



Astrometric Redshifts of Supernovae

Jaemyoung (Jason) Lee¹ , Masao Sako¹ , Richard Kessler² , and Alex I. Malz³

The LSST Dark Energy Science Collaboration

¹ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA; astjason@sas.upenn.edu² Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA³ McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Received 2024 May 4; revised 2024 October 8; accepted 2024 October 9; published 2024 December 12

Abstract

Differential Chromatic Refraction (DCR) is caused by the wavelength dependence of our atmosphere's refractive index, which shifts the apparent positions of stars and galaxies and distorts their shapes depending on their spectral energy distributions. While this effect is typically mitigated and corrected for in imaging observations, we investigate how DCR can instead be used to our advantage to infer the redshifts of supernovae from multiband, time-series imaging data. We simulate Type Ia supernovae in the proposed Vera C. Rubin Observatory Legacy Survey of Space and Time Deep Drilling Field, and evaluate astrometric redshifts. We find that the redshift accuracy improves dramatically with the statistical quality of the astrometric measurements as well as with the accuracy of the astrometric solution. For a conservative choice of a 5 mas systematic uncertainty floor, we find that our redshift estimation is accurate at $z < 0.6$. We then combine our astrometric redshifts with both host-galaxy photometric redshifts and supernovae photometric (light-curve) redshifts and show that this considerably improves the overall redshift estimates. These astrometric redshifts will be valuable, especially since Rubin will discover a vast number of supernovae for which we will not be able to obtain spectroscopic redshifts.

Unified Astronomy Thesaurus concepts: Type Ia supernovae (1728); Astrometry (80); Atmospheric refraction (115)

1. Introduction

Observations of Type Ia supernovae (SNe Ia) resulted in the unexpected but groundbreaking discovery of the accelerating universe (A. G. Riess et al. 1998; S. Perlmutter et al. 1999). While this discovery with only tens of SNe Ia is an extraordinary feat, subsequent efforts have allowed for increasingly precise constraints on the dark energy equation-of-state parameter w and the matter density Ω_m . Indeed, recent measurements from Pantheon+ (D. Brout et al. 2022) and the Dark Energy Survey (DES; DES Collaboration et al. 2024) now give roughly 0.15 uncertainty on the equation-of-state parameter w_0 for a flat w_0 CDM model from over 1000 SNe each. Combining supernovae measurements with Cosmic Microwave Background (CMB) and Baryon Acoustic Oscillations (BAO) measurements, which are largely independent from SNe Ia measurements, significantly improves the cosmological constraints.

Accurate redshifts are crucial for SN Ia cosmology. Up to this date, SN redshifts have relied on spectroscopic redshifts, either of the SN itself or from the host galaxy (D. A. Howell et al. 2005; M. Sako et al. 2018; C. Lidman et al. 2020; M. Smith et al. 2020). However, the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; Ž. Ivezić et al. 2019) will observe orders of magnitude more SNe than all of the SNe detected so far, making it no longer viable to perform spectroscopic follow-up for all of the detected SNe and their hosts.

Consequently, there have been numerous efforts to utilize SN photometry to accurately measure redshifts. Some of these efforts include using the peak flux in the observed bands as inputs for an empirical, analytic redshift estimator (Y. Wang 2006; Y. Wang et al. 2007, 2015), adding the redshift as an additional light-curve fitting parameter, also known as LCFIT+z (R. Kessler et al. 2010; N. Palanque-Delabrouille et al. 2010), and using a machine learning approach to learn the redshifts from light curves (H. Qu & M. Sako 2023). Many studies have sought to improve SN photometric redshift measurements by adding a host-galaxy photometric redshift (photo- z) prior, including three of the aforementioned works, as well as A. Mitra et al. (2023), where the resulting improvement in cosmological constraints were explored extensively.

While the impact of photometric redshift uncertainties on cosmology may not be severe (M. Dai et al. 2018; R. Chen et al. 2022; A. Mitra et al. 2023), better redshift measurements will lead to more constraining power in the cosmological parameter estimation by reducing the uncertainties in the redshifts used for cosmology. Here, we introduce another method to obtain redshift measurements using imaging data. This method utilizes the fact that sources at different redshifts (and hence different colors) are refracted by different amounts in our atmosphere. As illustrated in Figure 1, shorter-wavelength light is refracted more in our atmosphere than longer-wavelength light. This effect, known as Differential Chromatic Refraction (DCR; A. V. Filippenko 1982), is a result of the refractive index of our atmosphere being wavelength dependent.

Normally, DCR is an effect cosmologists want to mitigate for precision cosmology in both weak lensing (A. Alejandro Plazas & G. Bernstein 2012; J. E. Meyers & P. R. Burchat 2015; S. G. Carlsten et al. 2018) and SN Ia cosmology (J. Lee et al. 2023). However, M. C. Kaczmarszik et al. (2009) found that treating the observed DCR shifts in the Sloan Digital Sky Survey



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.




Figure 1. A schematic diagram showing how DCR affects light in the Earth’s atmosphere; shorter wavelengths are refracted more by the atmosphere, which means that, to an observer on Earth, a blue star appears to be higher up in the sky than a red star located at the same position in space. $R_{450\text{nm}}$ and $R_{650\text{nm}}$ denote the amount of refraction by our atmosphere at an air mass (AM) of 1.4 for 450 nm and 650 nm light, respectively.

(SDSS) u and g bands as additional colors in photometric redshifts improves quasar photometric redshifts, with the fraction of objects within $\Delta z = \pm 0.1$ of the spectroscopic redshifts increasing by 9% and catastrophic outliers reduced. As such, astrometric redshifts have the potential to strengthen photo- z constraints when combined with photo- z s. They also mention the possibility of applying these astrometric redshifts (astro- z s) to other sources with distinct spectral features such as Type Ia and Type II supernovae.

In this work, we develop the methodology to obtain astrometric redshifts for Type Ia supernovae. For this development, we added DCR effects to the SNANA⁴ (Super-Nova ANALysis; R. Kessler et al. 2009) simulation, and analyzed simulated SN Ia samples corresponding to the proposed Deep Drilling Fields (DDF) for LSST. We show astro- z constraints from only the DCR effect, utilizing multi-band, multi-epoch measurements, and we show how combining them with both host and SN photo- z s improves the overall redshift estimates.

The content of this work is as follows. In Section 2, we compute the amount of DCR shift that occurs depending on the air mass, source flux and filter response functions and describe how SNe at different redshifts are affected. Next, we describe the simulated LSST data set in Section 3. Then, we describe our methods to obtain redshift estimates from the simulated astrometry for each SN observation in Section 4. In Section 5, we show our results for LSST-like simulations in hypothetical (ideal) and realistic cases, as well as the improvement when combined with both host-galaxy and SNe photometric redshifts, and we extensively discuss our results. Finally, we end with a discussion and conclusion in Section 6.

2. Differential Chromatic Refraction

DCR is due to the dependency of our atmosphere’s refractive index on the source wavelength. The centroid shift caused by

DCR can be calculated as in A. Alejandro Plazas & G. Bernstein (2012):

$$\bar{R} = \frac{\int d\lambda R(\lambda, z_a) F(\lambda) \lambda S(\lambda)}{\int d\lambda F(\lambda) \lambda S(\lambda)}, \quad (1)$$

where \bar{R} is the altitude shift of the source toward the zenith, $R(\lambda, z_a)$ is the shift of each photon toward the zenith, z_a is the zenith angle, $F(\lambda)$ is the filter response function, and $S(\lambda)$ is the source spectral energy distribution (SED).⁵ $R(\lambda, z_a)$ can be calculated using

$$R(\lambda, z_a) \approx R_0 \tan(z_a), \quad R_0 = \frac{n^2(\lambda) - 1}{2n^2(\lambda)}, \quad (2)$$

with the air mass (AM), or the amount of air along the line of sight, being $\text{AM} = \sec z_a$ for $z_a < 80^\circ$. For an explicit expression of $n(\lambda)$, see A. V. Filippenko (1982). As apparent in Equation (1), the DCR altitude shift depends on the filter transmission function. It also depends on the source SED, which depends on both redshift and the epoch of the SN. The dependency on $R(\lambda, z_a)$ means that the shift becomes larger at higher AM.

Figure 2 shows that, as the SN redshift increases, the overlap of its SED with the filter function changes, so the DCR altitude shift changes with redshift. Although the amount of refraction that occurs due to DCR is usually between $43''$ and $46''$ for a given source viewed through the LSST $ugriz$ bands at an AM of 1.4, the observed shifts are much smaller because DCR also affects the reference stars that are used to determine the astrometric solution. Hence, we measure the shift with respect to the average-color star in a given CCD, which we take in this work to be a K5V star, adopted from the DES-SN5YR analysis (J. Lee et al. 2023). For LSST, we do not know beforehand the spectrum of the average-color star, but it is straightforward to use another spectral type as the reference star. Since the reference stellar type only determines the amplitudes of the measured DCR shifts, our subsequent results are not sensitive to the choice of reference stellar type unless its spectrum is significantly different from the one chosen here. Figure 3 shows the g -band altitude shift versus AM for an SN (with the light-curve stretch and color being 0) at rest-frame epoch $T_{\text{rest}} = 0$ for $z = 0$ and $z = 0.6$, where T_{rest} is the number of days from the time of peak brightness in the g band. Compared to the average-color star, the $z = 0$ SN is DCR-shifted upward (toward zenith) while the $z = 0.6$ SN is shifted downward (away from zenith) at $\text{AM} > 1.0$. We also show the expected increase in the DCR shifts at higher AM.

3. Data Set and Simulations

We perform our analysis on a simulated data sample generated assuming LSST survey characteristics. LSST is a ground-based survey using the 8.4 m Simonyi Survey Telescope in Chile with expected first light in 2025. Over the period of a decade, LSST will observe $18,000 \text{ deg}^2$ of the southern sky with its 3.2 gigapixel camera and a very wide 9.6 deg^2 field of view in six filters $ugrizY$ covering 320–1050 nm. We only use the DDF subsample (D. M. Scolnic et al. 2018) obtained by

⁴ <https://github.com/RickKessler/SNANA>

⁵ Note that the factor of λ is necessary in the integrands of both the numerator and denominator if $S(\lambda)$ is taken to be the spectral energy distribution. This takes into account that *photon-counting* detectors are used for our observations.

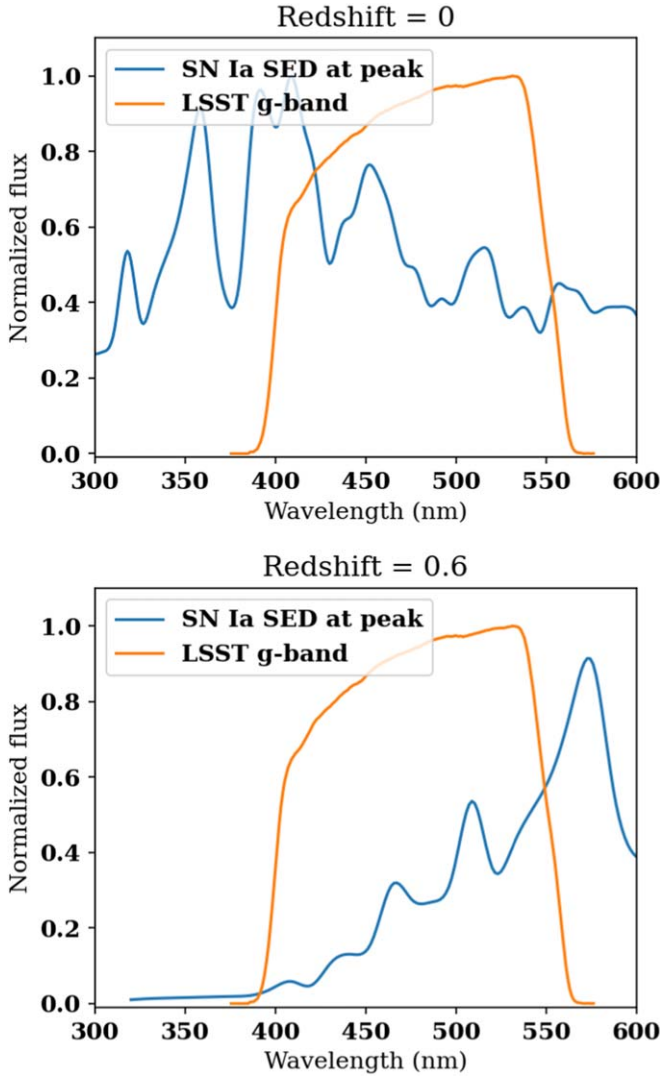


Figure 2. Example SN Ia SED at peak epoch overlaid on top of the LSST g band at $z = 0$ (top) and $z = 0.6$ (bottom). As shown in Equation (1), the resulting DCR shifts at the two redshifts are different since the overlap between the filter and the SED changes significantly. Because the product of the filter function and the SED, $F(\lambda)S(\lambda)$, is smaller for the $z = 0.6$ SN, the (absolute) DCR shift is also smaller, as shown in Figure 3.

coadding exposures from each of the bands taken within the same night.

We simulate LSST-like SNe Ia observations in the *ugrizY* bands using SNANA, which uses LSST-like SN Ia observing characteristics following the Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) data set (G. Narayan 2023), which is a sequel to the earlier PLAsTiCC data set (A. Bahmanyar & R. Biswas et al. 2018; R. Kessler et al. 2019b). The simulation includes realistic SN observations with a source model that generates SEDs for each epoch, galactic extinction that impacts the SED, and a noise model that accounts for the PSF, sky noise, and zero point as highlighted in R. Kessler et al. (2019a).

For our SED model in the simulation and analysis, we use the SALT3 model (W. D. Kenworthy et al. 2021) with wavelengths extended to the near-infrared (coverage up to 2000 nm) following J. D. R. Pierel et al. (2022). The SALT2 (J. Guy et al. 2007) and SALT3 models have a known artifact of producing negative SED fluxes in the ultraviolet region for




Figure 3. The resulting g -band DCR altitude shifts with respect to a K5V reference star for a $z = 0$ SN (blue) and $z = 0.6$ SN (orange) at different air masses (AMs), but both at peak epoch. As we expect, the shifts at the two redshifts are different and we see more shift at higher AM. Comparison with Figure 2 implies that the $z = 0$ SN is “bluer” (in the g band at least) than the reference star while the $z = 0.6$ SN is “redder.”

rest-frame epochs $T_{\text{rest}} < -15$ days. We hope that better-quality data in the future will fix this artifact, but here we simply force negative fluxes to be zero. We note that observations at $T_{\text{rest}} < -15$ days are rare occurrences; less than 1% of our simulated observations are at $T_{\text{rest}} < -15$ days. Host-galaxy properties are modeled as described in M. Lokken et al. (2023).

Some survey characteristics particularly important for this work include the filter response functions⁶ shown in Figure 4, and the simulated air mass as well as redshift distributions, shown in Figure 5. Because the DCR shifts are typically the largest in the u and g bands, observations in those bands will be especially helpful in measuring SN redshifts. The redshift distribution peaks around $z = 0.9$ without selection requirements (cuts; see end of Section 3), but at $0.6 < z < 0.9$ with cuts. Most of the observations are at $1.05 < \text{AM} < 1.25$, regardless of the selection cut, meaning that the vast majority of the SNe observations will result in some sort of shift away

⁶ <https://github.com/lst/throughputs/tree/1.9/baseline>




Figure 4. LSST *ugriz* filters (from left to right) according to their relative transmission and wavelength range. We do not show the *Y* band, as it is not used for the DCR calculations. Multiple detections with different bands are crucial for accurate astrometric redshifts.

from the reference star. Roughly 23% of the observations are at high AM (>1.4) where we see the most extreme shifts.

To simulate realistic astrometry, we added new functionalities in the SNANA simulations. We simulate true DCR shifts in the R.A. (α) and decl. (δ) directions by first calculating the DCR altitude shifts for the SN ($\bar{R}_{\text{ALT,SN}}$) and K5V star ($\bar{R}_{\text{ALT,K5V}}$) by performing the integral given in Equation (1) then adding $\Delta_{\text{DCR}}\alpha$ and $\Delta_{\text{DCR}}\delta$ to the reference star (K5V) R.A. and decl. after DCR, where

$$\begin{aligned}\Delta_{\text{DCR}}\alpha &= (\bar{R}_{\text{ALT,SN}} - \bar{R}_{\text{ALT,K5V}}) \frac{\sin q}{\cos \delta} \\ \Delta_{\text{DCR}}\delta &= (\bar{R}_{\text{ALT,SN}} - \bar{R}_{\text{ALT,K5V}}) \cos q\end{aligned}\quad (3)$$

and q is the parallactic angle. While Equation (3) is an approximation, the true shift calculated using spherical trigonometry results in a very similar answer. This is because the shifts considered are smaller than $1''$. We show in Appendix A a diagram of the celestial sphere for an observer in the Southern Hemisphere as well as detailed calculations of the shifts.

Next, because real observations include Poisson noise, the SN centroid accuracy depends on the S/N of the SN measurement. Thus, we add random Gaussian smearing as a function of $(\text{FWHM}_{\text{PSF}}/[\text{S/N}])$ for our statistical uncertainties. To parameterize the angular resolution, we use the GalSIM⁷ (B. T. P. Rowe et al. 2015) package to simulate images with *only* Poisson noise, calculate the standard deviation of R.A. and decl. shifts depending on the FWHM_{PSF} and S/N at multiple AMs, and fit a second-order polynomial of $(\text{FWHM}_{\text{PSF}}''/[\text{S/N}])$. Ideally, we would use GalSIM to simulate each SN observation including Poisson noise and DCR effects, as well as the host galaxy, and measure the centroid shift for each observation. However, this is computationally expensive, especially if we were to use the full scene-modeling pipeline that forward models the SN Ia and host galaxy as done for typical SN Ia surveys (P. Astier et al. 2006; D. Brout et al. 2019; P. H. Bernardinelli et al. 2023), which can take tens of hours per candidate. To reduce the computation time, we chose

to simulate the random shift arising from Poisson noise and DCR instead of simulating each image in this initial investigation. The quadrature sum of our fitted polynomial function for the standard deviation and a systematic floor gives the total uncertainties, $\sigma_{\text{tot}}^2 = \sigma_{\text{syst}}^2 + \sigma_{\text{stat}}^2$, where

$$\sigma_{\text{stat}} = 0.66 \left(\frac{\text{FWHM}_{\text{PSF}}''}{[\text{S/N}]} \right) + \frac{0.1}{1''} \left(\frac{\text{FWHM}_{\text{PSF}}''}{[\text{S/N}]} \right)^2, \quad (4)$$

with $\sigma_{\text{syst}} = 5$ mas. The simulated random R.A. and decl. shifts due to the total uncertainties are

$$\Delta_{\sigma,\text{tot}}\alpha = \sigma_{\text{tot}} \frac{\mathcal{N}_{\alpha}(0, 1)}{\cos \delta}, \quad \Delta_{\sigma,\text{tot}}\delta = \sigma_{\text{tot}} \mathcal{N}_{\delta}(0, 1), \quad (5)$$

where $\mathcal{N}_{\alpha}(0, 1)$ and $\mathcal{N}_{\delta}(0, 1)$ are random Gaussian deviates each with a mean of zero and standard deviation of 1. Combining everything, we have

$$\alpha_{\text{obs}} = \alpha_{\text{true}} + \Delta_{\text{DCR}}\alpha + \Delta_{\sigma,\text{tot}}\alpha, \quad (6)$$

and likewise for δ_{obs} . When the S/N is too low, σ_{stat} becomes larger than the largest DCR shifts we observe, or around $0''.1$ – $0''.2$, which typically occur at high-AM observations in the *ug* bands at $z < 0.1$. Thus, we only use DCR shifts when the S/N is larger than 3 ($\sigma_{\text{stat}} \approx 0''.23$ for a $1''$ PSF when $\text{S/N} = 3$). It is possible in theory to use all observations, including when $\text{S/N} < 3$, but this is likely not feasible for real observations, due to the difficulty of accurately fitting for centroids at such low S/N. Because Poisson noise causes the centroids to shift in random directions, using only $\text{S/N} > 3$ SNe for our analysis will not result in a bias.

To compare the simulated DCR measurement and the model DCR shift for a given epoch, we subtract the peak MJD of the SN from its observed date and take this value to be the observed epoch T_{obs} . To estimate peak MJDs, we use the “Fmax-clump” method as described in the SNANA manual, where peak MJD is defined to be the epoch of maximum flux in the sliding 50 days window containing the most detections. This method is robust because there is no fitting, and it ignores pathological fluxes that are far away from peak brightness. Next, we use the light-curve fitting module LCFIT+z (R. Kessler et al. 2010) with a weak cosmology prior as described in Appendix D, only for the purpose of obtaining light-curve fit peak MJD values that are input to the astro-*z* pipeline. Since we use the very high-cadence DDF only, the “Fmax-clump” method is adequate. However, for surveys with poorer cadence (e.g., WFD), a light-curve fit may be needed to get a better estimate of peak MJDs. Because the light-curve fitting is not always successful, this process leaves us with 13,735 out of 20,000 candidates, or 7804 out of the 8588 candidates that pass the selection cuts as described later in this section. The primary reason for unsuccessful light-curve fitting is not having observations before and after the peak epoch, while having low S/N is another reason. Thus, we combine the astro-*z* measurements obtained using LCFIT+z peak MJDs (when available) with “Fmax-clump” peak MJDs (when LCFIT+z peak MJDs are not available) when showing our default results in Section 5.2. We also show some performance metrics (i) with only “Fmax-clump” peak MJDs and (ii) assuming we know precisely the peak MJDs in Table 7.

Our simulation does not account for other wavelength-dependent effects in the telescope optics and detectors affecting SN centroid measurements. Some of these effects were investigated in detail for DECam in G. M. Bernstein et al. (2017, 2018).

⁷ https://galsim-developers.github.io/GalSim/_build/html/index.html




Figure 5. Redshift and air mass (AM) distributions for the LSST DDF simulations in SNANA. The peak redshift is around $z = 0.9$ for the No selection cut and the Epoch cut, but around $0.6 < z < 0.9$ for the Epoch + S/N cut, while the AM peaks between 1.05 and 1.10 for all three selection cuts. Higher-AM observations lead to better astro- z estimates.

The largest optics effect comes from refraction in the lens, shifting sources in different bands by different amounts, in particular near the boundaries of the focal plane. Such color-dependent radial displacements (also known as lateral color) are shown to be around $0''.050 \text{ mag}^{-1}$ and $0''.005 \text{ mag}^{-1}$ in the g and r bands in $g-i$ magnitudes, respectively, near the edges of the focal plane for DES. While J. Lee et al. (2023) mentions that these shifts are, at worst, similar to the DCR positional effects for DES, it is plausible to expect that at least some of these effects will be corrected for in LSST. For DECam, the uncertainty in the astrometric solution due to telescope optics such as unmodeled stray electric fields in the detectors as well as focal plane shifts between camera thermal cycles is about 3–6 mas on the focal plane (G. M. Bernstein et al. 2017), so our assumption of a $\sigma_{\text{syst}} = 5$ mas systematic floor is plausible. Therefore, the uncertainties we model in this study will account for the bulk of the uncertainties we will face in real observations.

Because we advocate the combination of photo- z s with our astro- z s, we apply selection cuts typically applied to photo- z s to our simulated samples when showing results for realistic simulations. Following R. Kessler et al. (2010), we first require that there are observations before $T_{\text{rest}} = -3$ days and after $T_{\text{rest}} = 10$ days when showing the results for all types of simulations. This means that, for each candidate we use, there are observations before and after the peak epoch. Given that we do not know T_{rest} beforehand, we relax the epoch cuts to $T_{\text{rest}}/(1 + z_{\text{max}})$ as in R. Kessler et al. (2010), or observations before $T_{\text{obs}} = -3/2.2$ days and after $T_{\text{obs}} = 10/2.2$, because our $z_{\text{max}} = 1.2$, allowing for more range in T_{obs} . For the same reason, we use peak MJDs from the “Fmax-clump” method to impose epoch cuts. Second, we require that there are at least three bands with detections of $S/N > 8$ for realistic simulations. The numbers of candidates that remain after selection cuts for the astro- z -only results are shown in Table 1. After imposing the epoch selection cut, we are left with 13,827 or 69.1% of the SNe Ia, while also imposing the S/N cut leaves us with 8,588 SNe Ia out of the simulated 20,000, or about 42.9%.

Finally, for the SN photo- z and its combinations with the other two redshift estimates where the light-curve fitting

Table 1
Number of Candidates That Pass the Given Selection Cut

Type	Number of Candidates
No selection cut	20,000
Epoch cut	13,827
Epoch + S/N cut (default)	8588
Epoch + S/N cut + AM cut	7528

Notes. Out of the 20,000 simulated candidates, the ones that pass the Epoch cut are used to show the results for Perfect and σ_{syst} -only simulations, while the default for realistic simulations is Epoch + S/N cut. The AM cut additionally requires that at least one observation for a candidate has an AM > 1.4 .

method LCFIT+ z is used, we take the fit probability (FITPROB), calculated by integrating the tail of the χ^2 distribution from the χ^2 of the light-curve fit and the number of degrees of freedom for each of the candidates (p -value), and require that $\text{FITPROB} \geq 0.01$. The number of candidates removed by this selection cut is less than 10% for all combinations, as shown in Table 5, but it noticeably improves the performance metrics.

4. Analysis Methodology

To measure the astro- z for a given SN, we minimize the χ^2 between the model and SNANA simulated DCR shifts. As introduced in Section 3, we use the extended SALT3 model to calculate the model DCR altitude shifts (MODEL) on a grid of (FILTER, z , EPOCH, AM, x_1 , c) with x_1 and c being the SALT3 light-curve parameters for stretch and color. The grids for each of the parameters are

$$\begin{aligned}
 z &: 0.00 \text{ to } 1.20, \text{ increments of } \Delta z = 0.01 \\
 \text{EPOCH} &: -18 \text{ to } 50, \text{ increments of } 1 \text{ day} \\
 \text{AM} &: 1.00 \text{ to } 3.00, \text{ increments of } 0.01 \\
 x_1 &: -3.0 \text{ to } 2.0, \text{ increments of } 0.5 \\
 c &: -0.30 \text{ to } 0.50, \text{ increments of } 0.05.
 \end{aligned} \tag{7}$$

As mentioned earlier in Section 3, we force the SED to be zero whenever the computed SALT3 fluxes are below zero when

calculating the DCR model shifts. These negative fluxes sometimes occur at early phases ($T_{\text{rest}} \leq -15$) in the ultraviolet, mostly at extreme color ($|c| \geq 0.2$) and stretch ($|x_1| \geq 2$). This can sometimes cause an issue, because the DCR shifts are infinite when the denominator of Equation (1) is zero. In such cases, we extrapolate the DCR shifts along the epoch axis using `scipy.interpolate`.

We measure the DCR shift for simulated data similarly to how we would for real data. First, we use the “observed” R.A. and decl. to compute the altitude of the SN at the telescope site, and subtract the altitude of the (σ_{tot}^2 -weighted) Y -band average R.A. and decl. for each candidate. This subtraction is needed because we cannot determine the exact coordinates of the K5V reference star, as all the astronomical objects in a given exposure are shifted by DCR as well. Hence, we take the Y -band average coordinates of the SN to be a proxy for the reference star coordinates, as the DCR shifts are smallest in the Y -band. While Poisson noise randomly scatters the Y -band positions, the average Y -band coordinates among multiple observations constitute a reasonable approximation. As only about 12,700 candidates have Y -band observations, we use the z -band average coordinates instead when we do not have Y -band observations for a given SN candidate. When we use the z -band average coordinates as the K5V reference star coordinates, we do not use the z -band DCR shifts when constructing our probability distribution functions, since the z -band average coordinates are sensitive to the DCR shifts in the z band. Further discussion on the choice of reference star is given in Appendix B.

The χ^2 between the model and the observed altitude shifts is defined as

$$\chi^2 = \sum_i \frac{(\Delta_{\text{OBS},i} - \Delta_{\text{MODEL},i})^2}{\sigma_{i,\text{total}}^2}, \quad (8)$$

where OBS indicates a real or simulated observation, i denotes each of the observations for a given SN, and $\sigma_{i,\text{total}}^2 = \sigma_{i,\text{stat}}^2 + \sigma_{i,\text{syst}}^2$, with $\sigma_{i,\text{stat}}$ given in Equation (4) and $\sigma_{i,\text{syst}} = 5$ mas as stated in Section 3. After the χ^2 is calculated for a given candidate over the three-dimensional grid of FILTER, x_1 , c , we calculate the posterior:

$$\mathcal{P}(z) = \int_{-0.3}^{0.5} \int_{-3.0}^{2.0} P(z, x_1, c) P(x_1) P(c) dx_1 dc, \quad (9)$$

where $P(z, x_1, c) \propto e^{-\chi^2/2}$ with the χ^2 being the sum over the $ugriz$ band χ^2 ($ugri$ when Y -band observations are absent) and $P(x_1)$ and $P(c)$ are priors for the light-curve stretch and color. $P(x_1)$ and $P(c)$ are obtained by binning asymmetric Gaussian distributions with the binning shown in Equation (7). The values with highest-probability, low-sided, and high-sided Gaussian widths are taken to be $(\bar{x}_1, \sigma_{x_1,-}, \sigma_{x_1,+}) = (-0.054, 0.043, 0.101)$ and $(\bar{c}, \sigma_{c,-}, \sigma_{c,+}) = (0.973, 1.472, 0.222)$ as in D. Scolnic & R. Kessler (2016), which are also used as the SNANA input. We do not need to normalize $\mathcal{P}(z)$, since we only need to compute the relative probabilities at each redshift.

We use $\mathcal{P}(z)$ to estimate the redshift in two different ways: (1) taking the redshift at the 50th percentile of the cumulative distribution function (hereafter, CDF50) and defining the 68% confidence interval as 16th to 84th percentiles, and (2) taking the redshift at the peak of the probability distribution function (hereafter, PDFPeak;

equivalent to the maximum likelihood) with uncertainties defined as the ± 34 percentiles around the peak.

There are advantages and disadvantages for either method. CDF50 gives better results when the posterior probabilities are multimodal without a distinct peak and the error bars are always within the range of redshifts we explore. However, when the χ^2 is unconstraining, the redshift is measured to be near the midpoint of the redshift range, meaning that the results are biased. On the other hand, PDFPeak can result in catastrophic outliers if the posterior has multimodal peaks or has similar values throughout all redshifts. Figure 6 shows examples for which (i) CDF50 gives a better redshift estimate (multimodal with similar probabilities), (ii) PDFPeak gives a better estimate (distinct peak but not large in amplitude), and (iii) both CDF50 and PDFPeak are able to determine the redshift very well (well-defined peak).

In addition to our astro- z results, we combine our astrometric redshifts with host-galaxy photometric redshifts (Host photo- z s) as well as SN Ia photometric redshifts (SN photo- z s) in Section 5. Henceforth, we denote the redshift measurements using the three methods as well as their combinations as shown in Table 2. z_{Host} are derived from ELASTiCC host photometry by using the `pzflow` package⁸ (J. F. Crenshaw 2021) as a generative model, trained on the CosmoDC2 simulations (D. Korytov et al. 2019), and are stored as quantiles corresponding to integrated (CDF) probabilities of [0%, 10%, ..., 100%]. The z_{SN} are obtained from the light-curve fitting method `LCFIT+z` (R. Kessler et al. 2010) using the $ugrizY$ bands as with our z_{DCR} simulations, where the redshift is fit simultaneously along with the light-curve parameters, i.e., time of maximum brightness t_0 , color c , stretch x_1 , and flux normalization x_0 .

When combining z_{DCR} with z_{Host} , we use the `qp` package⁹ (A. I. Malz et al. 2018) to reconstruct PDFs from the z_{Host} quantiles before multiplying the z_{DCR} and z_{Host} PDFs together. To combine either z_{DCR} or z_{Host} (or both) with z_{SN} , we provide the respective quantiles as priors in `LCFIT+z`. We do not determine $P(z_{\text{DCR}})$, $P(z_{\text{Host}})$, and $P(z_{\text{SN}})$ independently and combine the three PDFs for a joint PDF, because the SALT3 parameters such as x_0 , x_1 , and c and their uncertainties would be incorrectly associated with $P(z_{\text{SN}})$ only. Thus, the correct method is to combine $P(z_{\text{Host}})$ and $P(z_{\text{DCR}})$ and use this combined prior in `LCFIT+z` to obtain not only the combined redshift estimates but also the SALT3 parameters with correct uncertainties. When showing combined results, we impose the default selection cut described in Section 3 on the candidates with successful light-curve fits with z_{SN} and just the default selection cut when z_{SN} are not used.

Sometimes, the z_{DCR} and z_{Host} estimates cannot be determined (e.g., when the SN candidate has no z - or Y -band observations for z_{DCR}). In this case, we replace the PDF with a flat prior; if there is no z_{DCR} estimate for a candidate, only the z_{Host} prior is used and vice versa.

When combining two or more PDFs, sometimes the PDFs have negligible overlap, meaning that it is meaningless to combine them. Notable cases occur at $z_{\text{true}} < 0.4$, where z_{DCR} performs better than z_{Host} ; $z_{\text{DCR}+\text{Host}}$ performs worse than z_{DCR} by itself when z_{DCR} and z_{Host} are combined without considering the compatibility between z_{DCR} and z_{Host} PDFs. Hence, we impose an additional selection cut for $z_{\text{DCR}+\text{Host}}$ when z_{DCR} and

⁸ <https://github.com/jfcenshaw/pzflow>

⁹ <https://github.com/LSSTDESC/qp>




Figure 6. The relative posterior probabilities, true z , CDF50, and PDFPeak estimates for three cases highlighting the advantages and disadvantages for the two methods: (Left) CDF50 gives a better estimate (multimodal with similar peak amplitudes). (Middle) PDFPeak gives a better estimate (distinct peak but not large in amplitude). (Right) Both CDF50 and PDFPeak estimate the correct redshift.

Table 2
Denotation of Our Redshift Estimates Using Various Methods

Symbol	Measurement Method
z_{DCR}	from astrometry using DCR effect
z_{Host}	host photo- z PDF
$z_{\text{DCR+Host}}$	combined PDF from z_{DCR} and z_{Host}
z_{SN}	from LCFIT+ z , using flat redshift prior
$z_{\text{SN+Host}}$	from LCFIT+ z , using z_{Host} PDF as prior
$z_{\text{SN+DCR}}$	from LCFIT+ z , using z_{DCR} PDF as prior
$z_{\text{SN+DCR+Host}}$	from LCFIT+ z , using $z_{\text{DCR+Host}}$ PDF as prior

z_{Host} PDFs are incompatible. We assess the compatibility between the two PDFs using the Overlapping Index (OVL) between two PDFs, which is defined in the literature such as H. F. Inman & E. L. Bradley (1989), but originally postulated in a related form in K. Pearson (1894):

$$\text{OVL}(p_0, p_1) = \int \min(p_0(x), p_1(x)) dx, \quad (10)$$

where $p_0(x)$ and $p_1(x)$ are PDFs. Based on the performance metrics (discussed in the beginning of Section 5) for $z_{\text{DCR+Host}}$ compared to z_{DCR} and z_{Host} by themselves, we remove candidates for $z_{\text{DCR+Host}}$ when $\text{OVL} < 0.0344$, or when two normalized Gaussian distributions with the same standard deviations (σ) have means that are $3\sqrt{2}\sigma$ apart from each other. This choice eliminates only about 2.3% of the candidates that pass the Default selection cut, mostly at the lowest redshifts where incorrect z_{Host} PDFs can degrade $z_{\text{DCR+Host}}$ considerably, but significantly improves performance metrics at those redshifts. We note that removing candidates using an OVL cut may not be the most optimal way to treat the combination of incompatible PDFs, but it is sufficient for this analysis. While a small OVL can result from distributions with the same mean but very different widths, in our analysis, the use of OVL is (1) limited to a very small number of the candidates and (2) applied when combining with z_{DCR} and z_{Host} , neither of which have very narrow PDFs.

The LCFIT+ z implementation in SNANA produces a z_{SN} mean and uncertainty, not a PDF, and therefore an OVL value cannot be computed using z_{SN} . The $\text{FITPROB} \geq 0.01$ cut mentioned in Section 3 rejects most of the incompatible PDFs, and we believe that our treatment is sufficient for this initial investigation of combining SN Ia astrometric redshifts with

photometric redshifts. In Appendix C, we discuss the overlapping index and the treatment of z_{SN} PDFs in more detail.

Additionally, for $z_{\text{SN+DCR+Host}}$, we use either z_{DCR} alone or z_{Host} alone as priors when z_{DCR} and z_{Host} are incompatible with each other, depending on whether the overlapping index between z_{SN} and z_{DCR} or z_{Host} is higher. This choice further improves the performance metrics at low- z , as with $z_{\text{DCR+Host}}$ discussed earlier.

5. Results

In this section, we show the results of our z_{DCR} estimates for LSST-like simulations for: (i) Perfect ($\sigma_{\text{stat}} = 0$ and $\sigma_{\text{syst}} = 0$) and (ii) Realistic (with nonzero σ_{stat} and σ_{syst}) scenarios. In Appendix G, we show systematic effects only. For the Realistic case, we also show the results combined with z_{Host} as well as with z_{SN} obtained from LCFIT+ z . Along with the comparison between the estimated redshifts and true redshifts, we also show three different metrics depending on the case and the combination of PDFs: the mean bias (or binned residuals), outlier fraction, and the Median Absolute Deviation (MAD) deviation.

As is customary in photo- z literature (R. Kessler et al. 2010; J. Pasquet et al. 2019; H. Qu & M. Sako 2023), we define the residuals as

$$\Delta z \equiv \frac{z_{\text{estimate}} - z_{\text{spec}}}{1 + z_{\text{spec}}}, \quad (11)$$

where z_{estimate} is the estimated redshift and the bias is $\langle \Delta z \rangle$. z_{spec} refers to the spectroscopic redshift of the SN Ia or its host galaxy for an analysis using real observations, but is taken to be the true redshift from the simulations in this work. The uncertainty on $\langle \Delta z \rangle$ is defined as $\frac{\text{RMS}(\Delta z - \langle \Delta z \rangle)}{\sqrt{N_{\text{SN}}}}$, where rms is the root mean square and N_{SN} is the number of SN events in the sample or redshift bin. The outlier fraction η_x is defined as the fraction of candidates, where $|\Delta z| > x$: $x = 0.10$ is our default. Finally, the precision metric MAD deviation is defined as $\sigma_{\text{MAD}} = 1.4826 \times \text{Median}|\Delta z - \text{Median}(\Delta z)|$.

5.1. Perfect Simulations

We first show our results in Figure 7 for the Perfect case without any statistical or systematic uncertainties, assuming that our peak MJD measurements have no error and that we know both x_1 and c . This test validates the analysis and demonstrates the ideal performance of z_{DCR} for the LSST DDF




Figure 7. z_{DCR} vs. z_{spec} in the Perfect case for PDFPeak. As pointed out in Section 4, CDF50 results are equivalent to PDFPeak results for this case, and the 1σ uncertainties are zero. The right panels are the same as the left, but zoomed to $z < 0.6$. The outlier fraction $\eta_{0.10}$ is defined as the fraction where $|\Delta z| > 0.10$, with $\Delta z \equiv \frac{z_{\text{estimate}} - z_{\text{spec}}}{1 + z_{\text{spec}}}$. The imperfect z_{DCR} estimates are due to degeneracies in the DCR shifts at several redshifts, which is most apparent at $z_{\text{spec}} > 0.8$.

Table 3
Bias and Error on the Mean, $\eta_{0.10}$, and σ_{MAD} for Perfect and Realistic Simulations

Type	Bias (CDF50)	Bias (PDFPeak)	$\eta_{0.10}$ (CDF50)	$\eta_{0.10}$ (PDFPeak)	σ_{MAD} (CDF50)	σ_{MAD} (PDFPeak)
Perfect	-0.003 ± 0.000	-0.003 ± 0.000	0.014	0.014	0.008	0.008
Realistic	-0.048 ± 0.001	-0.024 ± 0.002	0.333	0.335	0.086	0.077

Note. We show in Appendix G that most of the degradation shown here for the Realistic case is caused by statistical uncertainties.

SNe Ia. Only the epoch selection cut is applied, since we are essentially assuming infinite S/N. As highlighted in Section 4, we set $\sigma_{\text{syst}} = 1$ mas in the χ^2 denominator to avoid infinite χ^2 (Equation (8)), while $\sigma_{\text{syst}} = 0$ in the simulation. The results we show here are equivalent for PDFPeak and CDF50.

In this perfect case, z_{DCR} give very precise estimates of the SN Ia redshifts for all redshifts, but especially for $z < 0.6$, where z_{DCR} almost always coincides with z_{spec} . When all redshifts are included, the outlier fraction $\eta_{0.10}$ is 1.35%, σ_{MAD} is 0.008, and bias is -0.003 ± 0.000 as shown in Table 3.

In Figures 8 and 9, the green X symbols show the binned bias, σ_{MAD} , and the outlier fractions versus z_{spec} . Again, the values shown for CDF50 and PDFPeak are equivalent. For the Perfect case, the bias is below 0.01 and $\eta_{0.10}$ is also below 5% in all redshift bins up to $z_{\text{spec}} = 1.2$.

Although the results shown here are for the Perfect case, there are some degeneracies and catastrophic outliers in Figure 7 (clumps or streaks), especially at the high redshifts. These artifacts are caused by the DCR shifts being degenerate at multiple redshifts at a given AM and observed epoch; fewer observations result in a higher probability that these degeneracies remain. Additionally, at higher redshifts, we typically do not have detections in the ug bands, due to low S/N (we only use detections where the coadded S/N > 3 , as mentioned in Section 3) where the DCR shifts are largest and therefore most informative.

5.2. Realistic Simulations

In Figures 10 and 11, we present our results for the Realistic case with both σ_{stat} , σ_{syst} , and realistic (mostly light-curve fit) peak MJD measurements. Additionally, we marginalize over

x_1 and c . After imposing the S/N selection cut on top of the epoch selection cut, we are left with about 42.9% of the simulated candidates (see Section 3).

Here, we find that the z_{DCR} estimates have degraded compared to the Perfect case, especially at high redshifts ($z_{\text{spec}} > 0.7$). CDF50 estimates very few SNe to be above $z = 0.8$, because the PDFs are nearly flat and the 50th percentile is clustered near the middle of the redshift range as shown in Figure 10 and reflected in the histograms of the residuals (Figure 11). PDFPeak and CDF50 show similar performance when all redshifts are included in terms of $\eta_{0.10}$ (33.5% and 33.3%, respectively), while the bias is lower for PDFPeak (-0.024 ± 0.002 compared to -0.048 ± 0.001 for CDF50). This suggests that, even at $z_{\text{spec}} > 0.8$, the high-redshift PDF peaks still provide some information, although by a modest amount.

In the $z_{\text{spec}} < 0.6$ panels of Figure 10, the accuracy is still relatively high compared to $z_{\text{spec}} > 0.6$, but CDF50 shows better performance. This is because CDF50 estimates the redshift to be $z \approx 0.6$ when the PDF is unconstraining, reducing catastrophic outliers when $z_{\text{spec}} < 0.6$. In Figures 8 and 9, the biases are close to zero at $z_{\text{spec}} < 0.7$ and the outlier fractions are below 20% at $z_{\text{spec}} < 0.6$ for both CDF50 and PDFPeak. The bias, however, deviates from zero at $z_{\text{spec}} > 0.7$ for both CDF50 and PDFPeak, but to a greater extent for CDF50 since most of the estimated redshifts are around 0.6. This is also the reason why the σ_{MAD} error bars are much larger for PDFPeak compared to CDF50; the CDF50 estimates are precise, but not accurate at high z . The diagonal streak (Figure 10) spanning $(z_{\text{spec}}, z_{\text{estimate}}) = (0.3, 0.5)$ to $(0.5, 0.3)$ could be due to the g -band DCR shifts typically being close to zero throughout all epochs around these redshifts, while there likely are not many

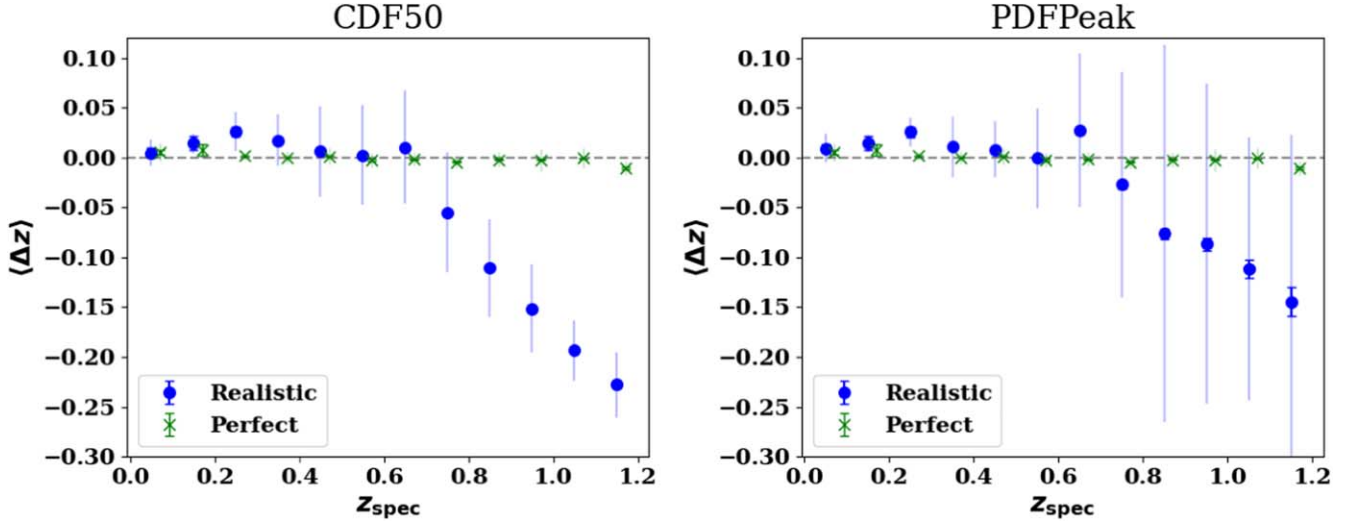


Figure 8. Binned residuals (bias), error on the mean (bold error bars), and the size of σ_{MAD} (light error bars) as a function of z_{spec} for Perfect and Realistic simulations for CDF50 (left) and PDFPeak (right). CDF50 and PDFPeak values are equivalent for Perfect simulations. For clarity, Realistic points (blue) are plotted at the midpoints of each redshift bin (0.05, 0.15, 0.25,...), while Perfect are displaced 0.02 to the right. The bias for the Realistic case noticeably deviates away from zero starting at $z_{\text{spec}} = 0.7$, especially for CDF50, where $z_{\text{DCR}} \approx 0.6$ for the unconstrained candidates. PDFPeak, on the other hand, displays significantly larger σ_{MAD} at higher z_{spec} . This is due to z_{DCR} being centered at 0.6 with little spread for CDF50, meaning that the redshift estimates are precise but not accurate.

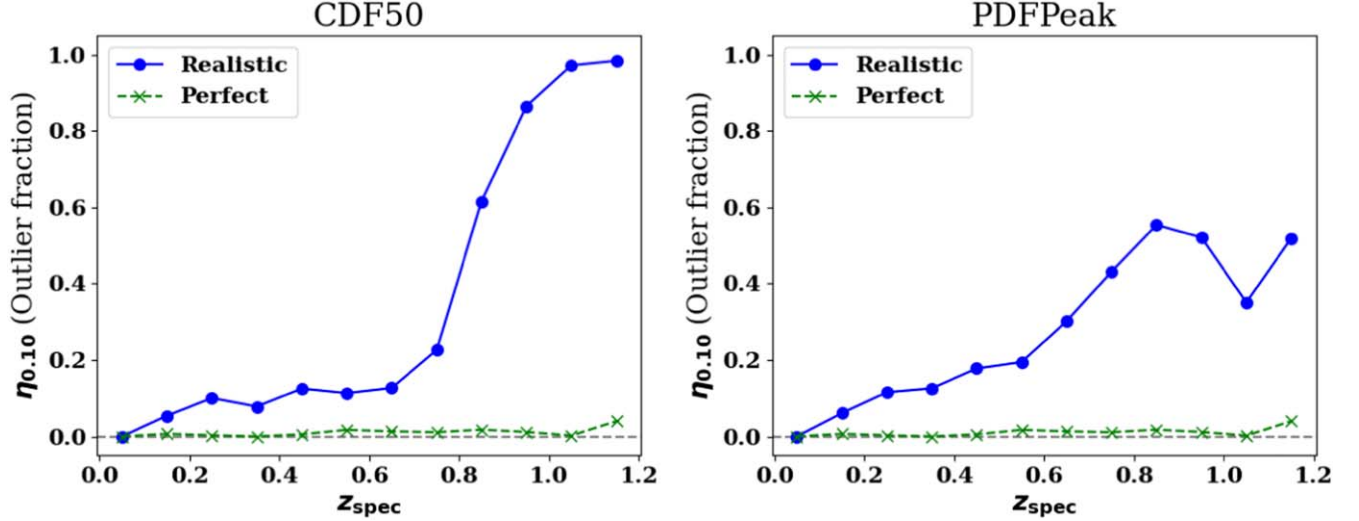


Figure 9. Binned outlier fractions as a function of z_{spec} for Perfect and Realistic simulations for CDF50 (left) and PDFPeak (right). The Perfect simulations show $< 5\%$ outlier fractions at all redshifts. For the Realistic simulations, the outlier fractions remain below 40% for both CDF50 and PDFPeak until $z_{\text{spec}} \approx 0.7$. At $z_{\text{spec}} \geq 0.8$, CDF50 estimates are mostly incorrect, as the PDFs are largely flat and unconstrained. Note that the outlier fractions for PDFPeak decreasing at the highest redshifts for the Realistic case is due to the possible estimated redshift range being limited from 0.0 to 1.2.

u -band observations, increasing room for degeneracies. In Appendix E.1, we further discuss the sensitivity of our results on various selection cuts.

5.2.1. Combining with z_{Host}

In this section, we discuss the performance of host-galaxy photometric redshifts (z_{Host}) combined with astrometric redshifts. Since we estimate z_{Host} using PDFs constructed from quantiles, we apply the CDF50 and PDFPeak methods the same way as for z_{DCR} . First, in Figure 12, we show z_{Host} as a function of z_{spec} for CDF50 and PDFPeak, with the same default selection cut imposed for the z_{DCR} -only simulations. Compared to z_{DCR} , the z_{Host} are much more accurate at higher redshifts but less accurate at lower redshifts. This is due to the distinct 4000 Å break for galaxies being in the u band at low redshifts, where the sensitivity is typically poor. Also, CDF50 gives more accurate estimates than

PDFPeak, with $\eta_{0.10} = 25.1\%$ for CDF50 and $\eta_{0.10} = 32.6\%$ for PDFPeak. As shown in Table 4, CDF50 performs better than PDFPeak in terms of bias and σ_{MAD} as well.

In Figure 13, we show the combined results, $z_{\text{DCR+Host}}$ versus z_{spec} . The estimated redshifts more tightly follow the true redshifts, with $\eta_{0.10}$ being 15.6% for CDF50 and 18.9% for PDFPeak, meaning that the outlier fractions are reduced by at least 10% for both CDF50 and PDFPeak compared to z_{Host} . The improvement is particularly striking at $z_{\text{spec}} < 0.6$, where the z_{Host} are not very constraining.

In Table 4, we also highlight the bias and σ_{MAD} for z_{DCR} , z_{Host} , and $z_{\text{DCR+Host}}$. We find that combining the two removes most of the bias for PDFPeak, and σ_{MAD} is considerably reduced for both CDF50 and PDFPeak.

Figures 14–15 show the mean binned residuals versus z_{spec} . We see large bias for z_{Host} at low redshifts for both CDF50 and

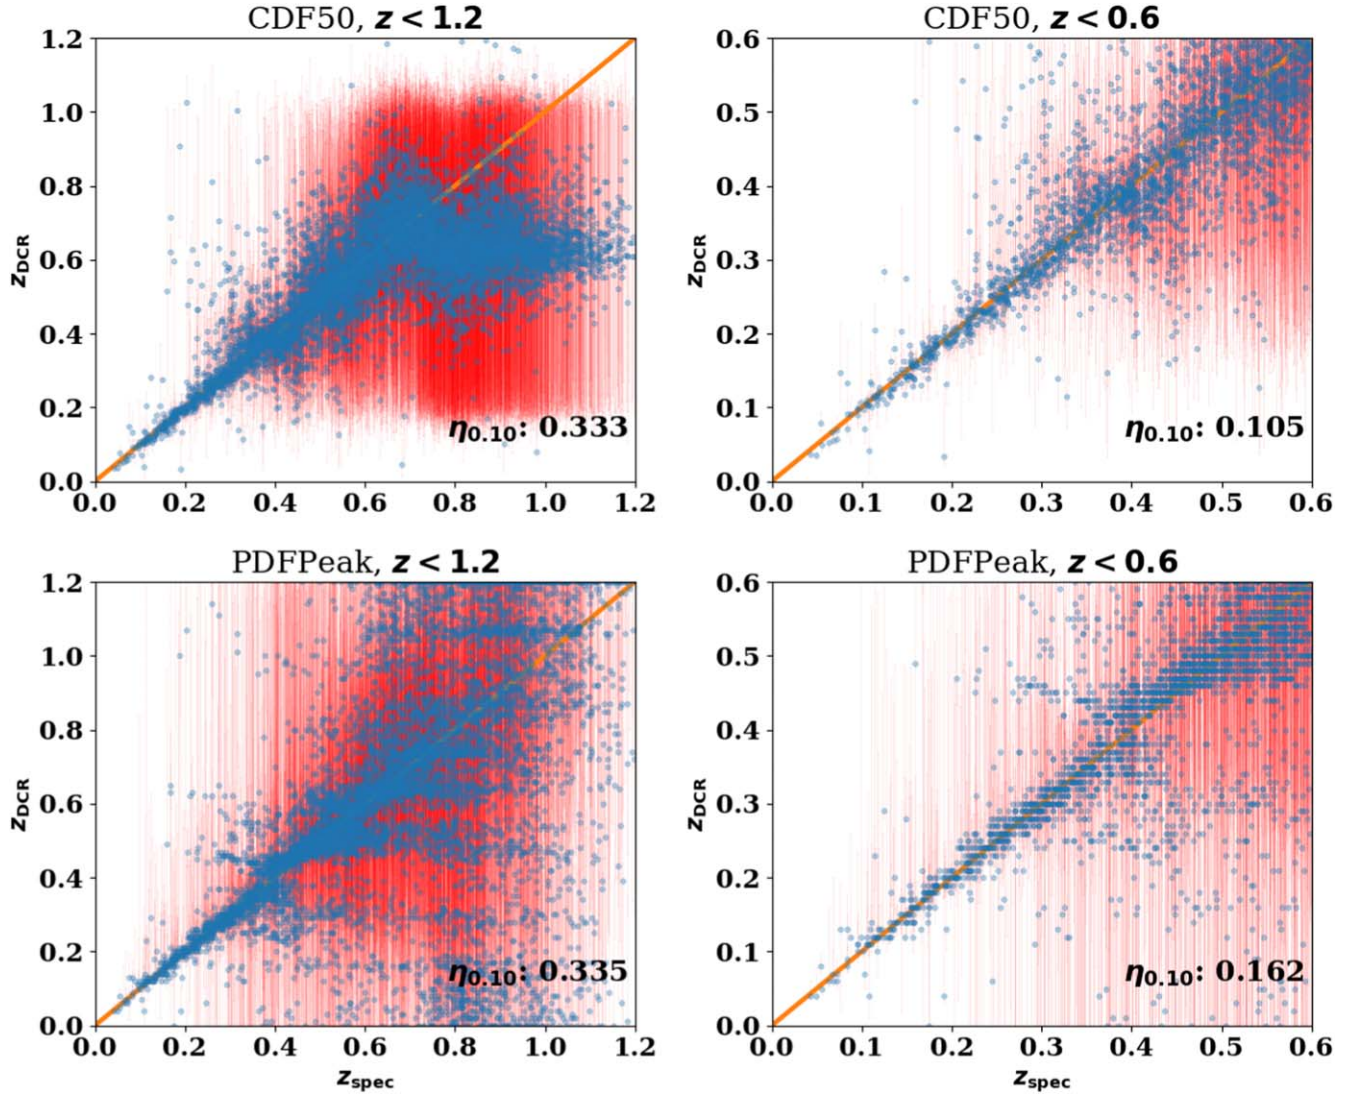


Figure 10. Similar to Figure 7, but for the Realistic case, with z_{DCR} vs. z_{spec} and 1σ error bars for the Realistic case for CDF50 (top panels) and PDFPeak (bottom panels). We see noticeable degradation in the z_{DCR} estimates, especially at $z_{\text{spec}} > 0.8$. The $z_{\text{spec}} < 0.6$ panels still show good agreement.




Figure 11. Histograms of the residuals Δz for the Realistic case for CDF50 (left) and PDFPeak (right). While the overall bias is significantly larger than in the Perfect case, we see little bias at $z_{\text{spec}} < 0.6$.

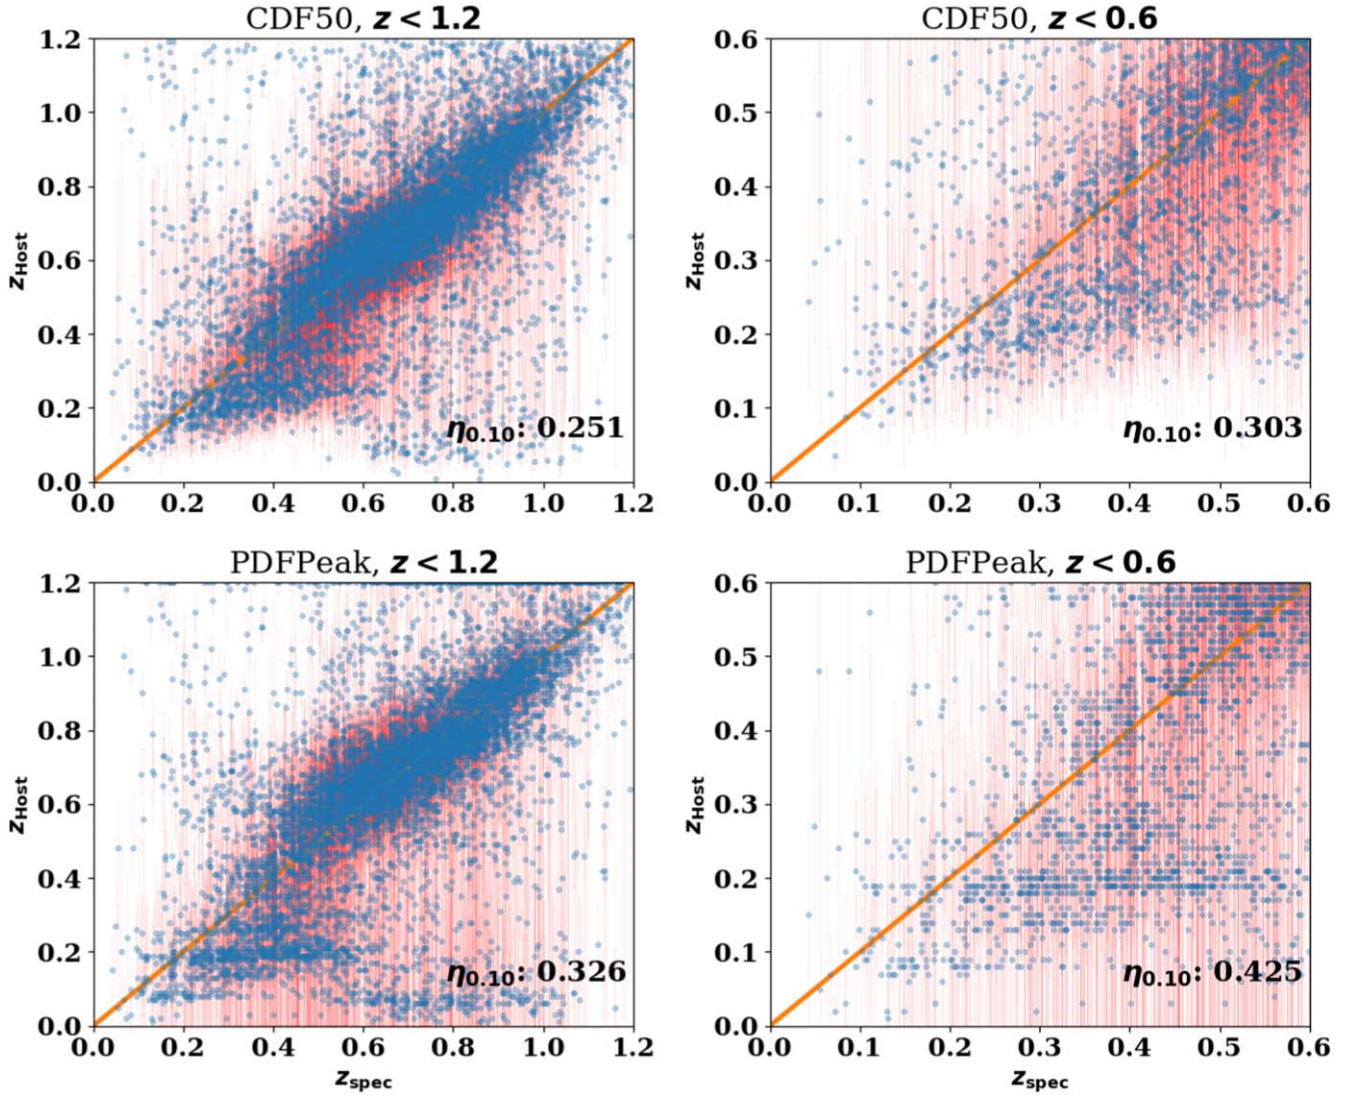


Figure 12. z_{Host} vs. z_{spec} and the 1σ error bars for CDF50 (top panels) and PDFPeak (bottom panels). For around 5% of the SNe that pass the Default selection cut, there are no z_{Host} estimates, so we assume flat PDFs as discussed in Section 4. While we do not show these candidates as data points and error bars in this figure, for clarity, they are always included for the performance metrics. We see that the z_{Host} are more accurate at higher redshifts, contrary to z_{DCR} . We also note that CDF50 performs better here, although this could be the result of the PDFs being constructed from quantiles with 11 bins.

Table 4
Bias and Error on the Mean, $\eta_{0.10}$, and σ_{MAD} for z_{DCR} only, z_{Host} only, and $z_{\text{DCR+Host}}$

Type	Bias (CDF50)	Bias (PDFPeak)	$\eta_{0.10}$ (CDF50)	$\eta_{0.10}$ (PDFPeak)	σ_{MAD} (CDF50)	σ_{MAD} (PDFPeak)
z_{DCR}	-0.048 ± 0.001	-0.024 ± 0.002	0.333	0.335	0.086	0.077
z_{Host}	-0.010 ± 0.001	-0.032 ± 0.002	0.251	0.326	0.067	0.085
$z_{\text{DCR+Host}}$	-0.016 ± 0.001	-0.013 ± 0.001	0.156	0.189	0.047	0.051

Note. Combining the two reduces the bias (PDFPeak) and σ_{MAD} , while the outlier fractions are reduced by at least 10%.

PDFPeak, little bias at the middle redshifts, and some bias at the higher redshifts. Combining with z_{DCR} reduces the bias, especially at the low redshifts. We note a similar trend for the outlier fractions as well; the z_{Host} outlier fractions are reduced significantly when combined with z_{DCR} at low redshifts.

5.2.2. Combining with z_{SN}

In this section, we show the redshift estimates for supernovae photometric redshifts (z_{SN}) using LCFIT+Z as well as z_{SN} combined with z_{Host} and z_{DCR} . Assessing the improvement

resulting from combining the z_{DCR} with $z_{\text{SN+Host}}$ is crucial to understanding how z_{DCR} can be used for upcoming surveys such as Rubin, where the number of supernovae detected will be far too large for spectroscopic redshift measurements.

We show in Table 5 the number of candidates remaining from the original 20,000 after requiring LCFIT+Z convergence which is about 61%–64% of the candidates, and then the default and FITPROB selection cuts, leaving of the candidates. In Figure 16, we show z_{SN} versus z_{spec} (top panels), $z_{\text{SN+Host}}$ versus z_{spec} (second panels), $z_{\text{SN+DCR}}$ versus z_{spec} (third panels), and $z_{\text{SN+DCR+Host}}$ versus z_{spec} (bottom panels) with the LCFIT+Z + Default selection

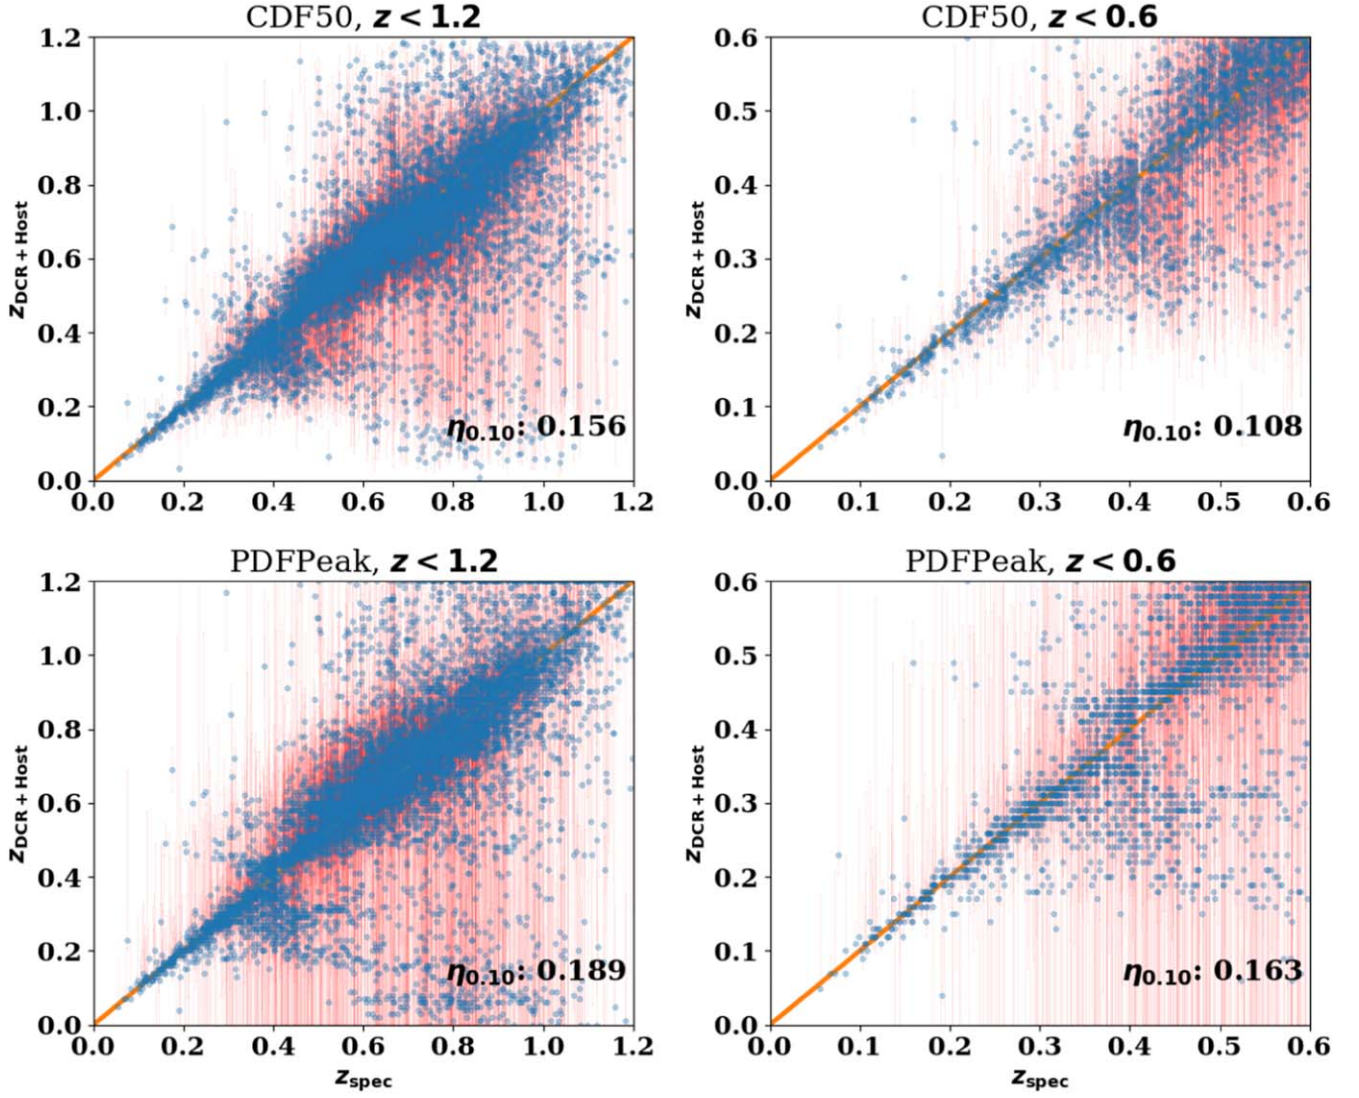


Figure 13. $z_{\text{DCR} + \text{Host}}$ vs. z_{spec} and the 1σ error bars for CDF50 (top panels) and PDFPeak (bottom panels). We see noticeable improvement from Figure 12, especially at $z_{\text{spec}} < 0.6$, where the z_{Host} values are not very constraining.

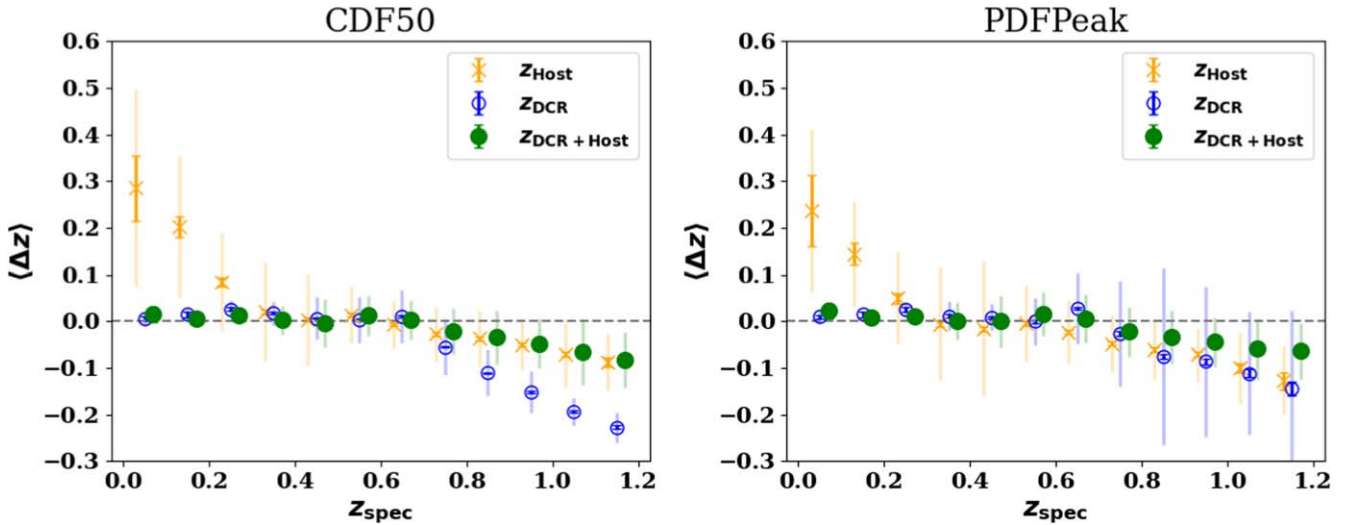


Figure 14. Binned residuals (bias), error on the mean (bold error bars), and the size of σ_{MAD} (light error bars) as a function of z_{spec} for z_{DCR} , z_{Host} , and $z_{\text{DCR} + \text{Host}}$ for CDF50 (left) and PDFPeak (right). Again, z_{DCR} is placed at the midpoints of each redshift bin, while z_{Host} and $z_{\text{DCR} + \text{Host}}$ are displaced to the left and right, respectively. Contrary to z_{DCR} , z_{Host} display more bias and larger error bars at lower redshifts, but combining with z_{DCR} alleviates this.

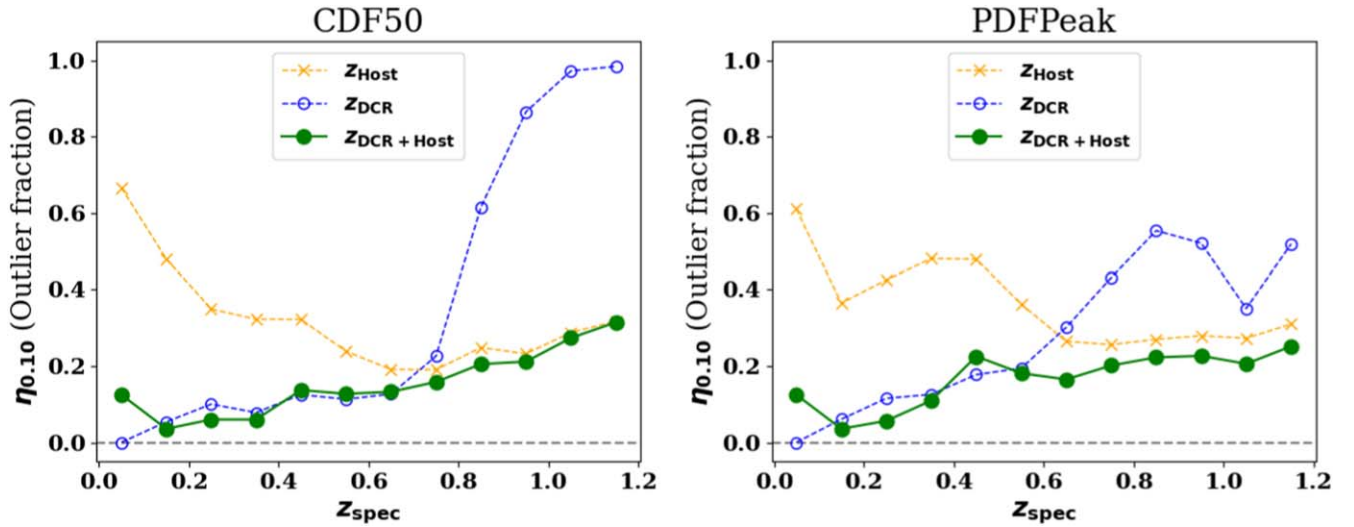


Figure 15. Binned outlier fractions as a function of z_{spec} for z_{DCR} , z_{Host} , and $z_{\text{DCR}+\text{Host}}$ for CDF50 (left) and PDFPeak (right). Combining with z_{DCR} significantly reduces the outlier fractions of z_{Host} at $z_{\text{spec}} < 0.6$ for both CDF50 and PDFPeak, while the high- z outlier fractions are not affected by much.

Table 5
Number of LCFIT+z Converged Candidates by Selection Cut

Type	Number of LCFIT+z Events with:		
	No Cuts	Default Cuts	Default + FITPROB Cuts
z_{SN}	12,537	7206	6577
$z_{\text{SN}+\text{Host}}$	12,397	7134	6585
$z_{\text{SN}+\text{DCR}}$	12,230	7596	7224
$z_{\text{SN}+\text{DCR}+\text{Host}}$	12,845	7534	7227

Note. Combining z_{DCR} recovers about 650 more candidates (10%) where the LCFIT+z values for photo- z s only do not converge.

cut. We would like to emphasize that adding z_{DCR} with the photo- z s results in a 10% increase in the number of candidates that pass the Default + FITPROB cuts.

As with z_{DCR} , z_{SN} have streaks of degeneracies at several locations and is more accurate at lower redshifts, although the performance is generally better with $\eta_{0.10}$ being 20.7% for all redshifts and 10.4% when limited to $z_{\text{spec}} < 0.6$. Combining with z_{Host} significantly improves the estimates at all redshifts, but the improvement is more pronounced at $z_{\text{spec}} > 0.6$, with $\eta_{0.10}$ reduced by about 14% for the all- z sample. This improvement is expected since z_{SN} and z_{Host} are independent estimates (as are the z_{DCR} values with z_{Host} , as discussed in Section 5.2.1). The third row of panels shows that adding z_{DCR} to z_{SN} considerably improves the estimates with $\eta_{0.10}$ reduced by about 6% for the all- z sample compared to z_{SN} only. Contrary to combining z_{Host} , the $z_{\text{SN}+\text{DCR}}$ improvement is more prominent at lower redshifts, with $\eta_{0.10}$ reduced by 8% when limited to $z_{\text{spec}} < 0.6$. Finally, combining z_{SN} with z_{Host} and z_{DCR} shows the full potential of nonspectroscopic redshift estimates, with $\eta_{0.10}$ decreasing to 6.0% for all redshifts and 1.9% when limited to $z_{\text{spec}} < 0.6$.

Table 6 shows the bias, $\eta_{0.10}$, and σ_{MAD} for the four combinations of redshift estimation methods involving z_{SN} . The measurements involving z_{SN} are unbiased, with the exception of $z_{\text{SN}+\text{DCR}}$. Combining z_{Host} with z_{SN} reduces σ_{MAD} to about two thirds of the z_{SN} -only value, while combining z_{DCR} with z_{SN} shows a similar impact. When all three methods are combined, the bias is nearly zero and σ_{MAD} is about 0.014.

There are a few points worth noting. First, while combining z_{SN} with z_{Host} reduces the outlier fraction by 14% as opposed to 6% with the addition of z_{DCR} , only about 37% of the true redshifts (for $z_{\text{SN}+\text{DCR}+\text{Host}}$) are below 0.6, where the z_{DCR} values are more informative. Because the z_{DCR} and z_{Host} each perform better at different redshifts, we show in Table 8 the bias, $\eta_{0.10}$, and σ_{MAD} for the different combinations at high z ($z_{\text{spec}} > 0.6$) and low z ($z_{\text{spec}} < 0.6$) separately. For z_{DCR} , z_{Host} , and $z_{\text{DCR}+\text{Host}}$, we show the PDFPeak values. Second, it may be counterintuitive that combining z_{SN} with z_{DCR} results in any noticeable improvement, since the information for both redshift estimates is based on the same SN Ia SED time series. We evaluate the independence of these two redshift estimates by computing the Pearson correlation coefficient between $z_{\text{SN}} - z_{\text{spec}}$ and $z_{\text{DCR}} - z_{\text{spec}}$, which we find to be 0.121. Such a low value shows that the two measurements are largely independent. This is because z_{DCR} relies on the product of the astrometric shifts due to DCR and the SED within a band (Equation (1)), while z_{SN} relies on only the SED integrated within a band. It is therefore reasonable that combining these two measurements results in a better redshift constraint.

Because z_{DCR} perform better at low z compared to z_{Host} and vice versa, we discuss our results in terms of low z and high z separately in Appendix E.2.

Finally, in Figures 17 and 18, we show the mean binned residuals and its σ_{MAD} error bars, as well as the outlier fractions by true redshift. Generally speaking, the smallest bias occurs at $z_{\text{spec}} < 0.2$ for $z_{\text{SN}+\text{DCR}}$, $0.2 < z_{\text{spec}} < 0.8$ for $z_{\text{SN}+\text{DCR}+\text{Host}}$, and $z_{\text{spec}} > 0.8$ for $z_{\text{SN}+\text{Host}}$. Similarly, the lowest outlier fractions occur at $z_{\text{spec}} < 0.8$ for $z_{\text{SN}+\text{DCR}+\text{Host}}$ and $z_{\text{spec}} > 0.8$ for $z_{\text{SN}+\text{Host}}$. For $z_{\text{SN}+\text{DCR}+\text{Host}}$, the outlier fractions are lowest at $0.0 < z_{\text{spec}} < 0.2$ (close to 0) and highest at $1.1 < z_{\text{spec}} < 1.2$ (still less than 20%). Contrary to when z_{DCR} values are combined with z_{Host} , combining z_{DCR} with z_{SN} noticeably degrades the performance at $z_{\text{spec}} > 0.8$.

In this section, we provided a variety of performance estimates for combinations of z_{DCR} , z_{Host} , and z_{SN} . While it is difficult to precisely pinpoint how much improvement will result from combining z_{DCR} with photo- z s for LSST, as the survey strategy has yet to be finalized, our methodology of obtaining z_{DCR} from astrometry is robust. Future work will

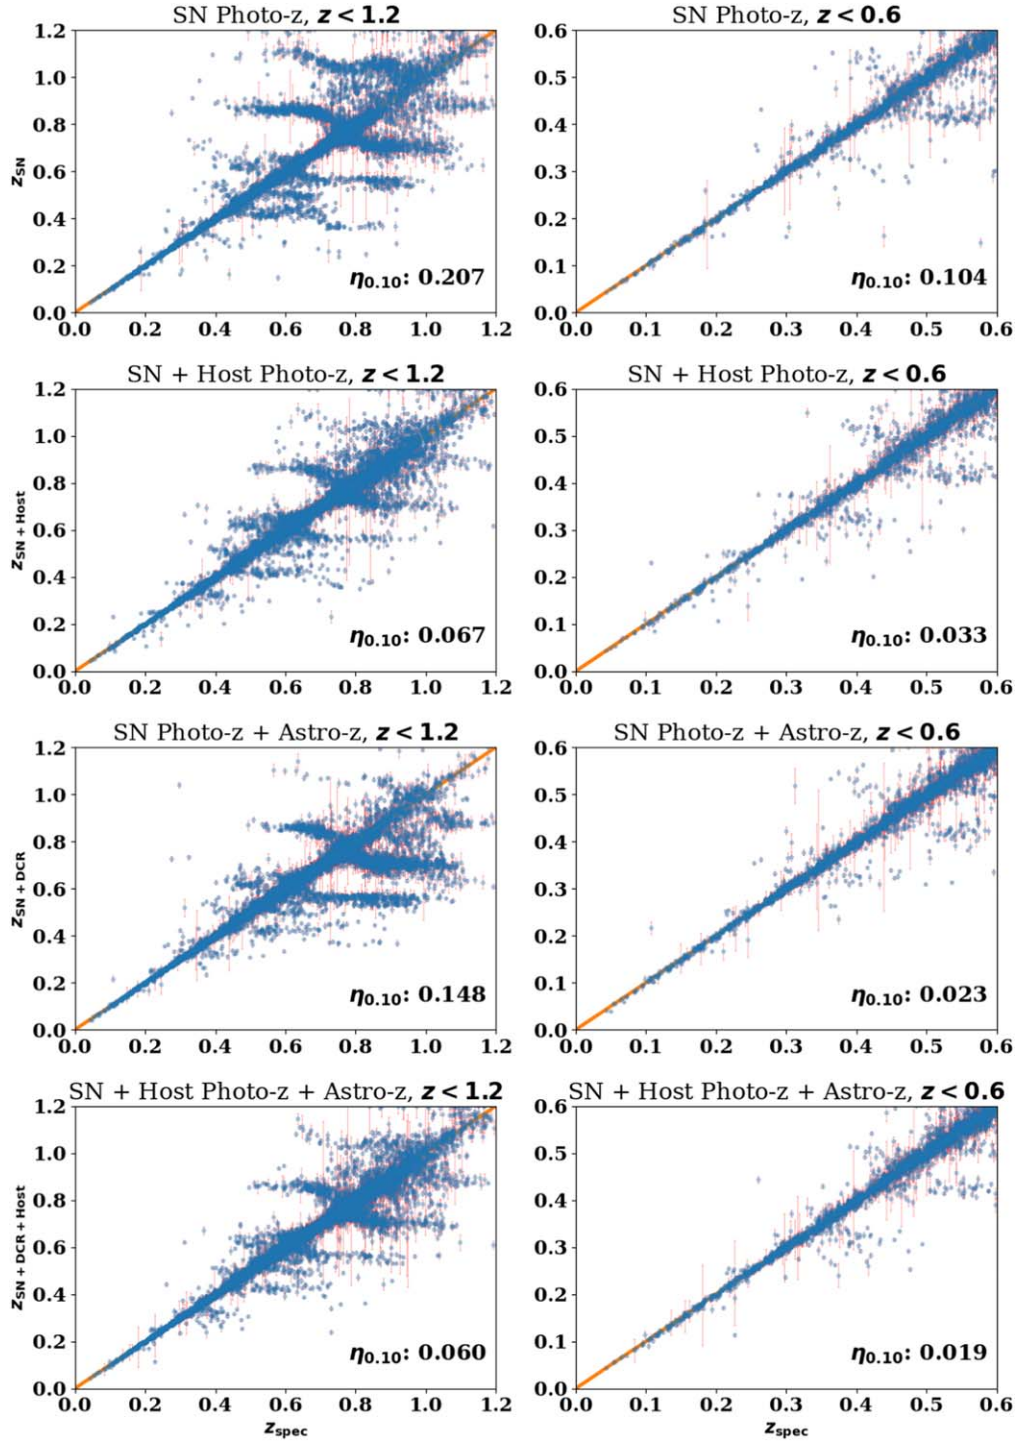


Figure 16. z_{SN} vs. z_{spec} (top row), $z_{\text{SN+Host}}$ vs. z_{spec} (second row), $z_{\text{SN+DCR}}$ vs. z_{spec} (third row), and $z_{\text{SN+DCR+Host}}$ vs. z_{spec} (bottom row). Combining z_{SN} with either z_{Host} or z_{DCR} significantly improves the estimates, more so at $z_{\text{spec}} > 0.6$ for z_{Host} and $z_{\text{spec}} < 0.6$ for z_{DCR} . Combining all three shows much better overall (throughout all redshifts) accuracy than any two combinations.

need to consider carefully how the z_{DCR} PDFs are combined with photo- z s, which could involve machine learning algorithms that utilize insight gained from the use of OVL and FITPROB, as well as diagnostics shown in Appendix F.

6. Discussion and Conclusion

In this paper, we outlined a method to obtain astrometric redshifts of SNe Ia from multiband, multipoch observations

using Differential Chromatic Refraction or DCR, which is caused by the wavelength dependence of our atmosphere's refractive index. Because the DCR shifts depend on the source SED, it can in principle be used to infer its redshift if a sufficient number of observations is available at moderately high air masses.

We evaluate our method on LSST-like simulations using an updated version of SNANA incorporating statistical and systematic uncertainties and detailed calculations of the DCR




Figure 17. Binned residuals (bias), error on the mean (bold error bars) and the size of σ_{MAD} (light error bars) as a function of z_{spec} for z_{SN} , $z_{\text{SN}} + \text{Host}$, $z_{\text{SN}} + \text{DCR}$, and $z_{\text{SN}} + \text{DCR} + \text{Host}$. Combining all three reduces the bias throughout all redshifts.

Table 6

Bias and Error on the Mean, $\eta_{0.10}$, and σ_{MAD} for z_{SN} , $z_{\text{SN}} + \text{Host}$, $z_{\text{SN}} + \text{DCR}$, and $z_{\text{SN}} + \text{DCR} + \text{Host}$

Type	Bias	$\eta_{0.10}$	σ_{MAD}
z_{SN}	0.007 ± 0.001	0.207	0.026
$z_{\text{SN}} + \text{Host}$	-0.000 ± 0.001	0.067	0.016
$z_{\text{SN}} + \text{DCR}$	-0.019 ± 0.001	0.148	0.017
$z_{\text{SN}} + \text{DCR} + \text{Host}$	-0.003 ± 0.001	0.060	0.014

Note. Combining all three methods significantly lowers the outlier fraction compared to z_{SN} only.

effect. We provide two types of point estimates: one where we take the 50th percentile of the cumulative distribution function of the posterior to be the estimated redshift (CDF50), and one where we take the peak of the posterior to be the estimated redshift (PDFPeak). We find that in the hypothetical cases of Perfect (no noise) simulations, z_{DCR} perform very well, although the estimate is not always accurate due to degeneracies. For the Realistic case, we find that our estimates at $z_{\text{spec}} > 0.7$ are substantially degraded, but the z_{DCR} estimates are still quite reasonable at $z_{\text{spec}} < 0.6$.

We then combine z_{DCR} with host-galaxy photometric redshifts as well as SN photometric redshifts using LCFIT + z . As z_{Host} are completely independent from z_{SN} and z_{DCR} , and z_{DCR} are close to independent from z_{SN} , which we confirmed by calculating the Pearson correlation coefficients between them, combining the different methods results in substantial improvement in the redshift estimates. Combining all three methods results in reduced bias, outlier fractions, and MAD deviation. The outlier fractions for z_{SN} is $\eta_{0.10} = 20.7\%$, while combining with both z_{Host} and reduces this to about 6.7% for the all-redshifts sample. Combining z_{DCR} on top of this further reduces $\eta_{0.10}$ to 6.0%, and results in about 10% more SN Ia candidates that pass selection cuts.

We believe that our analysis is important in several ways. First, this is the first demonstration of astrometric redshift measurements for SNe Ia. Contrary to M. C. Kaczmarczik et al. (2009), where the DCR shifts in the u and g bands were




Figure 18. Binned $\eta_{0.10}$ as a function of z_{spec} for z_{SN} , $z_{\text{SN}} + \text{Host}$, $z_{\text{SN}} + \text{DCR}$, and $z_{\text{SN}} + \text{DCR} + \text{Host}$. Combining all three lowers the outlier fractions throughout all redshifts, although $z_{\text{SN}} + \text{Host}$ show slightly lower outlier fractions at $z_{\text{spec}} > 0.8$.

treated as colors in the photo- z algorithm, we constructed and exploited the full PDFs. Astrometric redshifts are especially useful because they are independent of host-galaxy photometric redshifts, and nearly independent from SN photometric redshifts, allowing the combination to result in more accurate redshift measurements. z_{DCR} perform better at lower redshifts, complementing z_{Host} , which perform better at higher redshifts. This will also allow us to use more SNe Ia for cosmology compared to using just SN + Host photometric redshifts, as well as reduce uncertainties. While we are quite eager to see our new method used in a cosmology analysis, such an analysis is highly nontrivial and well beyond the scope of this work.

Although we use the coadded DDF for LSST in this initial analysis, we expect that astrometric redshifts will also be useful for the Wide-Fast-Deep (WFD) survey where the chances of obtaining spectroscopic redshifts are much lower. While astrometric redshifts are not expected to provide much information at the higher redshifts—especially for the WFD, due to low S/N and reduced cadence—we believe they will be similarly constraining, as shown in this work for $z_{\text{spec}} < 0.6$, which we have seen in preliminary analysis with the DES shallow fields. Although we found in Section 5.2 that restricting to candidates with at least one observation with $\text{AM} > 1.4$ does not improve z_{DCR} significantly for the DDF, a few intentional high-AM observations might be beneficial for the WFD fields, as the number of observations per candidate are expected to be lower than the DDF on average.

We plan to implement astrometric redshifts of SNe Ia into our SNe Ia analysis pipeline for LSST observations, although we will have to account for the additional uncertainties and bias coming from the telescope instrumental properties discussed in Section 3. In this work, we did not consider atmospheric variability, although its effect on DCR shifts is likely to be small and it is possible to incorporate the recorded atmospheric conditions for each observation. We recommend that LSST record the quantities required to calculate the atmospheric refraction index for each observation, such as air temperature, air pressure, and water vapor pressure. Ideally, the index of refraction would be calculated for each observation and stored for later access. We also did not include uncertainties arising from host-galaxy modeling, which will make analysis in real

observations more difficult. Additionally, we note that the performance of astrometric redshifts will depend sensitively on the detailed survey strategy such as the cadence and the air mass distribution of the observations, which have not yet been finalized for LSST. To prepare for using this new methodology in a future cosmology analysis on both simulations and real data, the DCR simulation tools have been integrated to SNANA, and an improved scene-modeling photometry pipeline is underway.

Some additional potential applications of this work include: utilizing the PSF shape changes due to DCR and wavelength-dependent seeing to estimate redshifts, as the 2D shape changes tend to be larger than the 1D DCR shifts, and extending astrometric redshifts to Type II supernovae. The former method is encouraging because the shape changes caused by wavelength-dependent effects result in more apparent magnitude shifts than the positional shifts due to DCR, but the shape changes will almost certainly be more difficult to measure than 1D positional shifts. Type II supernovae, like Type Ia supernovae, also display distinct emission lines, which result in substantially different DCR shifts depending on the redshift. Thus, we expect that astrometric redshifts for Type II supernovae will also be useful, as we have shown in this work for Type Ia supernovae. Another potential application could be to find mismatched host galaxies using $z_{\text{SN+DCR}}$. If the estimated $z_{\text{SN+DCR}}$ is very different from z_{Host} , the host galaxy may not have not been identified correctly. This could be useful whether we are using host galaxies for photometric redshifts or spectroscopic follow-up.

Acknowledgments

This paper has undergone internal review in the LSST Dark Energy Science Collaboration. We thank the internal reviewers Gautham Narayan and Robert Knop. We also thank the anonymous referee for catching an error in the equation used to calculate the DCR shifts in the initial version of this draft, as well as helping the authors improve the clarity of this work. LSST DESC acknowledges ongoing support from the Institut National de Physique Nucléaire et de Physique des Particules in France; the Science & Technology Facilities Council in the United Kingdom; and the Department of Energy, the National Science Foundation, and the LSST Corporation in the United States. LSST DESC uses the resources of the IN2P3/CNRS Computing Center (CCIN2P3-Lyon/VilleurbanneFrance) funded by the Centre National de la Recherche Scientifique; the Univ. Savoie Mont BlancCNRS/IN2P3 MUST computing center; the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under contract No. DE-AC02-05CH11231; STFC DiRAC HPC Facilities, funded by UK BIS National E-infrastructure capital grants; and the UK particle physics grid, supported by the GridPP Collaboration. This work was performed in part under DOE contract DE-AC02-76SF00515. J.L. and M.S. were supported by DOE grant DE-FOA-0002424 and NSF grant AST-2108094. The work of A.I.M. was supported by Schmidt Sciences.

Author Contributions Are Listed Below. J. Lee: led main analysis, wrote code and scripts for analysis, and wrote the manuscript. M. Sako: project initiator, supervision of lead graduate student, manuscript editing, administration, and management. R. Kessler: updated SNANA simulation to model atmospheric effects, light-curve fitting strategy, and manuscript editing. A.I. Malz: generated the photometric redshift PDFs.

Appendix A DCR Calculations

In this section, we describe more explicitly how we implemented the DCR R.A. (α) and decl. (δ) shifts into SNANA as mentioned in Section 3. This is similar to the case in the Appendix of J. Lee et al. (2023), but for an observer in the Southern Hemisphere. Figure 19 is a diagram of the celestial sphere for an observer (who would be in the center of the sphere) in the Southern Hemisphere, where Z is the zenith, (S)P and (N)P (denoted as just “P” from hereon) are the south and north celestial poles, and ϕ is the latitude of the observatory. Note that $\phi < 0$ for an observer in the Southern Hemisphere. This means that the angle between the South Pole and the observer’s south or the angle between the North Pole and the observer’s north is $|\phi|$. For an object located at X, a is the altitude, α and δ are the R.A. and decl., with $\delta < 0$ in this diagram, and α increases in the counterclockwise direction as viewed from above the North Pole. The angle XPZ or H , is the hour angle, and angle ZXP or q , is the parallactic angle, which are both typically defined using the north celestial pole.

In SNANA, we calculate the parallactic angles using the spherical sine and cosine laws:

$$\begin{aligned} \frac{\sin(q)}{\sin(90^\circ + |\phi|)} &= \frac{\sin(H)}{\sin(90^\circ - a)} = \frac{\sin(360^\circ - A)}{\sin(90^\circ - \delta)} \rightarrow \sin(q) \\ &= \frac{\cos(\phi)\sin(H)}{\cos(a)} = \frac{-\cos(\phi)\sin(A)}{\cos(\delta)} \\ \cos(90^\circ + |\phi|) &= \cos(90^\circ - a)\cos(90^\circ + |\delta|) \\ &\quad + \sin(90^\circ - a)\sin(90^\circ + |\delta|)\cos q, \\ \rightarrow \cos(q) &= \frac{\sin(\phi) - \sin(a)\sin(\delta)}{\cos(a)\cos(\delta)}, \end{aligned} \quad (\text{A1})$$

where we have defined $|\phi| = -\phi$ and $|\delta| = -\delta$, and A is the azimuth of the object.

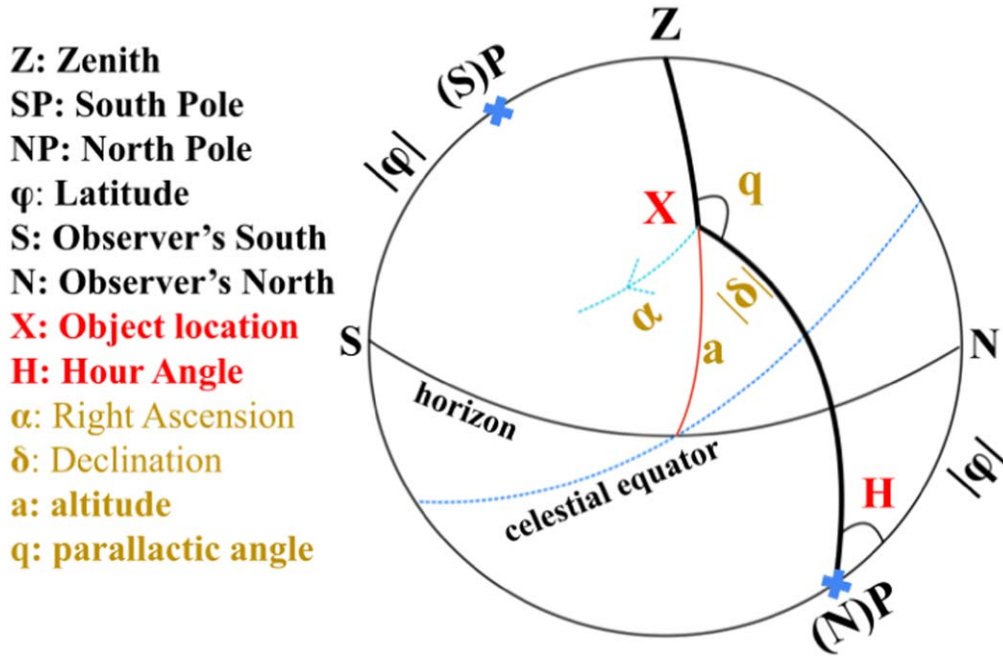
Now, consider a DCR altitude shift toward the zenith, Δa . In flat Euclidean geometry, we would calculate the α and δ shifts as $\Delta\alpha = \Delta a \sin(180^\circ - q) = \Delta a \sin(q)$ and $|\Delta\delta| = \Delta a \cos(180^\circ - q) = -\Delta a \cos(q)$. However, these relations do not necessarily hold on a sphere. To calculate the α and δ shifts properly, we can utilize Napier’s laws for right spherical triangles:

$$\begin{aligned} \sin(\Delta\alpha) &= \sin(180^\circ - q)\sin(\Delta a) = \sin(q)\sin(\Delta a) \\ \sin(|\Delta\delta|) &= \sin(x)\sin(\Delta a) = \tan(\Delta\alpha)\cot(180^\circ - q) \\ &= -\tan(\Delta\alpha)\frac{\cos(q)}{\sin(q)}, \end{aligned} \quad (\text{A2})$$

where the angle between $\Delta\alpha$ and $|\Delta\delta|$ is 90° , $|\Delta\delta|$ and Δa is $180^\circ - q$, and Δa and $\Delta\alpha$ is labeled as x . Since the shifts (altitude, R.A., and decl.) we are considering are well below 1° , $\sin(\Delta\alpha) \approx \tan(\Delta\alpha) \approx \Delta\alpha$, and Equation (A2) can be approximated as

$$\Delta\alpha \approx \sin(q)\Delta a, \quad |\Delta\delta| \approx -\Delta\alpha \frac{\cos(q)}{\sin(q)} = -\cos(q)\Delta a. \quad (\text{A3})$$

Finally, when we calculate the projected DCR shifts, we must account for the size of 1° becoming smaller as we move farther away from the celestial equator, and $\Delta\delta$ is clearly negative for



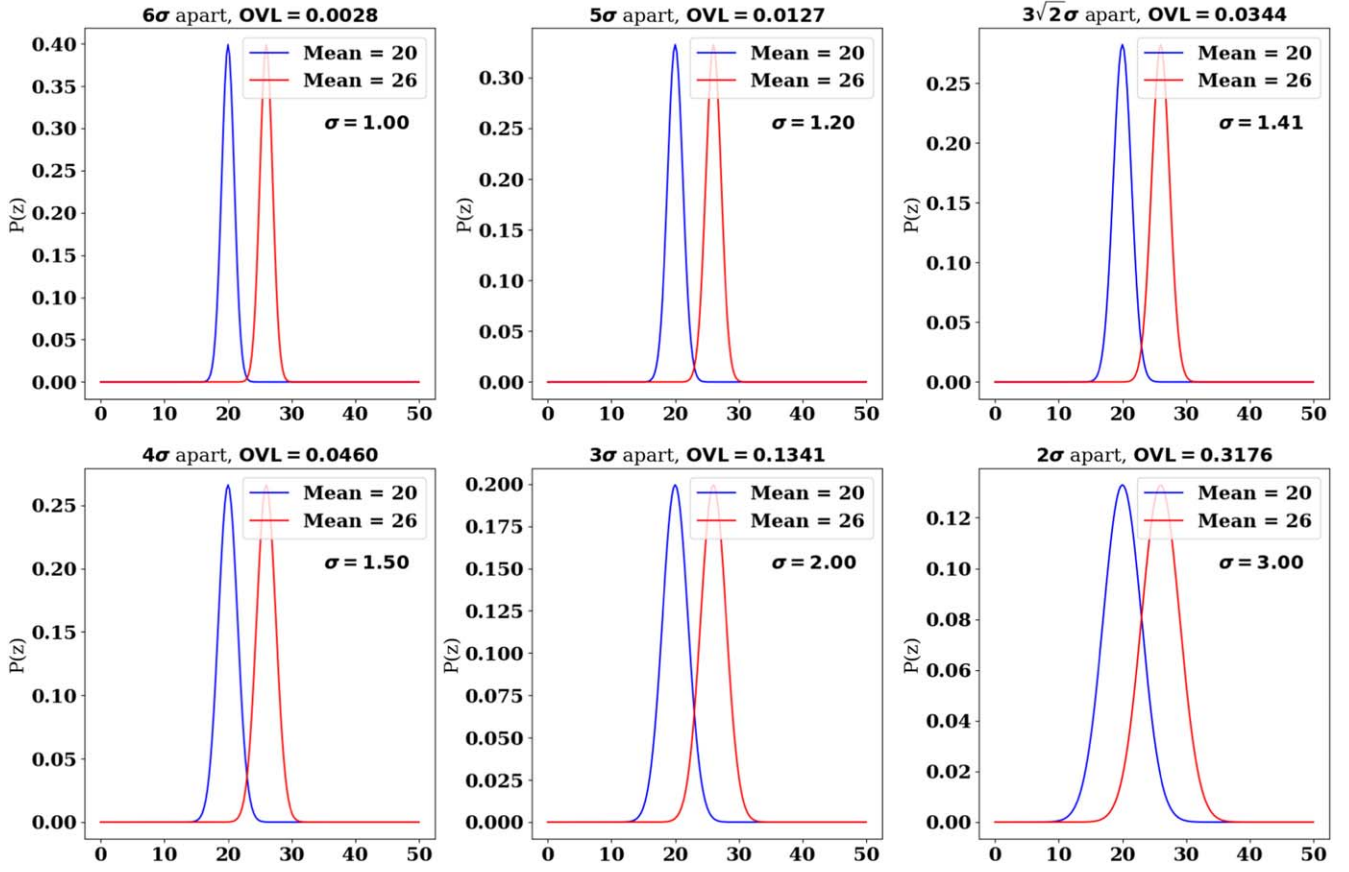


Figure 20. OVL depending on number of standard deviations of separation for two identical normalized Gaussian distributions with different means. Because the means of the two distributions are fixed, the larger the σ , the higher the overlap. We take the case when the means of the two distributions are $3\sqrt{2}$ apart or when $\text{OVL} > 0.0344$ as a threshold when combining z_{DCR} and z_{Host} .

While it is possible to compute the overlapping index assuming that z_{SN} PDFs are Gaussian (although this is not true, as shown in M. Sako et al. (2018)), we found that the OVL values between z_{SN} and z_{DCR} or z_{Host} are very small, due to the z_{SN} 1σ error bars from LCFIT+z being small, with a mean of 0.017 and median of 0.011. The LCFIT+z z_{SN} error bars are computed using the default MINUIT (F. James & M. Roos 1975) error option “Migration by Gradients” (MIGRAD). Better photo- z error estimates can be obtained, such as by using the “Minimization by Orthogonal Simplex” (MINOS) option in MINUIT, or MCMC fitting, but the fitting speed would be significantly slower.

Appendix D Cosmology Prior for Light-curve Fitting

In Section 3, we use a weak cosmology prior for LCFIT+z to obtain T_{obs} for Realistic z_{DCR} . Additionally, we use a cosmology prior in Appendix F to see improvement in the z_{SN} performance. Here, following R. Kessler et al. (2010) and A. Möller et al. (2024), we describe how the prior is applied.

We begin with the approximate fitted distance modulus, μ_{SALT3} :

$$\mu_{\text{SALT3}} = 30.0 - 2.5 \log_{10} x_0 + \alpha x_1 - \beta c, \quad (\text{D1})$$

with x_0 (amplitude), x_1 (stretch), and c (color) being SALT3 light-curve parameters. We use $\alpha = 0.14$ and $\beta = 3.2$ as our default. The difference between the fitted and theoretical

distance modulus, μ_{DIFF} , is given by

$$\mu_{\text{DIFF}} = \mu_{\text{SALT3}} - \mu_{\text{th}}(z_{\text{SN}}), \quad (\text{D2})$$

where μ_{th} is the theoretical distance modulus. The distance modulus uncertainty, σ_{μ} , can be computed using

$$\sigma_{\mu}^2 = \left(\frac{d\mu}{d\Omega_m} \Omega_{m,\text{ERR}} \right)^2 + \left(\frac{d\mu}{dw} w_{\text{ERR}} \right)^2, \quad (\text{D3})$$

where $\Omega_{m,\text{ERR}}$ and w_{ERR} are uncertainties for Ω_m and w , respectively. As mentioned in Appendix F, we use $\Omega_{m,\text{ERR}} = 0.03$ and $w_{\text{ERR}} = 0.1$. The reduced χ^2_{SALT3} is then given by

$$\chi^2_{\text{SALT3}} = \left(\frac{\mu_{\text{DIFF}}}{\sigma_{\mu}} P_{\mu\text{ERRSCALE}} \right)^2. \quad (\text{D4})$$

$P_{\mu\text{ERRSCALE}}$ is the distance modulus prior: a Gaussian profile with a standard deviation of $2\sigma_{\mu}$ when obtaining T_{obs} in Section 3, and $1\sigma_{\mu}$ when assessing improvement in the z_{SN} performance in Appendix F.

Appendix E Additional Results

E.1. Sensitivity of Astrometric Redshifts on Selection Cuts

We show in Table 7 the bias, $\eta_{0.10}$, and σ_{MAD} depending on the selection criteria: no selection cuts, default selection cut (42.9% of total candidates), and default + AM selection cut

Table 7
Bias and Error on the Mean, $\eta_{0.10}$, and σ_{MAD} for the Analysis Methods Indicated in the First Column

Type	Bias (CDF50)	Bias (PDFPeak)	$\eta_{0.10}$ (CDF50)	$\eta_{0.10}$ (PDFPeak)	σ_{MAD} (CDF50)	σ_{MAD} (PDFPeak)
No selection cut	-0.103 ± 0.001	-0.067 ± 0.001	0.577	0.444	0.139	0.123
Default	-0.048 ± 0.001	-0.024 ± 0.002	0.333	0.335	0.086	0.077
Default + AM cut	-0.041 ± 0.001	-0.020 ± 0.002	0.306	0.313	0.080	0.069
Fmax-clump	-0.048 ± 0.001	-0.025 ± 0.002	0.344	0.354	0.090	0.085
Ideal T_{obs}	-0.050 ± 0.001	-0.027 ± 0.002	0.320	0.324	0.083	0.072

Notes. With stricter selection cuts, we see improved metrics. Using peak MJDs from LCFIT+z (which is mostly the case for “Default”) improves performance compared to using “Fmax-clump.”

Table 8
Bias and Error on the Mean, $\eta_{0.10}$, and σ_{MAD} for Low z ($z_{\text{spec}} < 0.6$) and High z ($z_{\text{spec}} > 0.6$) for Various Combinations of the Three Methods Discussed in This Work

Type	Bias (Low z)	Bias (High z)	$\eta_{0.10}$ (Low z)	$\eta_{0.10}$ (High z)	σ_{MAD} (Low z)	σ_{MAD} (High z)
z_{DCR}	0.007 ± 0.002	-0.043 ± 0.002	0.162	0.437	0.033	0.123
z_{Host}	0.006 ± 0.003	-0.054 ± 0.002	0.425	0.268	0.123	0.065
$z_{\text{DCR+Host}}$	0.008 ± 0.001	-0.025 ± 0.002	0.163	0.203	0.041	0.055
z_{SN}	0.019 ± 0.002	0.001 ± 0.001	0.104	0.263	0.008	0.059
$z_{\text{SN+Host}}$	0.006 ± 0.001	-0.004 ± 0.001	0.033	0.086	0.008	0.023
$z_{\text{SN+DCR}}$	0.003 ± 0.001	-0.034 ± 0.001	0.023	0.227	0.007	0.041
$z_{\text{SN+DCR+Host}}$	0.003 ± 0.001	-0.007 ± 0.001	0.019	0.085	0.007	0.021

Notes. For z_{DCR} , z_{Host} , and $z_{\text{DCR+Host}}$, we show the PDFPeak values. z_{DCR} show good performance at low z , while z_{Host} show good performance at high z , which leads to $z_{\text{SN+DCR}}$ being significantly improved at low z compared to z_{SN} and $z_{\text{SN+Host}}$ being significantly improved at high z . $z_{\text{SN+DCR+Host}}$ show similar performance to $z_{\text{SN+DCR}}$ and $z_{\text{SN+Host}}$ at low z and high z , respectively.

(37.6% of total candidates), where the AM cut requires at least one observation at $\text{AM} > 1.4$ for a given candidate. As expected, the performance of z_{DCR} improves with more stringent cuts. We also see larger bias in CDF50 (compared to PDFPeak) for all three types of cuts, which is expected given that CDF50 estimates most $z_{\text{spec}} > 0.7$ SNe to be at $z_{\text{DCR}} \approx 0.6$. The outlier fractions are similar for CDF50 and PDFPeak, with the exception of “No selection cut,” where the bias for CDF50 is significantly higher, due to the inclusion of disproportionately more high- z SNe Ia. Additionally requiring the AM cut improves the estimates, but not significantly. While M. C. Kaczmarszik et al. (2009) suggest that a number of moderately high-AM observations will result in considerably more accurate redshift estimates for quasars, we note that, for the LSST DDF SNe Ia simulations, most of the candidates that pass the default selection cut also pass the AM cut, as there are multiple (usually over five) observations per SN. We also show the metrics depending on the choice of T_{obs} , and as expected, using peak MJDs from LCFIT+z results in some improvement from the “Fmax-clump” method, but less than the Ideal T_{obs} case, where we use the exact peak MJDs.

E.2. Low- z Versus High- z Improvements When z_{SN} Is Combined

In Table 8 (where we show the bias, $\eta_{0.10}$, and σ_{MAD} at low z and high z separately for PDFPeak when not combining with z_{SN}), both the bias and outlier fractions are much lower at low z than high z for z_{DCR} , while the outlier fractions are much lower at high z for z_{Host} as seen earlier. z_{SN} show larger bias at low z but a much smaller outlier fraction (10.4%) than at high z (26.3%), as with z_{DCR} . Combining z_{SN} with z_{Host} reduces both the bias and outlier fractions significantly at both low z and high z . While it is no surprise that the outlier fraction at high z is about a third of that for z_{SN} (8.6%), it is intriguing that the outlier fraction at low z is also

7% lower, at 3.3%, given that the z_{Host} -only outlier fraction at low z is 42.5%. Combining z_{SN} with z_{DCR} greatly reduces the bias and outlier fraction at low z , with the outlier fraction being 2.3%, but the bias at high z is only a little smaller than z_{DCR} at high z , and the outlier fraction at high z is reduced by a small amount (22.7%), as one would expect. When all three methods are combined, the biases at both low z and high z are close to zero, and the outlier fractions are similarly low at 1.9% and 8.5% for low z and high z , respectively, although using $z_{\text{SN+DCR}}$ at low z and $z_{\text{SN+Host}}$ at high z shows similar performance. A similar trend can be seen with σ_{MAD} as well; z_{DCR} has a much lower value at low z while z_{Host} has a lower value at high z , and hence combining each with z_{SN} displays lower (or at least similar) values at low z and high z , respectively. $z_{\text{SN+DCR+Host}}$ has the lowest σ_{MAD} values for both low z and high z .

Appendix F Diagnostics and Cosmology Prior

As discussed at the end of Section 3, we impose a $\text{FITPROB} \geq 0.01$ selection cut when combining with z_{SN} , which addresses some of the compatibility issues between the different redshift estimation methods. While this is sufficient for our initial analysis with simulations, we dive deeper and present some diagnostics as well as how much improvement can result from using a cosmology prior for LCFIT+z.

In Figure 21, we show some diagnostics histograms for $z_{\text{SN+DCR+Host}}$ (with the $\text{FITPROB} \geq 0.01$ cut and the compatibility cut for the $z_{\text{DCR+Host}}$ prior) for candidates that are within $\eta_{0.10}$ and outside $\eta_{0.10}$: FITPROB , prior- χ^2 , and host-galaxy log mass. For each of the panels, the blue points show the number of candidates within the x -axis bins for outliers, while the red histogram shows the distribution of nonoutlier (Not $\eta_{0.10}$)




Figure 21. Diagnostics for $z_{\text{SN}}+\text{DCR}+\text{Host}$. The blue points show the distributions for outliers, while the red histograms show the relative number of candidates for nonoutliers. The red bars have been scaled such that the area under the $\eta_{0.10}$ histograms and Not $\eta_{0.10}$ histograms are equal. We show the error bars for the $\eta_{0.10}$ points only because Not $\eta_{0.10}$ has many more candidates than $\eta_{0.10}$ and hence smaller error bars. As expected, outliers have lower FITPROB than nonoutliers, as well as slightly higher prior χ^2 . Additionally, outliers are more likely to occur in lower-mass host galaxies.

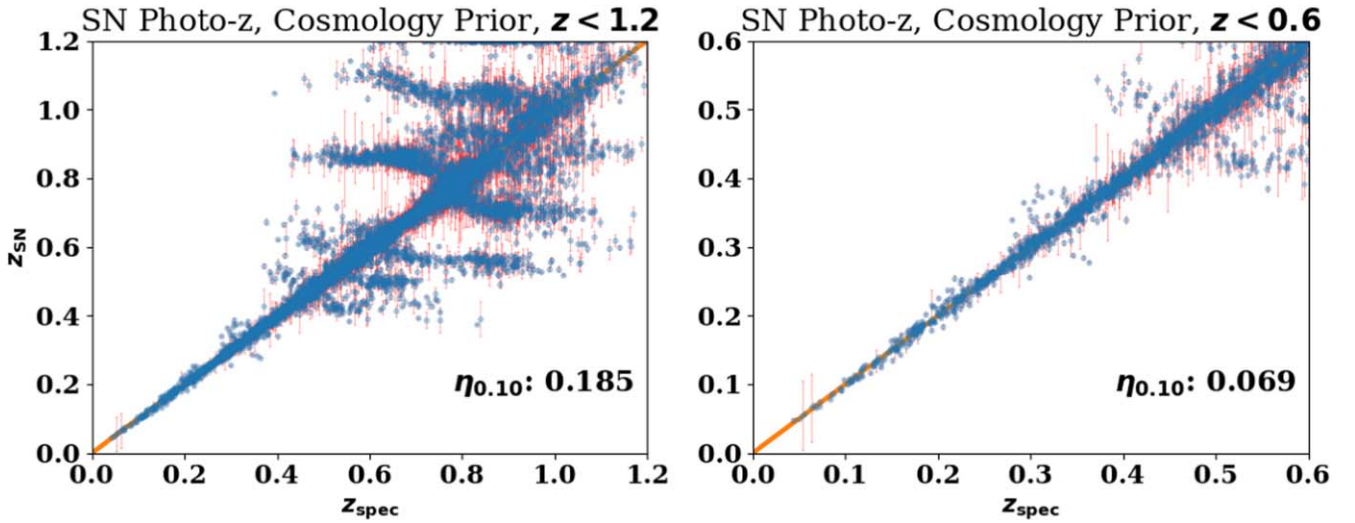


Figure 22. z_{SN} vs. z_{spec} with a weak cosmology prior for LCFIT+z. Compared to the top panels of Figure 16, where there is no cosmology prior, the estimates are closer to the true values throughout all redshifts, and the outlier fractions are lower.

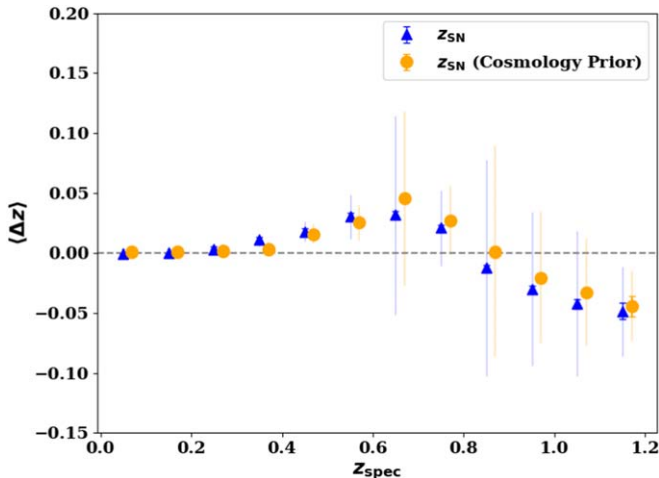


Figure 23. Binned residuals (bias), error on the mean (bold error bars), and the size of σ_{MAD} (light error bars) as a function of z_{spec} for z_{SN} , with and without cosmology priors. Using a cosmology prior generally tends to reduce the bias and σ_{MAD} except at $0.6 < z_{\text{spec}} < 0.8$.

candidates, scaled to have the same integral as the outlier distribution.

In the left panel, we see that $\eta_{0.10}$ has more candidates at lower FITPROB while Not $\eta_{0.10}$ displays the opposite, as expected. Adding useful information from FITPROB into a more complex machine learning algorithm could therefore improve the selection of events. In the middle, we see that the prior- χ^2 histograms are similar for $\eta_{0.10}$ and Not $\eta_{0.10}$, although $\eta_{0.10}$ is skewed toward slightly higher prior- χ^2 values. In the right panel, $\eta_{0.10}$ is shown to have more low-host-mass contribution compared to Not $\eta_{0.10}$. This could be due to low-mass host galaxies having lower S/N at high redshifts, degrading all three redshift estimates.

For all of the aforementioned results, we did not impose any cosmology prior, apart from determining T_{obs} using LCFIT+z for the Realistic case (Section 3). Here, we investigate potential improvement from using a weak cosmology prior for LCFIT+z. More detail on how the prior is applied is given in Appendix D. In Figures 22–24, we show the performance of z_{SN} only with a weak cosmology prior of $w = -1.0 \pm 0.1$ and $\Omega_m = 0.3 \pm 0.03$. In Figure 22, we see that, compared to the top

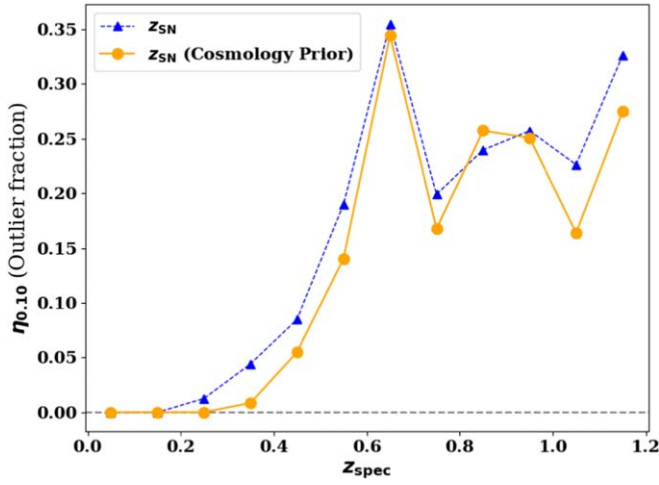


Figure 24. $\eta_{0.10}$ as a function of z_{spec} for z_{SN} , with and without cosmology priors. Similar to the bias and σ_{MAD} , a cosmology prior generally reduces $\eta_{0.10}$.

panel of Figure 16, the outlier fractions have noticeably decreased: $\eta_{0.10}$ is 0.207 and 0.185 without and with the cosmology prior, respectively, for all redshifts, while for $z < 0.6$ only, $\eta_{0.10}$ is 0.104 and 0.069, respectively. Note that the number of candidates that pass the selection cuts (LCFIT + z convergence + epoch and S/N cut + FITPROB ≥ 0.01 cut) have also increased by 7% compared to the case without the cosmology prior. In Figures 23 and 24, we show that the bias, $\eta_{0.10}$, and σ_{MAD} are generally slightly reduced, with the exception of some bins between $0.6 < z_{\text{spec}} < 0.9$. This additional analysis shows that, in practice, we can utilize a weak cosmology prior to obtain better estimates for both photometric and astrometric redshifts.

Appendix G Systematic Effects

Since we have assessed the Perfect and Realistic cases in Sections 5.1 and 5.2, here we show how sensitive our z_{DCR} results are to various levels of systematic effects in the absence of statistical uncertainties. The systematic uncertainties considered here include *unmodeled* telescopic effects as discussed in Section 3, and we assume that systematics leading to biases can typically be modeled and accounted for in a real survey. In Table 9, we show again the bias, $\eta_{0.10}$, and σ_{MAD} for the Perfect case as we saw in Section 5.1, as reference. We also show the metrics for a Perfect simulation with the following systematic

variations: $\sigma_{\text{syst}} = 2$ mas, $\sigma_{\text{syst}} = 5$ mas, $\sigma_{\text{syst}} = 10$ mas, c marginalized, x_1 marginalized, and realistic T_{obs} . For each systematic we tested, all the other conditions were taken to be the same as the Perfect case. As described in Section 4, CDF50 and PDFPeak results are the same when σ_{tot} should be zero; so for Perfect, c marginalized, x_1 marginalized, and realistic T_{obs} .

First, we show that $\sigma_{\text{syst}} = 5$ mas, which we took to be our default systematic floor throughout Section 5.2, only causes a slight degradation compared to the Perfect case, with $\eta_{0.10}$ being around 2.6% and σ_{MAD} being around 0.011 for both CDF50 and PDFPeak. CDF50 and PDFPeak show similar performance in terms of $\eta_{0.10}$ and σ_{MAD} , which indicates that the degradation primarily occurs at the highest redshifts where the PDFs tend to be flatter than at lower redshifts, meaning that systematic effects can smear them out enough to result in incorrect redshift estimates. Bias values are farther away from zero for CDF50 than PDFPeak, closely resembling the Realistic case.

Marginalizing over x_1 and c as well as taking $\sigma_{\text{syst}} = 2$ mas does not change the performance metrics much, suggesting that the level of degeneracies that result in incorrect redshift estimates at the high redshifts as shown in Figure 7 remain similar.

Using realistic T_{obs} (mostly using LCFIT+ z , with 11,595 candidates remaining after LCFIT+ z out of the 13,827 that pass the Epoch cut) causes a larger degradation, with $\eta_{0.10}$ at about 4.0%, or around two to three times that of the Perfect case, while σ_{MAD} also increases but to a smaller extent. The bias values are similar to the Perfect case, but have larger error bars. We find that the z_{DCR} performance with all of the systematics included in the Realistic case (using realistic T_{obs} , marginalizing over c and x_1 , and assuming $\sigma_{\text{syst}} = 5$ mas) but with $\sigma_{\text{stat}} = 0$ is considerably better than the Realistic case with $\eta_{0.10} = 6.9\%$ and $\sigma_{\text{MAD}} = 0.014$ for PDFPeak (similar values for CDF50), with the bias being -0.010 ± 0.001 and -0.004 ± 0.001 for CDF50 and PDFPeak, respectively, meaning that our z_{DCR} estimates are mostly limited by S/N.

As mentioned in Section 3, we note that our assumption of a $\sigma_{\text{syst}} = 5$ mas floor is reasonable and possible to achieve for LSST, given the DECam value of 3–6 mas (G. M. Bernstein et al. 2017). Additionally, we consider the z_{DCR} performance when σ_{syst} is 2 and 10 mas. In Table 9, we show that $\sigma_{\text{syst}} = 2$ mas results in $\eta_{0.10}$ being about half that of $\sigma_{\text{syst}} = 5$ mas, or similar to the Perfect case, while the bias and σ_{MAD} similarly improve slightly. $\sigma_{\text{syst}} = 10$ mas results in $\eta_{0.10}$ being about 3–4 times larger than the Perfect case at 6.9% and 4.4% for CDF50

Table 9
Bias, Outlier Fractions, and σ_{MAD} for Various Systematics

Type	Bias (CDF50)	Bias (PDFPeak)	$\eta_{0.10}$ (CDF50)	$\eta_{0.10}$ (PDFPeak)	σ_{MAD} (CDF50)	σ_{MAD} (PDFPeak)
Perfect	-0.003 ± 0.000	-0.003 ± 0.000	0.014	0.014	0.008	0.008
$\sigma_{\text{syst}} = 2$ mas	-0.007 ± 0.000	-0.003 ± 0.000	0.015	0.016	0.009	0.009
$\sigma_{\text{syst}} = 5$ mas	-0.010 ± 0.000	-0.004 ± 0.000	0.026	0.023	0.011	0.011
$\sigma_{\text{syst}} = 10$ mas	-0.019 ± 0.000	-0.006 ± 0.000	0.069	0.044	0.019	0.016
c marginalized	-0.007 ± 0.000	-0.004 ± 0.000	0.011	0.016	0.006	0.007
x_1 marginalized	-0.006 ± 0.000	-0.003 ± 0.000	0.012	0.014	0.007	0.008
realistic T_{obs}	-0.005 ± 0.000	-0.002 ± 0.000	0.040	0.040	0.010	0.010
All syst. for realistic	-0.010 ± 0.001	-0.004 ± 0.001	0.062	0.069	0.015	0.014

Notes. The impact of $\sigma_{\text{syst}} = 2$ or 5 mas, marginalization over c and x_1 is small, while using realistic T_{obs} results in more degradation. Larger σ_{syst} results in substantially larger degradation compared to other systematics. “All syst. for realistic” means using realistic T_{obs} , marginalizing over c and x_1 , and assuming $\sigma_{\text{syst}} = 5$ mas.

and PDFPeak, and the bias and σ_{MAD} also degrade substantially more compared to the other systematics considered here. The larger discrepancy in the metrics between CDF50 and PDFPeak for $\sigma_{\text{sys}} = 10$ mas suggests that the high- z estimates are affected disproportionately. The metrics we present here for larger σ_{sys} shows that limiting σ_{sys} to 5 mas or lower will be important when implementing z_{DCR} for LSST.

ORCID iDs

Jaemyoung (Jason) Lee  <https://orcid.org/0000-0001-6633-9793>

Masao Sako  <https://orcid.org/0000-0003-2764-7093>

Richard Kessler  <https://orcid.org/0000-0003-3221-0419>

Alex I. Malz  <https://orcid.org/0000-0002-8676-1622>

References

- Alejandro Plazas, A., & Bernstein, G. 2012, *PASP*, **124**, 1113
- Allam, T., Jr, Bahmanyar, A., Biswas, R., et al. 2018, arXiv:1810.00001
- Astier, P., Guy, J., Regnault, N., et al. 2006, *A&A*, **447**, 31
- Bernardinelli, P. H., Bernstein, G. M., Jindal, N., et al. 2023, *ApJS*, **269**, 18
- Bernstein, G. M., Abbott, T. M. C., Armstrong, R., et al. 2018, *PASP*, **130**, 054501
- Bernstein, G. M., Armstrong, R., Plazas, A. A., et al. 2017, *PASP*, **129**, 074503
- Brout, D., Sako, M., Scolnic, D., et al. 2019, *ApJ*, **874**, 106
- Brout, D., Scolnic, D., Popovic, B., et al. 2022, *ApJ*, **938**, 110
- Carlsten, S. G., Strauss, M. A., Lupton, R. H., Meyers, J. E., & Miyazaki, S. 2018, *MNRAS*, **479**, 1491
- Chen, R., Scolnic, D., Rozo, E., et al. 2022, *ApJ*, **938**, 62
- Crenshaw, J. F., 2021 jfcrenshaw/pzflow: v1.6.0, Zenodo, doi:10.5281/zenodo.4679913
- Dai, M., Kuhlmann, S., Wang, Y., & Kovacs, E. 2018, *MNRAS*, **477**, 4142
- DES Collaboration, Abbott, T. M. C., Acevedo, M., et al. 2024, *ApJL*, **973**, L14
- Filippenko, A. V. 1982, *PASP*, **94**, 715
- Guy, J., Astier, P., Baumont, S., et al. 2007, *A&A*, **466**, 11
- Howell, D. A., Sullivan, M., Perrett, K., et al. 2005, *ApJ*, **634**, 1190
- Inman, H. F., & Bradley, E. L. 1989, *CSTM*, **18**, 3851
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- James, F., & Roos, M. 1975, *CoPhC*, **10**, 343
- Kaczmarczik, M. C., Richards, G. T., Mehta, S. S., & Schlegel, D. J. 2009, *AJ*, **138**, 19
- Kenworthy, W. D., Jones, D. O., Dai, M., et al. 2021, *ApJ*, **923**, 265
- Kessler, R., Bernstein, J. P., Cinabro, D., et al. 2009, *PASP*, **121**, 1028
- Kessler, R., Brout, D., D'Andrea, C. B., et al. 2019a, *MNRAS*, **485**, 1171
- Kessler, R., Cinabro, D., Bassett, B., et al. 2010, *ApJ*, **717**, 40
- Kessler, R., Narayan, G., Avelino, A., et al. 2019b, *PASP*, **131**, 094501
- Korytov, D., Hearin, A., Kovacs, E., et al. 2019, *ApJS*, **245**, 26
- Lee, J., Acevedo, M., Sako, M., et al. 2023, *AJ*, **165**, 222
- Lidman, C., Tucker, B. E., Davis, T. M., et al. 2020, *MNRAS*, **496**, 19
- Lokken, M., Gagliano, A., Narayan, G., et al. 2023, *MNRAS*, **520**, 2887
- Malz, A. I., Marshall, P. J., DeRose, J., et al. 2018, *AJ*, **156**, 35
- Meyers, J. E., & Burchat, P. R. 2015, *ApJ*, **807**, 182
- Mitra, A., Kessler, R., More, S., & Hlozek, R. 2023, *ApJ*, **944**, 212
- Möller, A., Wiseman, P., Smith, M., et al. 2024, *MNRAS*, **533**, 2073
- Narayan, G. 2023, AAS Meeting, **55**, 117.01
- Palanque-Delabrouille, N., Ruhlmann-Kleider, V., Pascal, S., et al. 2010, *A&A*, **514**, A63
- Pasquet, J., Bertin, E., Treyer, M., Arnouts, S., & Fouchez, D. 2019, *A&A*, **621**, A26
- Pearson, K. 1894, *RSPTA*, **185**, 71
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *ApJ*, **517**, 565
- Pierel, J. D. R., Jones, D. O., Kenworthy, W. D., et al. 2022, *ApJ*, **939**, 11
- Qu, H., & Sako, M. 2023, *ApJ*, **954**, 201
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, **116**, 1009
- Rowe, B. T. P., Jarvis, M., Mandelbaum, R., et al. 2015, *A&C*, **10**, 121
- Sako, M., Bassett, B., Becker, A. C., et al. 2018, *PASP*, **130**, 064002
- Scolnic, D., & Kessler, R. 2016, *ApJ*, **822**, L35
- Scolnic, D. M., Lochner, M., Gris, P., et al. 2018, arXiv:1812.00516
- Smith, M., D'Andrea, C. B., Sullivan, M., et al. 2020, *AJ*, **160**, 267
- Wang, Y. 2006, *ApJ*, **654**, L123
- Wang, Y., Gjergo, E., & Kuhlmann, S. 2015, *MNRAS*, **451**, 1955
- Wang, Y., Narayan, G., & Wood-Vasey, M. 2007, *MNRAS*, **382**, 377