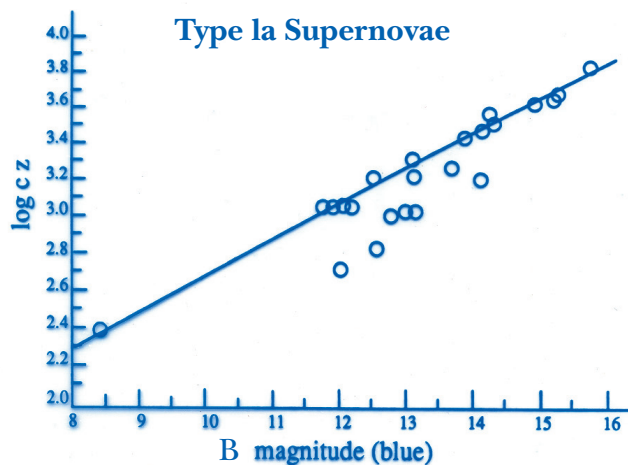


Figure 14. One approach to improving the standardization of the Type Ia supernovae suggested by David Branch's Oklahoma group (Branch *et al.* 1993, Vaughan *et al.* 1995) was to remove all the redder supernovae (plotted here as red-filled circles). This apparently removed both the dust-reddened supernovae and many peculiar Type Ia supernovae, thus greatly improving the dispersion about the Hubble line.

CALAN/TOLOLO: A NEW AND IMPROVED LOW-REDSHIFT DATASET

In the early 1990s, the Calan/Tololo supernova search joined in the game. This was a key moment for this field. I should particularly mention the names of Mario Hamuy, José Maza, Mark Phillips, and Nick Suntzeff who were leaders in this work, because they were responsible for setting the next stage. The dataset they developed contained relatively nearby supernovae that were just far enough out into the Hubble flow so that you could measure the relative distances very well (Hamuy *et al.*, 1993a, 1996). Using the exact same color cut from the Vaughn *et al.* paper, you could bring the brightness dispersion

down to about 18%, just by using this better measured dataset. Now, this was good enough for the cosmological measurement we wanted to make.

There was another approach to this that Mark Phillips proposed in the early 1990s: You could use the timescale of the supernova event, e.g. how rapidly it decayed from its peak brightness, to estimate what its brightness had been at the peak of its light curve (Fig. 15) (Phillips 1993). Mark saw this in the previously published supernova data, but even more clearly in the new Calan/Tololo data. Adam Riess then came up with a sophisticated statistical analysis of this, adding and subtracting template light curve shapes (Riess, Press and Kirshner 1995). Our group used a third method of capturing this supernova timescale that I had developed: Simply stretching or compressing the time axis of the light curve by a single “stretch factor” (Perlmutter *et al.* 1995a, 1997, 1999; see also Goldhaber *et al.* 1995, 2001).

Lightcurve Width-Luminosity Relation

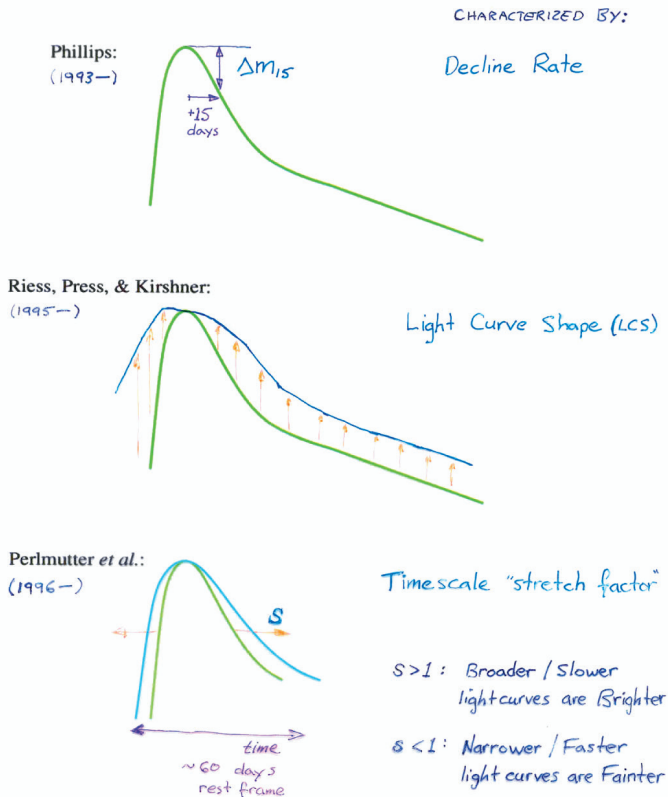


Figure 15. Three alternative approaches to using the time scale of the supernova event (or the shape of the light curve) as an indicator of how bright the supernova reached at peak. These methods followed Phillips' (1993) recognition that the faster the supernova's decline the fainter its peak magnitude.

With this approach, you could take the range of beautifully measured Calan/Tololo data that includes brighter/slower light curves as well as fainter/faster light curves, and for each one measure the stretch factor that would be needed to make them all match each other – and this stretch factor would predict the brightness of each one. In fact, if you “correct” each light curve by appropriately brightening the faster ones and dimming the slower ones while stretching or compressing their timescales, you can standardize all the light curves so that they all lie right on top of each other (Fig. 16). These techniques brought the field another step forward, further improving the brightness dispersion from 18% down to the 12-15% range.

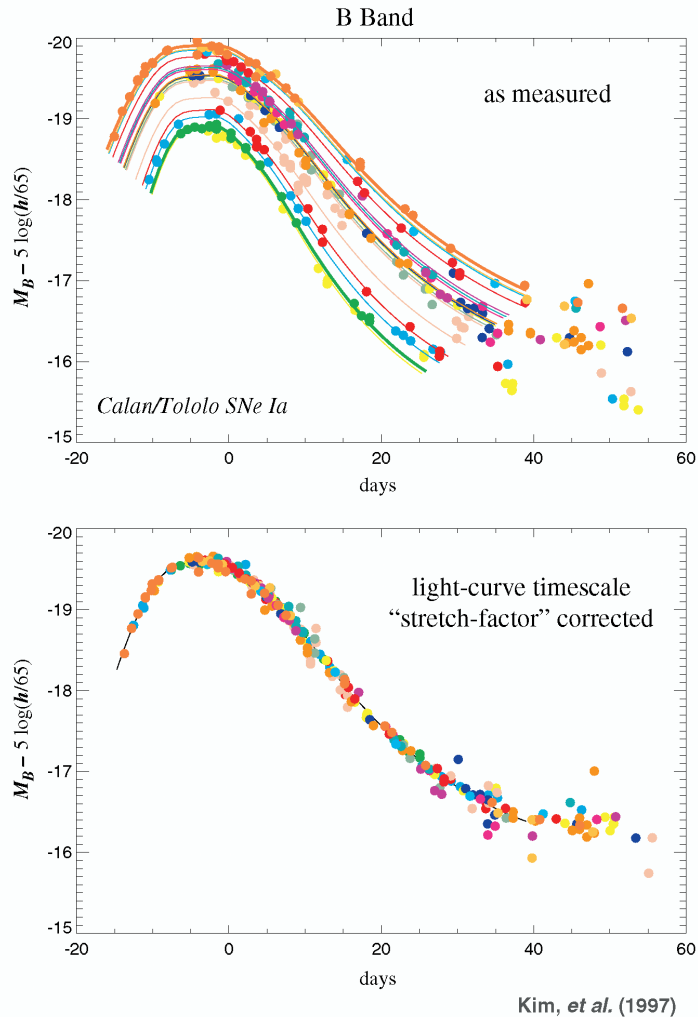


Figure 16. It is possible to “correct” each Type Ia supernova light curve by appropriately brightening the faster ones (low “stretch factor”) and dimming the slower ones (large “stretch factor”) while stretching or compressing their timescales. The lower panel shows the result of this, using a linear relation between the stretch of the light curve timescale and its peak luminosity.

With respect to the problem of dust dimming the supernova, I mentioned earlier that the color cut proposed by the Oklahoma group in the Vaughan *et al.* paper not only removed intrinsically redder supernovae, but also supernovae that were dimmed (and reddened) by dust. Another way to handle dust is to measure how red each supernova is, and then correct the brightness for that amount of dust. It was also found that such a color correction can account for some color-brightness relation intrinsic to the supernovae themselves, since it appears to follow a very similar trend to dust's (Tripp 1998, Tripp and Branch 1999). As I'll discuss later, once we had proven the batch-discovery-and-follow-up technique with that first batch of discoveries we were able to start getting the telescope time to make better supernova measurements, including the good color measurements needed for this purpose.

An important element of this dust story is that apparently most supernovae are not suffering much dimming by dust. At the time, you could already look at the range of relative brightnesses and colors of nearby supernovae and see that most supernovae were in a very narrow range – it would be spread more if dust were prevalent. And, of course, what you really care about if you are comparing nearby and distant supernovae is whether the range of supernova colors is the same for both groups, indicating essentially the same dust-dimming. So we developed two approaches to this color-study of dust: You could compare the nearby and distant distributions of supernova color, or you could correct each individual supernova's brightness using the color as an indication of how much correction is needed. (We used both of these approaches in the results that I'll be discussing, and we were thus able to robustly account for the dust with our measurements.)

There is also a third handle on dust that we began presenting at meetings, as an option for the future if we could find even more distant supernovae. At these much higher redshifts it becomes possible to differentiate the dimming due to dust from the cosmological effects. At great distances the supernova events are so far back in time you would not expect dust dimming to increase with distance the same way it does at closer distances. The farther back you go, the more the dimming should be due to the cosmology and not to dust (Fig. 17). In 1998 (Aldering *et al.* 1998), using the Hubble Space Telescope, we discovered, spectroscopically confirmed, and measured the light curve of a supernova with a redshift of 1.2. It is the first supernova ever found at a redshift above $z = 1$. We nicknamed it Albinoni (Fig. 17), and it was the kind of supernova needed to perform these tests. This approach was then taken to its conclusion by Adam Riess and colleagues, in some beautiful Hubble Space Telescope work (Riess *et al.* 2004).

Finally, there is the concern that the supernovae are not guaranteed to act the same over billions of years. Since we are looking so far back in time with these studies, the worry would be that the supernovae back then might not be the same as supernovae today. It would then be meaningless to compare the brightnesses if they were not identical. Our group realized that there was a nice route to addressing this problem: Just as the low-redshift supernova

sample had been separated into those hosted by elliptical galaxies or spiral galaxies, the same study could be performed for the high-redshift supernovae. Then the cosmology could be determined separately for the low- and high-redshift supernovae found in elliptical galaxies and for those in spirals (Fig. 18). These different host galaxy environments have very different histories, so if the cosmology results from these different environments agree we have a strong indication that the results are not strongly distorted by the environmental histories changing the behavior of the supernovae.

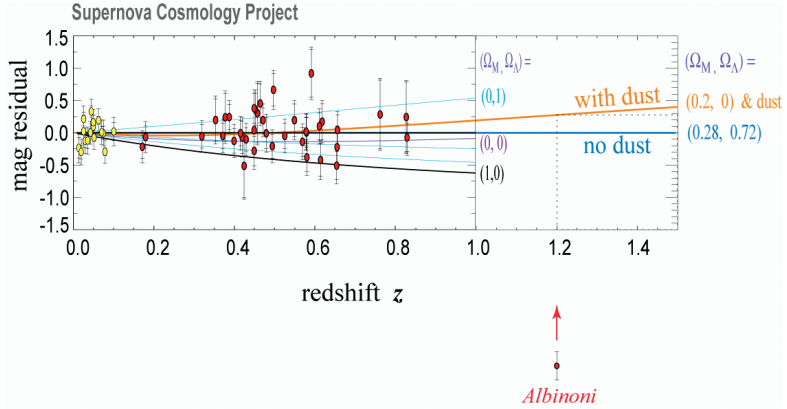


Figure 17. With the discovery (Aldering et al. 1998) of SN 1998eq (nicknamed “Albinoni”), we showed that it was possible to discover and spectroscopically confirm a Type Ia supernova well beyond redshift 1. We suggested that measurements at these very high redshifts would make it possible to separate dust from cosmology in the dimming seen on the Hubble diagram.

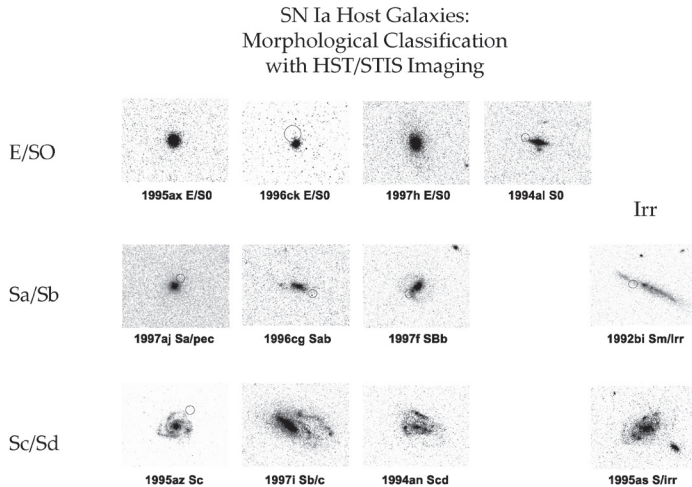


Figure 18. By performing the cosmology measurement with separate subgroups of supernovae found in different host-galaxy environments, it is possible to test that the measurements are not strongly distorted by the different evolutionary histories at low- and high-redshift. We proposed such tests in Perlmutter et al. (1995b, 1997, 1999) and implemented them in several stages, leading to the HST-morphology-based study shown here from our Sullivan et al. (2003) paper.

A NEW CONCERN

By mid-1994, we had answers in hand to the series of concerns listed in Fig. 5 about using distant supernovae for the cosmological measurement – but then we added one further new concern (Fig. 19). We began thinking about this when we were analyzing the cosmological implications of that very first high-redshift supernova we found in 1992. Of course, that supernova by itself did not give you a very definitive measurement. But, it happened to come out with a very low value for the slowing of the expansion of the Universe, and hence the mass density of the Universe responsible for this slowing. In fact, it was so low that we started to think about a mathematical term called the “cosmological constant” that Einstein had put into his equations describing the Universe’s expansion. Einstein early on rejected this term, once it was learned that the Universe was expanding. If it were there, however, it would have the effect of fighting against gravity’s slowing of the expansion. We realized that our first supernova’s very low values for the slowing could be due to some of this cosmological constant fighting against gravity – and then how could you tell what you were seeing? Is it less mass density, that is, less slowing due to gravity, or is it perhaps more of the cosmological constant? The concern was that we wouldn’t be able to tell these apart.

Why is the supernova measurement *not* easy?

a “new” problem

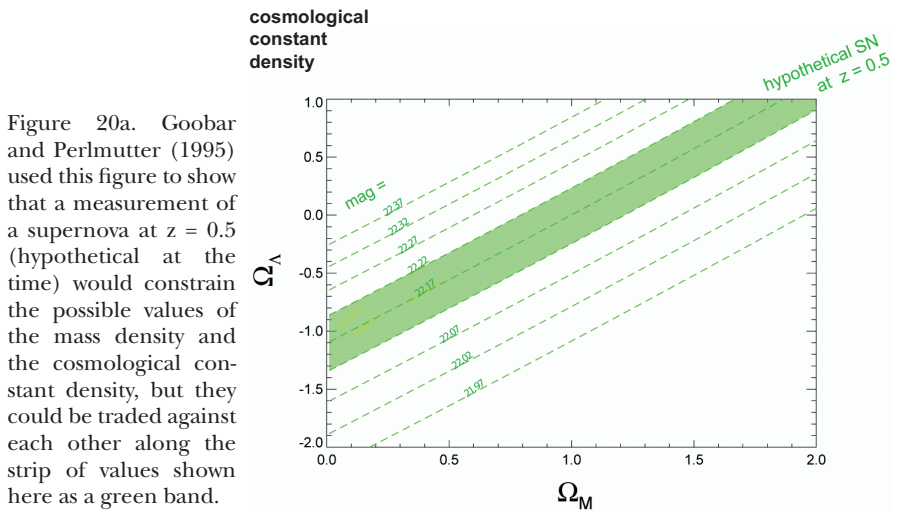
5. What if Einstein's "Cosmological Constant" (Λ) exists? It will fight against gravity due to mass (M) in the universe

-- how can you tell if there is less M or more Λ or vica versa?

Figure 19. By 1994 there was a new concern about the supernova measurement – in addition to the four problems described in Fig. 5.

The problem distinguishing between less mass or a stronger cosmological constant is illustrated in Fig. 20a (from a paper that Ariel Goobar and I wrote together). Say you have a supernova with redshift $z = 0.5$ (which we didn’t yet, so this was still hypothetical) and plot the combinations of mass density and the density of the cosmological constant consistent with its brightness. You get a steadily rising strip of possible values. With just this plot, you wouldn’t know if the true values describing our Universe were at the lower-left part of this strip, with low mass density and low cosmological constant density, or at the upper-right with high values for both densities. What Ariel and I realized was that if you go out to a larger range of redshifts, that is, if you study much

farther supernovae at redshifts as high as $z = 1$ and beyond, then this plotted strip of allowed cosmological values starts to rotate (Fig. 20b). Then the intersection of these strips for different redshift supernovae allows us to separate out these two effects – the mass density and the cosmological constant density on the expansion history of the Universe. Fig. 21 shows the plot we used in the paper to show that for a given mass density you could distinguish between a zero cosmological constant – no cosmological constant at all – and a significant cosmological constant. Interestingly, the values we chose for this example of a significant cosmological constant turned out to be very close to the final answer we found. Of course, we were expecting to find the zero cosmological constant result shown. (We pointed out in the paper that it would be easier to make the measurement if there were a cosmological constant, because the error bars are smaller in that region of this plot.)



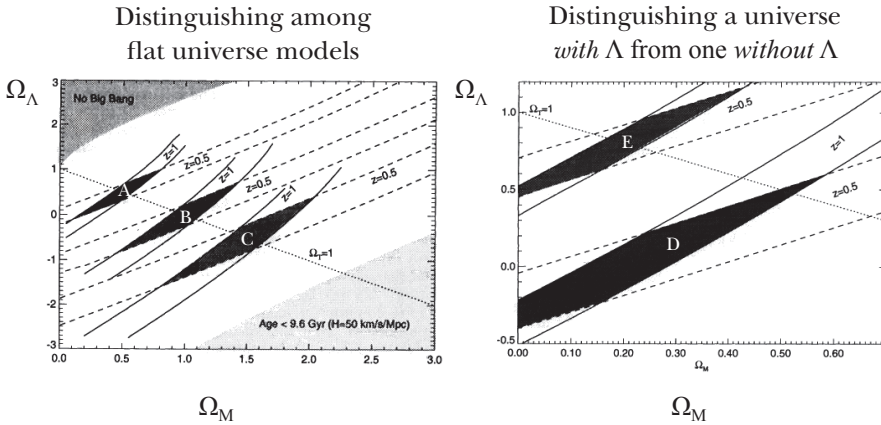


Figure 21. In Goobar and Perlmutter (1995), we gave a range of sample cosmologies to show how supernovae at, e.g., $z = 0.5$ and $z = 1$ could distinguish mass density from cosmological constant density. The examples in the right panel turned out to be particularly prescient, since they showed that in a universe with significant cosmological constant the smaller error bands make it easier to confirm the cosmological constant's existence.

AN UNEXPECTED OUTCOME

Now we were in a position at the end of 1994 to use the new on-the-rise/on-demand “batch” discovery strategy over and over again during each of the following semesters at the telescopes. It became a production job: We wanted to turn out enough of these supernovae and make these delicate high-precision measurements in multiple colors so that we could have the strong statistical sample needed to measure how much the Universe is slowing down. We were by now applying to telescopes all around the world to do this. We had to use telescopes to find the supernovae (at the Cerro Tololo 4-meter telescope in Chile), to follow the supernovae with spectroscopy (at the Keck 10-meter telescope in Hawaii), to follow the supernova with photometry (at the Isaac Newton telescope in the Canary Islands and the WIYN telescope in Tucson).

So it was a pretty dramatic scene during the weeks that we would conduct one of these supernova campaigns. One team would be flying down to Chile and then returning to Chile three weeks later for the “discovery images,” while another team back in Berkeley would be pulling the data in near-real-time over the then-fledgling internet to analyze the data. Meanwhile teams would head out to Hawaii, the Canary Islands, and Tucson, with email going back and forth updating everything we knew about each supernova in our batch of discoveries. By this time, the whole Supernova Cosmology Project (SCP) team was more than the three or four people that we had at the beginning. Fig. 22 shows much of the team at that time. I want to emphasize what a capable, creative, and dedicated group of people this was – and it made for a collaboration that was a great example of teamwork.



Figure 22. A collaboration meeting of the Supernova Cosmology Project in the late 1990s. *Top row:* Greg Aldering, Saul Perlmutter, Isobel Hook; *Second row:* Sebastian Fabbro, Alex Kim; *Third row:* Rob Knop, Pilar Ruiz-Lapuente; *Bottom row:* Peter Nugent, Ariel Goobar, Gerson Goldhaber. (*Not shown:* Reynald Pain, who was taking the picture.)

Semester by semester, we started to build up this increasingly larger sample of supernovae, covering an increasing redshift range. In Fig. 23 the color coding shows the supernova redshifts for each batch of supernova discoveries from a given semester's search-and-follow-up. First a half-a-dozen, and then a dozen, and then another dozen... and by 1997 we had enough supernovae in hand to get results that were statistically significant.

As we added the new batches of high-redshift supernovae to the Hubble diagram, publishing the results at each step, the history of the Universe's expansion slowly began to be apparent. The very first data (the red points around a redshift of $z = 0.4$ in Fig. 24a) appeared to favor a slowing Universe with no cosmological constant, but with only 7 supernovae the uncertainties were large. These were the very first high-redshift supernovae we had studied. After we had used them to show that the batch-discovery-and-follow-up method worked, we were able to request and obtain sufficient follow-up time on major telescopes to make more comprehensive measurements of the following batches of supernovae. So the next supernovae were all much better measured.

In a *Nature* article that appeared 1 January 1998, we then reported that even one very-well-measured supernova – it had Hubble Space Telescope observations – at twice the redshift (the red point at $z = 0.83$ in Fig. 24a) already began to tell a different story, pointing to a Universe with a cosmological constant! But the evidence really became strong almost immediately afterwards with 42 supernovae (the red points in Fig. 24b). Now there was a clear bulk of the supernova data indicating a Universe that is dominated by a cosmological constant, not ordinary matter.

Plotting the same data as an Expansion History of the Universe (Fig. 25a), we see that it doesn't match any of the entire range of possible slowing histories shown (the curved lines on the figure). Apparently we don't live in a Universe that is currently slowing in its expansion, but rather a Universe with one of the more interesting histories shown in the panel of Fig. 25b. Its expansion rate used to be slowing, but has been speeding up for the last half of its history, and presumably could speed up forever.




Figure 23. Each semester (shown in different colors) between 1994 and 1998 we employed the “batch search-and-follow-up” method to add another set of Type Ia supernovae. The distance (redshift) range increased, and we began to use the Hubble Space Telescope to follow the SNe to obtain the most precise measurements.

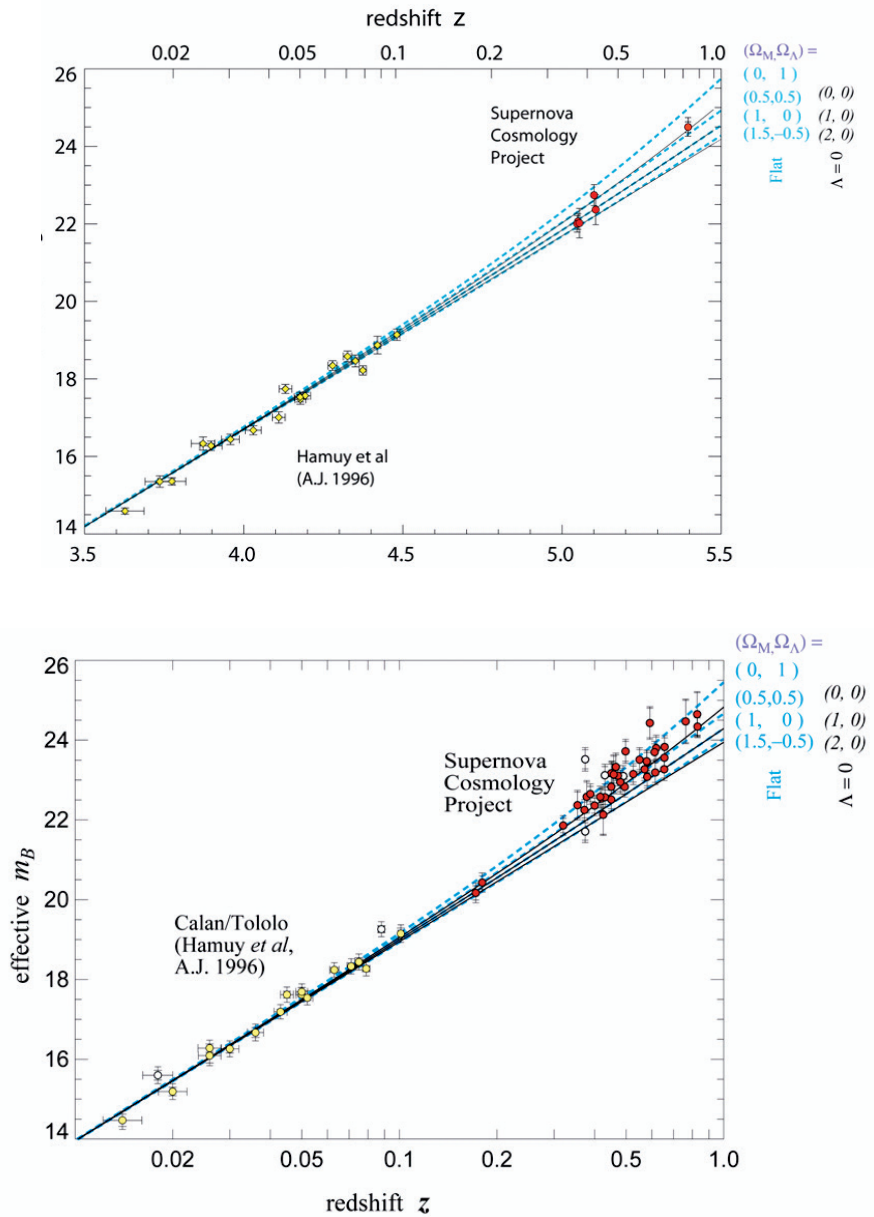


Figure 24. Fig. 24a top – Fig 24b bottom. As the SCP added the new batches of high-redshift supernovae to the Hubble diagram the history of the universe’s expansion slowly began to be apparent (Perlmutter et al. 1998). The very first data (the red points around a redshift of $z = 0.4$ in Fig. 24a) appeared to favor a slowing Universe with no cosmological constant, but with only 7 supernovae the uncertainties were large. Even one very-well-measured supernova – it had Hubble Space Telescope observations – at twice the redshift (the red point at $z = 0.83$ in Fig. 24a) already began to tell a different story. But the evidence really became strong with 42 supernovae (the red points in Fig. 24b). Now there was a clear bulk of the supernova data indicating a Universe that is dominated by a cosmological constant, not ordinary matter. Its expansion is apparently *speeding up*.





Figure 25. Fig. 25a top – Fig 25b bottom. Our supernova data clearly didn't fit with any of the decelerating options shown in Fig. 25a. To fit the data, we now had to add curves that are currently accelerating, as shown in the blue region of Fig. 25b. The best fit curve was decelerating for about the first seven billion years, and then accelerating for the most recent approximately seven billion years. This was the surprising result the supernovae were showing us. (Based on figure from Perlmutter 2003.)

Apparently we have a Universe that is dominated by some new ingredient, some previously unknown “Dark Energy” that makes the Universe expand faster and faster. It is so rare that you come across something that is not part of our current physics model! This is one of the best outcomes that you could ever get in a project like this. I feel very lucky to be able to work on this at all, because here was a project where any result you could discover would be exciting: We might have found that the Universe was infinite, or that it was finite and going to come to an end. Either of these results would have been great. Instead we found an answer that was even better than “great”, in that it was a surprise. This is not something that you can even wish for in science.

This outcome is a perfect example of how science can so often be two-headed. On the one hand, it was only possible to discover what came as a surprise to all of us because our field, physics, had already made such great progress in understanding the Universe. Less than a century ago we had no idea that there was more to the Universe than our own Milky Way. The immense size of the Universe, the fact that it is expanding, the fact that it is populated with such things as exploding stars – all this and more had to be discovered before we could do the work that led us to contemplate an unknown form of energy that accounts for more than two-thirds of everything there is.

It’s amazing how much we figured out, but on the other hand it’s amazing how big a mystery has opened up as a result, and how much we still have left to discover. One of the real pleasures of doing science – which will continue to be true, I believe, on any given day for the next few centuries – is that we have so much knowledge to build upon, yet there is still be so much for us to discover.

These two aspects of science remind us that science is a method, not a finished product. We don’t know where it will lead or what new, seemingly magical powers it will give us in the future. We never know whether what we find will turn out to be useful, but we do know that in the past, whenever we made a major step forward in our understanding of how the world works, we’ve ultimately been able to solve more problems, including very practical problems. I think that’s the only way we can proceed as basic scientists: We try to see what we can understand, and we hope it opens more possibilities for what we can do in the world.

SCENES FROM A DECADE: AN IMPRESSIONISTIC SKETCH

So far, you may have the impression of this science being a very cut-and-dried activity: We identify each problem, and solve it, and then see the results. But the actual experience of this work is completely different: It’s a non-stop whirlwind of activity and people. Unfortunately there are almost no photographs of these 10 year project to show this.

But a few years ago I tried to convey this with the following very fast, impressionistic, verbal sketch of scenes from the decade leading up to the discoveries. (I’ll supplement this with the few photos that we do have, Figs. 26–71.)

It begins with brainstorming in Berkeley with Carl Pennypacker (in 1987) as we first batted around hardware and software plans for a new high-redshift SN project in Rich Muller's group, which Rich soon embraced. And then the consequence: The mountaintop observatory cafeteria at Coonabarabran as Carl, graduate student Heidi Newberg, former-graduate student Shane Burns, and I got to know our pioneering Australia-based colleagues, Warrick Couch and Brian Boyle. They were installing and then using our weird crystal-ball of a wide-field corrector and camera at the AAT 4-meter telescope – which led to our first high-redshift (but unconfirmed!) SN.



Figure 26. Carl Pennypacker (left) and Saul Perlmutter.



Figure 27a and b. Rich Muller then and now.



Figure 28a and b. The F/1 wide-field CCD camera being mounted at the Anglo-Australian telescope.



Figure 29. (left to right) Heidi Marvin Newberg, Warrick Couch, Carl Pennypacker, and Shane Burns then.



Figure 30. Shane Burns now.




Figure 31. Brian Boyle.




Figure 32a and b. Warrick Couch then and now.



Figure 33. Heidi Marvin Newberg.



NOVEMBER
1989



JANUARY
1990

Figure 34a and b. The group's first (unconfirmed) high-redshift supernova, observed at the Anglo-Australian telescope.

Back in Berkeley, I have an image of Gerson Goldhaber overlaying transparencies with negative and positive images of fields full of galaxies – image analysis for the days when the computers were down! I should pause here to say that we are very sad that Gerson isn't here today – he died just over a year ago. He was a warm heart – and sharp eyes! – of our team. He was the host of the collaboration parties, and he should have been here to celebrate with us. We miss him.



Figure 35. Gerson Goldhaber.




Figure 36a–d. Transparencies showing successive images of distant galaxies with computer subtraction to isolate possible supernova.

In Australia, we had rain, rain, rain, and more clouds – and then the sunny relief of beautiful La Palma where our new Cambridge colleagues, Richard McMahon (working with Mike Irwin) studied the most distant quasars. Long nights debugging a new instrument for La Palma, and tense phone calls to the Isaac Newton telescope while the data was sent to Berkeley for analysis. And then, our first “official” high-redshift supernova – and a crucial La Palma spectrum from the Herschel telescope as Richard Ellis has by now joined the team (after being one of those Danish group pioneers).

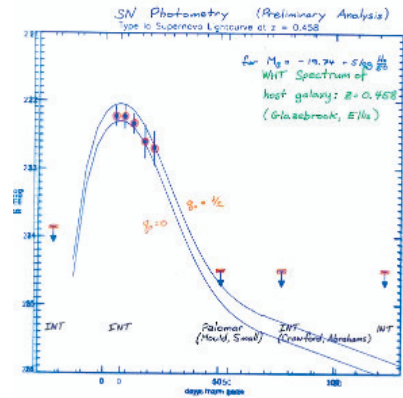


Figure 37. Anglo-Australian Telescope in Coonabarabran, Australia (Fig. 37a) and Isaac Newton Telescope on La Palma in the Canary Islands (Fig. 37b).



Figure 38a and b. Richard McMahon then and now.

Figure 39a. One of the preliminary graphs showing the brightening and fading light of the group's first official high-redshift supernova, SN 1992bi. The plot is on a transparency that is overlaid with another transparency showing the expected light curve, based on nearby Type Ia supernovae.






Figure 39b. Spectrum, obtained at the William Herschel telescope on La Palma, of the galaxy that hosted SN 1992bi.



Figure 40. Richard Ellis then and now.

By this time the Europeans had arrived full force in Berkeley – Ariel Goobar from Stockholm kicked it off by developing new analyses with our then grad student Alex Kim, and brainstorming with me about the cosmological measurements. A glimpse of Reynald Pain from Paris assessing the damage (and successes!) at the end of a complex telescope run – the first of the so-called “batch discoveries.” This epoch ends in my mind with a celebratory party at Gerson Goldhaber’s in the Berkeley hills, where we have a bottle of champagne for each of the half dozen SNe discovered in a batch.



Figure 41. Ariel Goobar (left) and Carl Penny-packer in 1998.



Figure. 41a. Ariel Goobar now.



Figure 41b. Alex Kim.



Figure 41d. Reynald Pain.

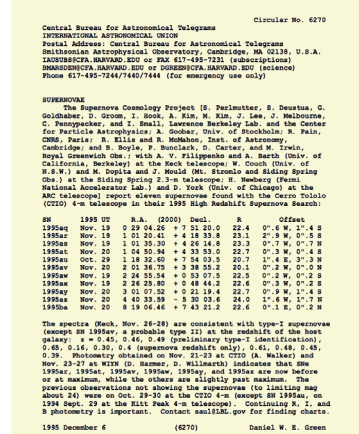


Figure 42. An IAU circular reporting the results of a “batch” SN discovery observing run.

The sociology changes a little as we move to mass production, with new outposts at telescopes around the world, typically manned by a lone team member tenuously connected by a stream of email, phone, and fax. Pilar Ruiz-Lapuente and Nic Walton are at La Palma, Chris Lidman is the voice at the VLT in Chile, Brad Schaefer at Kitt Peak. Larger expeditionary forces head to Cerro Tololo in Chile where all the SN discoveries now are generated (I picture Don Groom and Susana Deustua on one such trip), immediately followed by another team of us rushing with the new list of likely SNe to the oxygen poor mountaintop of the then new Keck Telescope. (I have a memory video-clip of Alex Filippenko – then on our team, his student Tom Matheson, and our new Cambridge PhD Isobel Hook and I crowding round the computer screen as SN after SN proved itself. And at the control center in Berkeley the graduate students working round the clock: I picture Matthew Kim and – from France – Sebastian Fabbro in the cramped room full of students, postdocs, and computers, where our former undergrad intern, Ivan Small, now presides over the growing search software.



Figure 43a and b. Pilar Ruiz-Lapuente then and now.



Figure 44. Telescopes around the world.



Figure 45a and b. Chris Lidman then and now.



Figure 46.
Brad Schaefer.



Figure 47.
Don Groom.



Figure 48.
Susana Deustua.



Figure 49.
Alex Filippenko.



Figure 50.
Tom Matheson.



Figure 51a and b. Isobel Hook then and now.



Figure 52.
Matthew Kim.



Figure 53.
Sebastien Fabbro.



Figure 54. Ivan Small.

I then imagine the calm of space as the Hubble Space Telescope quietly does its part of the job, but of course down on Earth the same flurry of humans – here Andy Fruchter and Nino Panagia – make it possible to use this robot effectively.



Figure 55. The Hubble Space Telescope.



Figure 56. Andy Fruchter.



Figure 57. Nino Panagia.

The last act begins with a view of the end-of-night clean up after the next collaboration party, but this time there are entire cases of bottles of champagne left un-drunk – we’re lightweights – all labeled with the names of the now-scores of new SNe to be analyzed. A fresh relief team of scientists is now on the field at Berkeley: Rob Knop, who thinks, types – and programs – faster than I talk, Peter Nugent, juggling SN theory and practice, and Greg Aldering pulling together all the strains of the analysis ...and the search.



Figure 58. Matthew Kim (*left*) and Rob Knop, relaxing afterwards.



Figure 59. Rob Knop.



Figure 60a and b. Peter Nugent then and now.



Figure 61a and b. Greg Aldering then and now.

A final push of analyses has all of our Berkeley-based pre-graduate-school interns working nights and weekends (parallel computing at its finest): First Julia Lee, and then Patricia Castro, Nelson Nunes, and Robert Quimby – all of whom continued careers in the field.




Figure 62a and b. Julia Lee then and now.




Figure 63. Nelson Nunes and Patricia Castro then (Fig. 63a), Nelson Nunes now (Fig. 63b) and Patricia Castro now (Fig. 63c).



Figure 64a and b. Robert Quimby then and now.

And we all fall gasping in a metaphorical heap with the surprising discovery about the Universe that you just heard about. ... And that's just our science team!

I think it's pretty clear from all this that the popular image of the lone scientist in a lab looks nothing like our experience: Science is – at least for us – an extremely social activity. This particular work was the product of an amazing community of scientists. In fact, between our two teams we include a large fraction, but not all, of the community of scientists studying supernova. I show here several of the key players that I mentioned in the talk, just *representing* the rest of this supernova community. It was an honor and a pleasure to work with all of you on this.



Figure 65. The Supernova Cosmology Project team in 2007.



Figure 66. The High-Z Supernova Search team in 2007.



Figure 67. José Maza.



Figure 68. Mario Hamuy.



Figure 69.
Craig Wheeler.



Figure 70.
Gustav Tammann.




Figure 71. David Branch.

A WORD OF GRATITUDE

None of this whirlwind of human choreography happens without the constant support of our families and friends, our teachers and mentors – and our staff at the universities, labs, and observatories – and also, in many cases, the really courageous administrators and funders who took risks on things when it wasn't obvious they were going to work. They are represented today by the family and friends who are here and we all thank you for helping – and putting up with – all this.

But the work is not done! We look forward to joining in with the next tag-teams of scientists as we delve into the mystery that we are currently calling Dark Energy.

And, finally, we are grateful for the Nobel Prize committees and foundation, who have found a way to encourage this human activity of science.

Thank you.

REFERENCES

1. G. Aldering *et al.*, *IAUC* **7046**, 1A (1998)
2. W. Baade, *Astrophys. J.* **88**, 285 (1938)
3. D. Branch, G. Tammann, *Annu. Rev. Astron. Astrophys.* **30**, 359 (1992)
4. D. Branch, D. Miller, *Astrophys. J. Lett.* **405**, L5 (1993)
5. D. Branch, A. Fisher, P. Nugent, *Astron. J.* **106**, 2383 (1993)
6. G. Goldhaber *et al.*, in: "Four Papers by the Supernova Cosmology Project", ed. S. Perlmutter, *Lawrence Berkeley Laboratory Report*, **LBL-38400** (1995); also in *Thermonuclear Supernovae*, ed. P. Ruiz-Lapuente, R. Canal, J. Isern, *NATO ASI Series C* (Kluwer, Dordrecht), Vol. **486** (1997)
7. G. Goldhaber *et al.*, *Astrophys. J.* **558**, 359 (2001)
8. A. Goobar and S. Perlmutter, *Astrophys. J.* **450**, 14 (1995)
9. M. Hamuy *et al.*, *Astron. J.* **106**, 2392 (1993a)
10. M. Hamuy *et al.*, *PASP* **105**, 787 (1993b)
11. M. Hamuy *et al.*, *Astron. J.* **112**, 2398 (1996)
12. E. Hubble, *Proc. Nat. Academy Sci.* **15**, 168 (1929)
13. A. Kim, A. Goobar, S. Perlmutter, *Pub. Astron. Soc. Pacific* **108**, 190 (1996)
14. C. T. Kowal, *Astron. J.* **73**, 1021 (1968)
15. C. Lidman *et al.*, *Astron. Astrophys.* **430**, 843 (2005)
16. B. Leibundgut, PhD Thesis, University of Basel (1988)
17. B. Leibundgut, *Astron. Astrophys.* **229**, 1 (1990)
18. D. Miller, D. Branch, *Astron. J.* **100**, 530 (1990)
19. R. Muller *et al.*, *Astrophys. J. Lett.* **384**, L9 (1992)
20. H. U. Nørgaard-Nielsen *et al.*, *Nature* **339**, 523, (1989)
21. P. Nugent, A. Kim, S. Perlmutter, *Pub. Astron. Soc. Pacific* **114**, 803 (2002)
22. N. Panagia, *Lecture Notes in Physics* (Springer Berlin/Heidelberg) **224**, 14 (1985)
23. S. Perlmutter *et al.*, in: "Four Papers by the Supernova Cosmology Project", ed. S. Perlmutter, *Lawrence Berkeley Laboratory Report*, **LBL-38400** (1995a); also in *Thermonuclear Supernovae*, ed. P. Ruiz-Lapuente, R. Canal, J. Isern, *NATO ASI Series C* (Kluwer, Dordrecht), Vol. **486**, (1997)
24. S. Perlmutter *et al.*, *Astrophys. J. Lett.* **440**, L41 (1995b)
25. S. Perlmutter *et al.*, *Astrophys. J.* **483**, 565 (1997)
26. S. Perlmutter *et al.*, *Nature* **391**, 51 (1998)
27. S. Perlmutter *et al.*, *Astrophys. J.* **517**, 565 (1999)
28. S. Perlmutter, *Physics Today* **56**, 53 (2003)
29. M. Phillips, *Astrophys. J. Lett.* **413**, L105 (1993)
30. A. Riess, W. Press, R. Kirshner, *Astrophys. J. Lett.* **438**, L17 (1995)
31. A. Riess *et al.*, *Astrophys. J.* **607**, 665 (2004)
32. M. Sullivan *et al.*, *Mon. Not. R. Astron. Soc.* **340**, 1057 (2003)
33. G. A. Tammann, B. Leibundgut, *Astron. Astrophys.* **236**, 9 (1990)
34. R. Tripp, *Astron. Astrophys.* **331**, 815 (1998)
35. R. Tripp, D. Branch, *Astrophys. J.* **525**, 209 (1999)
36. A. Uomoto, R. P. Kirshner, *Astron. Astrophys.* **149**, L7 (1985)
37. T. Vaughan, D. Branch, D. Miller, S. Perlmutter, *Astrophys. J.* **439**, 558 (1995)
38. J. C. Wheeler, R. Levreault, *Astrophys. J. Lett.* **294**, L17 (1985)