Cosmological Observations

Supernova surveys

There are different types of supernovae. **Type Ia** is the most luminous type.

Their observations allow us to estimate a cosmological function: the luminosity distance.

They are the most explored probe of the homogeneous Universe.

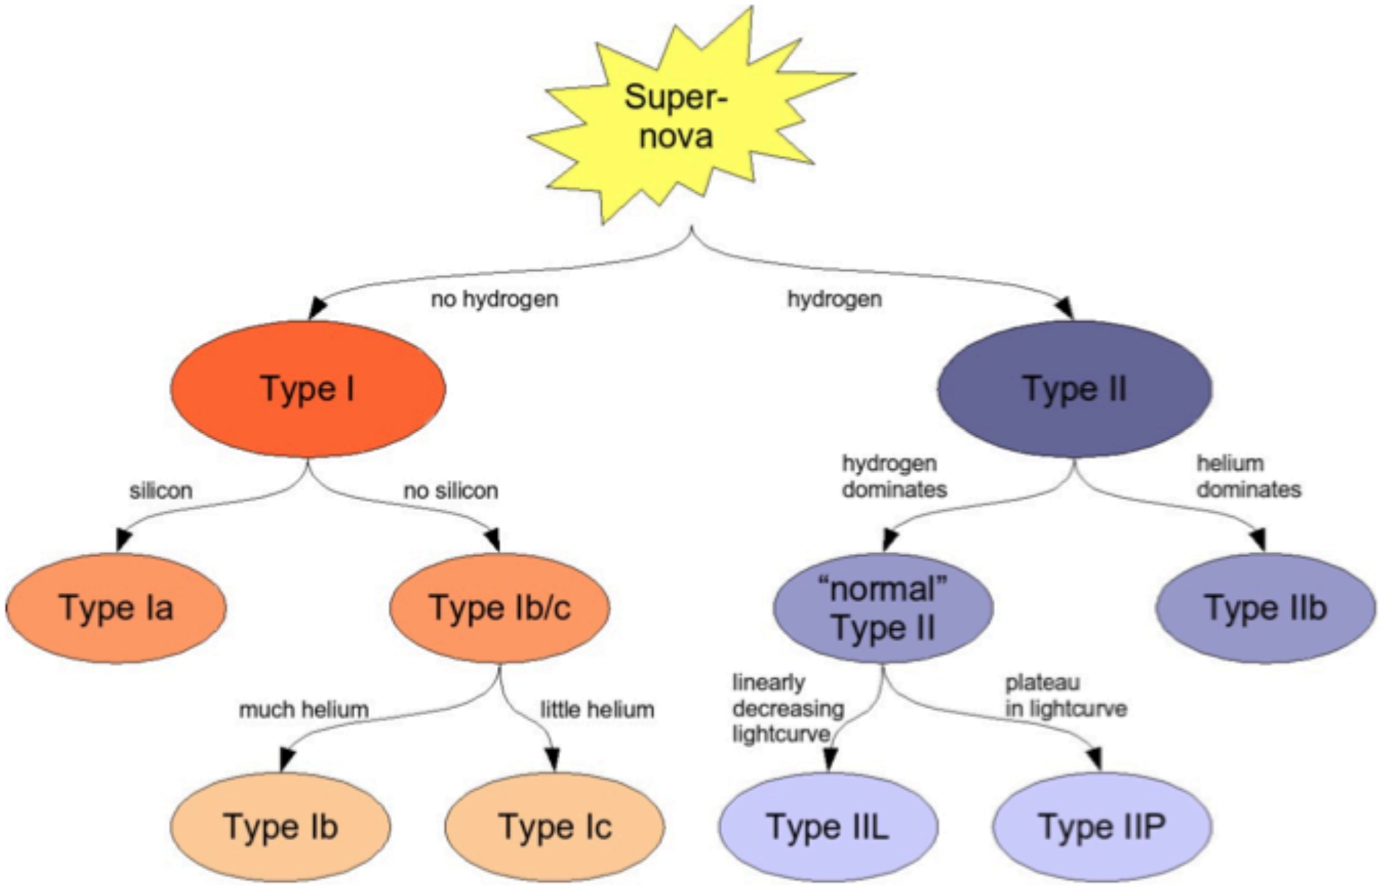

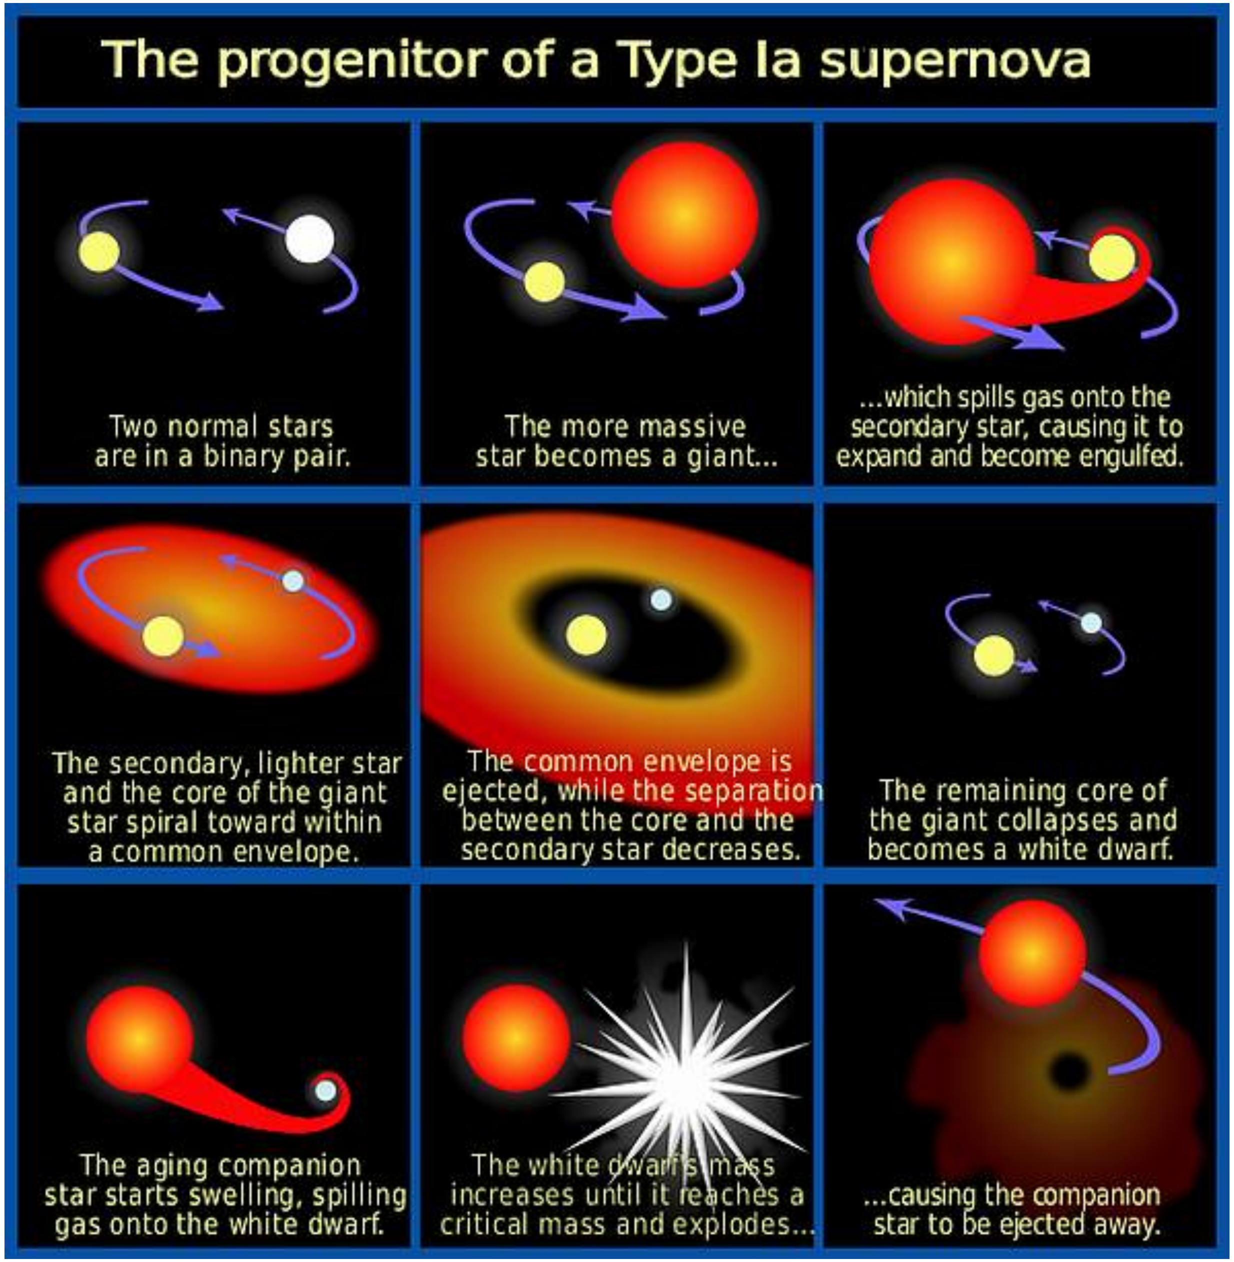

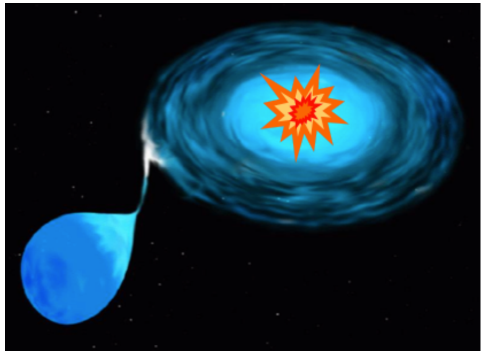

SNe Ia explosion formed by accretion onto a 1.4 M_Sun white dwarf (Chandrasekhar limit : the critical mass for a stable white dwarf)

WD no longer sustained by electron degeneracy pressure.

Becoming unstable, the white dwarf can either collapse further → forming a neutron star or a black hole

or it can explode → material in the star starts nuclear fusion producing metals, in particular a large amount of ^{56}Ni that decays through ^{56}Co to ^{56}Fe, **producing a very luminous explosion → M ~ -19.**

Other SNe types have different origins and do not require a binary system:

Sne Ib/Ic: collapse of stars with initial masses 8 - 30 M_Sun

Sne II: collapse of stars with initial masses > 30 M_Sun

Each type has a characteristic light curve

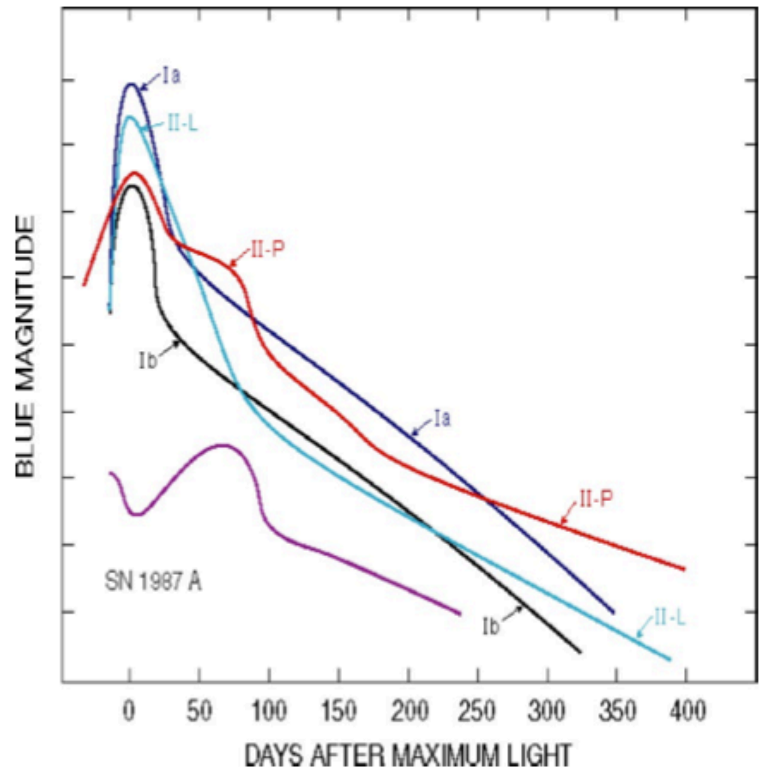

step 5 - **Constructing an estimator of the cosmological function**

The luminosity distance can be measured through the flux-luminosity relation. This relation is usually written as a difference of magnitudes, known as the **distance modulus.**

Flux is traditionally expressed in magnitudes, m.

Hipparcos separated the magnitude of visible stars in 6 qualitative classes. Since the eye sensitivity is roughly logarithmic → it turns out that apparently equal intervals are in reality equal ratios.

Stars of magnitude 5 have ~ 1% of the flux of stars of magnitude 1. This led to the modern (XIX century) definition of magnitude:

$$\frac{F_1}{F_2} = 100^{(m_2 - m_1)/5} \quad \rightarrow \quad m_1 - m_2 = -2.5 \log_{10}\left(\frac{F_1}{F_2}\right)$$

This is a **relative scale**. To define an absolute scale, the star Vega was chosen as reference → the flux of Vega corresponds to magnitude m=0:

$$m = -2.5 \log_{10} (F / F_{Vega})$$ apparent magnitude

The apparent magnitude depends on the distance, since $F \sim L / D^2$
(Notice that for large distances on an expanding spacetime this distance is the luminosity distance, since it needs to take into account the 'dilution of luminosity').

This leads to the definition of a distance-independent magnitude: the absolute magnitude, M: *the apparent magnitude an object would have if placed at a distance of 10 parsec.*

SNIa have M ~ -19 → they are very bright.

We can compare their magnitudes with the apparent magnitudes of the full moon (m = -12) or of the Sun (m = -26.7) → knowing the distance to the Sun, we find that $L_{SN} = 3 \times 10^9$ L_Sun

This implies that a SN Ia would appear as bright as the Sun if it was placed at a distance of D_L =31.6 pc ~ 100 lyr (this is the distance to some of the well-known night-sky stars in the Milky Way).

Luckily for us SN Ia are rare events - only 5 records in our galaxy in the past 1000 years → 1006, 1054, 1181, 1572, 1604

If the SN occurs in a nearby galaxy from the local group (D_L ~15 Mpc) → its apparent magnitude is m = 12 (this was the case of SN1987, a SN II that appeared in the LMC in 1987).

If the SN occurs at a distant galaxy (z = 1) → D_L ~ 6.7 Gpc (concordance model) → 450 x D_L (local galaxy) → its apparent magnitude is m = 25

In both cases, the apparent magnitude of the SN Ia is similar to the apparent magnitude of the whole galaxy

The difference between apparent and absolute magnitudes is a L/F ratio and **it is thus a direct measure of luminosity distance**.
This difference is known as the distance modulus:

μ = **m - M** = $-2.5 \log_{10} (F/F_{Vega}) + 2.5 \log_{10} (F_{10}/F_{Vega})$

$= -2.5 \log_{10} (F/F_{10})$

$= -2.5 \log_{10} [(L / D_L^2) / (L/10^2)]$

$= -5 \log_{10} (10/D_L)$

$= \mathbf{5\log_{10} (D_L) +25}$ (for D_L in Mpc, i.e.,we used 10 pc = 10^{-5} Mpc)

The goal of SN Ia cosmological surveys is to measure μ from the data, and then fit the cosmological predictions of μ (i.e. D_L) to the measured μ.

For this, we need **to define an estimator of μ from the observed quantities. An optimal estimator is one that gets an accurate (i.e, without bias) and precise (i.e. with high signal-to-noise ratio) measurement of μ.**

From images to the distance modulus estimator

i) Construct a survey to find SN Ia candidates

-- image the same part of the sky repeatedly looking for them
(monitoring required every few days)
-- subtract current image from image earlier to time to look for time variable event
(this can be a challenge with ground-based telescopes since smoothing from turbulent sky conditions may change from night to night)

In nearby galaxies

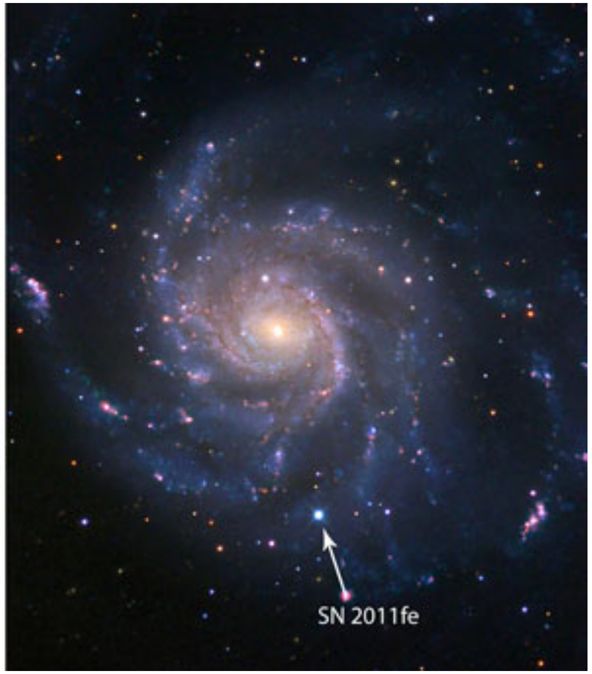

In distant galaxies

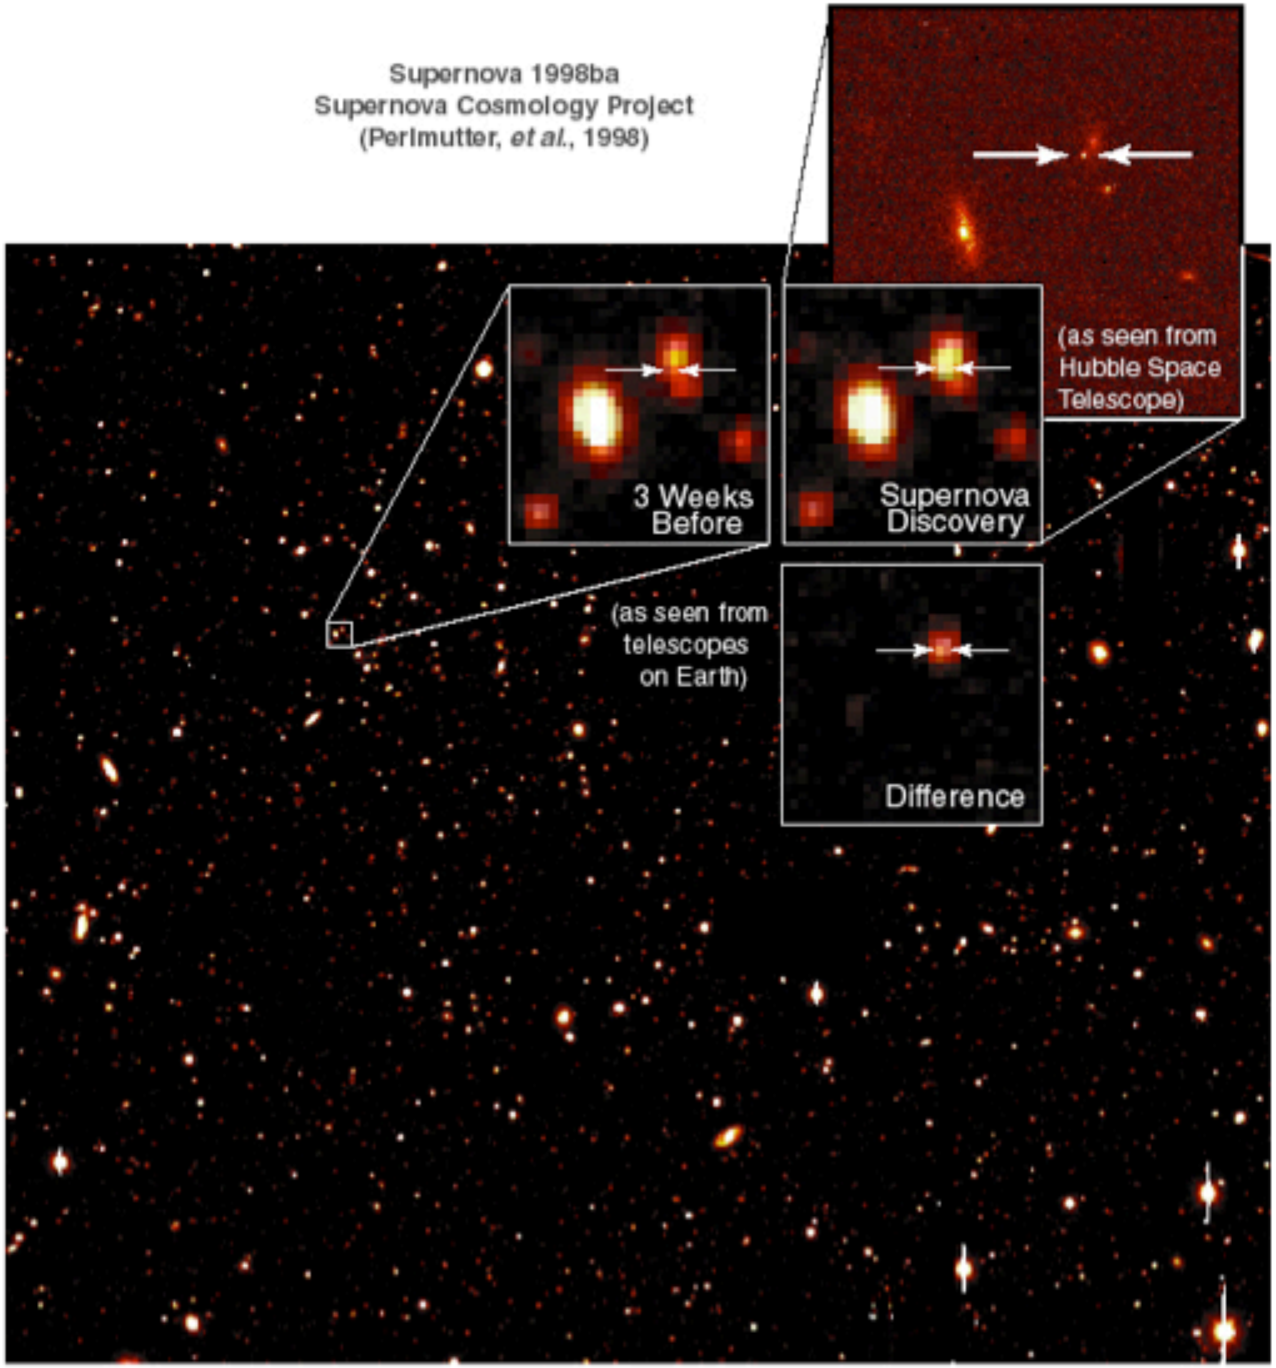

ii) Identify the SNe of type Ia and measure their redshifts

- Determine the SN type from its **spectrum**

(this is challenging for distant Sne since they have faint magnitudes ~24)

To distinguish SNe Ia from Sne Ib or Ic, need to check for the presence of Silicon lines (at ~6500 Angstroms rest-frame) → need spectroscopy in near-IR

- Use the spectrum also to **measure the redshift**
- Check if the **colors** of the source are consistent with being a SN Ia
- Check if the **light-curve** is consistent with being a SN Ia

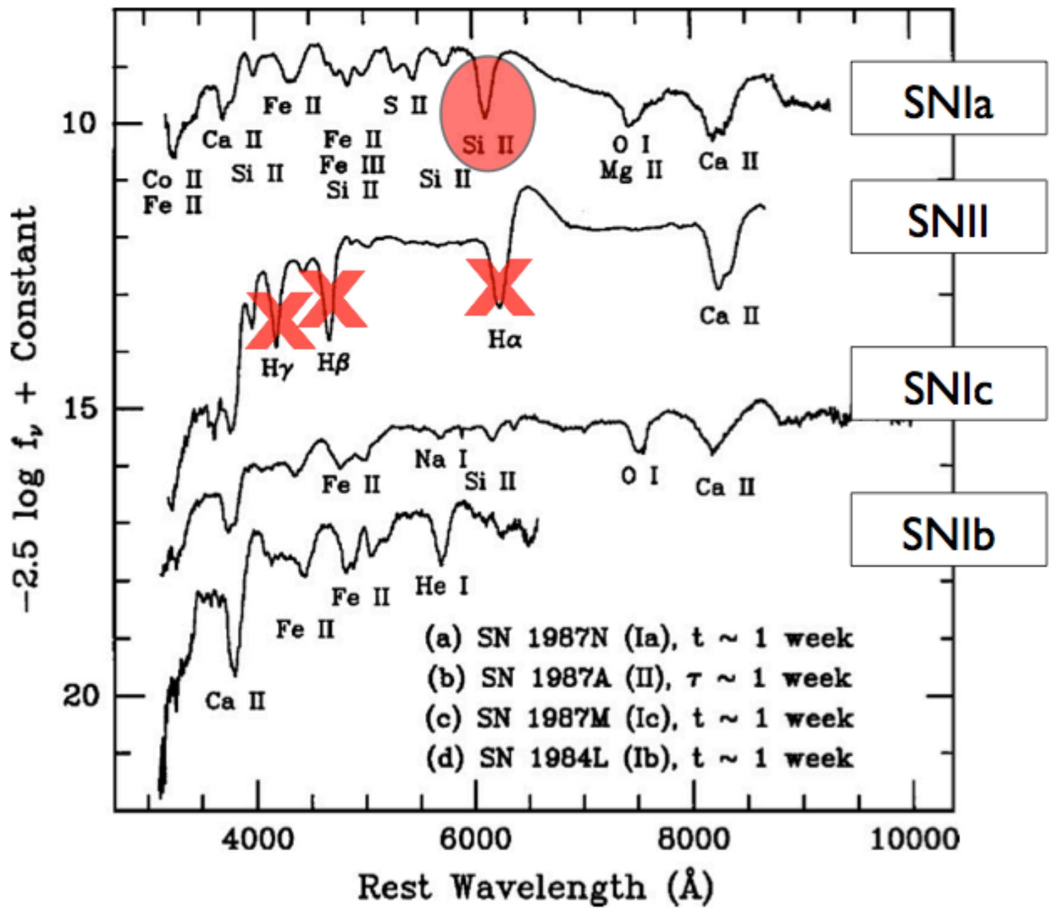

iii) Measure the SN flux and define the estimator of μ

Follow the SN event to get its peak amplitude and also get as much points in the light-curve as possible.

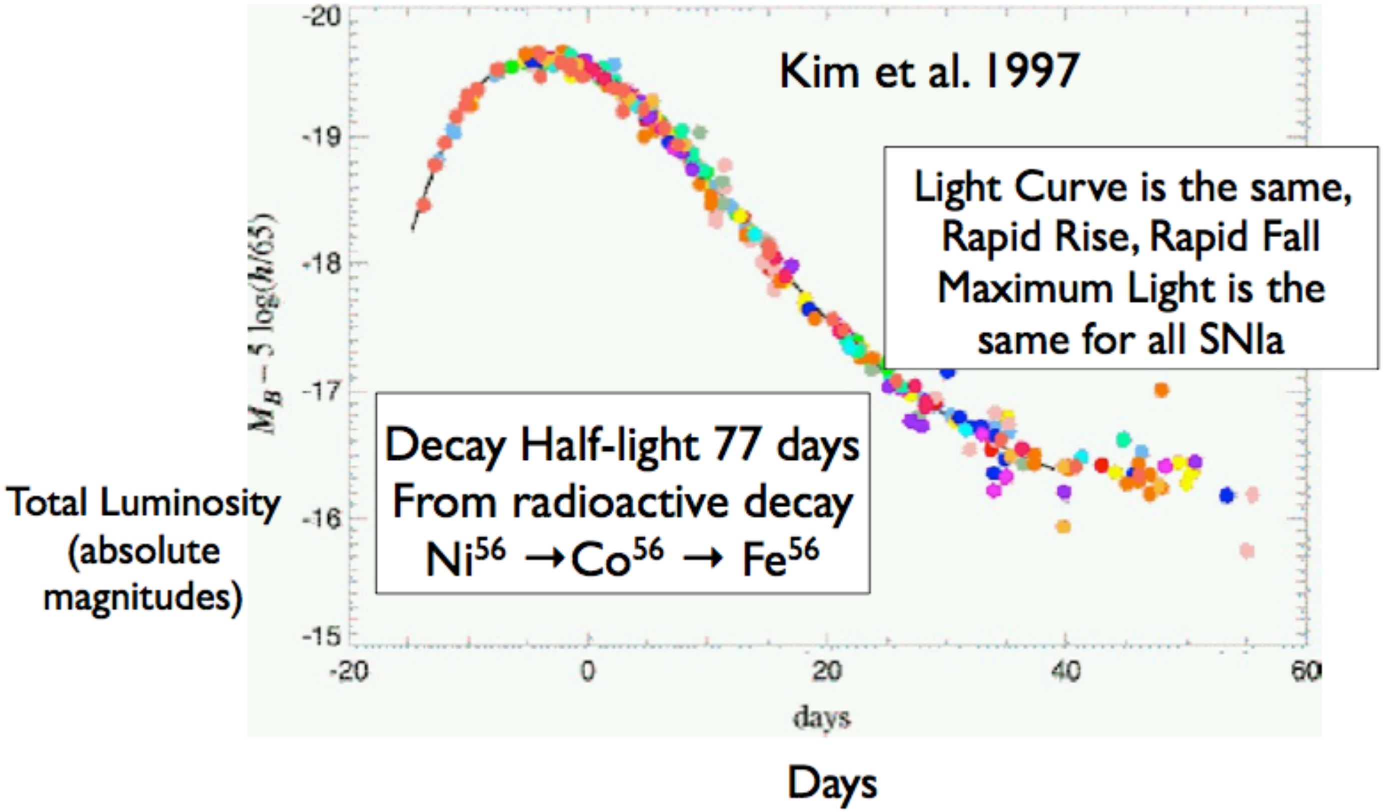

Since all SN Ia are formed in the same way, they are in first approximation assumed to be standard candles, i.e., all SN Ia would have the same light-curve: same absolute magnitude at the peak and same duration in time.

In this assumption, measuring the flux at the peak gives directly the distance, since the absolute magnitude is known:

The universal absolute magnitude is computed previously from observations of SNe in galaxies at known distances (from a **distance ladder** method). For this, we need to find galaxies that have simultaneously a SN and a Cepheid (or other calibrator). There are only 19 such Sne known up to now (so the uncertainty on M is large).

For those **calibration galaxies**, measuring the SN peak flux and knowing the distance → obtain the universal absolute magnitude (that would remain a fixed quantity for all subsequent observations of SN in galaxies with unknown distances, if M is universal).

→ with this information, the distance modulus of a given galaxy can be accurately computed from its definition: $\mu_i = m_i - M$

where the **magnitude m** is the measurement made from the flux of the galaxy obtained in this survey (with a certain uncertainty) and the **magnitude M** is a constant value known a priori from a distance ladder calibration previously made (also with an uncertainty).

So this expression $\hat{\mu}_i = m_i - M$ seems a good choice to be the estimator of the distance modulus of a galaxy from the data.

iv) Evaluate biases in the estimator

a) **Astrophysical effects**

In general, there are additional astrophysical effects that also contribute to the cosmological estimator. For example, there may be other effects contributing to the magnitude such that m-M is not just the cosmological contribution.

In SN surveys, an important effect to consider is the fact that **in reality the SNe are not standard candles,** i.e., that approximation is very weak**.**

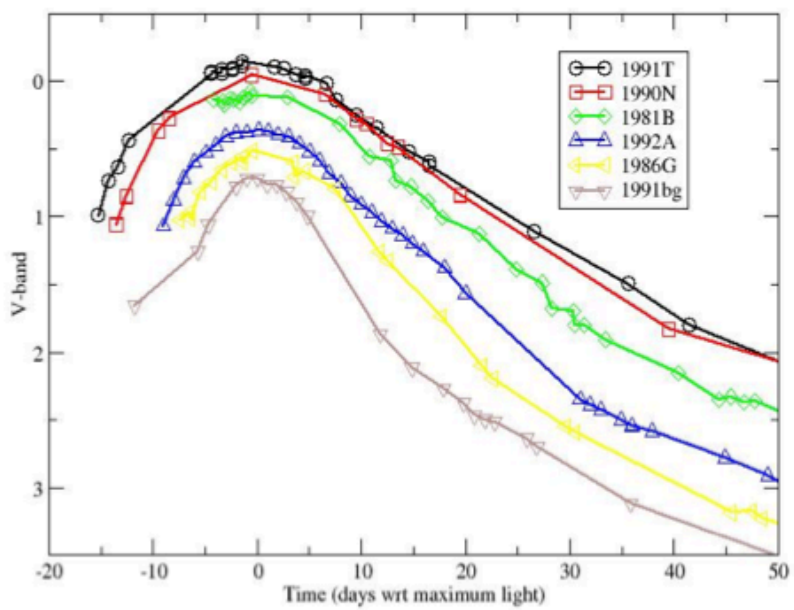

Each SN has its own light curve.

They are not universal after all → the value of M obtained for the control SN is not valid for all → each SN may have its own value of M.

In this way, if we apply the estimator **$\hat{\mu}_i$ = m_i - M,** we will get a wrong value for the distance modulus of the galaxy - a **biased** value, i.e., there will be a **systematic error** on all estimated distance modulus.

A more accurate (**unbiased**) estimator would be:

$$\hat{\mu}_i = m_i - M_i$$

Does this mean that we need to find the absolute magnitude of each SN? This is not possible. We only have very few SNe where M can be measured (the control ones).

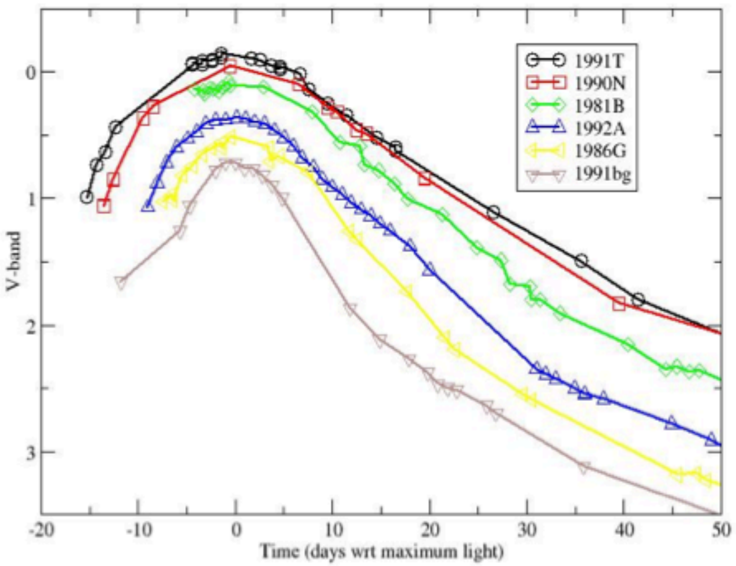

Fortunately, the SNe are not completely different. There is a correlation in the behaviour of their light-curves:

There seems to be a **shape- luminosity** relation in the light-curves → **the luminosities are indeed different but the peak amplitude depends on the decay time →** the brighter ones are systematically slower.

The existence of this (empirical) relation means that

the SNe are standardizable → their M values can be related with the standard "universal" M value.

So, if we apply a stretch factor to the light-curve, the peak will go up and reach the "universal light-curve" → the one that corresponds to the (**standard**) universal luminosity.

The question now is: if we would stretch the observed light-curve of a SN by a certain factor, this would correspond to change the luminosity by how much?

We do not know this! But the important point is that the impact of the stretching of the light-curve on the luminosity is the same for all SNe.

So we can **model** this effect with an arbitrary function of an arbitrary amplitude, and apply the model consistently to all SNe.

The **response of magnitude to stretch** is usually taken to be **linear**, i.e.,

$$\Delta M_i = M - M_i = \alpha (s_i - 1)$$

So, a galaxy that requires a stretch s for its light-curve to become identical to the standard one, has an absolute magnitude that differs ΔM_i from the universal value.
(s = 1 means no stretch, it is a SN already standard)

This means that the estimator of the distance modulus of a galaxy is not:

$\hat{\mu}_i = m_i - M$ → this gives a biased result;

it is also not:

$\hat{\mu}_i = m_i - M_i$ → this gives an unbiased result, but it is impossible to measure, so it cannot be an estimator;

but it is:

$\hat{\mu}_i = m_i - M + \alpha (s_i - 1)$ → this gives an unbiased result.

Notice that this method introduces one unknown parameter in the analysis:

the **stretch response parameter α**

Its value is unknown. How can we find out its value?

option 1:

If the luminosities (i.e., the absolute magnitudes) of some of the stretched SNe were known, that information could be used to calibrate the relation (i.e., to find the value of α).

However, this is not the case.

option 2:

An alternative would be to predict the absolute magnitudes of SNe from **astrophysical theory**.

However this cannot be done with enough precision, and it depends on many assumptions and astrophysical modelling (which would introduce additional astrophysical parameters, and would just move the problem to another place).

option 3:

The usual approach is to leave α as an additional free parameter of the model (introduced to model an extra effect), to be treated in the same way as the cosmological parameters.

This type of parameter is known as a nuisance parameter.

Despite the name, nuisance parameters are very important → without them the cosmological analysis would be biased → they model (and correct) a systematic effect.

In addition to the shape-luminosity relation, the light-curves also show a **color-luminosity relation,**

i.e., two SNe of different colours, if stretched by the same amount will not reach the same peak amplitude.

The bluer ones (the ones with higher amplitudes on the blue filters compared to redder filters) have larger luminosities.

So, for each measured SN, after stretching the light curve by a factor s, the amplitude needs to be further increased for it to match the standard light curve.

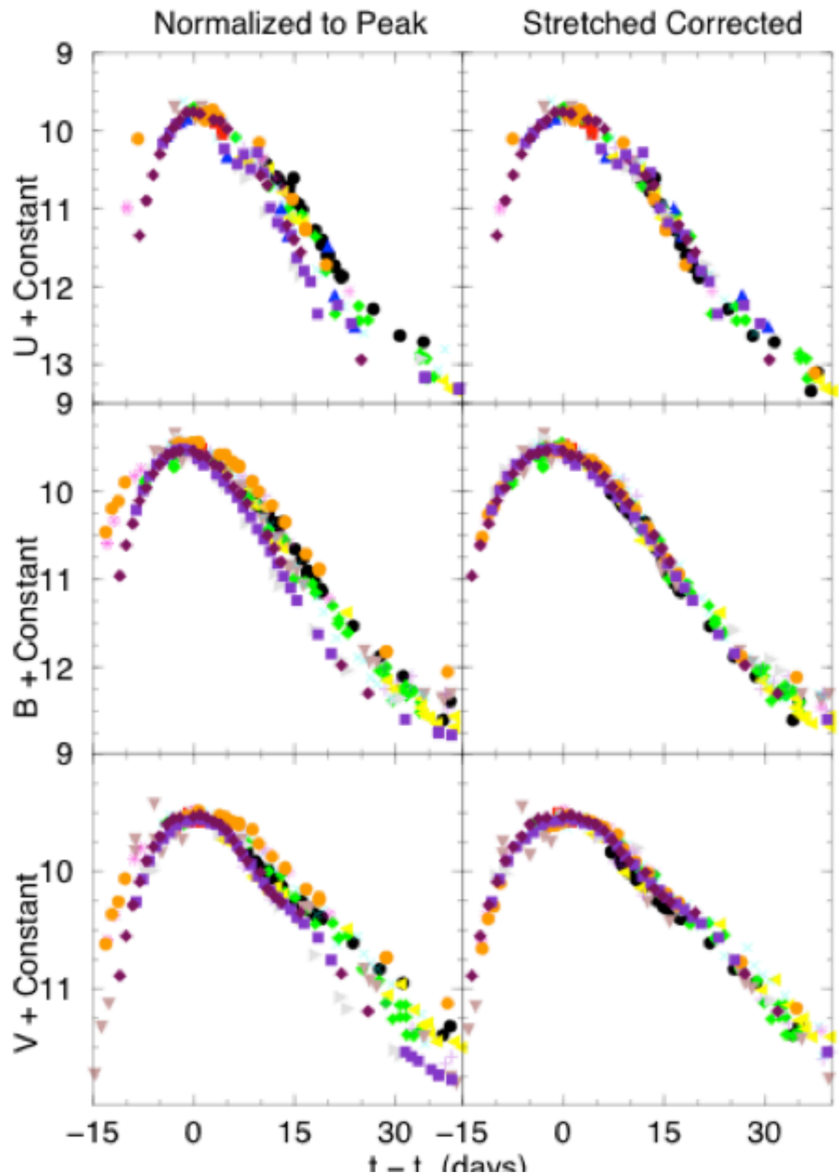

Moreover, by comparing the fluxes of one SN on different filters, sometimes it is found that the flux ratios are too different from the standard case to be explained just by this intrinsic color variation → differences also arise because of dust extinction in the galaxy host, which does not affect equally all bands.

So the amplitude of the light-curve is increased by a color factor, c, (that accounts both for color and for dust) in order to match the standard amplitude.

Once again, we need to find out what is the impact of this shift in the absolute magnitude.

This introduces another nuisance parameter : **the response factor β, that models a linear response:**

$$\Delta M_i = \beta c_i$$

The **estimator** of the distance modulus then becomes:

$$\hat{\mu}_i = m_i - M + \alpha\,(s_i - 1) - \beta\, c_i$$

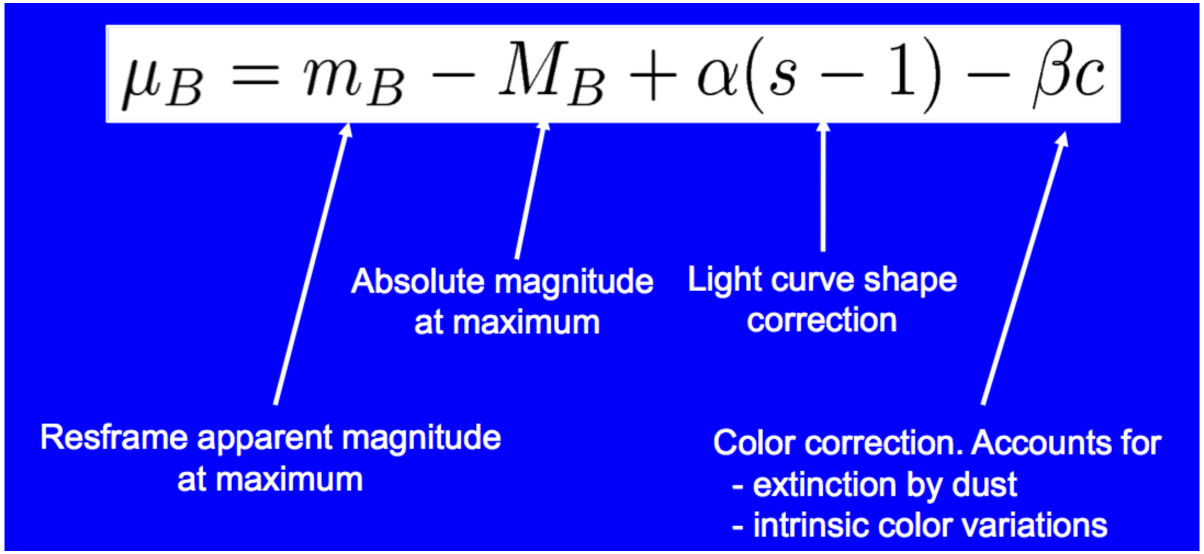

The estimator involves 3 measured quantities for each SN:

m_i : observed magnitude of each SN → i.e. the measured flux

s_i, c_i : stretch and color factors of each SN → measured from the light curves

Notice that the absolute magnitude of a galaxy with no stretch (s=1) and no color correction (c=0) is the reference value M, while the absolute magnitude of a galaxy with a corrected light curve is $M - \alpha (s_i - 1) + \beta c_i$

The estimator also involves 3 (global) model parameters:

M : reference absolute magnitude, known from the calibrators with an uncertainty → it may treated as a free parameter with a prior

α : response of magnitude to stretch → a free parameter

β: response of magnitude to color and dust → a free parameter

The nuisance parameters are not necessarily global across the whole sample of SN Ia. They may be different for SNe at different redshifts - evolution - or for SNe in different host galaxies - environment -

In this case, the analysis needs to be done with the SN separated in various sub-samples, with different parameter values in each.

b) **Instrumental effects**

The observational procedures in general introduce additional biases. In high-z measurements, **the use of an observing filter bias the flux-luminosity relation.**

Remember that the flux-luminosity relation, relates observed flux (in the observer's frame o) with the corresponding luminosity. But that luminosity is not the equal to the intrinsic luminosity (in the source rest-frame e), due to the expansion of the Universe (causing energy dilution and time dilation).

$$E_o = \frac{E_e}{1+z} \qquad \frac{\Delta t_o}{a(t_o)} = \frac{\Delta t_e}{a(t_e)}$$

This led to an additional factor of $(1+z)^2$ that is absorbed in the distance, leading to the definition of luminosity distance.

$$L_o = \frac{L}{(1+z)^2}$$

However, this reasoning is only valid when we consider the total luminosity emitted by the source (and the total measured flux), the so-called bolometric quantities.

In practice, the fluxes are measured within a filter → **only a part of the energy is detected** by the observer, i.e., an integral over the field spectral domain of a **flux density** per frequency bin ($F(v_0)\, dv_0$).

This flux corresponds to an emission of a different frequency in the galaxy's rest-frame ($L(v_e)\, dv_e$) due to the expansion redshift.

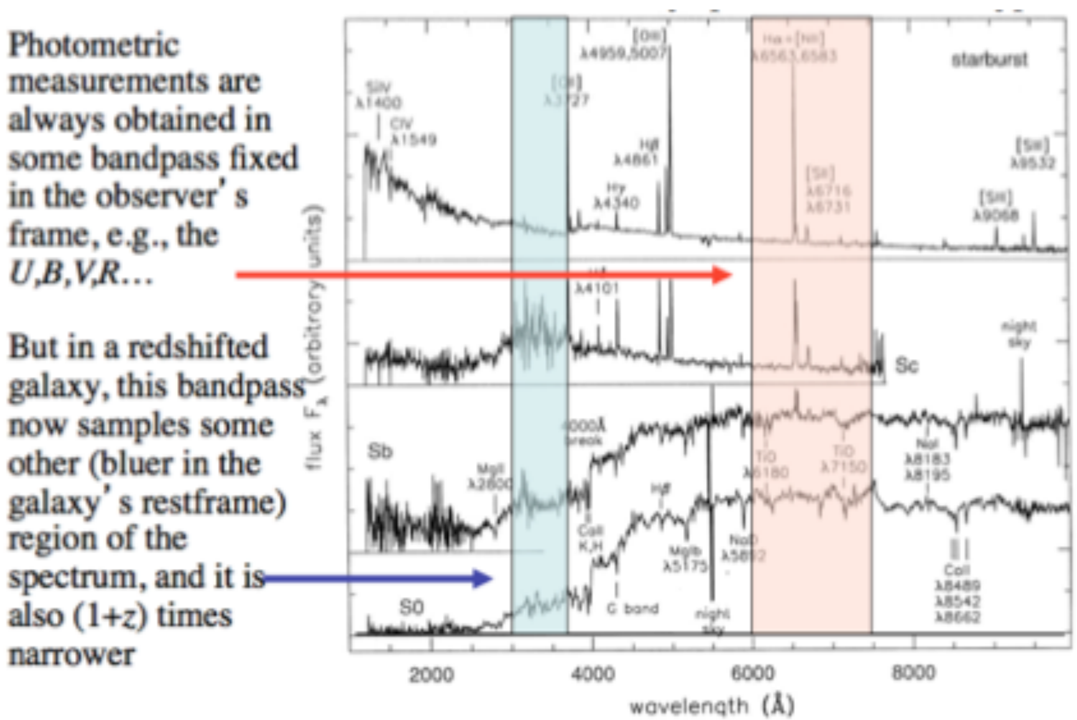

The flux on a detected frequency corresponds to an emission from a higher frequency

$$\nu_e = [1+z]\, \nu_o$$

The filter broadens at the observer's frame → the "unit" dv_0 shrinks

$$dv_e = (1+z)\, dv_0$$

The detected flux is an integration of the flux density and thus of the spectral energy density (SED) on a portion of the spectrum (i.e. within a filter, that provides a weighting function: the filter throughput),

This implies that the luminosity distance needs to be redefined, i.e., the $(1+z)^2$ factor valid for the bolometric case needs to be removed and replaced by a frequency-dependent function.

In practice, what is done is to keep the same luminosity distance and apply instead a correction to the flux-luminosity relation, i.e. to **unbias the luminosity distance**. This effectively introduces an extra term in the definition of the distance modulus, named the **K-correction**.

The K-correction is thus the difference between the observed magnitude for a source at redshift z *(for a specific filter and a specific source SED)* **and the magnitude that would be observed if there was no expansion.**

How can we compute the K-correction for a given SED and filter?

Let us consider the flux measured in a filter centered in ν_0, per unit frequency $d\nu_0$ (i.e., the **flux density**)

It was emitted in the rest-frame of the source as frequencies centered on a redshift ν_e , per unit frequency $d\nu_e$.

The standard bolometric flux-luminosity relation (already including the $(1+z)^2$ factor in the definition of luminosity distance) is:

$$F(\nu_0)\, d\nu_0 = \frac{L(\nu_e)\, d\nu_e}{4\pi D_L^2}$$

We now want to write the rest-frame luminosity in terms of the observed-frame luminosity. **For this, we need to introduce the two effects: frequency shift and filter broadening:**

$$F(\nu_0)\, d\nu_0 = \frac{L(\nu_0(1+z))\,(1+z)\, d\nu_0}{4\pi D_L^2}$$

The flux observed is the integral of this **flux density**, within the filter (or band).

$$\int_{band} F(\nu_o)\, d\nu_o$$

Using the flux-luminosity relation this is:

$$\int_{band} F(\nu_o)\, d\nu_o = \frac{\int L(\nu_o(1+z))\,(1+z)\, d\nu_o}{4\pi D_L^2}$$

If the luminosity density (the SED) is constant within the filter (e.g. in the case of a narrow filter), then the numerator just contains a $(1+z)^2$ factor that cancels out the one implicit in D_L^2 and we recover the standard relations

$$F_o = \frac{L_o}{4\pi D_C^2} \qquad F_o = \frac{L_e}{4\pi D_L^2}$$

In the more relevant case of a broad filter, we can multiply and divide the expression for the observed flux by the integral of L_o and write:

$$\int_{band} F(\nu_o)\, d\nu_o = \frac{\int L(\nu_o(1+z))\,(1+z)\, d\nu_o}{4\pi D_L^2}$$

$$= \frac{\int_{band} L(\nu_o)\, d\nu_o}{4\pi D_L^2} \; \frac{\int_{band} L(\nu_o(1+z))\,(1+z)\, d\nu_o}{\int_{band} L(\nu_o)\, d\nu_o}$$

↓ the previous result ↘ the correction =

The result is then: the correction due to the use of a filter is a factor (1+z), times the ratio f_ν between the integrated luminosities on the emission and observed bands, where

$$f_\nu = \frac{\int_{\nu_1}^{\nu_2} L[\nu(1+z)]d\nu}{\int_{\nu_1}^{\nu_2} L(\nu)d\nu}$$

This means that when we only measure part of the flux (using a filter) **the correct flux-luminosity relation is no longer**

F = L / (4π D_L^2) (the bolometric relation)

but

$$F = \frac{L}{4\pi D_L^2} f_\nu (1+z)$$

We can also write the flux-magnitude relation In terms of magnitude difference:

We start with the definition: $m - M = -2.5\log_{10}\left(\frac{F}{F_*}\right) + 2.5\log_{10}\left(\frac{F_{10}}{F_*}\right)$

where now $F = \frac{L}{4\pi D_L^2} f_\nu (1+z)$

For D = 10pc there is no correction (very low redshift) → z=0 and f_v = 1

$$\Rightarrow \quad m - M = -2.5 \log_{10} \left(\frac{k\, f_\nu (1+z)}{4\pi D_L^2} \, \frac{(10^{-5}\,\mathrm{Mpc})^2\, 4\pi}{k} \right) = 2.5 \log_{10} \left(\frac{D_L^2\,(10^5)^2}{(\sqrt{f_\nu})^2\,(1+z)} \right)$$

$$\Rightarrow \quad m - M = 5 \log_{10} D_L + 25 + 5 \log_{10} \left(\frac{1}{\sqrt{f}\,\sqrt{1+z}} \right)$$

$$\Rightarrow \quad m - M = 5 \log_{10} (D_L) + 25 - 2.5 \log_{10} \left(f_\nu (1+z) \right)$$

$$\boxed{m - M = 5 \log_{10}(D_L) + 25 - 2.5 \log_{10}(f_\nu) - 2.5 \log_{10}(1+z)}$$

The correction in the distance modulus is thus a constant shift in magnitude:

$$K_v(z) = 2.5 \log_{10} \left(f_v \, (1+z) \right)$$

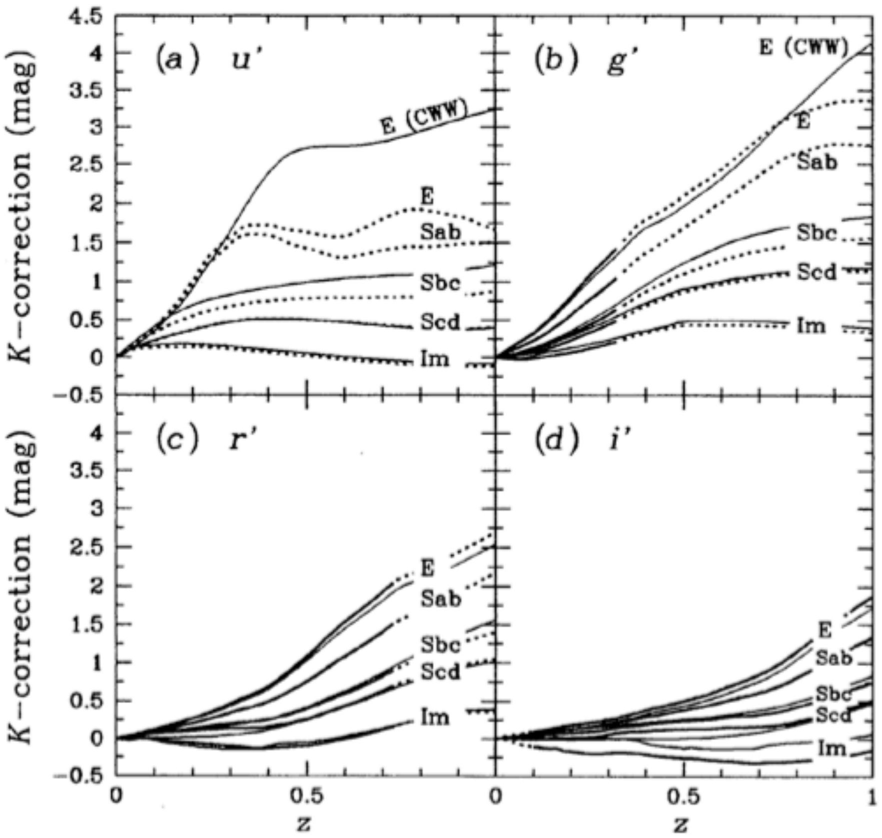

K-correction (in magnitude) for various filters and types of objects (spectra) as function of redshift.

The correction can be large, but it does not introduce modeling of an effect with a new free parameter.

It can be computed and directly applied. It depends on the observing filter, the redshift of the source, and the shape of the spectrum (SED).

We conclude that a **better estimator of the distance modulus is:**

$$\hat{\mu} = m - M + \alpha (s - 1) - \beta c - K_v(z)$$

c) **Other systematics**

Besides stretch, color and K-correction, there are many other possible sources of bias (also known as systematic effects) that impact the SN measurements:

- Peculiar velocities for low-z SN → the measured redshift are not only due to expansion → need to be corrected

- Contamination by core collapse SN for high-z SN → SN of other types mixed in the analysis

- Evolution of color-luminosity relation with redshift → the modeling of this bias should be done with $\beta(z)\,c$, and so β is no longer a single parameter

- Evolution of SNe with z → light-curves should not be matched to a single universal template

- Gravitational magnification → lensing effect changing the flux of the SNe

- Malmquist bias → a SN sample is biased towards the brightest objects, needs to correct the distance modulus by adding a magnitude of 1.38 σ^2_M

Assume we have a flux-limited sample of SNe and they have a distribution of absolute magnitudes $M_0 \pm \sigma_M$.

Then for high-z SNe the ones at the tail of the distribution may be outside of the limit → we observe a **biased sample, not representative of the full distribution → we lose systematically the faint objects, the sample is incomplete but not in a random way** → introduce a bias

We will think that those high-z SNe are brighter than they really are → we need to correct by **adding a magnitude value that increases with the width of the distribution** (the correction is computed from first principles, by integrating the SNe 'magnitude function').

(Note that if we lose objects in a random way - for example because of not observing the full sky - then there is no bias)

- Dependence on mass of the host galaxy → SNe appear systematically brighter when they are in massive galaxies by ~0.1 mag

Two possible simplest ways to proceed:

1) Add a further linear host term, H, to the analysis:

$$m_B = m_B - M_B + a(s-1) - bc + gH$$

– *Requires very precise measure of H, and robust errors*

2) Use two M_B – one for high-mass galaxies and one for low-mass

$$m_B = m_B - M_B^1 + a(s-1) - bc \quad \text{when } H < H_{\text{split}}$$

$$m_B = m_B - M_B^2 + a(s-1) - bc \quad \text{when } H > H_{\text{split}}$$

There are about 200 systematic effects identified in SNe analyses!

The dominant source of bias is the calibration of the universal absolute magnitude M (which depends on the distance ladder determinations and on the light-curves template-fitting).

Description	Ω_m	w	Rel. Area [a]
Stat only	$0.19^{+0.08}_{-0.10}$	$-0.90^{+0.16}_{-0.20}$	1
All systematics	0.18 ± 0.10	$-0.91^{+0.17}_{-0.24}$	1.85
Calibration	$0.191^{+0.095}_{-0.104}$	$-0.92^{+0.17}_{-0.23}$	1.79
SN model	$0.195^{+0.086}_{-0.101}$	$-0.90^{+0.16}_{-0.20}$	1.02
Peculiar velocities	$0.197^{+0.084}_{-0.100}$	$-0.91^{+0.16}_{-0.20}$	1.03
Malmquist bias	$0.198^{+0.084}_{-0.100}$	$-0.91^{+0.16}_{-0.20}$	1.07
non-Ia contamination	$0.19^{+0.08}_{-0.10}$	$-0.90^{+0.16}_{-0.20}$	1
MW extinction correction	$0.196^{+0.084}_{-0.100}$	$-0.90^{+0.16}_{-0.20}$	1.05
SN evolution	$0.185^{+0.088}_{-0.099}$	$-0.88^{+0.15}_{-0.20}$	1.02
Host relation	$0.198^{+0.085}_{-0.102}$	$-0.91^{+0.16}_{-0.21}$	1.08

So we finally found an estimator that should be better than the naïve one $\hat{\mu}_i = m_i - M$

This estimator is:

$$\hat{\mu}_i = m_i - M + \alpha (s_i - 1) - \beta c_i + \gamma H_i - K_v(z) + \text{other biases}$$

It takes into account the fact that

measured m = true $m + K_v(z)$

and

universal M = true $M + \alpha (s - 1) - \beta c + \gamma H$ + other biases

This estimator should give the true value of the distance modulus, i.e., if we average the measurements of N SNe Ia (at the same redshift), we should get the true μ :

(averaging over large N) → $\langle \hat{\mu} \rangle$ = true μ

an estimator with this property is called an **unbiased estimator.**

On the contrary, for the original estimator ^μ^= m - M,

(averaging over large N) → <^μ^> ≠ true μ

In that case the estimator is called a **biased estimator.**

Notice that the reason for averaging over many observations, is because the estimator has an uncertainty (error bars).

The **optimal** estimator should be **unbiased** (i.e., an estimator that provides **accurate** measurements) and at the same time should be measured with a high **signal-to-noise ratio** (i.e. an estimator that provides **precise** measurements).

v) Compute the uncertainty of the estimator

The values of the unbiased estimator averaged over the sample give the estimate for the quantity of interest (in our case $\mu(z)$).

But this is not enough to fully describe the measurement. **We also need to quantify the uncertainty of the estimator.**

This can be addressed in two different ways:

a) **Measuring the variance of the sample**

Consider a set of N_z SN Ia at the same redshift

All of these SNe should have the same value of μ_z.

However, the measured μ_{zi} for each SN will not be the same because the measurement process (including emission and correction factors) is a **random process**.

So, if we measure N_z SNe at the same redshift, they will constitute **N_z independent measurements of the same quantity**, where $\hat{\mu}_z$ is a **random variable** and each of the measured values μ_{zi}, is a realization of the probability distribution of the random variable $\hat{\mu}_z$.

If the distribution is Gaussian (which is always the case if N_z is large due to the **Central Limit Theorem**), the distribution is described by only two parameters (the lower order moments of a distribution). These are the mean, α, (i.e. the true value of μ), and the variance, σ^2

The larger the variance of a distribution, the most likely to observe a value μ_{zi} far from the true value α.

The σ of the distribution (i.e., the square root of the variance σ^2, also called the dispersion or the root-mean-square rms) is the error bar of the single observation μ_{zi}

This is the average error of each μ_{zi} measurement, but we are interested in the error of the ^μ_z^ estimate.

^μ_z^ is estimated from the individual measurements μ_{zi}. It is known that the maximum likelihood estimator of the mean of a Gaussian distribution is **the average of the realizations**:

$$\hat{\alpha} = \frac{1}{N} \sum_{i=1}^{N} \mu_i$$

↓ all at the same z

So, the uncertainty we are looking for is the error on the estimation of the mean.

The **variance of an estimator of a mean** is a well known result, and is given by:

$$\sigma_{\alpha}^{2} = \frac{1}{N(N-1)} \sum_{i=1}^{N} (\mu_i - \hat{\alpha})^2$$

Basically this is the variance of the random variable μ_{zi} divided by N_z

$$\sigma_{\alpha}^{2} = \frac{\sigma^2}{N}$$

This makes sense, since the estimated value of the mean will be closer to the true mean if we have a better sampling of the distribution (i.e., larger N_z).

The square root of this variance is the **uncertainty** associated with the estimator $\hat{\mu}_z$
→ **the error bar**.

This error is usually called the **statistical error** or also the **noise**.

The uncertainty introduced by the bias corrections, is in general not included in the variance of the random variable and needs to be taken into account separately as an extra contribution to the total error: the **systematic error**.

We just saw that **the uncertainty of ^μ_z^ decreases with the square root of the number of SNe observed at that z.**

If the variance of the random variable is large, then many measurements are needed (a large N_z) in order to obtain a small error on the mean, i.e, for the estimated value ^μ_z^ to be close to the true value μ_z.

In this case, we say that the estimator is **noisy** → it has a **low signal-to-noise ratio (S/N)**.

Note that a biased estimator is not necessarily noisy. On the contrary it can have a low statistical error if its measurements have little dispersion (but around a wrong value, since it is biased). The fact that the value is wrong means there is a large systematic error (it is biased).

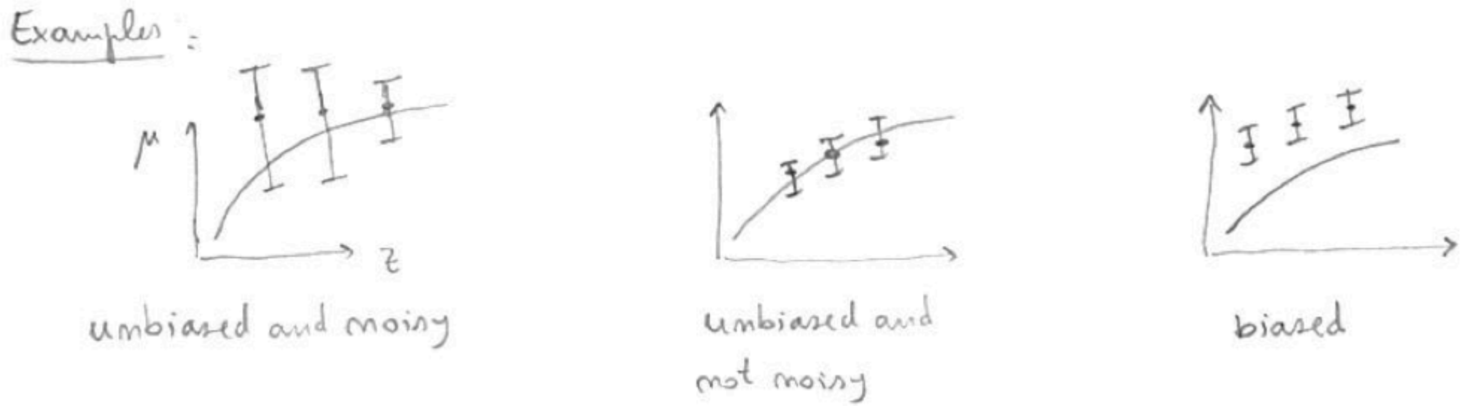

→ it is possible to have a high precision measurement with a low accuracy.

Consider now a set of N SNe Ia at various redshifts

For each one of the redshift bins (let us assume there are Nb bins), there is a different random variable μ_{zb}.

We measure N_{zb} SN for each of the bins zb, obtaining the various measurements μ_{zbi}.

We have then a vector of random variables, that is described by a multi-dimensional Gaussian (of dimension Nb).

The mean of a multi-dimensional Gaussian is a vector, that we want to estimate: ^(μ_{z1}, μ_{z2}, ... μ_{znb})^

The variance of a multi-dimensional Gaussian is a matrix, that consists of the variances of each random variable (which are the diagonal terms of the matrix), and the correlations between the various random variables (which are the off-diagonal terms of the matrix, also called the covariances) → this defines the **covariance matrix**.

The covariance matrix of a mean vector is computed similarly as in the case of a single random variable, but considering all correlations:

$$\sigma^2_{z_i z_j} = \frac{1}{N_i N_j} \sum_{k=1}^{N_i} \sum_{l=1}^{N_j} \left(\mu_k(z_i) - \langle \mu(z_i) \rangle_k \right) \left(\mu_l(z_j) - \langle \mu(z_j) \rangle_l \right)$$

This defines a Nb x Nb covariance matrix:

$$\sigma^2_{\mu\, z z'} = \begin{bmatrix} \sigma^2_{z_1 z_1} & \sigma^2_{z_1 z_2} & \sigma^2_{z_1 z_3} & \cdots \\ \| & \sigma^2_{z_2 z_2} & \sigma^2_{z_3 z_3} & \\ & & & \end{bmatrix}$$

The error bars for the measured function ^μ(z)^ are defined as the square root of the diagonal of the covariance matrix.

However the error bars do not quantify all the uncertainty of the measurement, for that we need the full covariance matrix.

In principle, the measurements of SNe at different redshifts should be uncorrelated → off-diagonal of the covariance matrix is zero and the diagonal contains all the information on the uncertainty of the estimator.

However some of the bias corrections introduce a correlation between redshift bins (for example the evolution effects). For that reason, the covariance matrix is in general not diagonal.

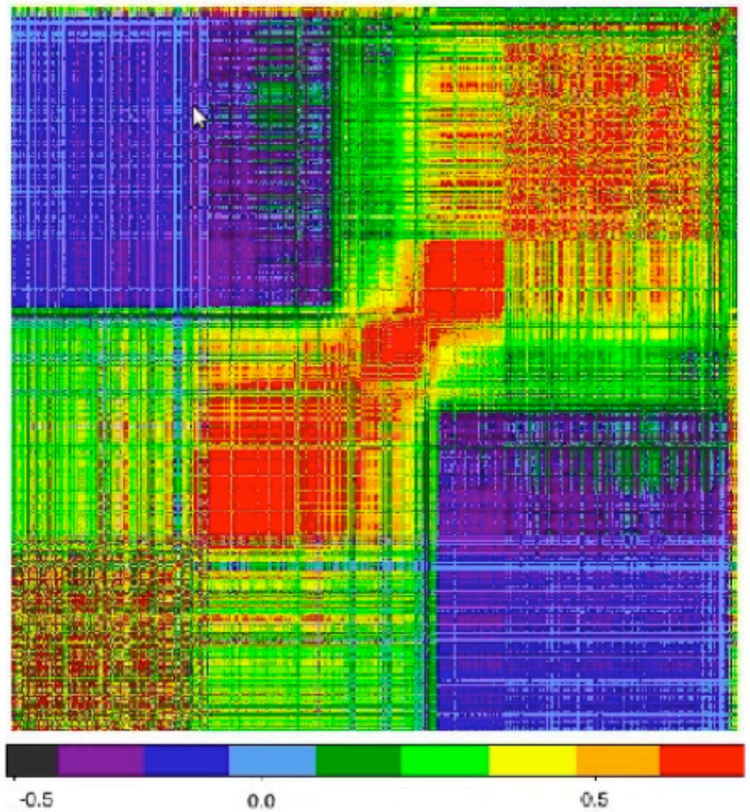

Using mock data

We just saw that we can compute the covariance matrix by computing the dispersion of the measurements.

However, this method does not give reliable results when the sample is small and is not a good representation of a complete sample → the results will suffer from sample variance.

In those cases it is better to measure the uncertainty of the estimator from simulations:

Simulate a random distribution of SN at various redshifts, with a distribution of fluxes and luminosities. From here we can build different realizations of SN samples, including noise, and measure the corresponding $\hat{\mu}(z)$ for each realization. The samples may be generated with bootstrapping methods. This will result in a set of **mock data**, that contains all the noise properties of the true data.

The uncertainty is then the covariance matrix of the mock $\hat{\mu}(z)$ data.

b) **Computing the variance of the estimator**

Instead of measuring the variance directly on the sample (or on mock data), the variance of the estimator can be analytically computed from its definition:

$$\sigma^2_{\mu_{zz'}} = \left\langle \left[\left(m_{z,i} - M + bias \right) - \left\langle m_{z,i} - M + bias \right\rangle_{N_z} \right] \left[\left(m_{z',j} - M + bias \right) - \left\langle m_{z',j} - M + bias \right\rangle_{N_{z'}} \right] \right\rangle_{ij}$$

This is the most consistent way to find the variance, since it is the formal definition and allows to consider not only the statistical error of the measurement (as in the first method) but also **the uncertainty of all terms contributing to the estimator**, such as:

- the uncertainty of the bias correction factors (the **systematic errors**)

- the intrinsic uncertainty of the true value. This is an important contribution for cosmological structure formation probes (not for the SN method), since the parameter values of the Universe are considered to be realizations of an unknown true value. So even the mean of a distribution has an intrinsic error (beyond the standard statistical error that decreases with the size of the sample). This contribution is called the **cosmic variance**.

The various contributions to the estimator uncertainty form the **error budget**.

Taylor expanding the formula of the variance of the estimator, we can write a linear expression for the variance, showing the explicit contributions of the error budget. This is the well-known **error propagation** formula. Assuming that all the terms are independent effects, the uncertainty of the estimator ^μ_z^ (for a given z) may be written as: (only a few terms are shown)

$$\sigma_\mu^2 = \left[\left(\frac{\partial \mu}{\partial m}\right)^2 \sigma_m^2 + \left(\frac{\partial \mu}{\partial M}\right)^2 \sigma_M^2 + \sigma^2_{\text{template fitting } s,c} + \sigma^2_{\text{spectroscopic identification}} \right]$$

(σ_m^2: uncertainty in the measured flux; σ_M^2: uncertainty in the intrinsic luminosity)

Note that this method allows us to compute the error of the estimator from the various statistical and systematic error contributions, but does not tell us how to compute those.

Computing the statistical error

Each of the error contributions for the budget need to be computed according to the physical process associated to that contribution.

For example, let us **consider the statistical error associated with the measurement of μ_z**, which propagates from the statistical error of the measurements of flux (magnitude), denoted in the error propagation formula by σ_m.

To compute it we need to realize that the measured signal is determined by the number of SN photons detected per pixel: it is a Poisson process.

The noise in a Poisson process **is the square-root of the number of detections.** So, if signal is N_photons_per_pixel

→ noise is sqrt (N_photons_per_pixel)

→ signal-to-noise ratio (S/N) is also sqrt (N_photons_per_pixel)

→ the dimensionless relative error is noise/signal, i.e., the inverse of the S/N, i.e.,

$$\sigma_m = 1/\text{sqrt}\,(N\_photons\_per\_pixel)$$

or, in percentage, the relative error is

$$\sigma_m = 100/\text{sqrt}\,(N\_photons\_per\_pixel)\ \ (\%)$$

We see that, for example, a S/N of 5 (also called a 5-sigma detection) means a relative error of 20%

We also see that the error decreases with the number of detected photons.
So in astronomical observations, **we can decrease the error by increasing the exposure time.**

When applying for telescope time, it is very important to compute in advance what is the needed exposure time.
This is determined by the S/N that we want to achieve, but also depends on the specific filter and telescope used.

"Exposure time calculators" (ETC) are codes that compute this for different observational configurations.

An ETC example

Assume we want to prepare the observation of one SN at redshift z=0.8 using the William Herschel 4m telescope at the Observatorio del Roque de los Muchachos (Canary islands).

The observation will be made in the B band and we want to obtain a S/N = 10 at the tail of the light curve

What is the exposure time that we need?

- The signal is determined by the number of SN photons detected per pixel.

Pointing to the host galaxy, we will detect N photons per pixel, approximately half of which come from the SN → signal is N/2

The noise depends on all the photons detected → noise is sqrt(N)

→ S/N = sqrt(N) / 2

We want S/N = 10 → **N = 400 photons per pixel**

- Now, consider the luminosity at the peak:

M_{SN} ~ -19 → $L_{SN} = 4 \times 10^9$ L_Sun = 1.5×10^{36} W= 1.5×10^{43} erg/s

We want a S/N of 10 at the end of the light-curve (~1 month after the peak) in order to have a good detection of the full light-curve. At the end of the light-curve, the luminosity is around 2 magnitudes fainter than at the peak → a factor of 6 in luminosity → L_tail = 2.5×10^{42} erg/s

- Consider now the distance to the SN

z = 0.8 → D_L = 5 Gpc (assuming the concordance model)

So the expected flux from the tail of the light-curve is:

$F = L / (4 \pi D_L^2) = 8.7 \times 10^{-16}$ erg/s/cm^2

This is the flux we will get in a square centimeter of the telescope

- Consider now the size of the telescope (diameter of 4 meters)

In 1 second, the telescope receives an energy of

$$8.7 \times 10^{-16} \times (400)^2 = 1.4 \times 10^{-10} \text{ erg from the SN}$$

However, the telescope optics and the CCD do not have 100% efficiency. Assuming the combined efficiency is only 30%

→ **the energy detected from the SN in 1 second is 4.2×10^{-11} erg**

- Consider now that we observe using the B filter,

The mean wavelength of the filter is 442 nm → frequency of 6.78×10^{14} Hz

This means that the mean energy of a photon detected in this filter is
$E = h\nu$ (where h is Planck's constant) $= 4.5 \times 10^{-12}$ erg

→ the telescope detects

$4.2 \times 10^{-11} / 4.5 \times 10^{-10}$ = **9.3 SN photons per second**

- Consider now that the size of the SN in the image is 4 pixels

→ in 1 second, we get 9.3 / 4 = 2.3 SN photons per pixel

Conclusion: knowing that we need 400 SN photons per second and per pixel, we need an exposure time of 400/2.3 = 170 s ~3 minutes

An exposure of 3 minutes in a 4m telescope, with these characteristics and in the B band, forms the image of a z=0.8 SN with a quality of S/N =10 at the tail of the light-curve

vi) The data vector

We have finally obtained the data vector, i.e., the distance modulus vector and its covariance matrix/error bars

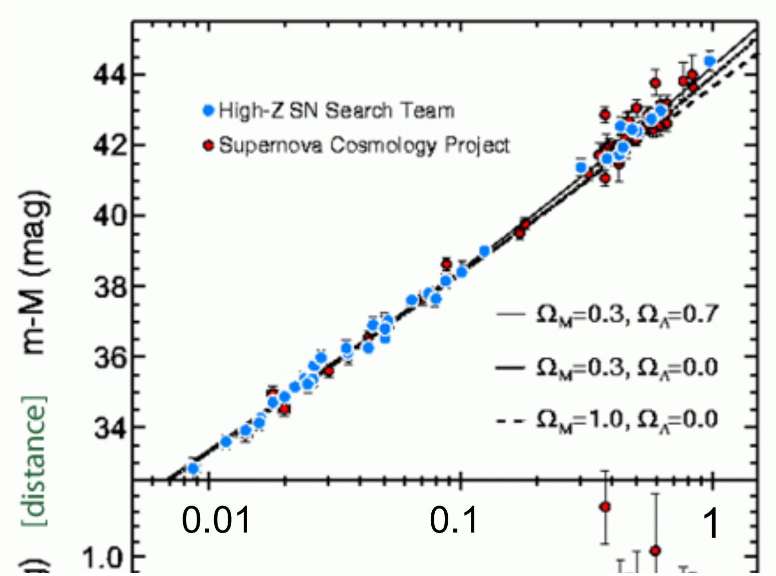

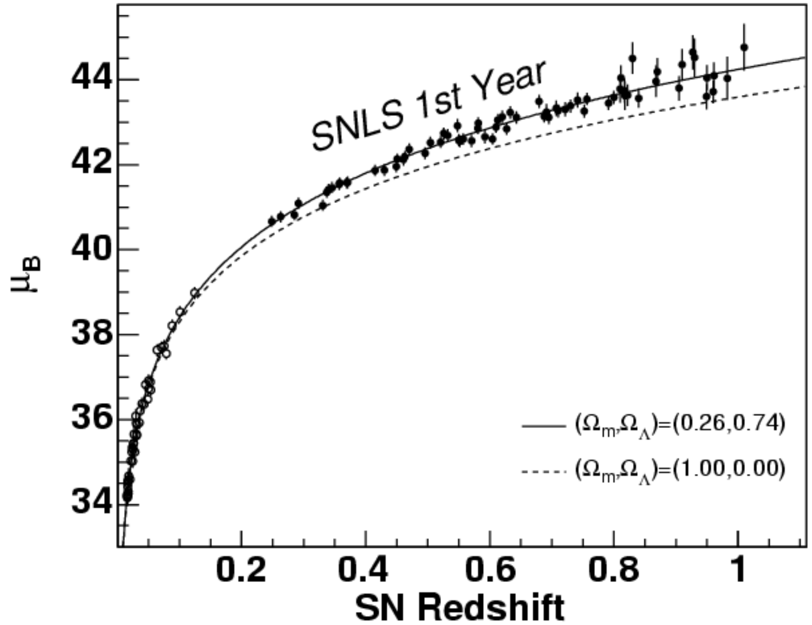

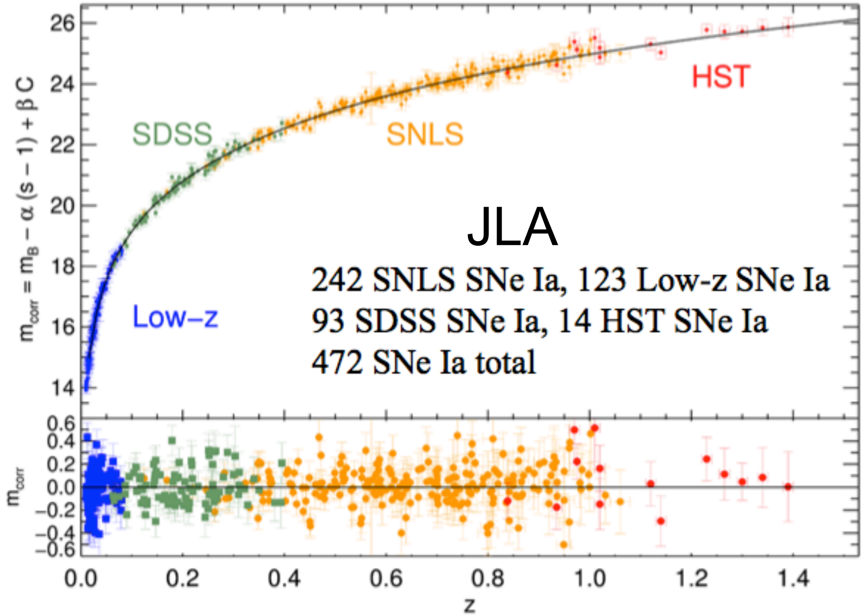

1998: Original sample	N ~ 40
2006: SNLS DR1	N ~ 100
2014: JLA compilation	N ~ 500
2018: Pantheon compilation	N ~ 1000

What do we do with the **data vector**?

We compare it with the theoretical computation, **“the theory vector”:**

$$\mu(z) = 5 \log_{10} [D_L (z; H0, \Omega, w)] + 25$$

and estimate the model free parameters in a **statistical inference analysis**, involving the computation of **likelihoods**.

Hopefully the data vector is measured with an **optimal estimator** **(accurate and precise)**

The estimator is unbiased and no extra uncertainty is introduced

(accurate and precise estimation of the cosmological parameters)

As we saw, in order to find an unbiased estimator of the distance modulus we needed to make a careful consideration of all sources of **systematics**.

We saw that in general **the systematic effects can be treated in different ways:**

a) **directly computed** (with perfect knowledge) and their values subtracted to the measurements

(e.g. the K-correction)

b) **modeled with a function** (perfectly known) that introduces additional **nuisance parameters** with unknown values

(e.g. the nuisance parameters)

The nuisance parameters can be **estimated** in the statistical inference analysis together with the cosmological parameters.

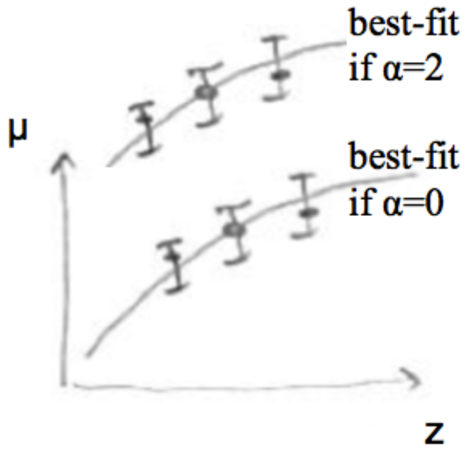

Depending on the value of α, the data points will be higher or lower in the μ(z) plot. → the best-fit cosmological model will depend on the value of α → **there is a degeneracy between α and the cosmological parameters.**

This makes it difficult to estimate simultaneously the nuisance and cosmological parameters with great precision.

- **marginalized**: usually there is no interested in finding the values of the nuisance parameters.
We only want **to estimate the cosmological parameters, but need to consider all the possible range of values of the nuisance parameters.**

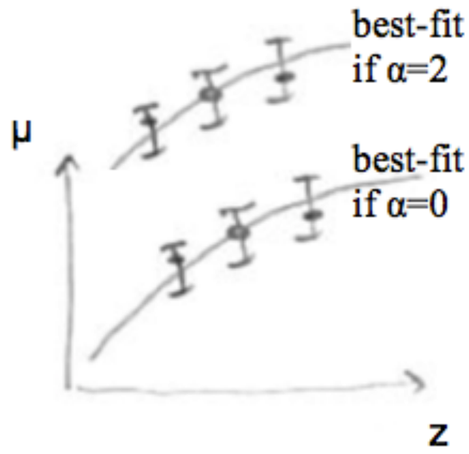

example: marginalize from α=0 to α=2

This implies that the estimate obtained for the cosmological parameters will have a larger uncertainty than if there were no nuisance parameters → **unbiasing the estimator results in decreasing the precision of the result** → but it increases its accuracy and the result is more reliable.

(accurate but not precise estimation of the cosmological parameters)

The presence of nuisance parameters worsens the constraints on the cosmological parameters.

In terms of Figure-of-Merit (FOM) - area of the confidence contours in the cosmological parameters space, where small area means strong constraints - this means that without the bias correction, the FoM will increase (is better) → but it is a wrong result, too optimistic.

c) **known unknowns**: **systematic error**

In reality, the computation of the corrections is not perfect, there is an uncertainty.

Also the astrophysical modeling may have uncertainties.

It may also happen that we are able to identify some systematics but are not able to compute them directly or to model them.

In these cases we are using a more or less biased estimator. To be able to make a meaningful analysis, in the presence of these **known unknowns**, we need to include additional uncertainties in the error bars: **systematic errors σ_{sys}** → the bias is replaced by an increase in the error bars.

In this case we can still get the right values inside our estimated interval, but it was just because we increased the uncertainty.

Notice that contrary to the statistical error (noise), systematic errors do not average out to zero with large N.

Notice also that if we underestimate the systematic uncertainty that should be allocated, we may end up with the wrong result

d) **unknown unknowns**: do nothing

In the case we are not aware that there are still additional effects and just use the biased estimator (without realising) without doing any correction or modelling, and do not even include systematic uncertainties,

then we can find precise results (if the estimator has a high signal-to-noise ratio) but they will be inaccurate.

Recent data: the Pantheon compilation

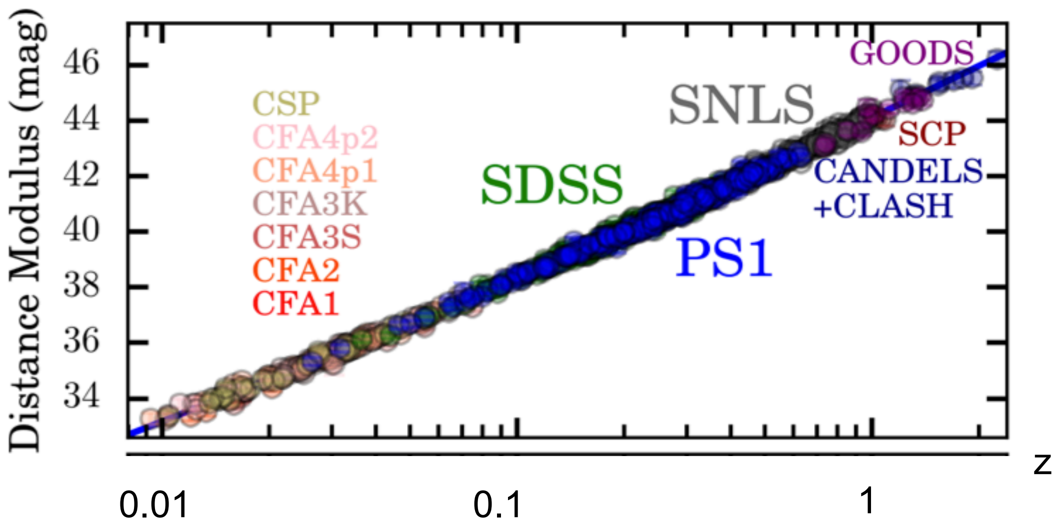

Sample	Number	Mean z
CSP	26	0.024
CFA3	78	0.031
CFA4	41	0.030
CFA1	9	0.024
CFA2	18	0.021
SDSS	335	0.202
PS1	279	0.292
SNLS	236	0.640
SCP	3	1.092
GOODS	15	1.120
CANDELS	6	1.732
CLASH	2	1.555
Tot	1048	

SNLS: SN Legacy Survey
SDSS: Sloan Digital Sky Survey
PS1: Pan-STARRS 1

Precision:

The large number of data points (N increased by a factor of 25) → leads to a factor of 5 improvement in precision.

Accuracy:

Systematic errors from bias corrections decrease because of better corrections, but do not decrease with N → a systematics floor may eventually be reached in future large surveys.

"Precision cosmology is hard, accurate cosmology is even harder"

- Michael Turner -