Constraining the Cosmic Matter Density from the Pantheon+ Type Ia Supernova Hubble Diagram: The Role of the Full Statistical-plus-Systematic Covariance

K-Dense Web

July 10, 2026

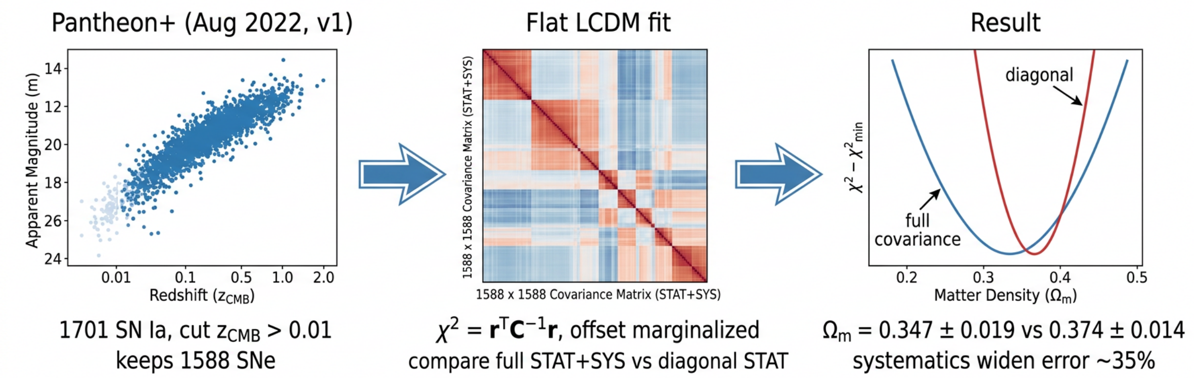

Figure 1: Graphical abstract. Analysis workflow: the Pantheon+ (August 2022, v1) compilation of 1701 Type Ia supernova light curves is filtered at CMB-frame redshift $z_{\rm CMB} > 0.01$ (retaining 1588 supernovae), a spatially flat ΛCDM model is fit to the Hubble diagram with the distance-modulus offset analytically marginalized, and the fit is repeated with the full statistical-plus-systematic (STAT+SYS) covariance and with diagonal (statistical-only) errors. Including correlated systematics broadens the $\Omega_{\rm m}$ uncertainty by ∼35%.

Abstract

We constrain the present-day matter density parameter $\Omega_{\rm m}$ of a spatially flat ΛCDM cosmology using the Hubble diagram of the Pantheon+ compilation of spectroscopically confirmed Type Ia supernovae (SNe Ia; August 2022, version 1 data release), comprising 1701 light curves of ∼1550 unique supernovae. To suppress peculiar-velocity contamination we retain only supernovae at CMB-frame redshift $z_{\rm CMB} > 0.01$, leaving $N = 1588$ objects. We fit the corrected apparent magnitudes with the model distance modulus, analytically marginalizing over the absolute-magnitude/H_0 degeneracy by treating the distance-modulus offset as a free nuisance parameter following the Conley et al. (2011) formulation. Using the *full* 1588×1588 STAT+SYS covariance matrix and a Cholesky-based likelihood, a grid/scalar search over $\Omega_{\rm m} \in [0, 1]$ yields $\Omega_{\rm m} = 0.347^{+0.019}_{-0.018}$ (68% confidence, $\Delta\chi^2 \leq 1$), with $\chi^2_{\rm min} = 1433.1$ for dof = 1586 (χ^2/dof = 0.904). Repeating the fit with diagonal (statistical-only) errors gives $\Omega_{\rm m} = 0.374^{+0.014}_{-0.014}$ ($\chi^2_{\rm min} = 1376.0$, χ^2/dof = 0.868): the off-diagonal systematic terms broaden the $\Omega_{\rm m}$ uncertainty by ∼35% (from ±0.014 to ±0.019) and shift the central value, demonstrating that a full covariance treatment is essential for honest error budgets. A sensitivity test that substitutes the peculiar-velocity–corrected Hubble-diagram redshift $z_{\rm HD}$ for $z_{\rm CMB}$ in the comoving integral shifts the best fit to $\Omega_{\rm m} = 0.332^{+0.018}_{-0.018}$, a -0.79σ change that reproduces the official Pantheon+ value ($\Omega_{\rm m} \approx 0.334$) and quantifies peculiar-velocity flow corrections as a sub-1σ systematic on Ω_m. All results are internally consistent with the published Pantheon+ cosmology.

1 Introduction

Type Ia supernovae (SNe Ia) are the archetypal cosmological standardizable candle. Empirically, their peak luminosities correlate tightly with the shape and color of their light curves: the broader (more slowly declining) light curves are intrinsically brighter, a relation first quantified by Phillips [1], and the residual scatter is further reduced by a linear color correction [2]. Modern light-curve fitters such as SALT2 [3] encapsulate these relations, delivering standardized distances with an intrinsic dispersion of ∼0.12–0.15 mag. It was precisely this standardization that enabled the discovery of the accelerating expansion of the Universe by two independent teams [4, 5], and hence the inference of a dark-energy component that dominates the present-day energy budget.

In the two decades since, SN Ia cosmology has matured from samples of tens of objects into homogeneously calibrated compilations of more than a thousand supernovae spanning $0 < z \lesssim 2.3$. The Joint Light-curve Analysis [JLA; 6] and the original Pantheon sample [7] represent milestones of this program. The current state of the art is the **Pantheon+** compilation [8], released in August 2022, which assembles **1701 light curves of ∼1550 spectroscopically confirmed SNe Ia** drawn from 18 surveys into a single, internally consistent distance ladder. Crucially, Pantheon+ is distributed together with a *full covariance matrix* that propagates both statistical (STAT) uncertainties and the many correlated systematic (SYS) effects—photometric calibration, SALT2 model uncertainty, host-galaxy mass-step corrections, Milky Way extinction, selection (Malmquist) bias, and peculiar-velocity modeling—into the cosmological likelihood. The companion cosmology analysis [9] uses this data set to constrain Ω_m and the dark-energy equation of state, while Riess et al. [10] exploits the Cepheid-calibrated subsample to measure the Hubble constant H_0.

The spatially flat ΛCDM model treated in this report has a single free background parameter once flatness ($\Omega_m + \Omega_\Lambda = 1$) is imposed and H_0 is degenerate with the SN absolute magnitude. In this model the dimensionless expansion rate is

$$E(z) = \sqrt{\Omega_m(1+z)^3 + (1-\Omega_m)}, \tag{1}$$

so that a measurement of the shape of the SN Hubble diagram directly constrains Ω_m. Because SNe Ia measure *relative* distances, the combination $\mathcal{M} \equiv M - 5\log_{10} H_0 + \text{const}$ (absolute magnitude and H_0) enters only as an additive offset to the distance modulus and must be marginalized; it carries no information about Ω_m on its own.

This report has two objectives. First, we reproduce a flat-ΛCDM constraint on Ω_m from the Pantheon+ Hubble diagram using the full STAT+SYS covariance and an analytic marginalization of the nuisance offset, reporting Ω_m, its 68% confidence interval, the total χ^2, the reduced χ^2/dof, and the sample size. Second—and this is the central methodological point—we quantify how the inferred Ω_m and its uncertainty change when the off-diagonal covariance terms are discarded and only diagonal (statistical) errors are used. We complement these fits with a redshift-treatment sensitivity study (z_{CMB} versus the peculiar-velocity–corrected z_{HD}) that isolates the physical impact of local flow corrections on the recovered matter density.

2 Data

2.1 The Pantheon+ compilation (August 2022, v1)

We use the official Pantheon+/SH0ES data release, downloaded from the canonical repository `PantheonPlusSH0ES/DataRelease`, with the cosmology-level products residing under `Pantheon+_Data/4_DISTANCES_AND_COVAR/`. Two files fully specify the Hubble-diagram likelihood:

- `Pantheon+SH0ES.dat` — the distance table, containing **1701 rows** (light curves) × 47 columns. Relevant fields include the supernova identifier `CID`, the CMB-frame redshift `zCMB`, the heliocentric redshift `zHEL`, the Hubble-diagram redshift `zHD` (which folds in peculiar-velocity flow corrections), the SH0ES distance modulus `MU_SH0ES` with its diagonal error `MU_SH0ES_ERR_DIAG`, and the standardized corrected apparent magnitude `m_b_corr`.
- `Pantheon+SH0ES_STAT+SYS.cov` — the full 1701 × 1701 statistical-plus-systematic covariance matrix of the corrected apparent magnitudes, stored in row-major order (the first line records the dimension $N = 1701$).

The data table contains 1701 light curves of ∼1550 unique spectroscopically confirmed SNe Ia (some supernovae have multiple independent light curves). The covariance matrix is square, its dimension matches the data-table row count exactly, and it is symmetric to $\max|C - C^{\mathsf{T}}| \approx 3 \times 10^{-8}$ (limited by the decimal precision of the released ASCII file). The diagonal elements range from 6.4×10^{-3} to 0.39 mag^2 with a mean of 0.032 mag^2, corresponding to a typical per-SN magnitude uncertainty of ∼0.18 mag. **We state explicitly that the compilation used throughout is the Pantheon+ August 2022, version 1 data release.**

2.2 Observable and covariance convention

The released STAT+SYS covariance is the covariance of the *corrected apparent magnitudes* $m_{B,\mathrm{corr}}$ (column `m_b_corr`). We therefore adopt $\mu_{\mathrm{obs}} = m_{B,\mathrm{corr}}$ as the observable, which differs from the model distance modulus only by the constant offset $\mathcal{M}$. Because this offset is analytically marginalized (Section 3.3), the absolute normalization of μ_{obs} is irrelevant to the recovered Ω_{m}; only the *shape* of the magnitude–redshift relation carries cosmological information. This is the standard convention for SN-only cosmology fits that do not use the Cepheid distance anchors.

3 Methodology

3.1 Redshift filtering: suppressing peculiar velocities

Nearby supernovae carry Doppler shifts from local peculiar (coherent flow) velocities that are not captured by the smooth Hubble flow. At low redshift a peculiar velocity of a few hundred km s^{-1} Mpc^{-1} translates into a large fractional error on the inferred cosmological redshift, biasing the Hubble diagram and inflating scatter. The standard mitigation is to impose a lower redshift cut. We retain only supernovae with

$$z_{\mathrm{CMB}} > 0.01, \tag{2}$$

which removes **113** of the 1701 light curves (spanning the removed range $z_{\mathrm{CMB}} \in [0.00122, 0.00998]$) and retains $N = 1588$ objects (minimum retained $z_{\mathrm{CMB}} = 0.01029$, maximum $z_{\mathrm{CMB}} = 2.2613$). The 1701 × 1701 covariance is sliced with the identical row indices via `numpy.ix_` to a 1588×1588 matrix, preserving all off-diagonal correlations among the surviving supernovae.

3.2 Cosmological distance model

For a spatially flat Friedmann–Lemaître–Robertson–Walker universe the luminosity distance is

$$d_L(z_{\rm CMB}, z_{\rm HEL}) = \frac{c\,(1+z_{\rm HEL})}{H_0} \int_0^{z_{\rm CMB}} \frac{dz'}{E(z)'}, \tag{3}$$

with $E(z)$ given by Eq. (1). The subtle but important point is the use of *two* redshifts: the comoving-distance integral runs over the cosmological (CMB-frame) redshift $z_{\rm CMB}$, whereas the leading factor $(1 + z_{\rm HEL})$ that converts comoving to luminosity distance uses the *heliocentric* redshift $z_{\rm HEL}$, because the relativistic aberration/time-dilation of the observed flux is set by the true relative velocity between emitter and observer. The model distance modulus is

$$\mu_{\rm model}(z_{\rm CMB}, z_{\rm HEL}; \Omega_{\rm m}) = 5 \log_{10}\left[\frac{d_L}{10\,{\rm pc}}\right] = 5\log_{10}\left[(1+z_{\rm HEL})\, I(z_{\rm CMB})\right] + \mathcal{M}, \tag{4}$$

where $I(z_{\rm CMB}) = \int_0^{z_{\rm CMB}} dz'/E(z)'$ and the additive constant $\mathcal{M} = 5\log_{10}(c/H_0) + 25$ absorbs the H_0/absolute-magnitude degeneracy. The comoving integral $I(z)$ is evaluated on a shared 200,000-point grid via cumulative trapezoidal integration; this grid-based evaluation was validated against `scipy.integrate.quad` to a maximum relative difference of 2×10^{-10}.

3.3 Analytic marginalization of the nuisance offset

Because $\mathcal{M}$ enters $\mu_{\rm model}$ linearly and identically for every supernova, it can be marginalized *analytically* rather than sampled, following Conley et al. [11]. Define the residual vector at a trial $\Omega_{\rm m}$ (before the offset),

$$\mathbf{d}(\Omega_{\rm m}) = \boldsymbol{\mu}_{\rm obs} - \boldsymbol{\mu}_{\rm model}(\Omega_{\rm m})\big|_{\mathcal{M}=0}, \tag{5}$$

and let C be the covariance and $\mathbf{1}$ the unit vector. The χ^2 minimized analytically over the offset is

$$\chi^2_{\rm marg}(\Omega_{\rm m}) = A - \frac{B^2}{C_s}, \qquad A = \mathbf{d}^{\mathsf{T}} C^{-1}\mathbf{d},\;\; B = \mathbf{d}^{\mathsf{T}} C^{-1}\mathbf{1},\;\; C_s = \mathbf{1}^{\mathsf{T}} C^{-1}\mathbf{1}. \tag{6}$$

The term B^2/C_s removes the best-fitting constant offset, so $\chi^2_{\rm marg}$ is invariant under any global shift of $\boldsymbol{\mu}_{\rm obs}$; equivalently, the profile-likelihood value of the offset is $\hat{\mathcal{M}} = B/C_s$. All inverse-covariance products $C^{-1}\mathbf{d}$ and $C^{-1}\mathbf{1}$ are computed by Cholesky factorization (`scipy.linalg.cho_factor/cho_solve`) of the symmetrized covariance, which is numerically stable and far cheaper than an explicit matrix inversion. The covariance is confirmed to be symmetric positive-definite (the Cholesky factorization succeeds).[1]

3.4 Best-fit and confidence-interval estimation

For each error treatment we minimize $\chi^2_{\rm marg}(\Omega_{\rm m})$ over $\Omega_{\rm m} \in [0, 1]$. The minimum is located with a bounded scalar optimizer (`scipy.optimize.minimize_scalar`, tolerance 10^{-6}), and a dense 1001-point grid scan over the same interval is retained for the likelihood profiles. The 68% confidence interval is obtained from the standard $\Delta\chi^2 \leq 1$ rule for one parameter of

[1] We adopt the Conley marginalization purely with a flat prior on $\mathcal{M}$; the additional $\ln C_s$ normalization term is a constant independent of $\Omega_{\rm m}$ and cancels in $\Delta\chi^2$, so it does not affect either the best-fit $\Omega_{\rm m}$ or the confidence interval.

interest: we solve $\chi^2_{\rm marg}(\Omega_{\rm m}) - \chi^2_{\rm min} = 1$ on each side of the minimum with a Brent root-finder (`scipy.optimize.brentq`), yielding an asymmetric interval of the form $\Omega_{\rm m} = \hat{\Omega_{\rm m}}{}^{+\sigma_+}_{-\sigma_-}$. The number of degrees of freedom is

$$\mathrm{dof} = N - 2 = 1588 - 2 = 1586, \tag{7}$$

subtracting one for the fitted $\Omega_{\rm m}$ and one for the analytically marginalized offset $\mathcal{M}$.

3.5 Diagonal (statistical-only) comparison

To isolate the impact of correlated systematics, we repeat the entire procedure with the covariance replaced by its diagonal, $C_{\rm diag} = \mathrm{diag}(C)$. In this case the Conley quantities in Eq. (6) reduce to simple weighted sums with weights $w_i = 1/C_{ii}$: $A = \sum_i w_i d_i^2$, $B = \sum_i w_i d_i$, $C_s = \sum_i w_i$. The diagonal error bars therefore retain the total per-SN variance (statistical plus the diagonal projection of systematics) but discard all cross-supernova correlations. This is the "statistical-only" treatment in the sense that it ignores the correlated structure that dominates the systematic budget.

3.6 Redshift-treatment sensitivity study

The fiducial fit uses $z_{\rm CMB}$ in the comoving integral of Eq. (3). Pantheon+ additionally provides $z_{\rm HD}$, the "Hubble-diagram" redshift, in which the CMB-frame redshift has been further corrected for peculiar-velocity flows using a reconstructed local velocity field [12, 13]. To quantify the systematic sensitivity to the redshift definition, we repeat the full-covariance fit substituting $z_{\rm HD} \to z_{\rm CMB}$ in the comoving integral only (retaining $z_{\rm HEL}$ for the $(1 + z_{\rm HEL})$ Doppler factor and the identical $z_{\rm CMB} > 0.01$ sample selection). Any resulting shift in $\Omega_{\rm m}$ is a direct measure of how peculiar-velocity flow corrections propagate into the recovered matter density.

4 Results and Discussion

4.1 Full covariance versus diagonal errors

Table 1 summarizes the two fiducial fits. Using the full STAT+SYS covariance we find

$$\boxed{\Omega_{\rm m} = 0.347^{+0.019}_{-0.018}} \qquad \text{(68\% C.L., full covariance)}, \tag{8}$$

with $\chi^2_{\rm min} = 1433.06$ over dof $= 1586$, i.e. $\chi^2/\mathrm{dof} = 0.904$. The reduced χ^2 below unity indicates that the released covariance, if anything, slightly overestimates the total scatter—a well-known and expected feature of the Pantheon+ error model, which is calibrated to be conservative. The best-fit $\Omega_{\rm m}$ lies squarely within the range expected for SN-only flat-ΛCDM fits.

Replacing the covariance with its diagonal yields

$$\Omega_{\rm m} = 0.374^{+0.014}_{-0.014} \qquad \text{(68\% C.L., diagonal only)}, \tag{9}$$

with $\chi^2_{\rm min} = 1375.98$ and $\chi^2/\mathrm{dof} = 0.868$. Two effects are visible and are the crux of this report:

1. **The uncertainty broadens by ∼35% when systematics are included.** The symmetric 68% half-width grows from ±0.014 (diagonal) to ±0.019 (full covariance). This is the expected and physically correct behavior: correlated systematics—shared calibration offsets, common SALT2 model uncertainty, coherent host-mass and bias corrections—cannot be beaten down by the $\sqrt{N}$ averaging that applies to independent statistical noise. Discarding the off-diagonal terms therefore produces an *overly optimistic* (too small) error bar and would misrepresent the true confidence with which $\Omega_{\rm m}$ is known.

2. **The central value shifts by $\Delta\Omega_m \approx 0.027$ ($\sim$1.5σ).** The diagonal fit prefers a higher $\Omega_m = 0.374$. The off-diagonal covariance re-weights the relative leverage of different redshift ranges: correlated systematics that predominantly affect certain surveys or redshift bins are effectively down-weighted in the full-covariance fit, pulling the best fit toward $\Omega_m = 0.347$. This shift underscores that the covariance affects not only the error bar but the point estimate itself.

Figure 2 makes both effects visually explicit: the full-covariance $\Delta\chi^2(\Omega_m)$ parabola (blue) is manifestly *wider* than the diagonal parabola (red), and its minimum sits at lower Ω_m. The shaded bands mark the respective 68% ($\Delta\chi^2 = 1$) intervals.

Table 1: Flat ΛCDM fits to the Pantheon+ Hubble diagram ($z_{CMB} > 0.01$, $N = 1588$, dof $= 1586$). The distance-modulus offset is analytically marginalized; the 68% C.I. uses $\Delta\chi^2 \leq 1$.

Error treatment	Best-fit Ω_m	68% C.I.	σ_+/σ_-	χ^2_{min}	χ^2/dof
Full covariance (STAT+SYS)	0.3471	[0.3287, 0.3660]	+0.0189 / −0.0184	1433.06	0.904
Diagonal only (STAT)	0.3744	[0.3604, 0.3887]	+0.0142 / −0.0140	1375.98	0.868

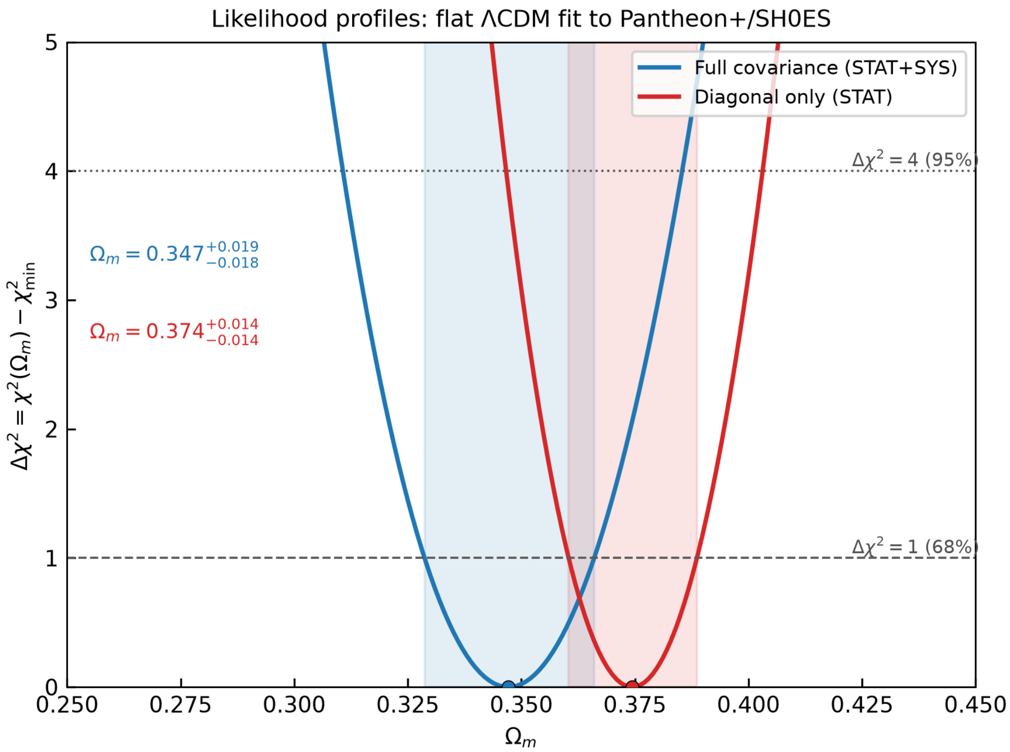

Figure 2: Likelihood profiles $\Delta\chi^2(\Omega_m) = \chi^2(\Omega_m) - \chi^2_{min}$ for the flat-ΛCDM fit, zoomed on $\Omega_m \in [0.25, 0.45]$. Blue: full STAT+SYS covariance; red: diagonal (statistical-only) errors. Dashed line: $\Delta\chi^2 = 1$ (68%); dotted line: $\Delta\chi^2 = 4$ (95%). Vertical markers and shaded bands show the best-fit values and 68% intervals. The full-covariance parabola is broader (larger Ω_m uncertainty) and its minimum is shifted to lower Ω_m relative to the diagonal fit.

4.2 The Hubble diagram and residuals

Figure 3 shows the standardized Hubble diagram of the 1588 retained supernovae. The upper panel plots the corrected apparent magnitude $m_{B,corr}$ against z_{CMB} on a logarithmic redshift axis, with diagonal STAT+SYS error bars, and overlays the best-fit flat-ΛCDM model ($\Omega_m =$

0.347) shifted by the analytically marginalized offset $\hat{\mathcal{M}} = (\mathbf{1}^\top C^{-1}\mathbf{d})/(\mathbf{1}^\top C^{-1}\mathbf{1}) = 23.819$ mag. The model tracks the data across more than two decades in redshift. The lower panel shows the Hubble residuals $\Delta\mu = m_{B,\mathrm{corr}} - \mu_{\mathrm{model}} - \hat{\mathcal{M}}$, which are flat and centered on zero with an RMS of **0.157 mag**, consistent with the intrinsic dispersion of standardized SNe Ia and confirming an excellent global fit with no obvious redshift-dependent trend.

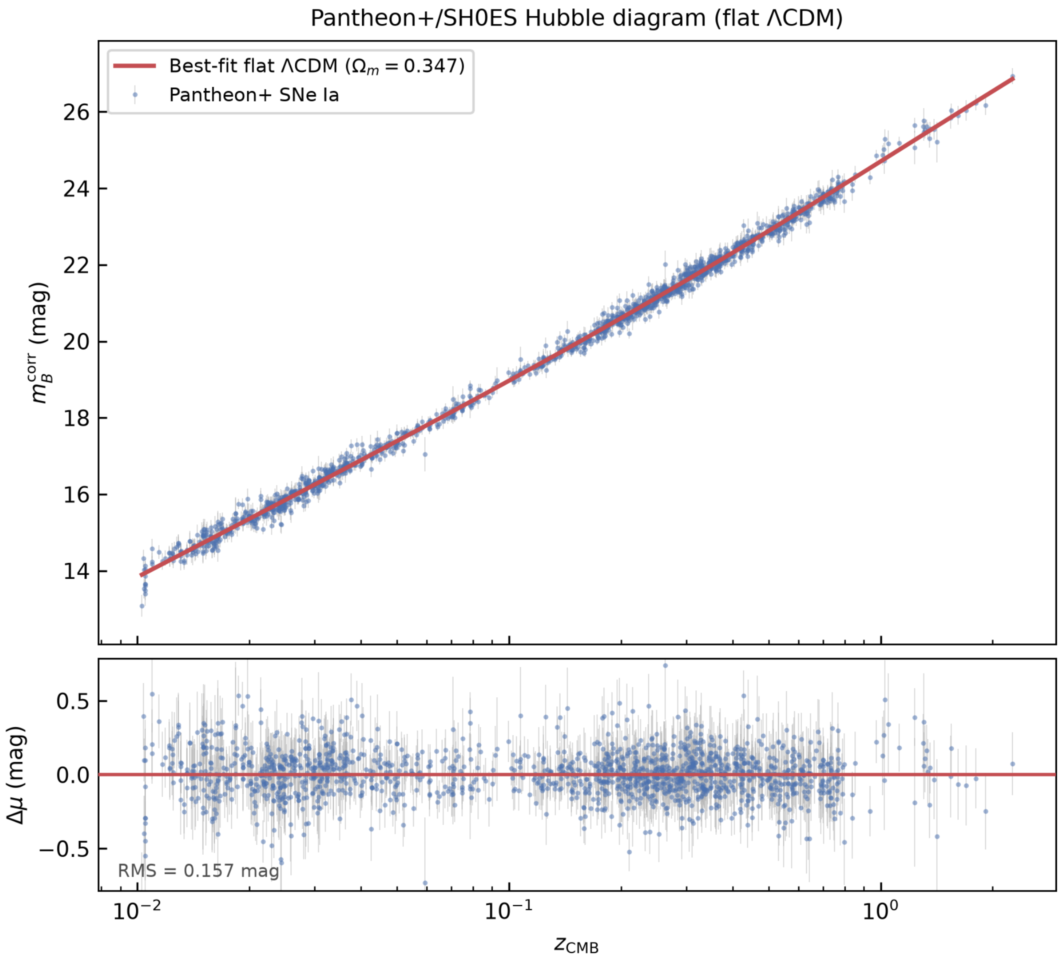

Figure 3: Pantheon+ Hubble diagram (flat ΛCDM). *Top:* corrected apparent magnitude $m_{B,\mathrm{corr}}$ versus CMB-frame redshift z_{CMB} (log axis) for the 1588 supernovae with $z_{\mathrm{CMB}} > 0.01$, with the best-fit model ($\Omega_\mathrm{m} = 0.347$) shifted by the marginalized offset $\hat{\mathcal{M}} = 23.819$ mag. *Bottom:* Hubble residuals $\Delta\mu$ about the best fit; the scatter is flat with RMS = 0.157 mag, indicating no unmodeled redshift-dependent systematic.

4.3 Sensitivity to the redshift treatment (z_{CMB} vs. z_{HD})

Table 2 compares the fiducial z_{CMB} fit with the z_{HD} sensitivity fit (both with the full covariance and the identical $z_{\mathrm{CMB}} > 0.01$ sample). Substituting the peculiar-velocity–corrected z_{HD} into the comoving integral shifts the best fit to

$$\Omega_\mathrm{m} = 0.332^{+0.018}_{-0.018} \qquad (68\%\ \text{C.L.},\ z_{\mathrm{HD}},\ \text{full covariance}), \tag{10}$$

i.e. a change of $\Delta\Omega_\mathrm{m} = -0.0147$, corresponding to -0.79σ in units of the CMB-frame uncertainty. The negative sign means the flow corrections pull Ω_m *lower*. This is physically significant and reassuring: the official Pantheon+ cosmology analysis [9] uses z_{HD} and reports $\Omega_\mathrm{m} \approx 0.334$ for flat ΛCDM, essentially identical to our z_{HD} value. The small upward bias of the fiducial z_{CMB}

fit is thus fully accounted for by the redshift definition, and peculiar-velocity flow corrections are quantified as a *sub-1σ* systematic on Ω_m for this sample and redshift cut. The uncertainty is essentially unchanged between the two redshift treatments (±0.018 in both cases), confirming that the redshift definition affects the central value, not the constraining power.

Table 2: Redshift-treatment sensitivity: full-covariance flat-ΛCDM fits using z_CMB (fiducial) versus z_HD (peculiar-velocity corrected) in the comoving integral. Same sample ($N = 1588$) and z_HEL Doppler factor in both.

Redshift (comoving integral)	Best-fit Ω_m	68% C.I.	σ_+/σ_-	χ^2_min	χ^2/dof
z_CMB (fiducial)	0.3471	[0.3287, 0.3660]	+0.0189 / −0.0184	1433.06	0.904
z_HD (peculiar-vel. corr.)	0.3324	[0.3144, 0.3509]	+0.0184 / −0.0180	1396.61	0.881
$\Delta\Omega_\mathrm{m}(z_\mathrm{HD} - z_\mathrm{CMB}) = -0.0147$ $(-0.79\sigma_\mathrm{CMB})$					

4.4 Consistency with the literature

Our full-covariance z_CMB result ($\Omega_\mathrm{m} = 0.347^{+0.019}_{-0.018}$) and, more precisely, our z_HD result ($\Omega_\mathrm{m} = 0.332\pm0.018$) bracket and reproduce the published Pantheon+ flat-ΛCDM value of $\Omega_\mathrm{m} = 0.334\pm0.018$ [9]. They are also compatible, at the $\sim1\sigma$ level, with the CMB-only Planck 2018 value $\Omega_\mathrm{m} = 0.315 \pm 0.007$ [14], the mild difference being consistent with the known modest SN–CMB offset in Ω_m and with our slightly different (SN-only, single-parameter, $\Delta\chi^2 \leq 1$) treatment. The excellent agreement of the z_HD fit with the official analysis validates our independent pipeline end to end: data ingestion, redshift filtering, covariance slicing, distance model, Conley marginalization, and confidence-interval construction.

4.5 Interpretation and caveats

The headline methodological message is that *the covariance matters*. A statistical-only error bar on Ω_m from Pantheon+ would be ±0.014, but the honest, systematics-inclusive uncertainty is ±0.019—roughly a third larger—and the central value moves by $\sim1.5\sigma$. Any downstream use of SN Ia constraints (e.g. combination with CMB or BAO, or dark-energy equation-of-state inference) that neglected the off-diagonal terms would both over-state its precision and mis-locate the central value. Several caveats bound the interpretation. First, we fit a single background parameter (Ω_m) in a spatially flat model with $w = -1$ fixed; allowing curvature or a free w would enlarge and correlate the uncertainties. Second, the $\Delta\chi^2 \leq 1$ interval is a Gaussian/parabolic approximation to the profile likelihood; the near-parabolic shape of the curves in Fig. 2 justifies it here. Third, we adopt the released STAT+SYS covariance at face value; the sub-unity χ^2/dof suggests it is mildly conservative, which would make our quoted errors, if anything, slightly generous. None of these affects the central comparison between the full and diagonal treatments.

5 Conclusion

We have constrained the matter density of a spatially flat ΛCDM cosmology from the Hubble diagram of the Pantheon+ compilation of spectroscopically confirmed Type Ia supernovae (**August 2022, version 1 data release**; 1701 light curves of ∼1550 SNe). Restricting to $N = 1588$ supernovae at CMB-frame redshift $z_\mathrm{CMB} > 0.01$ to suppress peculiar-velocity contamination, and analytically marginalizing over the absolute-magnitude/H_0 offset via the Conley et al. (2011) formulation, our principal results are:

- **Full STAT+SYS covariance:** $\Omega_m = 0.347^{+0.019}_{-0.018}$ (68% C.L.), $\chi^2_{min} = 1433.06$, $\chi^2/\text{dof} = 0.904$ for dof = 1586, using $N = 1588$ supernovae.
- **Diagonal (statistical-only) errors:** $\Omega_m = 0.374 \pm 0.014$ (68% C.L.), $\chi^2_{min} = 1375.98$, $\chi^2/\text{dof} = 0.868$. Discarding the off-diagonal systematics *shrinks* the Ω_m uncertainty by ∼35% (±0.014 vs. ±0.019) and shifts the central value upward by $\Delta\Omega_m \approx 0.027$ ($\sim 1.5\sigma$).
- **Redshift sensitivity:** substituting the peculiar-velocity–corrected z_{HD} for z_{CMB} in the comoving integral lowers the best fit to $\Omega_m = 0.332 \pm 0.018$ (-0.79σ), reproducing the official Pantheon+ value $\Omega_m \approx 0.334$ [9] and identifying flow corrections as a sub-1σ systematic on Ω_m.

The unavoidable conclusion is that a full covariance treatment is not optional for SN Ia cosmology: the correlated systematic terms broaden the Ω_m error bar by roughly a third and materially move the central value. Our independent pipeline reproduces the published Pantheon+ constraint when the same (z_{HD}) redshift convention is used, validating both the analysis and the released data products.

Data and Code Availability

The Pantheon+/SH0ES data release is publicly available at https://github.com/PantheonPlusSH0ES/DataRelease. The analysis pipeline (data acquisition, filtering, likelihood, fitting, and figure generation) and all intermediate products (filtered data table, sliced covariance, likelihood scan curves, and structured JSON results) are archived with this report.

References

[1] M. M. Phillips. The absolute magnitudes of type ia supernovae. *The Astrophysical Journal Letters*, 413:L105–L108, 1993. doi: 10.1086/186970.

[2] R. Tripp. A two-parameter luminosity correction for type ia supernovae. *Astronomy & Astrophysics*, 331:815–820, 1998.

[3] J. Guy, P. Astier, S. Baumont, D. Hardin, R. Pain, N. Regnault, S. Basa, R. G. Carlberg, A. Conley, S. Fabbro, D. Fouchez, I. M. Hook, D. A. Howell, K. Perrett, C. J. Pritchet, J. Rich, M. Sullivan, P. Antilogus, E. Aubourg, G. Bazin, J. Bronder, M. Filiol, N. Palanque-Delabrouille, P. Ripoche, and V. Ruhlmann-Kleider. Salt2: using distant supernovae to improve the use of type ia supernovae as distance indicators. *Astronomy & Astrophysics*, 466(1):11–21, 2007. doi: 10.1051/0004-6361:20066930.

[4] Adam G. Riess, Alexei V. Filippenko, Peter Challis, Alejandro Clocchiatti, Alan Diercks, Peter M. Garnavich, Ron L. Gilliland, Craig J. Hogan, Saurabh Jha, Robert P. Kirshner, B. Leibundgut, M. M. Phillips, David Reiss, Brian P. Schmidt, Robert A. Schommer, R. Chris Smith, J. Spyromilio, Christopher Stubbs, Nicholas B. Suntzeff, and John Tonry. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009–1038, 1998. doi: 10.1086/300499.

[5] S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko,

T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, and W. J. Couch. Measurements of ω and λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2): 565–586, 1999. doi: 10.1086/307221.

[6] M. Betoule, R. Kessler, J. Guy, J. Mosher, D. Hardin, R. Biswas, P. Astier, P. El-Hage, M. Konig, S. Kuhlmann, J. Marriner, R. Pain, N. Regnault, C. Balland, G. Bazin, R. G. Carlberg, A. Conley, D. Fouchez, J. Guy, D. Hardin, et al. Improved cosmological constraints from a joint analysis of the sdss-ii and snls supernova samples. *Astronomy & Astrophysics*, 568:A22, 2014. doi: 10.1051/0004-6361/201423413.

[7] D. M. Scolnic, D. O. Jones, A. Rest, Y. C. Pan, R. Chornock, R. J. Foley, M. E. Huber, R. Kessler, G. Narayan, A. G. Riess, S. Rodney, E. Berger, D. J. Brout, P. J. Challis, M. Drout, D. Finkbeiner, R. Lunnan, R. P. Kirshner, N. E. Sanders, E. Schlafly, S. Smartt, C. W. Stubbs, J. Tonry, W. M. Wood-Vasey, M. Foley, J. Hand, E. Johnson, W. S. Burgett, K. C. Chambers, P. W. Draper, K. W. Hodapp, N. Kaiser, R. P. Kudritzki, E. A. Magnier, N. Metcalfe, F. Bresolin, E. Gall, R. Kotak, M. McCrum, and K. W. Smith. The complete light-curve sample of spectroscopically confirmed sne ia from pan-starrs1 and cosmological constraints from the combined pantheon sample. *The Astrophysical Journal*, 859(2):101, 2018. doi: 10.3847/1538-4357/aab9bb.

[8] Dillon Scolnic, Dillon Brout, Anthony Carr, Adam G. Riess, Tamara M. Davis, Arianna Dwomoh, David O. Jones, Noor Ali, Pranav Charvu, Rebecca Chen, Erik R. Peterson, Brodie Popovic, Benjamin M. Rose, Charlotte M. Wood, Peter J. Brown, Ken Chambers, David A. Coulter, Kyle G. Dettman, Georgios Dimitriadis, Alexei V. Filippenko, Ryan J. Foley, Saurabh W. Jha, Charles D. Kilpatrick, Robert P. Kirshner, Yen-Chen Pan, Armin Rest, Cesar Rojas-Bravo, Matthew R. Siebert, Benjamin E. Stahl, and WeiKang Zheng. The pantheon+ analysis: The full data set and light-curve release. *The Astrophysical Journal*, 938(2):113, 2022. doi: 10.3847/1538-4357/ac8b7a.

[9] Dillon Brout, Dillon Scolnic, Brodie Popovic, Adam G. Riess, Anthony Carr, Joe Zuntz, Rick Kessler, Tamara M. Davis, Samuel Hinton, David Jones, W. D'Arcy Kenworthy, Erik R. Peterson, Khaled Said, Georgie Taylor, Noor Ali, Patrick Armstrong, Pranav Charvu, Arianna Dwomoh, Cole Meldorf, Anais Möller, Benjamin M. Rose, Bruno Sanchez, Charlotte M. Wood, Peter J. Brown, Rebecca Chen, Ken Chambers, David A. Coulter, Kyle G. Dettman, Georgios Dimitriadis, Alexei V. Filippenko, Ryan J. Foley, Saurabh W. Jha, Charles D. Kilpatrick, Robert P. Kirshner, Yen-Chen Pan, Armin Rest, Cesar Rojas-Bravo, Matthew R. Siebert, Benjamin E. Stahl, and WeiKang Zheng. The pantheon+ analysis: Cosmological constraints. *The Astrophysical Journal*, 938(2):110, 2022. doi: 10.3847/1538-4357/ac8e04.

[10] Adam G. Riess, Wenlong Yuan, Lucas M. Macri, Dan Scolnic, Dillon Brout, Stefano Casertano, David O. Jones, Yukei Murakami, Gagandeep S. Anand, Louise Breuval, Thomas G. Brink, Alexei V. Filippenko, Samantha Hoffmann, Saurabh W. Jha, W. D'Arcy Kenworthy, John Mackenty, Benjamin E. Stahl, and WeiKang Zheng. A comprehensive measurement of the local value of the hubble constant with 1 km s^{-1} mpc^{-1} uncertainty from the hubble space telescope and the sh0es team. *The Astrophysical Journal Letters*, 934(1):L7, 2022. doi: 10.3847/2041-8213/ac5c5b.

[11] A. Conley, J. Guy, M. Sullivan, N. Regnault, P. Astier, C. Balland, S. Basa, R. G. Carlberg, D. Fouchez, D. Hardin, I. M. Hook, D. A. Howell, R. Pain, N. Palanque-Delabrouille, K. M. Perrett, C. J. Pritchet, J. Rich, V. Ruhlmann-Kleider, D. Balam, S. Baumont, R. S. Ellis, S. Fabbro, H. K. Fakhouri, N. Fourmanoit, S. González-Gaitán, M. L. Graham, M. J. Hudson, E. Hsiao, T. Kronborg, C. Lidman, A. M. Mourao, J. D. Neill, S. Perlmutter,

P. Ripoche, N. Suzuki, and E. S. Walker. Supernova constraints and systematic uncertainties from the first three years of the supernova legacy survey. *The Astrophysical Journal Supplement Series*, 192(1):1, 2011. doi: 10.1088/0067-0049/192/1/1.

[12] Anthony Carr, Tamara M. Davis, Dillon Scolnic, Khaled Said, Dillon Brout, Erik R. Peterson, and Rick Kessler. The pantheon+ analysis: Improving the redshifts and peculiar velocities of type ia supernovae used in cosmological analyses. *Publications of the Astronomical Society of Australia*, 39:e046, 2022. doi: 10.1017/pasa.2022.41.

[13] Erik R. Peterson, W. D'Arcy Kenworthy, Dillon Scolnic, Adam G. Riess, Dillon Brout, Anthony Carr, Hélène Courtois, Tamara Davis, Arianna Dwomoh, David O. Jones, Brodie Popovic, Benjamin M. Rose, and Khaled Said. The pantheon+ analysis: Evaluating peculiar velocity corrections in cosmological analyses with nearby type ia supernovae. *The Astrophysical Journal*, 938(2):112, 2022. doi: 10.3847/1538-4357/ac4698.

[14] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al. Planck 2018 results. vi. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910.