

MY PATH TO THE ACCELERATING UNIVERSE

Nobel Lecture, December 8, 2011

by

ADAM G. RIESS

Johns Hopkins University, Baltimore, MD, USA, Space Telescope Science Institute, Baltimore, MD, USA.

INTRODUCTION

I think one of the most amazing facts about the universe is that it is expanding. I never would have guessed it. Even as an undergraduate, once I'd learned a little physics, I would have thought that the universe was eternal, static, and always in equilibrium. So in graduate school when I found out that the universe was expanding, I was awestruck. Then I learned if we could measure the expanding universe, the way we record the growth of a child with marks on a doorframe (Fig. 1), we could determine the *age* of the universe and predict its ultimate fate. This was staggering! I knew this is what I wanted to do. Since that time, charting the expanding universe to determine its nature has been my passion. Though I have to add: knowing what I know now, that the universe is not only expanding but also *accelerating*, I feel like King Alfonso X of Castile who saw Ptolemy's theory of the Cosmos and reportedly said "If the Lord Almighty had consulted me before embarking on creation thus, I should have recommended something simpler."

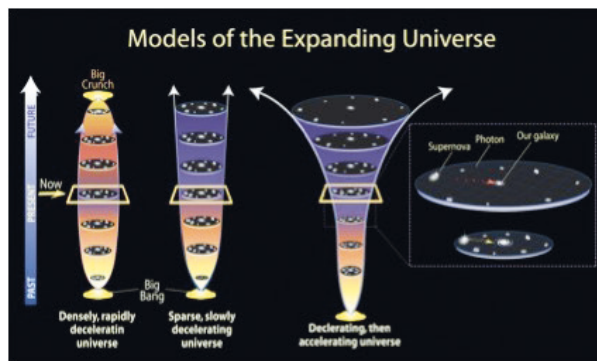


Figure 1. Models of the expanding Universe. We use the redshift and distance (brightness) of supernovae to measure the change in scale with time, respectively. Together these measure the expansion history.

THE EDUCATION OF A COSMOLOGIST

After junior year in high school, I spent a month that summer at the New Jersey Governor's School of Science. I went into the program thinking I was interested in genetic engineering but on a whim, I took a course on Einstein's theory of special relativity. When Dr. Jim Supplee of Drew University explained time dilation, length contraction, and the relativity of simultaneity I was hooked. From 1988 to 1992 I studied physics at the Massachusetts Institute of Technology. "Course 8", as physics was called at MIT, was demanding and MIT stretched my mind and capabilities more than any experience before or since. The hardest but most rewarding course was Junior Lab where we reproduced the great experiments of Nobel Prize winners from the first half of the twentieth century. I learned to love the process by which experimentalists and observers take raw measurements from a piece of equipment and convert them into basic statements of fact and tests of hypotheses. I would remain an ardent "data reductionist" for the rest of my career.

As I was finishing at MIT, Professor Robert Kirshner, then chairman of Harvard's astronomy department, called to tell me I was accepted to the astronomy graduate program and waitlisted for the physics program. He asked if I'd like to come visit and offered, as was standard practice, to pay my expenses. Given my proximity, he said, a subway ("T") token was in the mail to me! I rode two stops up the red line and was brought to the office of a model graduate student named Brian Schmidt. I would later learn critical lessons from him, but at the time we discussed the life of a graduate student. I decided to go to Harvard to earn a doctorate in astronomy. I knew next to nothing about astronomy and astrophysics but had a vague feeling that I wanted to learn more. That summer before grad school I worked on the MACHO Project at Lawrence Livermore National Laboratory where I briefly met Saul Perlmutter. So before I had even started grad school, I had already met many of the people who were going to make an enormous difference to the direction and shape of my career.

The only preliminary exam at Harvard at the time was based on Frank Shu's wonderful book "The Physical Universe", which I happily read from cover to cover. By the time I got to Chapter 14 on the expansion of the universe and Chapter 15 on gravitation and cosmology I knew what I wanted to work on: measuring the rate at which the universe expanded, or rather, the rate at which the expansion was decelerating. Cosmologists were then unable to measure that rate – and correspondingly, the universe's age – to better than a factor of two. Worse yet, their inaccurate measurements of the rate and the age meant, as Chapter 15 pointed out and recent work by Allan Sandage and others confirmed, that they also could not determine the mass density and fate of the universe. To solve both problems, they needed to gauge distances with precision across billions of light years. So far, their best indicator of distance was the brightest galaxies in clusters.

In the spring of 1993, I went to talk to Professor Robert Kirshner about choosing a research project. Kirshner and his student, Brian Schmidt, were finishing work on another distance indicator called core-collapse supernovae. He told me about some new work by Mark Phillips of Cerro Tololo Inter-American Observatory on a third distance indicator, another class of supernova called Type Ia supernovae. Type Ia supernovae were thought to arise when carbon-oxygen white dwarf stars accreted mass from a companion star, grew beyond the Chandrasekhar limit, and exploded. After their explosions, their light output rose, reached a peak and then declined. Because it was thought all explosions were the same, the peak brightness could be used to determine their distances: brightness falls off in a regular way with distance. And because Type Ia's are the most luminous of the common supernova types, peaking at 4 billion solar luminosities, the distances at which they can be seen are extremely high. But there was a catch. Some of these supernovae were not identical; some were *intrinsically* more luminous than others, causing astronomers to over- or under-estimate their distances. Mark found that the intrinsic luminosity at the peak appeared to correlate with the rate at which the light output declined after reaching peak. The more slowly their light declines, the more luminous they are. Thus the Type Ia's decline rate could be used to improve the precision of their distance estimates. Unfortunately the problem remaining was that variations in the amount of dust along the line of sight to a supernova changed its brightness and therefore degraded the precision of the distance estimates.

To understand these variables, a team led by Mario Hamuy, with Mark Phillips, Nick Suntzeff, Robert Schommer and Jose Maza in Chile were conducting the first large-scale program, the Calan/Tololo Survey, measuring the light curves of Type Ia supernovae. This survey would be groundbreaking, gathering the data that would provide the first proof that Type Ia supernovae would make excellent distance indicators.

Bob Kirshner was an expert in observing supernovae. He also had great common sense, a nose for what was important, and the ability to marshal the resources his students needed to succeed. He suggested we collaborate with another Harvard professor, William Press, an expert in developing algorithms – or “Numerical Recipes”, the title of his famous book – for analyzing complex data (Fig. 2). This first summer of work was slow going. Bill Press was away and I tried to glean from his latest paper “Interpolation, Realization, and Reconstruction of Noisy, Irregularly Sampled Data” (Rybicki and Press 1992) and from a preprint of Mark's (Phillips 1993) a way to make optimal use of the supernova's light curve – that is, the rise, peak, and decline of its light – to predict its true luminosity. At the time, the Calan/Tololo team had not yet finished its survey, so we could find little data with which to test new algorithms for predicting a supernova's luminosity, and even worse, most of that data was fatally contaminated. Bruno Leibundgut, an important future collaborator, had compiled an atlas of the historical observations going back to the 19th century of Type Ia supernova, most of them observed with photographic plates. But photographs, due to their analog nature, do

not allow one to accurately separate out the light of a fading supernova from the background light of its host galaxy. Thus the relative rates of rise and decline measured by my newly developed algorithms were less a clue about the supernova's luminosity than they were a measure of the error made in background subtraction. Before any progress could be made in improving measurements of the expansion rate, Type Ia's would have to be observed with the more precise, digital CCD detectors.

However, this time was well spent learning from Brian the techniques of using CCDs to measure the brightness of stars and I wore a furrow in the carpet between our offices. Later I shared a couple of observing runs at Mount Hopkins with Brian, Bob, and Peter Challis. Bob and I decided that part of my thesis would be to collect a large sample of Type Ia light curves, the first sample using CCDs on the skies over the northern hemisphere (the Cerro Tololo survey had been done in the skies of the southern hemisphere). I began searching the circulars from the International Astronomical Union for reports (most from amateurs) of Type Ia supernovae that had just exploded. Every time I found one, I needed follow-up observations of its light curves, so I would trade our group's future time on the Mt. Hopkins 1.2 meter telescope for present time from other observers, a typical cost of 30 minutes a night. We later formalized this time swap arrangement with the telescope's time allocation committee. This was a practical approach to collecting nearly nightly observations without having to live at the telescope.

After identifying a supernova candidate and confirming from its spectrum whether it was a young Type Ia (it often was) with the help of future collaborators Peter Garnavich and Pete Challis, I would contact the present observer to request the appropriate exposures, filters, and exposure times and I would provide a finder chart. The next day I would transfer the observations and calibration data from the computers at Mt. Hopkins to Harvard. Starting with SN 1993ac in October of 1993 and ending with SN 1996bv in November of 1996 I collected 22 well-observed Type Ia light curves, 1200 observations in all. A few years after the Calan/Tololo survey published their final sample of 29, we published ours, thus doubling the world's sample of high-quality, multi-filter CCD-measured light curves. These two samples would form the cornerstone of our team's discovery of acceleration.

At the end of that summer, when Bill Press returned, we worked together on a small sample of 8 well-observed supernovae, using the Rybicki and Press method to decouple the luminosity of a supernova from its distance in one passband – a narrow range of wavelengths – only. The virtue of this method, which we called the Light Curve Shape (LCS) method, was that it provided a true, empirical model which included the covariance of the data and model. So using the LCS method, we could use data from any point in the life of the supernova, and our error estimates would be more rigorous than prior approaches. When applied to the first ten SNe Ia from the Calan/Tololo survey, the LCS method reduced the scatter in the measurements of the expansion rate by more than half. It thereby reduced the number of supernovae required for a significant measurement of expansion or the

expected deceleration of expansion by a factor of 5.5. I was fortunate to be able to test the method on data taken by the Cerro Tololo team in the south, and I want to repeat here the gratitude I expressed to the Calan/Tololo Survey Team then. I quote the acknowledgment from the LCS paper written with Bob and Bill: “We are grateful to Mario Hamuy, Mark Phillips, Nick Suntzeff and the entire Calan/Tololo collaboration for the opportunity to study their outstanding data before publication.” This was a favor we were happy to be able to return by providing them early use of our own northern supernova sample a few years later for a paper written by Mark and the rest of the Calan/Tololo members.



Figure 2. Left: Left to Right, Bill Press, Adam Riess and Bob Kirshner at the Harvard Center for Astrophysics. Right: Application of the MLCS method to nearby supernovae to differentiate the 3 effects of distance, dimness and dust on the flux of supernovae, reprinted from Science Magazine in 1995.

Despite the improvements afforded by LCS, I knew we were still ignoring one of the biggest bugbears for accurately measuring distances in the Universe: the presence of dust. Dust in the galaxies hosting supernovae dims their light, fooling you into over-estimating their distances. Dust also reddens a supernova’s colors: the dust grains are of a size that scatters blue light more efficiently than red, leaving more red light to pass through and fooling you into thinking the supernova is redder than it is. Yet when previous workers, like Allan Sandage and Gustav Tammann tried to quantify the dimming from dust by measuring this reddening of the supernovae’s colors and assuming that they all had the same intrinsic color, the scatter in the distance measurements went up! As Bob said at the time “This is a bad sign that this is the right thing to do!”

At this point I had a new idea, probably the first I had had on my own. I could use the LCS method to disentangle the color intrinsic to the supernova (we learned that dimmer ones were intrinsically redder) from its extrinsic reddening by dust by using several different-colored filters. This

new Multicolor Light Curve Shape (MLCS) was a technique which could distinguish between the effects of distance, dust, and dimness on the brightness of Type Ia supernovae. With the MLCS and the new supernova samples, we could measure the expansion rate of the Universe to unprecedented precision, reducing the scatter from the narrow-waveband approach of LCS by 50% (Fig. 2). More importantly, MLCS could remove the presence of interstellar dust as a major source of uncertainty. By distinguishing an intrinsic and extrinsic change in supernova color it was further possible to make use of additional understanding (known in statistics as a “Bayesian Prior”) about the properties of dust to further improve the distance measurement. This addition would prove of particular importance in our later Nobel discovery: much of the color data we collected had a low ratio of signal-to-noise, a regime where such an approach excels and we used it for all our subsequent distance measurements.

As an aside, Robert J. Trumpler had first identified the importance of dust in determining the positions of stars in the 1930s, so I felt particularly honored to receive the 1999 PASP Trumpler Award for “the doctoral thesis judged particularly significant to astronomy”.

Using MLCS and the vastly improved sample of low- z supernovae we were able to make a number of additional, important measurements: the motions of the galaxies of our Local Group in relation to the universal standard frame of reference, the cosmic microwave background, the reddening properties of the dust in galaxies hosting supernovae and the linearity of the local expansion.

THE BIRTH OF THE HIGH-Z TEAM

In early March of 1994, I was still in graduate school and was on an observing run at the Multiple Mirror Telescope on Mt. Hopkins with Bob and Pete. While we were in the telescope control room, we received an excited call from Saul Perlmutter of Lawrence Berkeley Lab about a possible distant supernova that his team found. We obtained its spectrum for the Berkeley team immediately. The shift in a supernova’s spectrum due to the expansion of space gives its redshift (z) and the relation between redshift and distance is then used to determine the expansion rate of the Universe. Supernovae with greater redshifts and distances reveal the past expansion rate because their light was emitted when the Universe was younger. When compared to their nearby brethren they can measure how the expansion rate has changed over time. SN 1994G at $z = 0.425$, or about 5 billion light years away was the most distant Type Ia supernova known! This experience sparked an interest in all of us to fish for supernovae in the higher redshift waters. Later that year my thesis committee asked me to calculate the number of Type Ia supernovae we would need to observe at this record redshift if we added them to our nearby sample and then used MLCS to usefully measure how much the universe was decelerating. I made a table with the answer: a couple dozen should suffice.

About this time Brian began having discussions with Nick Suntzeff about

forming a new team to compete with Saul's. Brian's reasoning was that a team combining Calan/Tololo and Harvard groups, plus a few of their past members, would combine a solid understanding of supernovae and the critical low-redshift samples, and would therefore be able to hit the ground running. This newly formed "High-Z Team" (Nick was the principal investigator of the first proposal but Brian took the helm by 1996) quickly bought a set of custom filters to deal with the large redshifts of the targeted supernovae. Brian developed software whose importance is hard to overstate; it let us subtract the galaxy light to discover the supernovae while accounting for variations in the image quality. We were now in the race to measure the deceleration of the expansion rate and to predict the future of the Universe.

By April of 1995 Brian had found our team's first distant supernova, another record breaker at $z = 0.478$ and had measured its light output at various dates. I used MLCS to measure its distance and compared the result to different possible values for the universe's deceleration. I showed the result to the rest of the team and they chuckled nervously: that single supernova lay in the region of the diagram indicating the Universe was accelerating. We consoled ourselves that the error bar was big and one can be unlucky with a single object.

In the summer of 1995 a conference about Type Ia supernovae was held in Aiguablava, Spain, and members of both teams attended (Fig. 3). Saul's team, called the Supernova Cosmology Project (SCP), already had 7 high-redshift supernovae and was getting close to an answer. We had 1 high-redshift supernova and the hope that by the next year, we'd catch up to them.

In 1996, in the 6th cycle of Hubble Space Telescope operations, the director of the Space Telescope Science Institute, Robert Williams, awarded 28 orbits of his director's discretionary time on the Hubble Space Telescope to both teams to follow up the high redshift supernovae they found from the ground. The Supernova Cosmology Project was miffed that we were elevated in the process to equals, but Williams felt the Space Telescope should support competition for such an important problem.

That summer after finishing my PhD I began thinking about how well we could measure the deceleration if we followed different strategies. Would we be better off with, say, 3 measurements of 8 supernovae ... or 6 measurements of 4 supernovae? I wrote a Monte Carlo simulation (which allows the computer to repeat the experiment thousands of times) and found we should choose 6 measurements of 4 supernovae, and so choose quality over quantity. That spring we also had good luck with telescopes on the ground, finding 8 new SNe Ia out to $z = 0.62$.



Figure 3. Upper left: High-Z Team in Aspen. (Back row from left: Tonry, Suntzeff, Leibundgut, Filippenko and Hamuy. Front row: Jha, Riess, Schmidt and Kirshner.) Upper right: Riess, Goldhaber and Schmidt in Aiguablava. Lower left: Both Teams in Aiguablava. (From left: Pennypacker, Filippenko, Riess, Schmidt, Nugent and Suntzeff.) Lower right: High-Z Team observing at Keck Summit at fall 1996. (From left: Riess and Filippenko.)

By now, I had finished my PhD and needed to find a post-doc position. Unfortunately, our newly formed High-Z team did not have funds to hire a post-doc so I had to go out onto the job market. I received a firm offer from the Supernova Cosmology Project, communicated by the late Gerson Goldhaber, to work for them – an offer I nearly accepted until I came off the waitlist for a coveted Miller Fellowship at UC Berkeley, a position which offered a small research budget and full autonomy. At Berkeley was Professor Alex Filippenko who was an expert in supernova spectroscopy and an enthusiastic supporter of junior scientists. Moreover, he'd recently switched teams from the SCP to the High-Z Team and brought access to the large Keck Telescope with him. I happily went to Berkeley to work with Alex and continued working in the High-Z Team. Interacting with Alex, an energetic astronomer with great attention to detail, was a highlight of my time at Berkeley. Later that year the SCP began to discuss their initial results from their first seven supernovae. Their data supported a strongly decelerating universe with enough matter so that its mutual gravitation might even cause the future universe to re-collapse. They discussed this result at a conference Bob attended in 1996 in Princeton, and published it in 1997 [4]. The result was close to what was expected at the time, but it was clear that more data were badly needed.

By this time, our team was running into the usual academic management problem of having more chiefs than Indians. The original 9 members of the team had grown to ~15 including the additions of John Tonry, Christopher

Stubbs, Peter Garnavich, Craig Hogan and Peter Challis. And only a few of us were in the “sweet spot” of an academic career – that one fleeting intermediate post-doctoral stage, in which the scientist is fully up-to-speed on the subject but still has maximal time to devote to research. In astrophysicists’ collaborations, the rate-limiting step is usually not CPU power or telescope time; it’s the lack of human power.

We were having trouble analyzing our data, partly because we needed too much time to find and re-observe the supernovae; and partly because, unlike the Supernova Cosmology Project which was based in Berkeley, we were spread around the globe and couldn’t easily work together. By the beginning of 1997 we were already falling behind, collecting new data without yet analyzing the old. Brian suggested and the team agreed on a solution: offer individual junior members the opportunity and burden of pulling the past data together, whipping it into shape, drawing conclusions, and writing up the results. The reward was to be the leader of the study and the lead author on the published result. The junior members would thus gain the badly needed academic credentials needed to land a faculty job in exchange for their sweat equity. To jump start the process, junior members of the team (Brian, myself, Peter Garnavich, Pete Challis, Alejandro Clocchiatti, Al Diercks and David Reiss) got together in Seattle in March 1997 at the invitation of Chris Stubbs for a working summit we called “Reduce-Fest” to work on the software to analyze telescope images. We made some progress but in the end realized that our reference images weren’t good enough to let us reliably subtract host galaxy light from the images of the supernovae. Brian and I had vociferous arguments about the best way to measure the brightness of the supernovae when the image statistics were dominated by random fluctuations in the brightness of the night sky – statistics that in hindsight were integral to the future work. Brian and I argued a lot about science, but the arguments were always the good kind – wrestling with how to proceed on a thorny problem. When Brian moved to Australia these arguments occurred at strange hours for both of us and had large phone bills, paid by Alex.

Our team began to fear that the lack of a “first result” paper from the High-Z Team to match that already published by the Supernova Cosmology Project could limit our ability to compete successfully for telescope time. We knew we still had a lot of work to do to analyze the full sample of supernovae in hand, but we saw a way of showing progress. In 1997 we had used our Hubble Space Telescope time to follow 4 new high redshift supernovae, and the quality and homogeneity of the Hubble data allowed for a much faster analysis of those data. Peter Garnavich volunteered to play the “scout”, leading a quick march to analyze this Hubble data and to write a short letter about the results [1]. What Peter and our team saw in that first reconnaissance of the distant universe was already revealing. We could see that the universe was not decelerating strongly enough to re-collapse in the future.

Meanwhile, the Supernova Cosmology Project had changed their earlier conclusions based on their first 7 high redshift supernovae [3]. They now included their new highest redshift supernova which, owing to their Hubble

observations, was also their best measured. Both teams were now agreeing that the Universe was not decelerating enough to re-collapse. By the end of 1997, and at a dual press conference in January 1998 at the AAS, both the High-Z team and the Supernova Cosmology Project presented this result and hints of more results to come.

THE ACCELERATING UNIVERSE

Later in the year 1997 I began leading the effort to analyze the bulk of our team's data. Most of what I did was technical work. After using the new reference images to subtract the light of the galaxies, plus a software package Brian wrote, I measured the brightness for six of our supernovae relative to stars in the fields. Attendees of "Reduce-Fest" pitched in by re-measuring the single supernova they had worked on in Seattle. Next I calibrated the brightness of the reference stars against three nights of observations of standard stars. Because the supernovae had been monitored with different CCDs at different facilities I also needed to measure differences in CCDs' wavelength responses using the calibrated reference stars seen through thin clouds. A quick consultation with Nick Suntzeff was sufficient to conclude that intervening clouds were grey enough for this approach to work. I then set to work improving my light-curve measuring algorithm, MLCS. I needed to incorporate into it the new supernova light curves from my thesis and from the Calan/Tololo Survey, plus a few new improvements garnered in the prior two years. These improvements included a second order correlation now apparent between light curve shape and the supernova's luminosity, and the use of Monte Carlo simulations of the lines of sight through galaxies to improve upon the measurement of the dust between us and the supernova. This latter idea came from a new paper in 1997 by Hatano, Branch and Deaton. This new, improved version of MLCS ended up as the Appendix to the paper I was already writing [5].

If you are ever in a room full of physicists and astronomers and want to figure out which is which, ask each if they know how to calculate a "K-correction". Both will know how to do it, but only the astronomers will know it by its name. The need for this correction arises when we measure the distance to a supernova from its brightness because the way we see a supernova is also affected by the universe's expansion. Besides causing the redshift, cosmic expansion also dilates (expands) time intervals over which supernova light is collected, changes the size of the increments in brightness, and shifts the portion of the spectrum we observe. (A physicist would mutter about need for "relativistic corrections"). Brian and I iterated back and forth a few times before we were both comfortable applying these corrections. I later heard from members of the Supernova Cosmology Project that despite their own extensive study of these corrections, errors in making the corrections likely limited the accuracy of the analysis of their first seven supernovae. K-corrections are a tricky step.

By the fall I had run all the data through MLCS and later, with Mark

Phillips, through his own algorithm so we could compare results. Then, to reclaim a half dozen other supernovae previously abandoned because our observations didn't cover their full light curves, I used a new "Snapshot Method" I developed that year with Peter Nugent which used the spectrum to determine the supernova's progress along its incomplete light curve. I was keen to make use of every scrap of data we had, since I had begun to hear from Supernova Cosmology Project members (during games of touch football played in the muddy parks of Berkeley) that they now had about 40 supernovae at high redshift.

The fewer numbers of high-redshift supernova that our team had did have compensations. First, our team was able to muster a larger sample at low redshift by adding to the 17 supernovae from the Calan/Tololo Survey that both teams used, another 17 from my thesis and my Snapshot paper. Second, because of our extensive color measurements and use of the Bayesian prior, our high redshift supernovae had half the scatter of the other teams' sample. In all, during that fall of 1997, a lot of new and important developments for our supernova cosmology work came together and critically increased the credibility of our subsequent discovery.

I had measured the supernovae used to determine the universe's expansion rate in the present and in the past, and could then measure how that rate changed over the last few billion years. Knowing how the expansion decelerated – a quantity called q_0 – I could predict the amount of mass the universe must have – a quantity called Omega matter, or Ω_M . The higher the universe's mass the more its gravity pulls against its expansion, and the more the universe decelerates. The equation is surprisingly simple, $q_0 = \Omega_M / 2$. But what I initially measured and wrote in my lab notebook in the fall of 1997 was stunning! The only way to match the change in the expansion rate I was seeing was to allow the universe to have a "negative" mass. In other words, up-ending the equation, the Universe wasn't decelerating at all – it was accelerating (Fig. 4)!

That simple equation assumes that matter is the only important component of the Universe. So initially, I hadn't considered any force besides the gravity of matter, and now my computer programs were telling me that only an imaginary negative mass could match the apparent acceleration and cause the reverse of attractive gravity. The more complete equation for the deceleration parameter, q_0 , in Einstein's Theory of General Relativity is $q_0 = \Omega_M / 2 - \Omega_\Lambda$ where Ω_Λ , or Omega lambda, is the energy density of empty space. Einstein called the Λ component the cosmological constant, and in effect it's a repulsive gravity. A contemporary particle physicist would call it vacuum energy, that is "the zero point energy summing all possible particles in the vacuum", and then would complain that calculating it yields a nonsensically enormous answer.

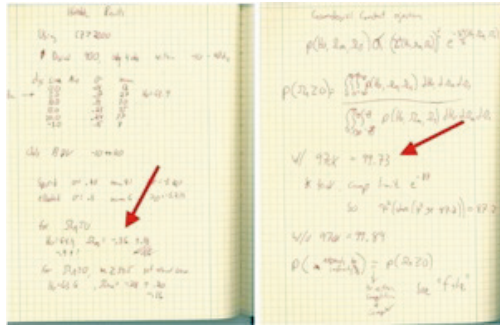


Figure 4. The author's lab notebook. Left: After performing various tests of the data I decided to analyze the expansion rate data in terms of the mass density for the Universe it suggested. The answer I got, -0.36 ± 0.18 made no sense, unless the Universe was accelerating! Right: A few days later I calculated the significance of the cosmological constant to 99.7% – 99.8% confidence no matter what the mass density. If the mass density was even the smallest conceivable amount, ~ 0.2 , the confidence rose to 4–5 sigma.

Since there is no such thing as negative mass, in something like confusion and desperation (in the same spirit Einstein had introduced it long ago), I re-introduced the famous cosmological constant to the equation and immediately found that its repulsive gravity could explain the acceleration I was seeing (Fig. 4). Its presence was significant in a statistical sense as well as an absolute sense – in fact the Universe was 70% in this form alone! This was remarkable and even my modest experience told me that such “discoveries” are usually the result of simple errors. So I spent a couple of weeks double checking my results but could find no errors. Then I thought hard about unexpected astrophysical sources of contamination to the observations. Here my thesis on recognizing and correcting for the effects of interstellar dust was helpful. Although dust in the high redshift galaxies could mimic the effect of acceleration and dim the light of distant supernovae, my use of MLCS and Mark Phillips’s use of his own dust-correcting algorithm made this unlikely. With growing confidence in the results, I first told Brian who spot-checked the final calculation. At the beginning of January 1998 he wrote me that he was getting the same answer. Later, the media quoted him saying: “My own reaction is somewhere between amazement and horror.”

I was also able to rule out a number of other concerns. One was that high redshift supernovae, born when the universe was younger, might be somehow different. But a comparison of the distances from nearby supernovae in old (elliptical) and young (spiral) galaxies limited the possible size of the difference to less than a third of the size of the acceleration signal. Moreover, my MLCS and Mark’s own algorithm gave the same result: the light curves and spectra of the nearby and high redshift supernovae were indistinguishable. Another concern was that a kind of exotic type of dust in the host galaxies could have had unusually large grains that would not redden the supernova light and could therefore go undetected. I calculated that the low scatter in the high redshift supernovae limited that kind of dust to an insignificant amount. A third concern was a well-known bias: astronomers tend to find

the brightest objects of a class preferentially. Brian did a simulation showing that this effect was insignificant too. We even calculated possible but unlikely explanations – the presence of a local void in the Universe, contamination of the sample by other supernova types and an effect known as gravitational lensing – and ruled out the significance of all of them.

Coincidentally, another exciting event was occurring in my life. Nancy Joy Schondorf and I were married on January 10, 1998, the best day of my life. We planned a honeymoon to Hawaii, after the next supernova observing run on the Big Island.

Meanwhile, the rest of the team did their own spot-checking of the results and more thinking and they too could not find any mistakes. It is amusing to look back at the emotions expressed about the result in the emails between team members over a couple of days in early January of 1998:

A. Filippenko, Berkeley, CA, 1/10/1998 10:11 am:

“Adam showed me fantastic plots before he left for his wedding. Our data imply a non-zero cosmological constant! Who knows? This might be the right answer.”

B. Leibundgut, Garching, Germany, 1/11/1998 4:19 am:

“Concerning a cosmological constant I’d like to ask Adam or anybody else in the group, if they feel prepared enough to defend the answer. There is no point in writing an article, if we are not very sure we are getting the right answer.”

B. Schmidt, Australia, 1/11/1998 7:13 pm:

“It is true that the new SNe say that [the cosmological constant] is greater than zero... How confident are we in this result? I find it very perplexing...”

M. Phillips, Chile, 1/12/1998 04:56 am:

“As serious and responsible scientists (ha!), we all know that it is FAR TOO EARLY to be reaching firm conclusions about the value of the cosmological constant...”

R. Kirshner, Santa Barbara, CA, 1/12/1998 10:18 am:

“I am worried. In your heart you know [the cosmological constant] is wrong, though your head tells you that you don’t care and you’re just reporting the observations... It would be silly to say ‘we MUST have a nonzero [cosmological constant]’ only to retract it next year.”

J. Tonry, Hawaii, 1/12/1998 11:40 am:

“...who remembers the detection of the magnetic monopole and other gaffes?...on the other hand, we should not be shy about getting our results out with appropriate disclaimers...”

A. Filippenko, 1/12/1998 12:02 pm:

“If we are wrong in the end, then so be it. But at least we ran in the race.”

A. Riess, Berkeley, CA, 1/12/1998 6:36 pm:

(Sent on the eve of our honeymoon over understandably icy stares from my wife!)

“The results are very surprising, shocking even. I have avoided telling anyone about them because I wanted to do some cross checks (I have) and I wanted to get further into writing the results up before [the other team] got wind of it... ..The data require a nonzero cosmological constant! Approach these results not with your heart or head but with your eyes. We are observers after all!”

A. Clocchiatti, Chile, 1/13/1998 7:30 am:

“If Einstein made a mistake with the cosmological constant... Why couldn't we?”

N. Suntzeff, Chile, 1/13/1998 1:47 pm:

“I really encourage you [Adam] to work your butt off on this. Everyone is right. We need to be careful and publish good stuff with enough discussion to make it believable to ourselves. ... If you are really sure that the [cosmological constant] is not zero – my god, get it out! Take responsibility as organizer and grunt. I mean this seriously – you probably never will have another scientific result that is more exciting come your way in your lifetime.”

Nick was absolutely right.

On February 20 we had a team-wide teleconference to decide whether to go forward with the paper I was writing, a draft of which I had earlier sent around. We decided to proceed. Somehow – I'm still not clear how – Jim Glanz of Science Magazine got wind of our result, interviewed many of us, and broke the story. Alex Filippenko on our team discussed the results at the UCLA Dark Matter conference in Los Angeles back to back with the Supernova Cosmology Project team, with both teams now claiming to see acceleration. I finished our paper with long nights and with crucial assistance from team members and on March 13, I submitted it: “Observational Evidence From Supernovae for an Accelerating Universe and a Cosmological Constant” (Fig. 5). It was accepted in May 6 [5]. The Supernova Cosmology Project published the same conclusion 9 months later [2]. Together the two team's conclusion became the “Breakthrough of the Year” in 1998 of Science Magazine (Fig. 5).

OBSERVATIONAL EVIDENCE FROM SUPERNOVAE FOR AN ACCELERATING UNIVERSE
 AND A COSMOLOGICAL CONSTANT

ADAM G. RISS,¹ ALISH V. FILIPPENKO,¹ PETER CHALLINOR,² ALESSANDRO CIACCIARETTI,¹ ALAN DEBBES,¹
 PETER M. GARAVITO,¹ RAY L. GRIFFITHS,¹ CHRIS J. HOLLAND,¹ SUZUKI IMAI,¹ ROBERT P. KILBINGER,¹
 B. LEBRONNQUET,¹ M. M. PHILLIPS,¹ DAVID RISS,¹ BRIAN F. SCHMIDT,^{1,3} ROBERT A. SCHMIDT,¹
 E. CHRIS SMITH,^{1,4} J. SPRINGEL,¹ CHRISTOPHER STUBBS,¹
 NICHOLAS B. SANDERS,¹ AND JOHN TONER^{1,5}
 Received 1998 March 11; revised 1998 May 8




Figure 5. The High-Z Team discovery paper [5] and the Hubble diagrams over-plotting the supernovae measured by both teams.

EXTRAORDINARY CLAIMS REQUIRE EXTRAORDINARY EVIDENCE

Our finding, that the Universe was currently accelerating, immediately suggested a profound conclusion: the universe’s cosmic energy budget is dominated by a type of smoothly distributed “dark energy”. Dark energy, a new component of the universe with negative pressure, causes the repulsive variety of gravity to dominate over matter’s attractive gravity. The cosmological constant would be dark energy’s poster child, the one with the pedigree, which came from Einstein. Was cosmic acceleration caused by dark energy the correct interpretation of the supernova data, or did we make a mistake in the interpretation of the supernova data? Past attempts to use distant objects (like the brightest cluster galaxies) to measure the change in the expansion rate had been foiled by evolution of the object’s intrinsic luminosity. So we needed a good test of cosmic speedup. As scientists often say, extraordinary claims require extraordinary evidence.

The situation circa 2000 was well summarized in the popular book by Donald Goldsmith “The Runaway Universe”, published that year. The essence of what we had learned from high-redshift supernovae in 1998 was that they were 20% fainter than we expected. We had assumed the usual rule that “fainter means farther” and had come to our conclusion of acceleration and dark energy. In 1999, Anthony Aguirre of the Harvard CfA wrote a series of bracing papers pointing out the supernova evidence could be mimicked by gray dust – mythical stuff akin to the Loch Ness Monster or Bigfoot of astronomy. His idea was that, in addition to the familiar, interstellar dust

grains that redden and dim supernova light, intergalactic space was full of dust “needles” one-tenth of a millimeter long and one-tenth that length in width. Aguirre calculated rigorously that such needles could absorb light equally well at all visible colors making them look “gray”. Gray dust would be nearly impossible to detect. Such a possibility seemed outlandish, but so of course did dark energy. Aguirre showed that at the time no astronomical observation ruled out the possibility of gray dust between galaxies in amounts sufficient to explain away the evidence for acceleration and dark energy. When trying to discriminate between improbable options Occam’s Razor becomes a blunt tool.

Fortunately a good test of competing hypotheses had just become available. If distant supernovae appeared faint not because the universe was accelerating but because of some astrophysical cause – like a pervasive screen of gray dust making supernovae appear dim, or past supernovae being born dimmer – then the extra dimming seen in distant supernovae would be expected to steadily increase as the distance to the supernovae increases. That is, if we were to look twice as far through a uniform screen of gray dust, we should see twice as much dimming. But if instead dark energy had begun increasing the size of the universe only recently, then the more distant supernovae should bear witness of the universe at an earlier stage in which dark energy was subordinate to matter. At that earlier stage, matter’s attractive gravity would have been decelerating the expansion and forming structures like galaxies and clusters. While acceleration increases distances and dims the light of supernovae, deceleration does the opposite, shortens distances and brightens light. The cumulative effects of both phases would likely look different than the effects of either gray dust or evolution if we measured supernovae beyond a redshift of one. As Goldsmith wrote in 2000:

“Gray dust and systematic differences can mimic the effects of a nonzero cosmological constant with high precision only so long as we examine distant supernovae within a relatively constricted range of distances... [I]f however astronomers observe distant supernovae over a much larger range...then cosmological models allow astronomers to disentangle all other effects from the crucial one: the acceleration produced by a nonzero cosmological constant... astronomers must therefore not rest on their current supernova assets... they must push their frontiers farther into space. Only then can they eliminate the possibility that gray dust has fooled them and show that the runaway universe deserves general acceptance.”

This was a call to arms we had already accepted!




Figure 6. Supernovae observed with Hubble at $z > 1$ confirm the result. By observing the transition from acceleration to deceleration (looking back in time) the High-Z Team papers could rule out simple astrophysical dimming (dust or evolution) as an alternative to acceleration with a mixed dark matter and dark energy Universe.

Unfortunately, finding supernovae so far back and away is difficult because they are so faint – a 60-watt light bulb held at a distance of 400,000 miles, twice the distance of the Moon. Five years of attempts by both teams had demonstrated that such faint supernovae were too hard to find reliably from the ground. So far, each team had found only one supernova each at $z = 1.2$, the Supernova Cosmology Project in 1998 and the High-Z Team in 1999. The Hubble Space Telescope could see more deeply than any ground-based observatory, but the small field of view of its main camera, Wide Field Planetary Camera 2 (WFPC2) made finding supernovae unlikely unless they occurred in the younger universe in dramatically higher numbers. In 1997 High-Z Team members Ron Gilliland and Mark Phillips made a reconnaissance of the high-redshift universe by re-observing the Hubble Deep Field with WFPC2 two years after it was first observed. They discovered two supernovae, one of which, SN 1997ff, appeared to be a high-redshift Type Ia because it was in an old, elliptical host. While interesting, this single observation did not allow measurements of the supernova’s light curve or colors to estimate its distance, and thus could not provide the test of the gray dust (or evolution) dimming. However, in 2001, after I had moved to STScI, I began looking for additional observations of SN 1997ff and encountered incredible serendipity!

The PI, Rodger Thompson, of the first near-infrared camera on Hubble, NICMOS, had used his camera to stare deeply into the Hubble Deep Field. By unbelievable good fortune, SN 1997ff had been observed for a couple of weeks shortly after its discovery in the very corner of the NICMOS field. In fact it was so close to the corner that it fell in and out of the frame through the intentional jiggling of the field (a technique for reducing the pixeliza-

tion of the image). I spent a number of months retrieving the color and light curve data needed to determine the supernova's distance and redshift. SN 1997ff was a Type Ia supernova at $z = 1.7$, by far the greatest distance ever observed. More importantly, it was about 60% brighter than we would have expected for the gray dust (or evolution) scenario. This was a good sign that the Universe had once been decelerating and that supernovae were faithfully tracing the history of cosmic expansion.

However, we did not want to hang such an important conclusion on a single, serendipitous supernova. We really wanted to find more. This became a possibility in 2001 when a new camera was due to be placed on Hubble, the Advanced Camera for Surveys (ACS), built by Holland Ford of Johns Hopkins University. ACS would improve Hubble's "discovery space" (area times depth) by an order of magnitude. I led a proposal in 2001 to use ACS to find and measure half-a-dozen Type Ia supernovae at z 's greater than 1 by "piggybacking" on a survey to collect images of distant galaxies, the GOODS Survey led by Mauro Giavalisco. The idea was that the GOODS Team would reimage a field composed of 15 ACS pointings every 45 days. Our new team, the Higher-Z Team (with members Tonry, Filippenko, Kirshner Leibundgut, Challis and Jha from the old High-Z Team and new members Casertano, Strolger, Ferguson, Giavalisco, Mobasher and Dickinson) would subtract these images to find fresh supernovae. Unlike previous ground-based surveys, we would already have a good idea of the type and redshift of the supernovae from color measurements obtained by the GOODS Team. If we found a supernova at z greater than 1 with the right colors to be a Type Ia supernova, we could interrupt Hubble's schedule to re-observe it with ACS and NICMOS. Our proposal was accepted, but that did not guarantee our plan. The NASA crew of the Space Shuttle flight STS 109 first had to install ACS! One of the greatest privileges I have ever had as a scientist was to present the science case for the observations to the astronauts before their trip to Hubble. These folks (John Grunsfeld, Mike Massimino, Scott Altman, Jim Newman, Duane Carey, Nancy Curry, Richard Linnehan) were real life heroes who risked their lives to refurbish Hubble. Their successful mission in March 2002 brought Hubble again to the peak of its capabilities and allowed us to undertake the proposed investigation.

Compared to our searches using ground-based telescopes, using Hubble to find and follow supernovae had advantages and disadvantages. On the positive side: the weather was always good in space, the image quality always razor sharp and moonlight never got in the way of our observations. On the negative side: the Hubble doesn't have the protection of Earth's atmosphere and magnetic fields that ground-based telescopes do, so cosmic rays strike Hubble about 100 times more frequently. A cosmic ray strike can look quite similar to a supernova: they both appear as a new source not in the prior image. Because cosmic rays affect only about 2% of the pixels in a Hubble image, and because the odds of consecutive strikes is small, most astronomers can distinguish a real astronomical source from a cosmic ray strike by obtaining a second image. However, because a supernova could appear on

any pixel of the ACS's 16 million pixels, we determined we needed 4 consecutive images to rule out the fluke of even 3 consecutive strikes. Another challenge was imposed by the way Hubble is scheduled: observing schedules are uploaded to the telescope weekly. We could alter Hubble's schedule only on a Tuesday by noon, in advance of the next schedule upload. It was ironic that a supernova's light could travel for 9 billion years but needed to reach Hubble shortly before Tuesday to be of any use! We found that if you look for a supernova only on the weekends, you could naturally meet this requirement.

Our year-long program from 2002 to 2003 to measure supernovae on Hubble was highly successful. We found 6 Type Ia supernovae at redshifts over 1.25 [6]. They allowed us to rule out gray dust and evolution and to clearly determine that the universe was decelerating before it began accelerating (Fig. 6). In physics, a change in the value or sign of deceleration (which results from a change in force) is called a jerk. So when we announced this result in a conference in 2003 we described the change from the universe's past deceleration to its present acceleration as evidence of recent "cosmic jerk". I saw Dennis Overbye from the New York Times in the front row as I discussed the result and asked him to please not run a picture of me next to a headline "Cosmic Jerk Discovered" – to no avail.

Over the next two years we continued using Hubble to collect Type Ia supernovae at redshifts greater than 1 and by 2007 we published a sample of 23. Not only did this data appear to confirm that supernovae were faithfully communicating the history of cosmic expansion; they could also be used to help determine if the properties of dark energy had been changing over the last 10 billion years. So far those properties seem unchanged, adding credence to Einstein's cosmological constant. But to quote Edwin Hubble from "The Realm of the Nebulae":

"With increasing distance, our knowledge fades, and fades rapidly. Eventually, we reach the dim boundary – the utmost limits of our telescopes. There, we measure shadows, and we search among ghostly errors of measurement for landmarks that are scarcely more substantial. The search will continue. Not until the empirical resources are exhausted, need we pass on to the dreamy realms of speculation."

Although Hubble did not know about acceleration, his description fits our work. Today we seek to learn more about the cause of cosmic acceleration by refining a wide range of cosmological measurements. Many feel this challenge is one of the greatest for cosmology and fundamental physics and I agree. Since about 2003, WMAP data, measurements of Baryon Acoustic Oscillations (BAO), large-scale structure, weak lensing and the integrated Sachs Wolfe effect also give strong evidence for dark energy in ways independent from supernovae. The BAO technique alone now provides independent confirmation of the recent acceleration. For my own work I have been focusing on improving the measurements of the present expansion rate, also known as the Hubble constant, because knowing this to percent level precision would significantly aid the present quest. Already we have improved the determination of the Hubble constant by a factor of

3 to 3.5%. This, combined with the WMAP measurements of the cosmic microwave background are good enough to measure the properties of dark energy about as well as high redshift supernovae – to around 10% – and to provide another independent check of the results. An ambitious goal will be to achieve about a 1% measurement. I don't expect to win another Nobel Prize from this new work, but it should keep me out of trouble for a while.

I want to close by expressing my deep gratitude to the amazing people and facilities I have been honored to work with. To my colleagues on the High-Z and Higher-Z Teams, it has been a pleasure to share this scientific adventure with you. I thank those who built the instruments and facilities at CTIO to allow us to find supernovae (Bernstein and Tyson) and the Calan/Tololo Survey for helping ground and inspire the subsequent work. I thank the men and women who have helped make the Hubble Space Telescope the premier scientific instrument of our time and the astronauts who risked their lives to maintain it. Most of all I thank my family, my wife Nancy and my children for help keeping me sane and for reminding me that the world down on Earth is at least as interesting as the universe around it.

REFERENCES

1. P. M. Garnavich *et al.*, *Astrophysical Journal*, **493**, L53 (1998).
2. S. Perlmutter, *Astrophysical Journal*, **517**, 565 (1999).
3. S. Perlmutter *et al.*, *Nature*, **392**, 311 (1998).
4. S. Perlmutter *et al.*, *Astrophysical Journal*, **483**, 565 (1997).
5. A. Riess *et al.*, *Astronomical Journal*, **116**, 1009 (1998).
- A. Riess *et al.*, *Astrophysical Journal*, **607**, 665 (2004).