

Quantifying concordance of correlated cosmological data sets

Marco Raveri¹, Georgios Zacharegkas², and Wayne Hu³

¹*Center for Particle Cosmology, Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

²*Kavli Institute for Cosmological Physics, Department of Astronomy & Astrophysics, The University of Chicago, Chicago, Illinois 60637, USA*

³*Kavli Institute for Cosmological Physics, Department of Astronomy & Astrophysics, Enrico Fermi Institute, The University of Chicago, Chicago, Illinois 60637, USA*



(Received 21 January 2020; accepted 5 May 2020; published 21 May 2020)

We develop estimators of agreement and disagreement between correlated cosmological data sets. These account for data correlations when computing the significance of both tensions and excess confirmation while remaining statistically optimal. We discuss and thoroughly characterize different approaches commenting on the ones that have the best behavior in practical applications. We complement the calculation of their statistical distribution within the Gaussian model with one estimator that takes non-Gaussianities fully into account. To illustrate the use of our techniques, we apply these estimators to supernovae measurements of the distance-redshift relation, absolutely calibrated by the local distance ladder. The suite of best estimators that we discuss finds results that are in excellent agreement between estimators and find no indications of significant internal inconsistencies in this data set above the 1% probability threshold. This shows the robustness of local determinations of the Hubble constant to features in the distance-redshift relation.

DOI: [10.1103/PhysRevD.101.103527](https://doi.org/10.1103/PhysRevD.101.103527)

I. INTRODUCTION

The remarkable ability of the Λ CDM model to explain a wide range of observations, such as the spectrum of the fluctuations in the cosmic microwave background, observations of gravitational lensing and the clustering of galaxies, has made it widely accepted as the standard cosmological model.

However, despite its successes, as the precision of different experiments has increased, so too has the statistical significance of discrepancies between their inferences of the Λ CDM model parameters (for a recent review see [1]). Such discrepancies deserve close attention since they may hint at the existence of new physical phenomena or to the presence of residual systematic effects that are not yet understood.

In parallel with increased experimental precision, cosmological data sets have become increasingly complex to the point that understanding whether different probes agree or not requires the use of dedicated statistical tools [2–27].

In this paper, we develop concordance and discordance estimators (CDEs) for data sets that are correlated. This allows us to study both internal consistency of a data set and the mutual consistency of different correlated data sets. In the former case splitting a data set in different parts naturally leads to correlated data pieces. The latter case, even though present cosmological measurements are only weakly correlated, will become increasingly important in

the future as correlations will become more relevant, reflecting the fact that different experiments will be measuring different properties of the same underlying sky.

We focus on extending the CDEs introduced for uncorrelated data sets [20] to correlated data sets. In particular, we discuss estimators quantifying parameter shifts between two correlated data sets and goodness-of-fit loss when two data sets are joined together. We analyze these estimators under the Gaussian linear model (GLM), assuming Gaussianity of the data and model parameters. We comment on mitigation against non-Gaussianities that is built in some of them in practical applications.

To complement and check these results we also discuss a purely Markov chain Monte Carlo (MCMC) approach, which uses the full distribution of parameters, as inferred from MCMC samples, to compute the probability of a parameter shift. This technique fully takes into account parameter non-Gaussianities and any nonlinear aspect of the model.

When considering correlated data sets we follow two different strategies. In the first approach we consider the two disjoint data sets and build estimators working on their separate parameter inference though properly including their correlation. In the second one, we always consider the joint data set but fit the two parts of the data set with different cosmological parameters. This approach was also employed in [28–35]. We call the first strategy the data split

one while we refer to the second one as the parameter split strategy. The data and parameter split techniques are equivalent when data sets are uncorrelated but different when they are correlated.

Overall we find that the parameter split methodology is more convenient in practice. All estimators that we consider, in this case, can be easily obtained from the posterior distribution while the same is not true for the data split estimators. Moreover the different parameter split estimators should agree if non-Gaussianities are negligible, providing essential cross-checks of the reported results.

As a demonstration of our methodology we apply these estimators to the Pantheon type Ia supernova (SN) sample [36] calibrated with measurements of the Hubble constant from [37]. We choose to split the SN data into two subsets at the redshift values of $z = 0.3$ and $z = 0.7$, loosely corresponding, respectively, to the time of dark energy–dark matter equality and the time at which cosmic acceleration begins.

We find that the Λ CDM model provides a good fit to these data with the exception of the $z > 0.7$ SN measurements for which the agreement seems to be too good at the 94% confidence level. We also find excellent agreement between the results of different parameter split estimators, regardless of mild non-Gaussianities in the SN posterior that are effectively mitigated. As reported by different estimators the first SN split at $z = 0.3$ is in good agreement with the Λ CDM model while the second split, at $z = 0.7$, results in parameters that are too close to each other at about 97% probability. Furthermore, we do not find significant indications of differences in the estimates of the Hubble constant between different SN redshift splits indicating that its direct measurement is robust against split in the SN catalog and to features in the SN distance-redshift relation.

This paper is organized as follows. In Sec. II we briefly review the Gaussian linear model. In Sec. III we describe the SN data set that we employ in this work as an illustration of our statistical techniques. Section IV includes a quantification of the importance of SN data correlations and their impact on cosmological parameters. The properties and differences of the two approaches that we follow, data splits and parameter split, are discussed in Sec. V, before a detailed discussion of the two methodologies separately in Sec. VI and Sec. VII, respectively. We summarize our conclusions in Sec. VIII.

Details of our techniques are presented in a series of Appendices. In Appendix A we discuss the canonical correlation analysis for quantifying the impact of correlations. In Appendix B we provide a worked pedagogical example that clarifies the differences between data splits and parameter split. In Appendices C and D we derive the exact distributions of the goodness-of-fit loss statistic in the cases of data and parameter splits, respectively. In Appendix E, we generalize our discussion of splitting the data into an arbitrary number of subsets.

II. THE GAUSSIAN LINEAR MODEL

In this section we gather some basic definitions that we will later use throughout the paper. For an in-depth discussion of the Gaussian linear model (GLM) we refer the reader to [20].

We denote the multivariate Gaussian distribution in N dimensions with mean $\bar{\theta}$ and covariance \mathcal{C} as $\mathcal{N}_N(\theta; \bar{\theta}, \mathcal{C})$. For a given data set, D , described by a model \mathcal{M} that depends on a set of parameters θ , the posterior probability distribution of the parameters is given by

$$P(\theta|D, \mathcal{M}) = \frac{\mathcal{L}(\theta)\Pi(\theta)}{\mathcal{E}}, \quad (1)$$

where the likelihood is the probability of the data at any given choice of parameters $\mathcal{L}(\theta) = P(D|\theta, \mathcal{M})$ and any prior knowledge is encoded in $\Pi(\theta)$. The normalization of the posterior, $\mathcal{E} \equiv P(D|\mathcal{M})$, is the evidence that provides the probability distribution of the data given the model \mathcal{M} .

In this section, we assume that the prior distribution is Gaussian in the model parameters, $\Pi(\theta) = \mathcal{N}_N(\theta; \theta_\Pi, \mathcal{C}_\Pi)$, with mean θ_Π and covariance \mathcal{C}_Π . As discussed in [20] this is a good choice to use in practice as it allows us to treat Gaussian priors on nuisance parameters exactly and models the most relevant features of informative flat priors: the scale of the prior and its central value.

We further assume that the likelihood is a Gaussian distribution in data space, $\mathcal{L}(\theta) = \mathcal{N}_d(x; m, \Sigma)$ and we denote by d the number of data points x and Σ their covariance matrix. The mean of the distribution is given by the model prediction, $m(\theta)$.

The GLM assumes that one can linearly expand the model prediction, $m(\theta)$, around a given parameter point. Since we are working with Gaussian priors, for simplicity, in the following we assume that the linear model expansion point is the prior center θ_Π and we can write

$$m(\theta) \approx m_\Pi + M(\theta - \theta_\Pi), \quad (2)$$

where $m_\Pi \equiv m(\theta_\Pi)$ and $M \equiv (\partial m / \partial \theta)|_{\theta_\Pi}$ is the Jacobian of the transformation between parameter and data space.

Given the model prediction m_Π , the residual of a randomly chosen data point x , henceforth $X \equiv x - m_\Pi$, can be projected onto a component along the linear model, $\mathbb{P}X$, and another component orthogonal to it, $X - \mathbb{P}X = (\mathbb{I} - \mathbb{P})X$. The projector can be thought of as a two-step process. The first step is to construct the linear combinations of data, namely $\tilde{M}X$, that give the parameter estimates: $\theta - \theta_\Pi = \tilde{M}X$, where $\tilde{M} = \mathcal{C}M^T\Sigma^{-1}$, with

$$\mathcal{C} = \langle (\theta - \theta_\Pi)(\theta - \theta_\Pi)^T \rangle = (M^T\Sigma^{-1}M)^{-1}, \quad (3)$$

as the parameter covariance or inverse Fisher matrix. Then, as a second step, given the parameter estimates we transform back into data space using the Jacobian,

$M(\theta - \theta_{\Pi}) = M\tilde{M}X \equiv \mathbb{P}X$. Thus $\mathbb{P} \equiv M\tilde{M}$ is the full projector:

$$\mathbb{P} = MCM^T\Sigma^{-1}, \quad (4)$$

and $(\mathbb{I} - \mathbb{P})$ is its complement.

In the GLM the maximum likelihood (ML) is given by

$$\begin{aligned} \ln \mathcal{L}_{\max} = & -\frac{1}{2}X^T(\mathbb{I} - \mathbb{P})^T\Sigma^{-1}(\mathbb{I} - \mathbb{P})X \\ & -\frac{d}{2}\ln(2\pi) - \frac{1}{2}\ln(|\Sigma|), \end{aligned} \quad (5)$$

where we used $|\cdot|$ to denote the determinant of a matrix. Notice that the first line of Eq. (5) contains all the pieces that depend on the data while the second one contains normalization constants that are often neglected. The parameters corresponding to the maximum likelihood model are given by

$$\theta_{\text{ML}} = \tilde{M}(x - m_{\Pi} + M\theta_{\Pi}). \quad (6)$$

Over realizations of data, the maximum likelihood parameters are distributed as $\mathcal{N}_N(\theta; \theta_{\Pi}, \mathcal{C})$. These expressions refer to the true maximum likelihood of a model and they should be obtained without reference to the prior; θ_{Π} appears here due to the assumption that the prior mean coincides with the GLM expansion point.

The maximum posterior (MAP) parameters combine the ML parameters with the prior:

$$\theta_p = C_p(C_{\Pi}^{-1}\theta_{\Pi} + C^{-1}\theta_{\text{ML}}), \quad (7)$$

where $C_p^{-1} = C_{\Pi}^{-1} + C^{-1}$. Under the GLM the maximum posterior parameters are distributed as $\mathcal{N}(\theta; \theta_{\Pi}, C_p)$.

Within the GLM the probability of the data, i.e., the evidence, is Gaussian distributed for the Gaussian priors that we consider and is given by $\mathcal{E} = \mathcal{N}_d(x; m(\theta_{\Pi}), \Sigma + MC_{\Pi}M^T)$.

We define all the statistics, Q , that we discuss in this paper to follow the convention that if $P(Q > Q^{\text{obs}})$ approaches zero then the observed value lies in the tail of the distribution that we would associate with a tension; if it approaches one the observed value would be in the tail associated with excess confirmation.

III. SN DATA SET AND MODEL

As an example case we study the internal consistency of the Pantheon type Ia supernovae (SN) sample [36] to redshift splits under the Λ CDM model.

The Pantheon collaboration provides measurements of the SN magnitude, corrected for stretch, color, etc., relative to a fiducial absolute magnitude $\mathbf{m} - \mathfrak{M}_{\text{fid}}$, with its covariance Σ_{SN} . $\mathfrak{M}_{\text{fid}} = -19.34$ is predetermined by a fit under the assumption of a Hubble constant of $H_0^{\text{fid}} = 70$ for

definiteness (see e.g., [38]).¹ Here and throughout H_0 is quoted in units of $\text{km s}^{-1} \text{Mpc}^{-1}$, whereas $c = 1$ in general formulas. The likelihood is then analytically marginalized over the true absolute magnitude \mathfrak{M} when considering the distance modulus $\mathbf{m} - \mathfrak{M}$. Hence the reference Pantheon SN likelihood is not Gaussian in $\mathbf{m} - \mathfrak{M}$. Some of the methods that we discuss rely on Gaussianity of the likelihood in data space and in addition we wish to explore compatibility of H_0 determinations between subsets of the data. To achieve this we introduce the SN absolute magnitude \mathfrak{M} as an additional model parameter. We take the SN-based measurements of the Hubble constant in [37], of $H_0^{\text{SN}} = 74.03 \pm 1.42$, to infer a measurement of $\mathfrak{M} = \hat{\mathfrak{M}} \pm \sigma_{\mathfrak{M}}$ as

$$\begin{aligned} \hat{\mathfrak{M}} &= 5 \log_{10} \frac{H_0^{\text{SN}}}{H_0^{\text{fid}}} + \mathfrak{M}_{\text{fid}}, \\ \sigma_{\mathfrak{M}} &= \frac{5}{\ln 10} \frac{\sigma_{H_0^{\text{SN}}}}{H_0^{\text{SN}}}, \end{aligned} \quad (8)$$

that we use as a prior on the additional model parameter \mathfrak{M} . The SN data likelihood, as provided by the Pantheon collaboration, remains a Gaussian distribution,

$$\mathcal{L}_{\text{SN}} = \mathcal{N}\left(\mathbf{m} - \mathfrak{M}_{\text{fid}}; 5 \log \frac{d_L}{10 \text{ pc}} + \mathfrak{M} - \mathfrak{M}_{\text{fid}}, \Sigma_{\text{SN}}\right), \quad (9)$$

but the relationship to the luminosity distance d_L comes through the likelihood for the absolute magnitude data,

$$\mathcal{L}_{\mathfrak{M}} = \mathcal{N}(\mathfrak{M}; \hat{\mathfrak{M}}, \sigma_{\mathfrak{M}}^2), \quad (10)$$

such that

$$\mathcal{L} = \mathcal{L}_{\text{SN}} \times \mathcal{L}_{\mathfrak{M}}. \quad (11)$$

This obviates the need to marginalize \mathfrak{M} when considering cosmological constraints. Note that by inferring $\mathcal{L}_{\mathfrak{M}}$ from the determination of H_0 from [37] rather than directly calibrating \mathfrak{M} in the process of Pantheon data reduction, we force the Pantheon data set as a whole to return the same mean value for H_0 but allow for nontrivial consistency tests with subsets of the data.

As a function of SN redshift, the flat Λ CDM model for d_L is

$$d_L(z) = \frac{1+z}{H_0} \int_0^z \frac{dz'}{\sqrt{\Omega_m(1+z')^3 + 1 - \Omega_m}}. \quad (12)$$

¹Note that <https://github.com/dscolnic/Pantheon> provides $\mathbf{m} - \mathfrak{M}_{\text{fid}}$ whereas <https://github.com/cmbant/CosmoMC> provides \mathfrak{M} from which $\mathfrak{M}_{\text{fid}}$ can be extracted.

We place flat priors on the range of the two cosmological parameters $H_0 \in [40, 100]$ and $\Omega_m \in [0, 1]$. These control respectively the amplitude and shape of $d_L(z)$.

The Pantheon SN sample covers the redshift range $z \in [0.01, 2.26]$ with 1048 SN distance measurements. Given this many data points it is important to fix the data splits with a solid *a priori* criterium to avoid look-elsewhere type corrections to statistical significance that are hard to quantify.

We therefore choose two relevant physical times in the SN redshift range with which to cut the SN sample in two:

- (i) $z_{\text{cut}} = 0.3$ approximately the time of dark matter/dark energy equality;
- (ii) $z_{\text{cut}} = 0.7$ approximately the time at which cosmic acceleration begins.

These two data splits are very different and complementary in ways that allow us to explore the most common situations that one might face in practical applications. The first split has almost equal weight in both parts with 630 SN below $z = 0.3$ and 418 data points above. The second split is heavily weighted toward the first part of the data set, with 924 SN below $z = 0.7$, compared with 124 SN above $z = 0.7$.

When analyzing the splits separately, the absolute magnitude measurement would be applied to each half separately. When joining the two data splits, we need to take into account this double counting of data. This is equivalent to introducing two absolute magnitude measurements that are fully correlated between the splits. As we shall see, this provides an extreme, albeit trivial, example of fully correlated data points between data sets and their impact.

In all the following, cosmological predictions for the Λ CDM model are obtained with the CAMB [39] code. The parameter posterior distributions are obtained with Markov chain Monte Carlo (MCMC) sampling with the CosmoMC [40] code and their analysis largely relies on the GetDist code [41].

IV. IMPACT OF CORRELATIONS

The stronger the correlation between two data sets the more crucial it is to use statistical tools which take these correlations into account.

We can see in Fig. 1 the impact that data correlations have on the joint SN posterior. In both panels we show the full results taking into account all correlated modes, while we also show the results when we neglect the correlation between the data sets of the SN split, while keeping the correlation information within each subset of the split. As we can clearly notice, the posterior is influenced in two different ways: the peak of the distribution shifts; also, the variance changes and looks more constraining when we neglect correlations, because we are not considering the part of the information in the two data sets that is redundant.

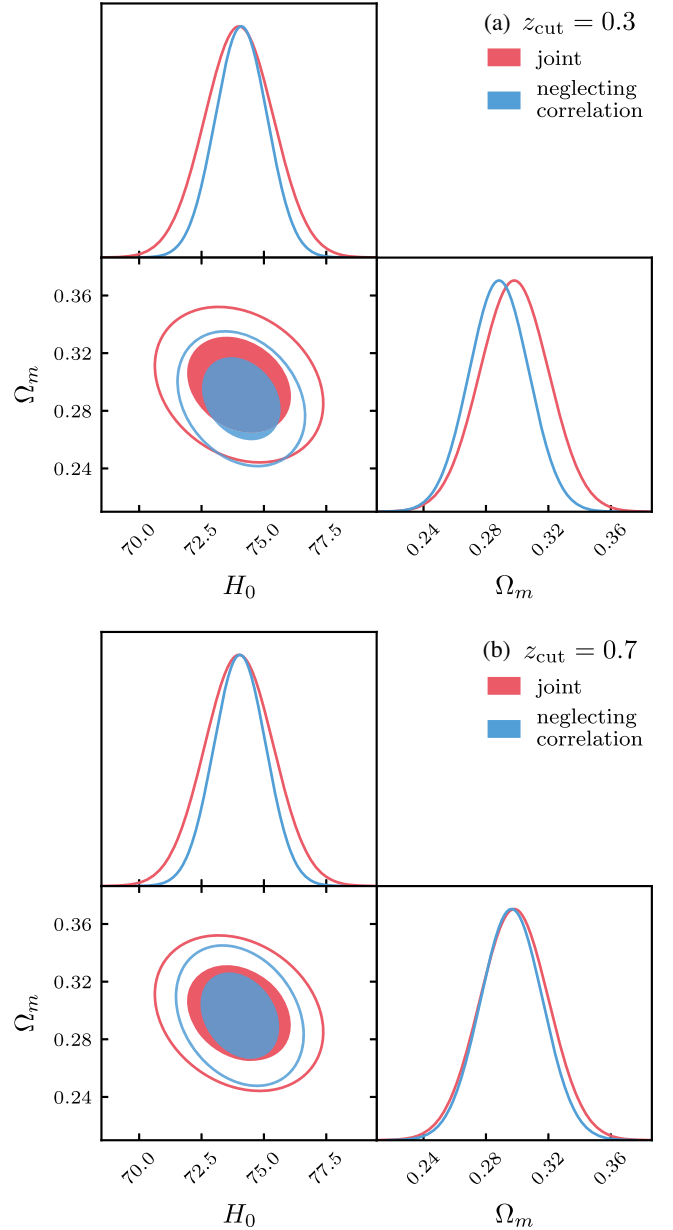


FIG. 1. The joint marginalized posterior of the full Pantheon SN data set compared to the results obtained neglecting the correlation between the $z > z_{\text{cut}}$ and $z < z_{\text{cut}}$ SN measurements for the two SN splits that we consider, in panels (a) and (b) respectively. The filled contour corresponds to the 68% CL region while the continuous contour shows the 95% CL region.

In multiple dimensions the correlation strength can be quantified by means of the canonical correlation analysis (CCA) [42]. CCA allows us to understand the change in parameter variances as we summarize here and fully discuss in Appendix A.

The two data splits that we consider have an almost completely correlated data mode due to the common calibration while the second correlated mode has a correlation coefficient of approximately $\rho_{12} = 0.5$ for both splits.

The maximum error that we would make on the determination of the parameter variance, with respect to the full joint estimate, if we were to neglect correlations is discussed in Appendix A and is bounded by

$$1 - |\rho_{12}| \leq \left(\frac{\sigma_J^2}{\sigma_{\text{no corr}}^2} \right) \leq 1 + |\rho_{12}|, \quad (13)$$

where $\sigma_{\text{no corr}}^2$ is the wrong estimate of the variance of one parameter obtained neglecting correlations, σ_J^2 is the correct estimate of the variance, accounting for all data correlations, and ρ_{12} is the maximum correlation coefficient.

In Fig. 1 we can clearly see that neglecting the duplication of the SN amplitude leads, when joining the two SN splits, to a false $\sqrt{2}$ improvement in the determination of the Hubble constant. This is because the two determinations from the splits are nearly fully correlated because of the absolute magnitude calibration that they share. The second correlated mode is then responsible for the residual underestimate of the error in the determination of Ω_m at about the 20% level and reflects correlations in the measurement of SN magnitudes which resemble that parameter.

We can understand the shift in the posterior in Fig. 1 as a difference in the SN maximum posterior due to data correlations. Underlying this difference is the impact of correlations on the likelihood at a given parameter point. As discussed in [20] the value of the likelihood at MAP, $\hat{\theta}_p$, can be used as a goodness-of-fit measure to test the consistency of a data set with the model at hand. The MAP measure for goodness of fit is then given by

$$Q_{\text{MAP}} \equiv -2 \ln \mathcal{L}(\hat{\theta}_p) + 2 \left\langle \ln \mathcal{L}(\hat{\theta}_p) \right\rangle_D + \langle Q_{\text{MAP}} \rangle_D, \quad (14)$$

where the average is over data realizations. For Gaussian priors Q_{MAP} is distributed as a sum of Gamma distributed variables which can be (conservatively) approximated by a chi-square distribution of $d - N_{\text{eff}}$ degrees of freedom, $Q_{\text{MAP}} \sim \chi^2(d - N_{\text{eff}})$, where

$$N_{\text{eff}} = N - \text{tr}[\mathcal{C}_{\Pi}^{-1} \mathcal{C}_p] \quad (15)$$

is the number of effective parameters that are being constrained by the data over the prior, with N being the total number of model parameters. Notice that there might be cases where the data covariance matrix is singular. In these cases one needs to compute the number of data points as the rank of the covariance, $d = \text{rank}(\Sigma)$.

The results of the application of this goodness-of-fit statistics to the SN data are summarized in Table I. As we can see the full SN catalog is a reasonably good fit. Notice that the joint data set contains two fully correlated measurements of \mathfrak{M} so that the data covariance is singular and the number of data points is computed as the rank of the

TABLE I. The results of the application of the Q_{MAP} goodness-of-fit statistics to the SN data sets that we consider. The first column specifies the data set that is considered, the second one the value of likelihood at maximum posterior. The third column shows the number of effective parameters, as in Eq. (15), highlighting that one parameter is constrained by the measurement of the absolute magnitude of SN. The fourth column reports the number of data points used, highlighting the extra measurement of the SN absolute calibration. The fifth column shows the probability to exceed the Q_{MAP} goodness-of-fit statistic.

Data	$-2 \ln \mathcal{L}_{\text{MAP}}$	N_{eff}	d	$P(Q_{\text{MAP}} > Q_{\text{MAP}}^{\text{obs}})$
$z < 0.3$	633.63	$1.95 + 1$	$630 + 1$	43.0%
$z > 0.3$	388.63	$1.96 + 1$	$418 + 1$	82.9%
$z < 0.7$	926.92	$1.99 + 1$	$924 + 1$	44.8%
$z > 0.7$	99.64	$1.49 + 1$	$124 + 1$	93.6%
Joint	1026.86	$1.99 + 1$	$1048 + 1$	65.8%

data covariance. If we were to neglect correlations both the position of the maximum posterior and the likelihood value at MAP would change. In the case of the $z_{\text{cut}} = 0.3$ split these changes are at about $2\Delta \ln \mathcal{L} = 0.2$, which corresponds to about half a sigma shift, which we see in the posterior of Fig. 1(a). We also find that the corresponding results for the $z_{\text{cut}} = 0.7$ split are smaller. Neglecting correlations, in both cases, has also the effect of overestimating the number of degrees of freedom of the Q_{MAP} distribution. These would be $d_{\text{SN}} + 2$ given that the full correlation between the \mathfrak{M} measurements is neglected and the covariance matrix becomes nonsingular. In this case then, the goodness-of-fit test is mostly insensitive to the presence of correlations between the two SN splits, given the small change in the likelihood and the large number of data points.

We conclude this section by discussing the implications of the goodness-of-fit values shown in Table I. In the $z_{\text{cut}} = 0.3$ case, as with the joint data set, both splits contain enough SN measurements to measure the two relevant cosmological parameters. The fit to the data is reasonably good and does not indicate the presence of tension or confirmation of high statistical significance. When considering the split of the data at $z_{\text{cut}} = 0.7$, the vast majority of the data points falls in the low-redshift subset. This means that the low-redshift part measures both parameters better than the prior while the high redshift one starts being influenced by the prior, as reflected in N_{eff} . The first part of the data provides a reasonably good fit, while the second one leans toward the fit being too good at a probability level that, however, does not exceed 95%.

V. DATA SPLITS AND PARAMETER SPLITS

We can split a correlated data set by taking partitions of the joint data set, that we denote as D_J , in two pieces that we indicate with D_1 and D_2 . Hereafter we denote quantities that refer to the joint data set with J and quantities referring

to the first and second data sets with the subscript 1 and 2 respectively. In Appendix E we discuss the natural generalization to an arbitrary number of splits. We indicate the joint $D_J = D_1 \cup D_2$ data covariance as

$$\Sigma_J \equiv \begin{pmatrix} \Sigma_1 & \Sigma_{12} \\ \Sigma_{21} & \Sigma_2 \end{pmatrix}, \quad (16)$$

which is in general not block diagonal. Since the full covariance has to be symmetric then $\Sigma_1 = \Sigma_1^T$, $\Sigma_2 = \Sigma_2^T$ and $\Sigma_{21} = \Sigma_{12}^T$. Notice that Σ_1 , Σ_2 and Σ_J have all to be symmetric and positive definite.

Within the GLM this data separation projects onto parameter space through the single and joint Jacobian matrices that are related by $M_J^T = (\partial m_1^T / \partial \theta, \partial m_2^T / \partial \theta) \equiv (M_1^T, M_2^T)$. The GLM estimate of the ML parameters in this case is given by

$$\begin{aligned} \theta_{aS}^{\text{ML}} &= C_{aS} M_a^T \Sigma_a^{-1} (x_a - m_{\Pi a} + M_a \theta_{\Pi}), \\ \theta_J^{\text{ML}} &= C_J M_J^T \Sigma_J^{-1} (x_J - m_{\Pi J} + M_J \theta_{\Pi}), \end{aligned} \quad (17)$$

with $a \in \{1, 2\}$. Hereafter we denote quantities that are obtained within the single analysis of the split data sets with S . The maximum posterior parameters are obtained by adding on top of these estimates the prior, as in Eq. (7). As we can see the inference of the parameters for the separate data splits depends only on the given data set, while their joint inference is influenced by the complementary set, through the correlation between the two. For this reason it is not possible, in presence of data correlations, to reconstruct the joint ML parameters as a linear combination of parameter quantities that live in the single parameter space.

We can still, however, compute the covariance between different data split parameters both at the ML and MAP level as

$$\begin{aligned} \text{cov}(\theta_{1S}^{\text{ML}}, \theta_{2S}^{\text{ML}}) &= C_{1S} M_1^T \Sigma_1^{-1} \Sigma_{12} \Sigma_2^{-1} M_2 C_{2S}, \\ \text{cov}(\theta_{1S}^p, \theta_{2S}^p) &= C_{p1S} C_{\Pi}^{-1} C_{p2S} \\ &\quad + C_{p1S} M_1^T \Sigma_1^{-1} \Sigma_{12} \Sigma_2^{-1} M_2 C_{p2S}. \end{aligned} \quad (18)$$

As we can see these depend on both parameter space and data space quantities to take into account that the single parameter covariances do not include correlation contributions.

As an alternative strategy we can think of the split as originating in parameter space, describing the two parts of the joint data set with different sets of parameters and always fitting the joint likelihood. We denote with the subscript C quantities that refer to this strategy of parameter duplication and, for example, we work with a $2N$ parameter vector that is defined by $\theta_C \equiv (\theta_{1C}, \theta_{2C})^T$. Similarly, we can define the duplicated prior parameter vector by $\theta_{\Pi C} = (\theta_{\Pi}, \theta_{\Pi})^T$. One subtlety is that our null hypothesis is still that the data is drawn from the prior distribution of a

single parameter. We shall see that this difference between the split analysis and statistical properties of the data causes minor issues when counting these parameters if they are partially, but not fully constrained, by the prior.

In the GLM the block structure of the covariance in Eq. (16) then projects on the two parameter copies with the Jacobian given by

$$M_C = \begin{pmatrix} \partial m_1 / \partial \theta_1 & \mathbb{O} \\ \mathbb{O} & \partial m_2 / \partial \theta_2 \end{pmatrix} = \begin{pmatrix} M_1 & \mathbb{O} \\ \mathbb{O} & M_2 \end{pmatrix}. \quad (19)$$

The maximum likelihood estimate of the copy parameters is given by Eq. (6) and explicitly reads

$$\begin{aligned} \theta_C^{\text{ML}} &= \begin{pmatrix} \theta_{1C}^{\text{ML}} \\ \theta_{2C}^{\text{ML}} \end{pmatrix} \\ &= C_C M_C^T \Sigma_J^{-1} (x_J - m_{\Pi C} + M_C \theta_{\Pi C}), \end{aligned} \quad (20)$$

where we have used the definition of the parameter copies covariance $C_C^{-1} = M_C^T \Sigma_J^{-1} M_C$ and we have defined the duplicate prior center model prediction $m_{\Pi C} = (m_{\Pi}, m_{\Pi})^T$. The parameter copy ML is then Gaussian distributed, over the space of data draws, with $\theta_C^{\text{ML}} \sim \mathcal{N}(\theta_C; \theta_{\Pi C}, C_C)$. The maximum posterior parameters are obtained by adding copies of the Gaussian priors on top of the ML result. We write explicitly the block structure of the parameter copies posterior covariance as

$$C_{pC} \equiv \begin{pmatrix} C_{p1C} & C_{p12C} \\ C_{p21C} & C_{p2C} \end{pmatrix}, \quad (21)$$

which allows us to write the posterior estimate for the first parameter copy as

$$\begin{aligned} \theta_{1C}^p &= \theta_{1C}^{\text{ML}} - C_{p1C} C_{\Pi}^{-1} (\theta_{1C}^{\text{ML}} - \theta_{\Pi}) \\ &\quad - C_{p12C} C_{\Pi}^{-1} (\theta_{2C}^{\text{ML}} - \theta_{\Pi}), \end{aligned} \quad (22)$$

and similarly for the second parameter copy.

As we can see the GLM posterior distribution for one of the parameter copies is related to the parameters of the other copy in two ways: first indirectly in the ML estimate and then directly at the MAP level. In particular, at the ML level, the parameters of one copy are related to the parameters of the other because of the shared data in Eq. (20). This is a natural consequence of the parameter duplication technique: since we always fit the joint data set, split parameters are influenced by data in both splits through their correlation.

In the posterior, there is a direct coupling between the ML and posterior estimators of the copies. In Eq. (22) this coupling is mediated by C_{p12C} in the last term. In the limit where the data is uncorrelated this term would vanish and, therefore, the sets would not be able to communicate with each other; we would, therefore, retrieve the expressions in

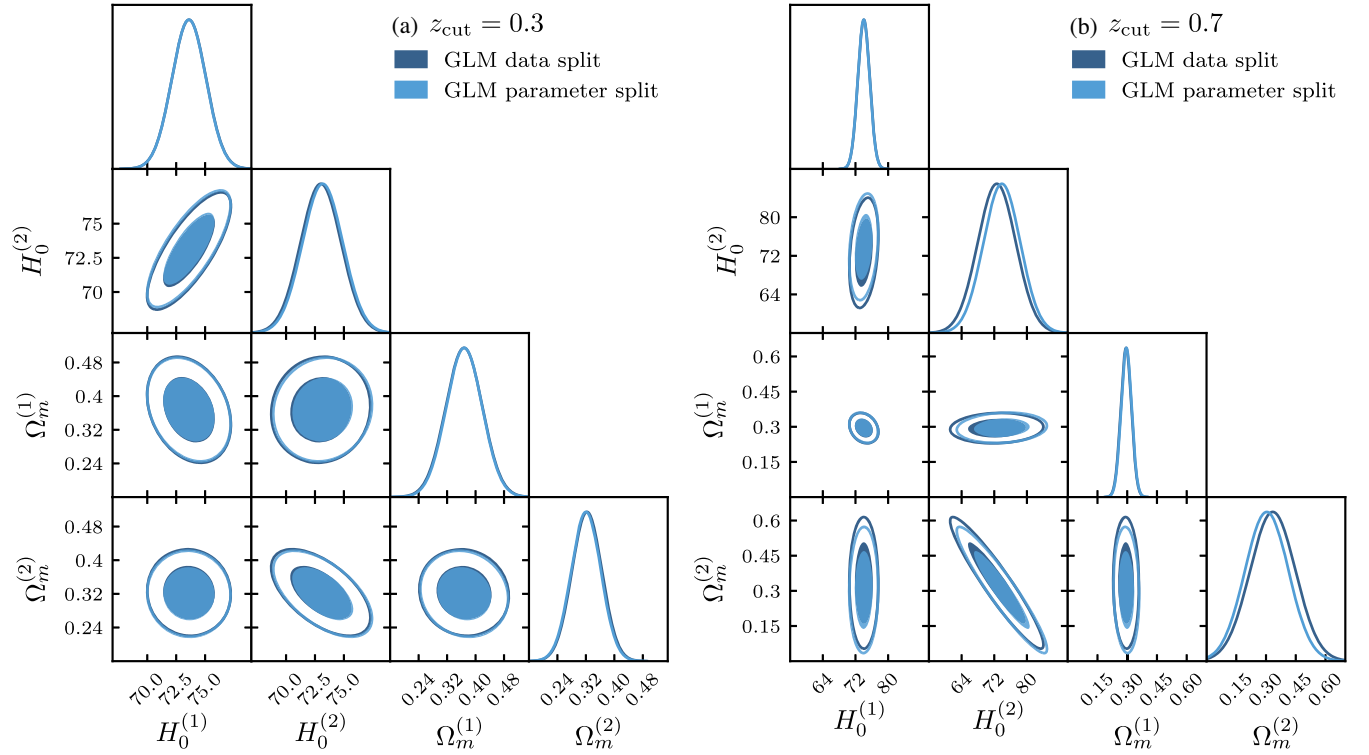


FIG. 2. The GLM joint marginalized posterior parameter distribution for SN data splits and parameter splits. Different panels show different SN redshift cuts. The filled contour corresponds to the 68% CL region while the continuous contour shows the 95% CL region.

[20], which would also be the same as the corresponding expressions under the data split methodology. With correlated copy parameters, the maximization of the *joint* posterior depends on the ML values of each. Contrast this with the case of the separate parameters of the data split. Even though the ML values are still correlated according to Eq. (18), maximization of the posterior for each parameter is performed without reference or knowledge of its complement.

The joint parameter results can be viewed as a subspace of the parameter copies where all the copies have the same value. We define the projection matrix on this subspace as

$$D_C^T = (\mathbb{I}_{N \times N}, \mathbb{I}_{N \times N}). \quad (23)$$

When expressed as D_C , it is known as the design matrix, which takes a single set of N parameters and produces two separate parameters, i.e., the $2N$ copy parameters, to be estimated by the data. Notice that the relation between the joint and copy Jacobian is given by $M_J^T = D_C^T M_C^T$. The joint parameter covariance is a linear combination of the copy parameter covariance given by $C_J^{-1} = D_C^T C_C^{-1} D_C$. The linear combination of the copy parameter estimates that forms the joint parameter estimate is

$$\theta_J^{\text{ML}} = C_J D_C^T C_C^{-1} \theta_C^{\text{ML}}. \quad (24)$$

Likewise $D_C C_J D_C^T C_C^{-1}$ is the parameter projector that projects the copy parameters onto the space where they are the same $(\theta_{1C}^{\text{ML}}, \theta_{2C}^{\text{ML}})^T \rightarrow (\theta_J^{\text{ML}}, \theta_J^{\text{ML}})^T$.

The physical meaning of the parameter estimates between the data split and parameter duplication approaches in principle differs in the presence of data correlations. In Appendix B we present a simple example which illustrates these differences. Here we would like to emphasize that the data split strategy provides us with two distributions $P(\theta_{1S}|D_1)$ and $P(\theta_{2S}|D_2)$ that are interpreted as the posterior for the parameters of one data set with no knowledge of the other. On the other hand, the second strategy provides the joint parameter distribution $P(\theta_{1C}, \theta_{2C}|D_J)$ for both data sets. When marginalized over one of the parameter copies this gives $P(\theta_{1C}|D_J)$ which is the distribution of the parameters of the first data set, given that the full data set has been measured.

In our example, we duplicate all the SN cosmological parameters but we do not duplicate the parameter describing the absolute magnitude. This would be fully correlated, since the corresponding measurements in the joint data set are fully correlated, closely matching the example case discussed in Appendix B. As a consequence duplicating the SN calibration parameter and keeping track of its correlations gives the same results as not duplicating it and we omit its duplication for simplicity. We have checked that all the results that we discuss are unchanged.

We find that the difference between the two ways of estimating parameters from the split SN data is minor. In Fig. 2 we show the GLM prediction of the parameter posterior in the two split senses that we have discussed.

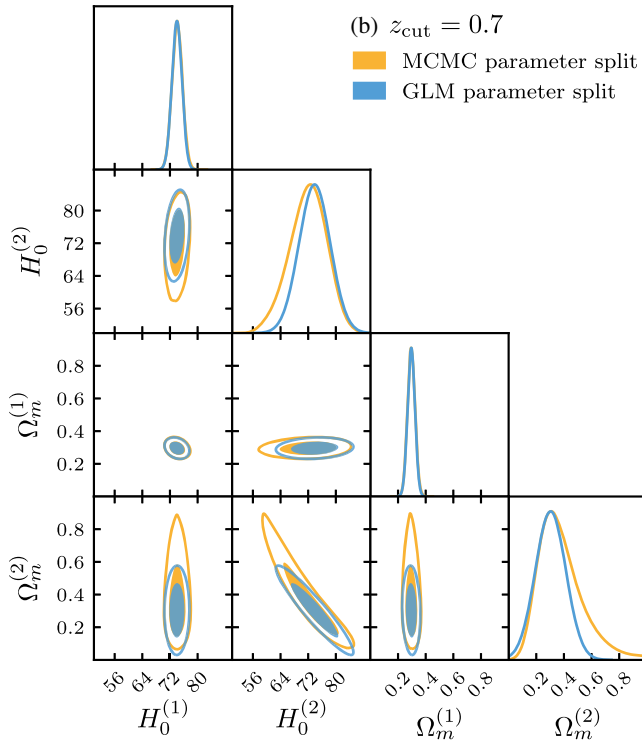
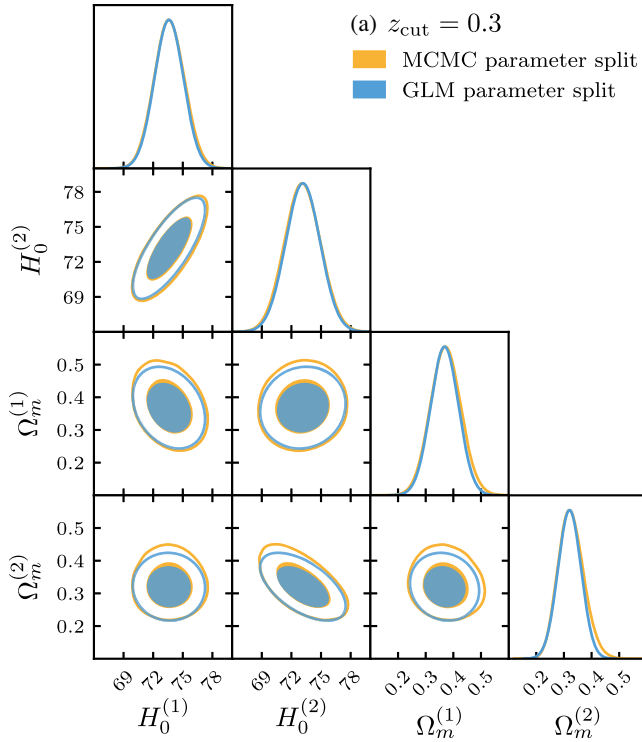


FIG. 3. The comparison of the GLM and MCMC joint marginalized posterior parameter distribution for SN parameter splits. Different panels show different SN redshift cuts. The filled contour corresponds to the 68% CL region while the continuous contour shows the 95% CL region.

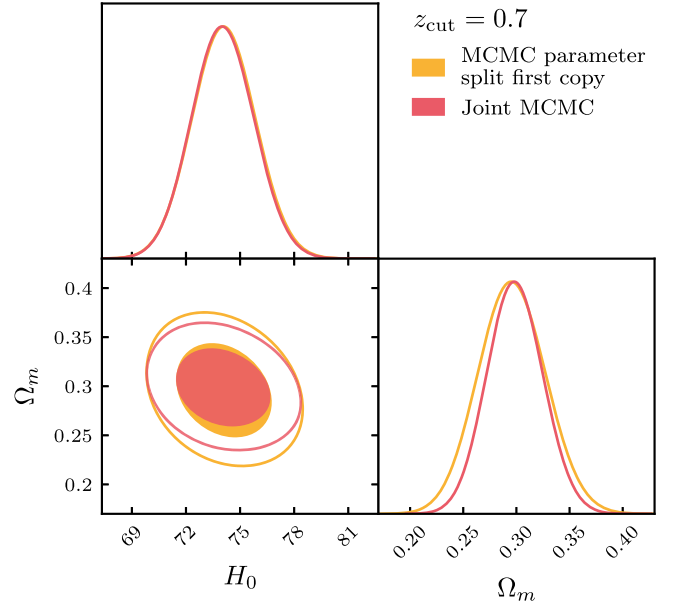


FIG. 4. The comparison of the MCMC joint marginalized posterior parameter distribution for the full SN data set and the low redshift end of the $z_{\text{cut}} = 0.7$ split. The filled contour corresponds to the 68% CL region while the continuous contour shows the 95% CL region.

In the $z_{\text{cut}} = 0.3$ case we observe very little difference between the two parameter estimates. This shows that, even though data correlations are relevant, for this data split, the two techniques do not strongly differ. The $z_{\text{cut}} = 0.7$ case, on the other hand, shows somewhat more differences. In this case the high z part of the SN catalog is significantly weaker than the low z part and correlated data modes become more relevant making the two techniques more different. As a result we see a shift in the MAP estimate and a small decrease in their covariance. In this case, in fact, the weaker split leverages the correlation with the strongest one to achieve slightly smaller error bars and parameter estimates that are closer to the ones of the strongest split. We refer the reader to Appendix B for an in-depth discussion of these effects, explained through a simple example.

Next we consider the extent to which the GLM model works overall to describe the redshift splits that we consider. To this end, notice that the joint parameter distribution cannot be obtained with standard parameter estimation techniques in the data split case, for which we use only the GLM, while it can be straightforwardly obtained in the parameter split case. We find that in the $z_{\text{cut}} = 0.3$ case the GLM works remarkably well, as shown in Fig. 3 where we compare the GLM posterior to the MCMC posterior. Both ends of the split have comparable constraining power and contain enough SN measurements to constrain both amplitude and shape parameters. The high redshift end of the split has fewer SN

measurements and hence shows hints of a slight non-Gaussian decay of the probability tail of $\Omega_m^{(2)}$.

The $z_{\text{cut}} = 0.7$ split on the other hand is different. The high redshift part of the data set contains few SN and the amplitude/shape degeneracy is far less constrained. As we can see in Fig. 3, especially for the joint $\Omega_m^{(2)}$ and $H_0^{(2)}$ distribution the GLM contours are clearly different from the MCMC ones. In particular we see high nonlinearities in the model (i.e., a markedly “banana” shaped degeneracy) that, when marginalized over, result in significant skewness of the 1D posteriors. The $z_{\text{cut}} = 0.7$ case will then serve as a good example of how some tension estimators have built in strategies to mitigate these types of non-Gaussianities.

Finally we remark that, in contrast, the posterior of the low redshift end and the joint SN posterior, shown in Fig. 4, is very close to Gaussian. Both redshift splits, in fact, contain a large number of SN that is sufficient to shrink the measured errors so that the model does not explore its nonlinear part. For these two posteriors we expect the GLM to work well in the $z_{\text{cut}} = 0.7$ case too.

VI. DATA SPLIT CDES

In this section we discuss CDEs in case of data splits and show their application to the SN example. Specifically, in Sec. VI A we present parameter shift statistics and in Sec. VI B we discuss goodness-of-fit loss.

A. Parameter shifts

Given two data sets we can compute the difference between the parameters obtained by considering the two data sets alone: $\Delta\theta_S \equiv \theta_{1S}^p - \theta_{2S}^p$. Within the GLM this is Gaussian distributed and it can be shown that its expectation value over data realizations is zero. To form the optimal quadratic form to detect shifts in parameters,

$$Q_{\text{DM}}^S \equiv \Delta\theta_S^T [\mathcal{C}(\Delta\theta_S)]^{-1} \Delta\theta_S, \quad (25)$$

we need to compute the parameter difference covariance $\mathcal{C}(\Delta\theta_S)$. For a discussion of optimal quadratic forms see Appendix D in [20]. Within the GLM the shift covariance can be obtained starting from the covariance in data space and results in

$$\begin{aligned} \mathcal{C}(\Delta\theta_S) &= C_{p1S} + C_{p2S} - C_{p1S} C_{\Pi}^{-1} C_{p2S} - C_{p2S} C_{\Pi}^{-1} C_{p1S} \\ &\quad - C_{p1S} M_1^T \Sigma_1^{-1} \Sigma_{12} \Sigma_2^{-1} M_2 C_{p2S} \\ &\quad - C_{p2S} M_2^T \Sigma_2^{-1} \Sigma_{21} \Sigma_1^{-1} M_1 C_{p1S}. \end{aligned} \quad (26)$$

As we can see this expression agrees with [20] in the limit of uncorrelated data sets. It cannot be, however, expressed in terms of parameter space quantities only when data correlations are present. In this case the parameter shift covariance depends on both the parameter and data covariance that are connected through the model

Jacobian to account for the fact that data correlations are omitted from the single parameter estimates.

In addition to this, we can also write parameter shifts in update form, by comparing the parameters of one posterior (for simplicity 1 here) to the joint parameter determination: $\Delta\theta_S^U \equiv \theta_{1S}^p - \theta_J^p$. This is, again, Gaussian distributed with zero mean and covariance:

$$\begin{aligned} \mathcal{C}(\Delta\theta_S^U) &= C_{p1S} + C_{pJ} - C_{p1S} C_{\Pi}^{-1} C_{pJ} - C_{pJ} C_{\Pi}^{-1} C_{p1S} \\ &\quad - C_{p1S} M_1^T \Sigma_1^{-1} (\Sigma_1, \Sigma_{21})^T \Sigma_J^{-1} M_J C_{pJ} \\ &\quad - C_{pJ} M_J^T \Sigma_J^{-1} (\Sigma_1, \Sigma_{12}) \Sigma_1^{-1} M_1 C_{p1S}. \end{aligned} \quad (27)$$

This agrees with [20] in the limit of uncorrelated data sets, where we recover $\mathcal{C}(\Delta\theta_S^U) = C_{p1S} - C_{pJ}$, but becomes significantly more complicated in general due to the presence of data correlations. We denote with

$$Q_{\text{UDM}}^S \equiv (\Delta\theta_S^U)^T [\mathcal{C}(\Delta\theta_S^U)]^{-1} \Delta\theta_S^U \quad (28)$$

the optimal data split parameter shift statistics in update form. Under the GLM, both Q_{DM}^S and Q_{UDM}^S are chi-squared distributed with the number of degrees of freedom equivalent to the rank of their respective covariance matrix.

In case of uncorrelated data sets the statistical significance of parameter shifts in update form is the same as the statistical significance of the difference between $\theta_1^p - \theta_2^{\text{ML}}$ since these two quantities are related by a linear transformation. However, in the presence of data correlations this is not the case since the single parameters do not contain the information on the data correlation that is contained in the joint parameter determination. In other words, it is not possible to write the update parameter shift as a linear combination of the shift in the two single parameters. Hence, we would expect to see some differences between the two estimates, related to the presence of correlated data and parameter modes.

From the previous discussion it appears clear that using the optimal, inverse covariance weighted, CDEs for data split parameter shifts is challenging in the presence of data correlations. Their covariances cannot be written in parameter space and depend on both the posterior and data covariance. These can be related to each other by projection operations involving derivatives of the observables that are cumbersome to compute accurately. These considerations limit the applicability of these methods in practice.

The SN case is, however, simple enough that we can successfully apply these estimators within the GLM. In the remainder of this section we present the challenges in applying them to the SN case and we comment on the results.

The SN data Jacobian, M_J , is estimated numerically by linear finite differences computed around the best fit of the joint SN data set. The finite difference parameter step is computed such that it would correspond to a SN chi-square

difference of one, ensuring that the derivatives are estimated on the scale at which they are relevant and are not contaminated by numerical noise. We assume that the model is fully linear so that the joint Jacobian determines the single data split Jacobian.

All the other quantities that are needed to compute Q_{DM}^S and Q_{UDM}^S are estimated from the GLM. A numerically challenging aspect of computing Q_{DM}^S and Q_{UDM}^S is identifying directions that can contribute to parameter shifts and those that do not. The latter parameter combinations can be either prior constrained or fully correlated, as can be seen from Eqs. (26) and (27). In practice, due to numerical noise, the parameter shift covariances are never exactly zero along these directions.

In the uncorrelated case this problem is solved, at least for parameter shifts in update form, by computing the quadratic form using the Karhunen-Loeve (KL) decomposition of the covariances involved, as discussed in [20]. In this case we select the directions that are used to compute the two parameter shift estimators based on the KL decomposition of the shift covariance matrices and the parameter covariance of the most constraining of the two data sets. Once the KL decomposition is performed the spectrum of the KL eigenvalues can be examined to understand if there is a clear separation of modes with KL eigenvalues very close to zero and directions that are significantly different from zero. This strategy also avoids problems with parameters having different units since the KL modes are invariant under changes of parameter basis. This also results in a wide separation between directions that can and cannot contribute to a shift making it easier to identify and remove the latter. Once the directions that cannot contribute any shift are isolated and removed the parameter shifts and their covariance are both projected on the other directions and Q_{DM}^S and Q_{UDM}^S are computed. The number of degrees of freedom of the two tests is given by the number of KL modes that are retained. In the SN example, this number is 2, since the absolute magnitude constraint does not differ between the sets.

In Table II we show the results of the application of the data split parameter shift statistics. Notice that with the data split strategy some quantities entering in the calculation of the results cannot be obtained from MCMC sampling so we estimate the results with the GLM only.

When we consider parameter differences in update form we always pick the low redshift SN cut as the base parameters for the update since, among the two possibilities, it contains a larger number of SN measurements and is hence more Gaussian.

As we can see the statistical significance of the reported results differs for the two estimators, as expected because of non-negligible data correlations. The $z_{\text{cut}} = 0.3$ split, in particular is found to be in agreement in both cases, with slightly different statistical significance. The second SN split, at $z_{\text{cut}} = 0.7$, on the other hand, is found to have

TABLE II. Results of the application of the data split parameter shift estimator in normal, Q_{DM}^S , and update form, Q_{UDM}^S , to the SN split that we consider. The first column shows the SN split that is being considered while the second column reports the value of the Q_{DM}^S and Q_{UDM}^S parameter shifts statistics. The third column shows the number of degrees of freedom of the two statistics and the fourth column the probability to exceed the measured value of Q_{DM}^S and Q_{UDM}^S . All quantities used to compute the results in this table are obtained with the GLM.

(a) Data split GLM difference in means			
Redshift cut	Q_{DM}^S	d.o.f.	$P(Q_{\text{DM}}^S > Q_{\text{DM obs}}^S)$
$z_{\text{cut}} = 0.3$	4.94	2	8.5%
$z_{\text{cut}} = 0.7$	0.11	2	94.5%
(b) Data split GLM update difference in means			
Redshift cut	Q_{UDM}^S	d.o.f.	$P(Q_{\text{UDM}}^S > Q_{\text{UDM obs}}^S)$
$z_{\text{cut}} = 0.3$	3.52	2	17.2%
$z_{\text{cut}} = 0.7$	0.53	2	76.8%

parameters that are too close, with respect to their covariance, at 94.5% probability in normal form while in agreement in update form. Since both results are computed within the GLM, and the prior is only weakly informative, the difference between the two estimates is given by different weighting of correlated data modes and reflects the fact that, in the presence of data correlations, Q_{DM}^S and Q_{UDM}^S are not expected to give the same results.

B. Goodness-of-fit loss

In addition to shifts in parameters we can use, as a CDE, the statistics of the ratio of the joint and single likelihoods at maximum posterior, Q_{DMAP}^S [20]. In the case where we consider data split we refer to this estimator as

$$Q_{\text{DMAP}}^S \equiv 2 \ln \mathcal{L}_1(\theta_{p1}^S) + 2 \ln \mathcal{L}_2(\theta_{p2}^S) - 2 \ln \mathcal{L}_J(\theta_p^S). \quad (29)$$

This quantifies goodness-of-fit loss as it corresponds to the degradation of the performances of the model when fitting two data sets jointly vs fitting the joint data. When two data sets are considered separately the model can invest all its parameters in improving the fit to data. On the other hand, when the two data sets are joined, the parameters have to compromise between the two and the joint fit will be worse. However note that in the correlated case the two data sets are not independent so that the joint likelihood is not the product of the two independent likelihoods regardless of the parameters. Consequently Q_{DMAP}^S is not necessarily positive definite, complicating its interpretation as a goodness-of-fit loss.

Even computing the statistics of Q_{DMAP}^S for correlated data sets, within the GLM, proves extremely hard in case of

data set splits. In Appendix C we report its statistics and further elaborate on the technical difficulties in practically computing it. Overall the algebraic expressions involved in its calculation are defined, as it happens for data split parameter shifts, in terms of quantities living in both parameter and data space.

Despite their complicated nature, in Appendix C we provide the full expressions that are necessary in order to compute the exact distribution of the GoF loss statistic with the data split technique. We have, furthermore, made use of these expressions to show that the distribution would not be, necessarily, well approximated with a simple chi-square distribution.

VII. PARAMETER SPLIT CDES

In this section we follow the strategy of quantifying agreement and disagreement by duplicating model parameters. We first go through the analytic aspects of calculating the CDEs and then report the results of applying them to the SN example that we consider in this work. In Sec. VII A we present parameter shift statistics, in Sec. VII B we discuss exact Monte Carlo parameter shift statistics, while in Sec. VII C we consider goodness-of-fit loss.

A. Parameter shifts

We first consider the difference between the duplicate parameter posteriors, denoted by $\Delta\theta_C \equiv \theta_{1C}^p - \theta_{2C}^p$, in the case of parameter splits. To form the optimal estimator of the significance of the shifts, we construct the quadratic form:

$$Q_{\text{DM}}^C \equiv (\Delta\theta_C)^T [\mathcal{C}(\Delta\theta_C)]^{-1} \Delta\theta_C \quad (30)$$

using their covariance to weight shifts in different parameter space directions. In this case the covariance reads

$$\mathcal{C}(\Delta\theta_C) = \mathcal{C}_{p1C} + \mathcal{C}_{p2C} - \mathcal{C}_{p12C} - \mathcal{C}_{p21C}. \quad (31)$$

Notice that, in the uncorrelated limit $\mathcal{C}_{pijC} = \mathcal{C}_{piC} \mathcal{C}_{\Pi}^{-1} \mathcal{C}_{pjC}$ for $i, j \in [1, 2]$. Furthermore, Eq. (31) implies that in the case of parameter duplication we can express the covariance of the parameter shifts using just the results from the MCMC chains. This is not true in the case of data splits, however, where the expression of the covariance includes terms related to the data covariance which account for the correlations.

We then calculate the covariance of parameter shifts in update form using one of the two parameter copies, namely θ_{1C}^p , and the parameters inferred from the joint data set, θ_J^p . Therefore, defining $\Delta\theta_C^U = \theta_{1C}^p - \theta_J^p$, the covariance of parameter shift in update form is written as

$$\mathcal{C}(\Delta\theta_C^U) = \mathcal{C}_{p1C} - \mathcal{C}_{pJ}. \quad (32)$$

which is invariant in form to the one without correlations, which is discussed in [20].

We denote with

$$Q_{\text{UDM}}^C \equiv (\Delta\theta_C^U)^T [\mathcal{C}(\Delta\theta_C^U)]^{-1} \Delta\theta_C^U \quad (33)$$

the optimal parameter-split parameter shift statistic in update form. Notice that, since Eq. (32) is invariant in form with respect to the uncorrelated case considered in [20], we can compute Q_{UDM}^C by means of the KL decomposition to filter out modes that are not improved by the data over the prior and hence subject to sampling noise. Under the GLM, both Q_{DM}^C and Q_{UDM}^C are chi-squared distributed with number of degrees of freedom equal to the rank of their covariances.

The statistical significance of the two Q_{DM}^C and Q_{UDM}^C estimators is the same for the maximum likelihood parameters while it might differ at the maximum posterior level in case of partially informative priors. This difference stems from the fact that the update form of parameter shifts contains only one copy of the prior in the joint, whereas in the single parameter shift the prior is applied once to each data set. Therefore, θ_J^p cannot be formed from a linear combination of θ_{iC}^p . We can instead define a joint parameter estimate that is so constructed

$$\tilde{\theta}_J^p = \tilde{C}_{pJ} D_C^T \mathcal{C}_{pC}^{-1} (\theta_{1C}^p, \theta_{2C}^p)^T \quad (34)$$

with covariance $\tilde{C}_{pJ}^{-1} = \mathcal{C}_{pJ}^{-1} + \mathcal{C}_{\Pi}^{-1} = D_C^T \mathcal{C}_{pC}^{-1} D_C$, so that

$$\theta_{1C}^p - \tilde{\theta}_J^p = \tilde{C}_{pJ} D_C^T \mathcal{C}_{pC}^{-1} (\mathbb{O}, \theta_{1C}^p - \theta_{2C}^p)^T, \quad (35)$$

where the vector \mathbb{O} has length N_p . This clearly shows that the statistical significance of $\theta_{1C}^p - \tilde{\theta}_J^p$ is the same as $\theta_{1C}^p - \theta_{2C}^p$ since the two are related by a linear, invertible transformation. We can then write the update parameter difference as

$$\Delta\theta_C^U = (\theta_{1C}^p - \tilde{\theta}_J^p) + (\tilde{\theta}_J^p - \theta_J^p), \quad (36)$$

which, in the uncorrelated limit reduces to $\theta_{1C}^p - \theta_{2C}^{\text{ML}}$. This agrees with the discussion in [20] of their Eq. (47). More generally, the additional difference can be computed from \mathcal{C}_{pJ} and \mathcal{C}_{Π} and can cause $\Delta\theta_C^U$ to be larger than the difference implied by $\theta_{1C}^p - \theta_{2C}^p$ since the Gaussian priors in each copy tend to bring the posteriors closer together. Note that for flat, range bound, priors as in our SN example the two copies do not lead to a stronger joint prior so that $\tilde{\theta}_J^p = \theta_J^p$.

Furthermore the difference between Q_{DM}^C and Q_{UDM}^C becomes relevant only if there is a non-negligible shift along partially prior constrained directions since the two estimators agree in the fully data and prior constrained limits.

TABLE III. Results of the application of the parameter split parameter shift estimator in normal, Q_{DM}^{C} , and update form, $Q_{\text{UDM}}^{\text{C}}$, to the SN split that we consider. The first column shows the SN split that is being considered while the second column reports the value of the Q_{DM}^{C} and $Q_{\text{UDM}}^{\text{C}}$ parameter shifts statistics. The third column shows the number of degrees of freedom of the two statistics and the fourth column the probability to exceed the measured value of Q_{DM}^{C} and $Q_{\text{UDM}}^{\text{C}}$. All quantities used to compute the results in (a) and (b) are obtained from the GLM while in (c) and (d) they are obtained from the means and covariances estimated by an MCMC.

(a) Parameter split GLM difference in means			
Redshift cut	Q_{DM}^{C}	d.o.f.	$P(Q_{\text{DM}}^{\text{C}} > Q_{\text{DM}}^{\text{C}} \text{ obs})$
$z_{\text{cut}} = 0.3$	4.51	2	10.5%
$z_{\text{cut}} = 0.7$	0.03	2	99.0%
(b) Parameter split GLM update difference in means			
Redshift cut	$Q_{\text{UDM}}^{\text{C}}$	d.o.f.	$P(Q_{\text{UDM}}^{\text{C}} > Q_{\text{UDM}}^{\text{C}} \text{ obs})$
$z_{\text{cut}} = 0.3$	4.52	2	10.4%
$z_{\text{cut}} = 0.7$	0.03	2	99.0%
(c) Parameter split MCMC difference in means			
Redshift cut	Q_{DM}^{C}	d.o.f.	$P(Q_{\text{DM}}^{\text{C}} > Q_{\text{DM}}^{\text{C}} \text{ obs})$
$z_{\text{cut}} = 0.3$	4.63	2	9.9%
$z_{\text{cut}} = 0.7$	0.37	2	83.2%
(d) Parameter split MCMC update difference in means			
Redshift cut	$Q_{\text{UDM}}^{\text{C}}$	d.o.f.	$P(Q_{\text{UDM}}^{\text{C}} > Q_{\text{UDM}}^{\text{C}} \text{ obs})$
$z_{\text{cut}} = 0.3$	4.75	2	9.3%
$z_{\text{cut}} = 0.7$	0.06	2	96.9%

It is clear at this point that making use of the parameter split methodology provides some advantages compared to the data splitting method. Equations (31) and (32) for the covariances for the parameter split statistics should be compared with Eqs. (26) and (27) for data split statistics. Crucially the former can be simply calculated from parameter covariances whereas the latter require manipulations of the data covariance. We can also therefore check the GLM results using parameter covariances taken from the MCMC chain when evaluating Eqs. (30) and (33).

We now discuss the results obtained from applying the parameter shift estimators to the SN data considered in this work. We note here that when we consider parameter differences in update form we always use the low-redshift data part to compare with the joint data set since, compared with the other data subset, it contains a larger number of SNe and is therefore more Gaussian.

The summary of our results is presented in Table III. We show in the Table the results obtained by doing a full GLM calculation of all covariances and parameter values and the results obtained using the parameter mean and covariances

from the MCMC sampling. As we can see in Table III from degree of freedom counting, the number of effective data constrained parameters in this case is the same as in the data-split applications in Sec. VI (summarized in Table II), as expected.

We can further see that, in the $z_{\text{cut}} = 0.3$ case, the parameter shift estimates in both standard form and update form agree very well within the GLM since the prior is only very weakly informative. The difference in mean result also qualitatively agrees with the result of the data split calculation in Sec. VI.

The results obtained from the MCMC are slightly different from the GLM one because the parameter centers and covariances are computed from the samples and are influenced by slight non-Gaussianities in the distribution while the GLM parameters are obtained within the linear model.

In the $z_{\text{cut}} = 0.7$ case the two GLM results agree, as they should given the weak priors, but point toward parameters that are too close to each other in units of their covariances. Even though the high redshift part of this split is partially prior constrained in the tails of the posterior, the means and covariances are not substantially influenced by the prior.

As we can further see, the MCMC results, in the $z_{\text{cut}} = 0.7$ case, are significantly different reflecting the fact that non-Gaussianities are more relevant in this case. In particular the Q_{DM}^{C} MCMC result is sensibly lower in statistical significance. Of the various estimators for the parameter means and covariances entering into Q 's, this is the only one that utilizes those of the high redshift part directly, rather than evaluated at a position that is influenced by the stronger low redshift data. In Fig. 3, we see that the slowly decaying tails of $H_0^{(2)}$, $\Omega_m^{(2)}$ increase the MCMC covariance, separate the means, and miss the fact that the posterior peaks are anomalously close given their local curvatures. In update form, the impact of the long tails is mitigated since it focuses on the peak region that is consistent with both data sets. We shall see in Sec. VII B, this means that the $Q_{\text{UDM}}^{\text{C}}$ MCMC results are more accurate than the purely GLM ones even in such a non-Gaussian case.

We also notice that the overall results in this case are qualitatively different from the data split ones that are less statistically significant. This might happen because one part of the split is much weaker than the other and the parameter split estimates in this case are more heavily influenced by the strongest data set, as discussed in Sec. V. When the data split estimates and parameter split estimates differ they might point toward a problem in the covariance matrix rather than the parameter mean. This effect is compatible with our goodness-of-fit results, in Sec. V, which showed that the high redshift part of the $z_{\text{cut}} = 0.7$ split leans toward a fit which is too good, at about the same statistical significance.

While our general rule is to compute the update parameter difference by updating the stronger with the

weaker data set, for the $z_{\text{cut}} = 0.3$ case their strengths are nearly equal. We have checked that reversing the ordering to update the high redshift data with the low redshift data does not change the statistical significance appreciably in this case as expected.

The last aspect that we can quantify is the error in the assessment of statistical significance that we could have made if we were to neglect the correlation between the two data sets.

For the $z_{\text{cut}} = 0.3$ split both Q_{DM} and Q_{UDM} would largely misestimate statistical significance resulting in a probability to exceed of 23% and 15% respectively. Notice that the update form mitigates the error since all correlations are still accounted for in the joint estimate. In the $z_{\text{cut}} = 0.7$ case, on the other hand, the two estimators neglecting correlations would estimate a probability to exceed of 99% and 56% which is again largely wrong.

B. Monte Carlo exact parameter shifts

Having an MCMC parameter estimation in the case of parameter duplication presents us with the additional possibility of computing parameter shifts as a Monte Carlo integral, as we discuss in this section.

We first consider the joint posterior probability density function of the two parameter copies $P(\theta_{1C}^p, \theta_{2C}^p)$. We can then calculate the distribution of their difference, denoted by $\Delta\theta_C \equiv \theta_{1C}^p - \theta_{2C}^p$, as the N -dimensional convolution integral:

$$P(\Delta\theta_C) = \int_{V_p} P(\theta_{1C}^p, \theta_{1C}^p - \Delta\theta_C) d\theta_{1C}^p, \quad (37)$$

over the whole parameter space volume V_p . Note that this equation is general and describes the probability to observe a parameter shift $\Delta\theta_C$ without assuming the parameters to be independent. In the limit of uncorrelated data sets the joint probability distribution in the above expression reduces to $P(\theta_{1C}^p, \theta_{2C}^p) = P_1(\theta_{1C}^p)P_2(\theta_{2C}^p)$.

To compute the statistical significance of a shift in parameters we then evaluate the integral:

$$S = \int_{P(\Delta\theta_C) > P(0)} P(\Delta\theta_C) d\Delta\theta_C, \quad (38)$$

where the volume of integration is defined as the region of parameter space where the probability to get a shift $\Delta\theta_C$ is above the isocontour of no shift, $\Delta\theta_C = 0$.

To form the MCMC chain of parameter differences in the case of correlated data sets we can take, sample by sample, the difference between the first and second copy of the parameters, without changing the weights of the samples. The result would be the MCMC estimate of the convolution integral in Eq. (37). Since the parameter duplication chain is run to convergence in the full $2N$ dimensional space the parameter difference chain is appropriately sampled.

Once we have the samples from the parameter difference probability we can compute the integral in Eq. (38) with a mixture of kernel density estimation (KDE) and MCMC techniques.

The probability of a difference in parameter, for every sample in the difference chain, is estimated through KDE with a Gaussian smoothing kernel that uses the scaled parameter difference covariance. The smoothing kernel is explicitly given by

$$K(\Delta\theta_{1C}, \Delta\theta_{2C}) = \mathcal{N}(\Delta\theta_{1C}; \Delta\theta_{2C}, \Lambda\mathcal{C}(\Delta\theta_C)\Lambda), \quad (39)$$

where $\Lambda = \text{diag}(\sqrt{\lambda})$ is a scaling matrix with λ as the smoothing scaling parameter.

We fix this parameter using Silverman's rule of thumb [43] to

$$\sqrt{\lambda} = \left(\frac{4}{n_s(N+2)} \right)^{\frac{1}{N+4}}, \quad (40)$$

where n_s is the number of samples in the chain and N is the number of parameters.

For a given MCMC sample j the KDE probability of a shift is given by

$$P(\Delta\theta_{jC}) = \frac{1}{\sum_{i=1}^{n_s} w_i} \sum_{i=1}^{n_s} w_i K(\Delta\theta_{jC}, \Delta\theta_{iC}), \quad (41)$$

where w_i denotes the weights of the samples and given that the smoothing kernel is normalized. Equation (41) is also computed for the zero shift so that the MCMC estimate of the integral in Eq. (38) is given by the number of samples that have a KDE probability of shift above the KDE probability of zero over the total number of samples.

This approach has several advantages. First, the combination of MCMC and KDE makes the estimate weakly sensitive to the choice of the smoothing kernel. The amount of over/under smoothing that the kernel might be doing is balanced by the fact that that would also happen for the zero shift estimate and would drop in the difference. In other words we never just use directly the probability of a zero shift, as obtained from Eq. (41), which would largely depend on the smoothing kernel in general, but rather compute how many samples from the distribution are above that probability. The second advantage is that this parameter shift estimate is now completely accounting for all possible non-Gaussianities in the parameter posterior.

The challenge in using this estimate is that, for statistically significant results, the estimate is likely to be noisy due to the fact that the MCMC chain would have very few samples in the tail of the distribution.

This sampling error can, however, be estimated in two ways. The first is given by a shot noise estimate, by taking the square root of the number of MCMC samples in the smallest probability tail to account for both tensions and

TABLE IV. Results of the application of Monte Carlo exact parameter shift statistics, with the parameter split methodology, to the SN split that we consider. The first column shows the redshift of the data split that we consider. The second column shows the result of Eq. (38), measuring the significance of tension or confirmation using exact MCMC techniques. The reported uncertainty is an estimate of the sampling error on a given quantity.

Redshift cut	S
$z_{\text{cut}} = 0.3$	$9.9 \pm 0.4\%$
$z_{\text{cut}} = 0.7$	$96.4 \pm 0.9\%$

confirmation results. The second is estimated as the variance of the result across multiple MCMC chains of the same distribution. In this case we have n_{chains} chains and we compute the shift probability for each of them and then estimate the error as the ratio of the variance across chains, weighted by the number of chains, $\sigma_S^2 = \text{var}(S)_{\text{chains}}/n_{\text{chains}}$, since the fiducial result uses all of them and given that different chains are independent. The two error estimates are usually in good agreement for well converged chains.

In Table IV we show the results of the MCMC calculation applied to the considered SN splits. Since the SN parameter space is only two dimensional we can also show, in Fig. 5, the posterior distribution of the difference in parameters and use that to check the reliability of these estimates.

As we can see both results reported in Table IV match very well the posterior distribution even though the significance of the shift is not directly computed from that posterior estimate. We highlight that this graphical test, which is impossible in higher dimensions, still depends on

the KDE smoothing that, in this case, is set to be the optimal one as described in [41].

As we can further see, the $z_{\text{cut}} = 0.3$ results match the GLM result for Q_{DM}^C and Q_{UDM}^C in Table III. This is expected since we have shown that the parameter posterior, in this case, is very close to Gaussian, as also highlighted by the difference posterior in Fig. 5.

In the case of the $z_{\text{cut}} = 0.7$ split the MCMC result agrees, within sampling errors, with the parameter update result. This case is, in fact, the most non-Gaussian that we consider, as can also be seen in Fig. 5, so Q_{DM}^C is expected to misestimate statistical significance.

We can also see from Table IV that the estimated sampling errors in the $z_{\text{cut}} = 0.7$ case are higher than the ones of the $z_{\text{cut}} = 0.3$ case. Both chains were run to comparable convergence but the former result is higher in statistical significance. This means that the chain contains less sample in the tail, hence increasing the error estimate. A smaller error could be achieved by running the second chain longer, at the expense of possibly significant computational resources.

These results also show that mitigation of non-Gaussianities by parameter update statistics computed from the MCMC samples, as in Table III(d) is working as expected and the two results are compatible within sampling errors.

In any case, when non-Gaussianities are suspected to be relevant, and there is reason to believe that their mitigation from the parameter shift in update form is not enough, the results can be checked with the MCMC techniques that we have just shown. We, however, highlight that reaching an acceptable noise level in the MCMC estimate, for statistically significant results, requires very long chains to accurately sample the tails of the distribution.




FIG. 5. The joint posterior of parameter split parameter differences for the two SN split that we consider. The filled contour corresponds to the 68% CL region while the continuous contour shows the 95% CL region. The dashed lines represent the position of zero shift while the dashed contour shows the probability level that intersects zero, as reported in Table IV.

C. Goodness-of-fit loss

The last CDE that we discuss is goodness-of-fit loss with the parameter split approach. In contrast to the data split case, when considering parameter splits Q_{DMAP}^C becomes easy to compute, as we discuss below.

At maximum likelihood level the statistics of goodness-of-fit loss is chi-squared distributed as a consequence of the fact that the parameter copies decompose the joint parameter estimate, as we show in Appendix D.

At the posterior level the goodness-of-fit loss statistics is defined as

$$Q_{\text{DMAP}}^C \equiv 2 \ln \mathcal{L}_J(\theta_C^p) - 2 \ln \mathcal{L}_J(\theta_J^p). \quad (42)$$

In Appendix D we discuss in detail its exact distribution as a linear combination of chi-squared variables. In practice the distribution of Q_{DMAP}^C can be approximated by that of a single chi-squared distribution, matching the mean of the exact distribution, with degrees of freedom:

$$\langle Q_{\text{DMAP}}^C \rangle = N_{\text{eff}}^C - N_{\text{eff}}^J - \text{tr}[\mathcal{C}_{\text{IC}^1 \mathcal{C}_{pC}}^{-1} (\mathbb{I}_{2N} - \mathcal{C}_{\text{IC}^1 \mathcal{C}_{pC}}^{-1} \mathbb{J}_{2N})], \quad (43)$$

where we have defined $\mathbb{J}_{2N} \equiv D_C D_C^T - \mathbb{I}_{2N}$.

As we can see, the statistics of Q_{DMAP}^C can be easily computed from the posterior MCMC samples. In the uncorrelated case it also reduces to the Q_{DMAP} statistics discussed in [20]. Compared to the uncorrelated case, we notice that in the correlated case there is an extra term that is present in the mean of the exact Q_{DMAP}^C distribution, as shown in Eq. (43), in addition to the difference in the number of effective parameters. Notice that this term vanishes for fully data or prior constrained directions. Its appearance is associated with the mismatch of assuming the data is drawn from a single parameter and prior but analyzed with split parameters and independent priors.

The results of the goodness-of-fit loss estimator for our SN analysis are shown in Table V.

As we can see, for both SN split, results are in very good agreement with both the results for parameter shifts in update form in the previous sections and the MCMC exact calculations.

In Appendix D we show that the exact distribution, in the SN case, is indeed very well approximated with a chi-square distribution and that results are largely unaltered.

We conclude the section discussing the misestimate of statistical significance that would be made if correlations between data sets were neglected. In both cases this would lead to large differences in the results with a probability to exceed the value of Q_{DMAP}^C at the 80% and 25% level for the $z_{\text{cut}} = 0.3$ and $z_{\text{cut}} = 0.7$ splits respectively. In this example the correlation would be playing a crucial role in identifying a statistically significant discrepancy that would not be identified if correlations were not properly accounted for.

TABLE V. Results of the goodness-of-fit loss statistic, using parameter duplication, to the SN split considered in this work. The first column shows the redshift of the data spit that we consider. In the second column we present the results for Q_{DMAP}^C . The third column shows the number of degrees of freedom of the Q_{DMAP}^C statistic, while the fourth column shows the probability to exceed the measured value of Q_{DMAP}^C .

Redshift cut	Q_{DMAP}^C	d.o.f.	$P(Q_{\text{DMAP}}^C > Q_{\text{DMAP}}^C \text{ obs})$
$z_{\text{cut}} = 0.3$	4.55	1.94	90.24%
$z_{\text{cut}} = 0.7$	0.03	1.58	3.7%

VIII. CONCLUSIONS

We have introduced, thoroughly discussed, and illustrated with supernovae data, a set of estimators of agreement and disagreement between cosmological data sets in the presence of non-negligible data correlations.

In particular we have explored two different approaches. The first corresponds to considering separately different correlated data sets and building estimators based on the separate inference of both, while keeping track of data correlations in assessing agreement or disagreement. We called this a data split approach. A complementary approach, which we refer to as parameter split, consists in splitting the parameter space, duplicating all relevant parameters, always analyzing the joint data set.

Both strategies are equivalent in the limit of vanishing data correlations but differ otherwise, as we have thoroughly explored. Namely, we have studied and characterized the distribution of parameter shifts estimators and goodness-of-fit loss estimators with both strategies, discussing differences, their strengths and weaknesses.

We have found that in practical applications the parameter split strategy is easier to implement since it allows us to compute the statistical significance of both tensions and excess confirmation from the output of standard parameter estimation pipelines.

The parameter split strategy also provides a suite of estimators that have complementary properties. This includes a parameter shift estimator in update form that is mitigated against possible non-Gaussianities of the parameter distributions, while maintaining the ease of application of a Gaussian estimator. This can be complemented, as we have shown, by a fully MCMC estimator that quantifies parameter shifts regardless of the Gaussianity of the parameter distribution, at the expense of computational power due to the necessity of sampling well the tail of different distributions. Goodness-of-fit loss with parameter duplicates further provides a check that the reported results are reliable in a completely different way.

In cases where the parameter posterior is Gaussian and contains parameter space directions that are either fully constrained by the data or the prior the three estimators are expected to give the same assessment of statistical

significance, providing an essential cross-check of the validity of these assumptions. When this is true, the different estimators are also optimal having minimum variance among all the possible estimators that one can define.

We have found that combining these estimators in the following recipe works well to assess the compatibility of two different correlated data-sets. First parameter shifts in update form, Eq. (33), can be used to single out the parameter space directions that contribute the most to the result and aid its physical interpretation. This first estimate of statistical significance can be checked with the MCMC estimator of Eq. (38) to provide a safeguard against significant non-Gaussianities in the posterior. In case of a Gaussian posterior the results can be confirmed with the estimator of goodness-of-fit loss, Eq. (42).

We have applied, following our discussion, all estimators to the Pantheon SN data set split at two cosmologically relevant redshifts, $z_{\text{cut}} = 0.3$ and $z_{\text{cut}} = 0.7$, roughly corresponding to the times of DE-DM equality and the redshift at which cosmic acceleration begins.

We have shown that different data split estimators are not expected to recover the same results in the presence of data correlations, even when the Gaussian approximation for the parameter posterior works well. On the other hand we have found that the parameter split estimators all recover results that are in good agreement, as it is expected, when model parameters are either fully constrained by the data or the prior, as in the SN cases that we consider.

The parameter split estimators report that the two ends of the SN catalog, split at $z_{\text{cut}} = 0.3$, agree well and show no statistically significant indication of tensions nor excess confirmation. On the other hand the two parts of the high redshift split, at $z_{\text{cut}} = 0.7$, report excess confirmation at about 96% probability. As we have discussed, this could be related either to the covariance of the low-high redshift SN being misestimated, or errors reported too conservatively. The latter explanation seems consistent with goodness-of-fit results indicating that the high redshift end of the $z_{\text{cut}} = 0.7$ SN split seems too good of a fit to the Λ CDM model, at about the 94% confidence level.

The SN constraints on the shape of the distance-redshift relation are one of the reasons why late times resolutions of the Hubble constant tensions are not viable [44–46] and it is hence important to understand the source of the excess goodness of fit in the high redshift part of the Hubble diagram that we report finding.

As the accuracy and complexity of different cosmological measurements grows and in preparation for the analysis of the current and future surveys it is important to solve remaining outstanding issues. In particular the impact of non-Gaussianities on the behavior of different estimators needs to be understood and fully taken into account. In addition our work opens the possibility of performing extensive tests of internal consistency of a single data set by

splitting it into different parts that would naturally be correlated. This raises the issue of look-elsewhere corrections for multiple tests being performed on the same data that needs to be fully quantified. Nonetheless, this work provides important building blocks in this construction by providing estimators of agreement and disagreement between correlated cosmological data sets and quantifiable tests of non-Gaussianity in parameter posteriors.

ACKNOWLEDGMENTS

We thank Dillon Brout, Rick Kessler, and Dan Scolnic for useful discussions. M. R. was supported in part by NASA ATP Grant No. NNH17ZDA001N, and by funds provided by the Center for Particle Cosmology. W. H. and M. R. were supported by U.S. Department of Energy Contract No. DE-FG02-13ER41958 and the Simons Foundation. Computing resources were provided by the University of Chicago Research Computing Center through the Kavli Institute for Cosmological Physics at the University of Chicago.

APPENDIX A: CANONICAL CORRELATIONS

The strength of the correlation between two data sets corresponds to the magnitude of the cross covariance Σ_{12} block in suitable units of Σ_1 and Σ_2 , which in general have dimensions $d_1 \times d_2$, $d_1 \times d_1$ and $d_2 \times d_2$ respectively. In one dimension ($d_1 = d_2 = 1$) this would be quantified by the Pearson correlation coefficient:

$$\rho_{12} = \frac{\sigma_{12}^2}{\sigma_1 \sigma_2}, \quad (\text{A1})$$

where we used the notation $\Sigma_1 = \sigma_1^2$, $\Sigma_2 = \sigma_2^2$, $\Sigma_{12} = \sigma_{12}^2$ for one-dimensional quantities. The quantity ρ_{12} is bounded to be in $[-1, 1]$ with zero meaning absence of correlation, while 1 and -1 indicate complete correlation and anti-correlation respectively.

In multiple dimensions the equivalent procedure is often referred to as canonical correlation analysis (CCA) [42] that we now discuss.

The idea is to look for the optimal linear combination of the two data vectors that maximize the correlation between them. If we take two vectors in the space of data, x_1 and x_2 , we can build the quantity

$$\rho_{12} = \frac{x_1^T \Sigma_{12} x_2}{\sqrt{x_1^T \Sigma_1 x_1} \sqrt{x_2^T \Sigma_2 x_2}}, \quad (\text{A2})$$

which we seek to maximize with respect to x_1 and x_2 .

It can be shown that x_1 , x_2 and ρ_{12} are the solutions of the two eigenvalue problems:

$$\begin{aligned}\Sigma_1^{-1}\Sigma_{12}\Sigma_2^{-1}\Sigma_{21}x_1 &= \rho_{12}^2x_1, \\ \Sigma_2^{-1}\Sigma_{21}\Sigma_1^{-1}\Sigma_{12}x_2 &= \rho_{12}^2x_2.\end{aligned}\quad (\text{A3})$$

We refer to the solutions of Eq. (A3) as the CCA modes and $-1 < \rho_{12} < 1$ as their correlation coefficient. Notice that the sign of ρ_{12} is arbitrary and corresponds to a convention for the relative sign of x_1 and x_2 . If given data modes (x_1, x_2) are positively correlated, then $(-x_1, x_2)$ are negatively correlated. When considered this way, as a pair of data vectors spanning the joint space, the CCA modes are equivalent to the Karhunen-Loeve (KL) modes of Σ_J and $\text{diag}(\Sigma_1, \Sigma_2)$, which clarifies their implications for parameter estimation. These KL modes, e^K , are solutions to the generalized eigenvalue problem,

$$\begin{pmatrix} \Sigma_1 & \Sigma_{12} \\ \Sigma_{12} & \Sigma_2 \end{pmatrix} e^K = \lambda_K \begin{pmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{pmatrix} e^K, \quad (\text{A4})$$

where K indexes the modes. The KL modes form a complete and statistically independent basis for the joint data in that their amplitudes have no covariance for both Σ_J and $\text{diag}(\Sigma_1, \Sigma_2)$. The KL eigenvalue therefore is the ratio between the variances of these mode amplitudes with and without the Σ_{12} correlations. Unlike an ordinary eigenvector decomposition, these modes are not orthogonal in the Euclidean sense, but rather orthogonal under the metrics provided by the covariance matrices.

The relationship to the CCA modes is that for each one of the $\min(d_1, d_2)$ unique $|\rho_{12}|$ eigenvalues with CCA solutions x_1 and x_2 there are two KL modes with $e^K = (x_1, x_2)$ and $(-x_1, x_2)$, and $\lambda_K = 1 \pm |\rho_{12}|$. The remaining $|d_1 - d_2|$ modes are uncorrelated, with $\lambda_K = 1$, and have support only across the larger of the two data sets.

The impact of correlations on parameters is bounded by the largest correlation or the pair of modes for which $\lambda_K = 1 \pm \max(|\rho_{12}|)$. This maximal effect occurs if the parameter of interest is exactly one of these two KL mode amplitudes, where the parameter variance would be misestimated by this factor, with the sign distinguishing an under and over estimate respectively. Notice that since the maximum possible correlation is itself bounded by $|\rho_{12}| \leq 1$ then $0 \leq \lambda_K \leq 2$, meaning that the ratio of errors is likewise strictly bounded between 0 and $\sqrt{2}$. Neglecting correlations can make parameter errors at most infinitely overestimated or underestimated by $\sqrt{2}$. The former occurs when such correlations allow a zero noise measurement of a parameter. The latter occurs when correlations make the information in the two data sets completely redundant. We illustrate these ideas with a simple example in Appendix B.

For a more general parameter of interest which is not exactly a KL mode itself, we can compute the impact of correlations by summing the parameter information in each mode independently. Given the sensitivity per KL mode to a parameter vector θ as

$$c^K \equiv \frac{\partial m}{\partial \theta} e^K = M^T e^K, \quad (\text{A5})$$

the Fisher matrix is

$$\mathcal{F}_{ij} \equiv \sum_K \frac{1}{\lambda_K} c_i^K c_j^K, \quad (\text{A6})$$

whereas falsely neglecting the correlations would give

$$\tilde{\mathcal{F}}_{ij} \equiv \sum_K c_i^K c_j^K, \quad (\text{A7})$$

where the index K runs over the KL modes. Note that if the KL modes with low correlation dominate the information on a given parameter, the impact of correlations on parameter variances decreases from the extreme of $1 \pm \max(|\rho_{12}|)$. The Fisher estimate of the parameter covariance is then $\mathcal{C} = \mathcal{F}^{-1}$.

In the remainder of this section we comment on the correlated data modes of the two SN splits that we consider.

In Fig. 6 we show the spectra of the correlation coefficients for the two splits. As we can see the spectra are similar and contain about ten correlated data modes, while the remaining ones are nearly uncorrelated. In both cases, the first mode is completely correlated with $|\rho_{12}| = 1$, and corresponds to the redundant measurement of the SN absolute calibration that is present in both splits. The second mode has a correlation coefficient of $|\rho_{12}| = 0.52$ for the $z_{\text{cut}} = 0.3$ split and $|\rho_{12}| = 0.51$ for the $z_{\text{cut}} = 0.7$ one and corresponds to the first genuine SN data correlation. In the same figure we also show the spectrum




FIG. 6. The spectrum of the absolute value of the canonical correlation coefficients for the two SN data split considered. The continuous and dashed lines show the results for the unbinned and binned SN data respectively.

of canonical correlations for the redshift-binned SN measurements. As we can see there are fewer modes, corresponding to fewer data points, but the correlation coefficients are qualitatively unaltered showing that the correlations that we see are not due to noise in the covariance matrix but rather come from shared systematic correlations.

In this regard, the CCA analysis shows that correlations must be kept since there are several linear combinations of the data whose errors would be severely misestimated otherwise. On the other hand this does not necessarily mean that cosmological parameters will be equally affected and we now want to quantify the impact that these data correlations have on the determination of cosmological parameters.

In the Λ CDM model SN magnitudes at different redshifts depend on cosmology through two parameters, Ω_m and H_0 . The effect of variations of these parameters on the SN magnitude is shown in Fig. 7. As we can see these describe variations in the amplitude and shape of $m - \mathfrak{M}_{\text{fid}}$ as a function of redshift. In addition to these parameters we have another parameter that describes the absolute magnitude of the SN with its corresponding data constraint which is required to make inferences from the distance modulus $m - \mathfrak{M}$.

The inferred errors on the parameters $\Omega_m, H_0, \mathfrak{M}$ are therefore correlated regardless of whether the SN magnitude measurements are themselves correlated. In particular,




FIG. 7. The Jacobian of the SN magnitude with respect to different cosmological parameters. Parameters are scaled to have unit variance so that the Jacobian is directly in magnitude units. Different lines correspond to different parameter variations as shown in the legend.

the absolute magnitude calibration is degenerate with H_0 and both of these parameters are mildly degenerate with Ω_m , as can also be seen in Fig. 1.

To make the impact of correlations on parameters clear, and mostly unaltered by marginalization, we now perform a principal component (PC) analysis on the SN covariance. Note that we compute the principal components of the covariance after transforming H_0 to $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$ to have a dimensionless set of parameters with comparable scalings.

The parameter coordinates of the SN principal components are reported in Table VI. As we can see the first PC is mostly influenced by changes in Ω_m . The second PC is not cosmologically interesting but is a highly constrained direction where changes to both \mathfrak{M} and cosmological parameters make all SN brighter or dimmer. The third PC corresponds to the direction along the $\mathfrak{M} - H_0$ degeneracy where measurements of \mathfrak{M} determine H_0 .

We can now employ Eq. (A6) to consider the cumulative impact of the KL modes in the PC parameter space, after ordering them by absolute value of correlation coefficient and such that odd modes are correlated and even modes are anticorrelated. Since we are using the PCs, the entries of the Fisher matrix provide a good representation of the inverse PC variance.

In Fig. 8 we can see the Fisher matrix entries as we sum different KL modes, for the two cases where we keep and neglect the data correlations. Both curves are shown in units of the full Fisher matrix including correlations.

As we can see, for the first PC and for both redshift splits, the dominant contribution to the difference between the two results comes from the third KL mode. The third mode is correlated with a correlation coefficient of about $\rho_{12} = 0.5$ and this corresponds, as in Eq. (A6), to an underestimate of the variance by a factor of 1.5. This fractional underestimate is diluted somewhat by the sum of higher KL modes which have smaller ρ_{12} , especially for the $z_{\text{cut}} = 0.7$ case where the lower redshift side can measure shape changes from Ω_m on its own. The net result is that the variance of PC1 is underestimated by a factor of 1.4 for $z_{\text{cut}} = 0.3$ and 1.2 for $z_{\text{cut}} = 0.7$. This is also consistent with the underestimate of the variance of Ω_m shown in Fig. 1 when we neglect correlations between the two data set splits.

TABLE VI. The cosmological parameter combinations defining the principal components (PC) of the SN covariance.

Principal component	$\Delta\Omega_m$	Δh	$\Delta\mathfrak{M}$
PC 1	0.99	-0.13	0.06
PC 2	-0.14	-0.94	0.32
PC 3	-0.02	0.33	0.94




FIG. 8. The amplitude of the Fisher matrix elements corresponding to the three parameter principal components in units of their full data set values. The parameter combinations defining the three principal components are shown in Table VI. All curves are obtained by summing the first i KL modes ordered by absolute value of correlation coefficient and such that odd modes are correlated and even modes are anticorrelated. The two colors correspond to the cases where we include or exclude the correlation between SN at $z > z_{\text{cut}}$ and $z < z_{\text{cut}}$, as shown in the legend.

The dominant contribution to the difference in results for the second PC comes from the fourth mode which is anticorrelated with a correlation coefficient of about $\rho_{12} = -0.5$ corresponding to an overestimate of the PC variance of about 1.5. This well constrained mode hardly influences cosmological results, which marginalize over \mathfrak{M} .

The third PC gets most of its contribution from the first correlated mode. This is due to the shared absolute magnitude calibration which gives $\rho_{12} = 1$ and hence a factor of 2 underestimate. Since the third PC involves the $H_0 - \mathfrak{M}$ degeneracy, it also explains the underestimate of the variance of H_0 shown in Fig. 1.

We can now look at the most relevant correlated data modes that are shown in Fig. 9. These are obtained from the full SN data set, so that the discreteness of the data is evident, but exhibits far smoother trends in redshift than the scale of the individual redshifts themselves. The data split, in fact, highlights coherent effects across the redshift sample.

Comparing Figs. 9 and 7 we can see that the third KL mode is qualitatively very similar to the effect of changing Ω_m , possibly with a small amplitude component, and so it is not surprising that the data correlation corresponding to this mode reflects almost entirely on the parameter variance. The fourth mode differs from the third one by a sign flip in the high redshift part and looks less like a smooth change in $\mathbf{m} - \mathfrak{M}_{\text{fid}}$, especially for $z_{\text{cut}} = 0.7$.

Overall we see that the CCA decomposition provides a powerful tool for quantifying and understanding the impact of data correlations on parameter estimation. When the high ranked modes resemble the desired parameters themselves, this impact is maximal.




FIG. 9. The third and fourth most correlated modes for the two SN splits that we consider. Different colors represent different modes, as shown in the legend. The vertical line shows the redshift where we split the SN data.

APPENDIX B: PEDAGOGICAL EXAMPLE OF DATA VS PARAMETER SPLITS

In this Appendix we illustrate the difference between splitting data and splitting parameters in the case of correlated data sets with a pedagogical example that can be fully treated analytically.

Take a d dimensional data vector x_i split as $(x_1 \dots x_{d_1})$ and $(x_{d_1+1} \dots x_{d_1+d_2})$ with d_1 and d_2 elements respectively. Let $\langle x_i \rangle = 0$ and define a data covariance matrix to have all uncorrelated elements but for the last (x_{d_1}) and first (x_{d_1+1}) data points of the two sets which have a correlation coefficient R :

$$\Sigma = \begin{pmatrix} \mathbb{I} & \mathbb{O} & \mathbb{O} & \mathbb{O} \\ \mathbb{O} & 1 & R & \mathbb{O} \\ \mathbb{O} & R & 1 & \mathbb{O} \\ \mathbb{O} & \mathbb{O} & \mathbb{O} & \mathbb{I} \end{pmatrix}, \quad (\text{B1})$$

where, hereafter, \mathbb{I} and \mathbb{O} denote the identity and zero matrices of appropriate dimensions respectively. In the language of CCA, there is only one pair of correlated modes proportional to $(\mathbb{O}, 1, \pm 1, \mathbb{O})$ with $\rho_{12} = \pm R$.

We use a simple model with a parameter that controls the mean of the data sets. The joint model Jacobian is then $M_J^T = (\mathbb{I}, 1, 1, \mathbb{I})$ while the parameter split Jacobian is given by

$$M_C^T = \begin{pmatrix} \mathbb{I} & 1 & 0 & \mathbb{O} \\ \mathbb{O} & 0 & 1 & \mathbb{I} \end{pmatrix}. \quad (\text{B2})$$

Following Eq. (6), the maximum likelihood parameter split estimates of the parameters are therefore

$$\begin{aligned} \theta_{1C} &= \frac{1}{\alpha} \left[d_2 \sum_{i=1}^{d_1} x_i + \left(\sum_{i=d_1+2}^{d_1+d_2} x_i - (d_2-1)x_{d_1+1} \right) R \right. \\ &\quad \left. - (d_2-1) \sum_{i=1}^{d_1-1} x_i R^2 \right], \\ \theta_{2C} &= \frac{1}{\alpha} \left[d_1 \sum_{i=d_1+1}^{d_1+d_2} x_i + \left(\sum_{i=1}^{d_1-1} x_i - (d_1-1)x_{d_1} \right) R \right. \\ &\quad \left. - (d_1-1) \sum_{i=d_1+2}^{d_1+d_2} x_i R^2 \right], \end{aligned} \quad (\text{B3})$$

where we have defined $\alpha \equiv d_1 d_2 - (d_1-1)(d_2-1)R^2$. The parameter covariance matrix is then given by Eq. (3):

$$C_C = \frac{1}{\alpha} \begin{pmatrix} d_2 - (d_2-1)R^2 & R \\ R & d_1 - (d_1-1)R^2 \end{pmatrix}. \quad (\text{B4})$$

These results can now be compared to the data split estimators of the parameters:

$$\begin{aligned} \theta_{1S} &= \frac{1}{d_1} \sum_{i=1}^{d_1} x_i, \\ \theta_{2S} &= \frac{1}{d_2} \sum_{i=d_1+1}^{d_1+d_2} x_i, \end{aligned} \quad (\text{B5})$$

with covariance matrix given by

$$C_S = \frac{1}{d_1 d_2} \begin{pmatrix} d_2 & R \\ R & d_1 \end{pmatrix}. \quad (\text{B6})$$

As we can see, all results coincide in the uncorrelated limit, $R \rightarrow 0$, while generally differing when data correlations are present. In particular, in case of the parameter copies, Eq. (B3), each parameter estimate depends on the full data set, even the uncorrelated pieces of the complementary set, where the weights are proportional to the correlation coefficient R . This is because those uncorrelated data still inform the mean of the correlated data point. This example illustrates the fundamental difference between the two statistics: the data of set 2 influence the parameters of set 1 and vice versa for parameter splits but not for data splits.

Similarly we can compute the parameter estimate for the joint data set:

$$\begin{aligned} \theta_J &= \frac{1}{\alpha_J} \left[(1+R) \sum_{i=1}^{d_1+d_2} x_i - R(x_{d_1} + x_{d_1+1}) \right], \\ \alpha_J &= (1+R)(d_1 + d_2 - 2) + 2 \end{aligned} \quad (\text{B7})$$

with covariance

$$C_J = \frac{1+R}{\alpha_J}, \quad (\text{B8})$$

and explicitly verify that we can decompose the joint parameter determination as a linear combination of the two parameter copies, as in Eq. (24) while the same does not apply for the data split parameter determinations.

It is now instructive to consider the two extreme cases of fully correlated and anticorrelated data sets, corresponding respectively to $R = 1$ and $R = -1$. In both cases the data covariance, Eq. (B1), becomes singular, with one data combination fixed with zero variance. In the $R = 1$ case the difference between x_{d_1} and x_{d_1+1} has zero variance and hence $x_{d_1} = x_{d_1+1}$. In the $R = -1$ case the sum of x_{d_1} and x_{d_1+1} is fixed, so that $x_{d_1} = -x_{d_1+1}$.

The different parameter estimators, discussed above, then respond differently in these two cases, depending on how the correlated mode projects on the parameters of

the model. The data split parameters remain unaltered and do not respond in a particular way to the extreme correlation, since the presence of that correlation is ignored in the parameter fit in the first place.

From Eq. (B3), the parameter split parameters for $R = 1$ become equal:

$$\theta_{1C} = \theta_{2C} = \theta_J = \frac{1}{\alpha} \left(\sum_{i=1}^{d_1} x_i + \sum_{i=d_1+2}^{d_1+d_2} x_i \right), \quad (\text{B9})$$

with $\alpha = d_1 + d_2 - 1$, the number of independent data points. It follows that the difference of the two copy parameters is fixed to zero while their sum has variance $2/\alpha$. In this case the fully correlated data acts as a bridge so that the best parameter estimator of each is always the joint estimator that uses all of the data optimally.

In the opposite case, when $R = -1$, the two copy parameters are given by

$$\theta_{1C} = -\theta_{2C} = \frac{1}{\alpha} \left(\sum_{i=1}^{d_1} x_i - \sum_{i=d_1+2}^{d_1+d_2} x_i \right) \quad (\text{B10})$$

while the joint parameter estimate is $\theta_J = 0$. This reflects the fact that a shift in θ_J reflects a shift in the mean of all points whereas a noise fluctuation can only shift the difference between the correlated points, not their sum. The sum of the two copy parameters is then fixed to zero while their difference has variance $2/\alpha$. The former represents a parameter that can be measured free of noise in the joint case when including correlations and is a simple example of saturating the $\lambda_K = 0$ KL bound discussed in the previous section below Eq. (A4).

As we can see, in these two extreme cases, if the two data sets share some linear combination of their data that can be measured free of noise, then information is fully shared between the parameter splits in a manner that depends on the projection of this linear combination onto parameter space. The case of partial correlation is analogous but in that case the two parameters are likewise no longer fully correlated.

To further clarify the difference between the statistics for finite R , let us take the simplest example where $d_1 = 1$ and $d = d_2 + 1$. The estimator of parameter split parameter difference is then

$$\Delta\theta_C \equiv \theta_{1C} - \theta_{2C} = x_1 - \frac{1 + (d-2)R}{d-1} x_2 + \frac{R-1}{d-1} \sum_{i=3}^d x_i, \quad (\text{B11})$$

$$\text{var}(\Delta\theta_C) = \frac{(1-R)(d + (d-2)R)}{d-1},$$

while the data split estimator is

$$\Delta\theta_S \equiv \theta_{1S} - \theta_{2S} = x_1 - \frac{1}{d-1} \sum_{i=2}^d x_i,$$

$$\text{var}(\Delta\theta_S) = \frac{d-2R}{d-1}. \quad (\text{B12})$$

Even though $\Delta\theta_C \rightarrow 0$ as $R \rightarrow 1$, its variance does as well since the fluctuations in the uncorrelated data drop out of the difference. Thus for finite R , $\Delta\theta_C$ may be significantly anomalous even though its magnitude is much less than $\Delta\theta_S$. Nonetheless, for a given data realization, the two would report different statistical significance in general.

To understand the difference in significance let us illustrate this with a simple example. Suppose the anomalous aspect of the data were an extreme fluctuation in the value of x_1 itself. For a typical realization of x_2 , this would appear as an anomalous, but different, value for $\Delta\theta_C$ and $\Delta\theta_S$. However the realization of x_2 contains fluctuations from both the correlated noise and the uncorrelated noise. A rare fluctuation in the uncorrelated piece will separate the significance of the two statistics. For example if the uncorrelated piece separated x_2 from x_1 more than expected given R , it would affect the parameter split estimator more than the data split. As $R \rightarrow 1$, the former depends mainly on $x_1 - x_2$, which is controlled by the uncorrelated piece of the noise rather than the anomalous value of x_1 itself. In this sense, it is more important for the parameter split statistics that correlations are modeled accurately than it is for the data split statistics.

In this simplified case it is also easy to write the correlation between the two parameter shift estimates:

$$\text{corr}(\Delta\theta_C, \Delta\theta_S) = \sqrt{\frac{(1-R)(d + R(d-2))}{d-2R}}, \quad (\text{B13})$$

which shows that in the $R = 0$ case the two estimators are completely correlated and become gradually uncorrelated as R increases. As discussed above, the correlation of a single data point reduces the parameter split difference whereas the data split difference still fluctuates because of all of the uncorrelated data, thereby decorrelating the two estimators.

APPENDIX C: EXACT DISTRIBUTION OF GOODNESS-OF-FIT LOSS STATISTIC WITH DATA SPLIT

In this section we discuss in detail the exact distribution of the Q_{DML}^S and Q_{DMAP}^S estimators, with the data split methodology.

We first consider the ratio of maximum likelihoods of the joint data set and the two subsets. By direct calculation it can be shown that, up to constant offsets that is irrelevant to the calculation of statistical significance this ratio can be written as the quadratic form in the data:

$$\begin{aligned}
Q_{\text{DML}}^S &= X_J^T \left[(\mathbb{I}_J - \mathbb{P}_J)^T \Sigma_J^{-1} (\mathbb{I}_J - \mathbb{P}_J) - \begin{pmatrix} (\mathbb{I}_1 - \mathbb{P}_1)^T \Sigma_1^{-1} (\mathbb{I}_1 - \mathbb{P}_1) & \mathbb{O} \\ \mathbb{O} & (\mathbb{I}_2 - \mathbb{P}_2)^T \Sigma_2^{-1} (\mathbb{I}_2 - \mathbb{P}_2) \end{pmatrix} \right] X_J \\
&\equiv X_J^T A_{\text{DML}}^S X_J,
\end{aligned} \tag{C1}$$

where X_J is the full data vector, distributed according to the evidence of the joint data set, and A_{DML}^S is the matrix that defines Q_{DML}^S . The indices 1 and 2 denote the two subsets of the data after we split the joint set.

Through explicit computation, the form of the joint projector is given by

$$\mathbb{P}_J \equiv \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix}$$

with

$$\begin{aligned}
p_{11} &= M_1 C_J (M_1^T - M_2^T \Sigma_2^{-1} \Sigma_{21}) K_1^{-1}, \\
p_{12} &= M_1 C_J (M_2^T - M_1^T \Sigma_1^{-1} \Sigma_{12}) K_2^{-1}, \\
p_{21} &= M_2 C_J (M_1^T - M_2^T \Sigma_2^{-1} \Sigma_{21}) K_1^{-1}, \\
p_{22} &= M_2 C_J (M_2^T - M_1^T \Sigma_1^{-1} \Sigma_{12}) K_2^{-1},
\end{aligned}$$

where we have defined

$$\begin{aligned}
K_1 &\equiv \Sigma_1 - \Sigma_{12} \Sigma_2^{-1} \Sigma_{21}, \\
K_2 &\equiv \Sigma_2 - \Sigma_{21} \Sigma_1^{-1} \Sigma_{12}.
\end{aligned}$$

In order to calculate the distribution of Eq. (C1) we follow the procedure discussed in Appendix A of [20] and compute the eigenvalues, λ , of $A_{\text{DML}}^S \mathcal{S}_J$, where \mathcal{S}_J is the covariance for the joint distribution of the data which, for Gaussian priors, is $\mathcal{S}_J = \Sigma_J + M_J C_{\Pi} M_J^T$. This allows to decompose Q_{DML}^S in the following way:

$$Q_{\text{DML}}^S = \sum_i \lambda_i \mathcal{U}_i^2, \tag{C2}$$

where each $\mathcal{U}_i \sim \mathcal{N}_J(x_J; \mathbb{O}, \mathbb{I})$ so that Q_{DML}^S is a weighted sum of chi-squared variables.

By direct calculation we have

$$\mathcal{A}_{\text{DML}}^S \equiv A_{\text{DML}}^S \mathcal{S}_J \equiv \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{C3}$$

where

$$\begin{aligned}
A &= [\Sigma_1^{-1} M_1 C_1 - K_1^{-1} (M_1 - \Sigma_{12} \Sigma_2^{-1} M_2) C_J] M_1^T, \\
B &= -K_1^{-1} (M_1 - \Sigma_{12} \Sigma_2^{-1} M_2) C_J M_2^T \\
&\quad - (\mathbb{I}_1 - \Sigma_1^{-1} M_1 C_1 M_1^T) \Sigma_1^{-1} \Sigma_{12}, \\
C &= -K_2^{-1} (M_2 - \Sigma_{21} \Sigma_1^{-1} M_1) C_J M_1^T \\
&\quad - (\mathbb{I}_2 - \Sigma_2^{-1} M_2 C_2 M_2^T) \Sigma_2^{-1} \Sigma_{21}, \\
D &= [\Sigma_2^{-1} M_2 C_2 - K_2^{-1} (M_2 - \Sigma_{21} \Sigma_1^{-1} M_1) C_J] M_2^T.
\end{aligned}$$

An analytic solution to the above eigenvalue problem is not easily obtained, but can be obtained numerically to evaluate the exact distribution of Q_{DML}^S . We highlight that, similarly to what happens for data split parameter shifts, the calculation of the statistics involves quantities that are defined both at the parameter space and data space level.

Note that the expressions we derived above reduce to the corresponding ones in [20] in the limit of uncorrelated data sets.

Furthermore we can notice that the quadratic form defined by Q_{DML}^S is not necessarily positive definite. This is a consequence of the fact that the projector on the joint parameter space is not a subspace of the span of the single data set projector. This severely limits the possibility of approximating Q_{DML}^S with a chi-squared distribution, which is positive definite, especially for events in the confirmation tail that would be very close to $Q_{\text{DML}}^S = 0$.

In addition, the fact that Q_{DML}^S is not chi-squared distributed means that correlated data fluctuations are not optimally weighted.

We then consider the ratio of likelihoods at maximum posterior (DMAP) in the data split case. To do so we add the extra terms that transform ML into MAP so that the matrix that controls Q_{DMAP}^S is given by

$$\begin{aligned}
\mathcal{A}_{\text{DMAP}}^S &= \mathcal{A}_{\text{DML}}^S + \Sigma_J^{-1} M_J C_J C_{\Pi}^{-1} C_{pJ} M_J^T \\
&\quad - \begin{pmatrix} \tilde{M}_1^T C_{\Pi}^{-1} C_{p1} M_1^T & \tilde{B} \\ \tilde{B} (1 \leftrightarrow 2) & \tilde{M}_2^T C_{\Pi}^{-1} C_{p2} M_2^T \end{pmatrix},
\end{aligned} \tag{C4}$$

where, for compactness, we have defined $\tilde{B} \equiv \tilde{M}_1^T C_{\Pi}^{-1} C_{p1} C_1^{-1} C_{p1} (C_{\Pi}^{-1} \tilde{M}_1 \Sigma_{12} + M_2^T)$.

These results, for both Q_{DML}^S and Q_{DMAP}^S , can be used to compute the respective exact distributions. The trace of these distributions coincides with the results obtained in the uncorrelated case, but we notice that it is problematic to approximate them with simpler distributions because both of them are not positive definite.




FIG. 10. Spectrum of the exact distribution of Q_{DML}^S in panel (a) and Q_{DMAP}^S in panel (b). Different lines correspond to different SN split, as shown in the legend. Modes that have eigenvalues below 0.05 are not shown.

The fact that both Q_{DML}^S and Q_{DMAP}^S are not positive definite means that there are aspects of the data where the joint likelihood is better than the product of the separate likelihoods. This can never happen for uncorrelated data sets and is a consequence of the presence of correlated data modes. In particular, the data modes that are fit by the model separately are the ones that would contribute to the positive definiteness of the above statistics, since they can zero out different data fluctuations, whereas the correlated data modes that are left out can contribute negatively. Fitting the data jointly, however, always takes correlated modes into account, so that the contribution to the chi-square from them is considered.

We close this Appendix by discussing the application of these exact statistics to our SN example. In Fig. 10 we show the eigenvalues of the matrix in Eq. (C3) for Q_{DML}^S and the eigenvalues of the matrix in Eq. (C4) for Q_{DMAP}^S . This clearly shows that both estimators are not positive definite which means that an approximation of the exact distribution with a χ^2 distribution would not be appropriate. We can also see that the number of nonzero components largely exceeds the number of parameters as a consequence of the fact that for this data split estimator data and parameter modes are mixed. Notice also that both distributions are similar since the prior is not informative.

APPENDIX D: EXACT DISTRIBUTION OF GOODNESS-OF-FIT LOSS STATISTIC WITH PARAMETER SPLIT

In this section we discuss in detail the exact distribution of the Q_{DML}^C and Q_{DMAP}^C estimators, as presented in Sec. VII C, and how they can be approximated.

We first consider the statistics of ML ratios (DML). We focus on the distribution of the DML statistic between the joint chain and the one with the duplicated parameter space. To do so, we begin by considering the ML parameter split determination, $\theta_C^{\text{ML}} = (\theta_{1C}^{\text{ML}}, \theta_{2C}^{\text{ML}})^T$, and the joint ML parameters, θ_J^{ML} . We then use them to define the difference in joint log likelihood at the ML point as

$$Q_{\text{DML}}^C \equiv -2 \ln \mathcal{L}_J(\theta_J^{\text{ML}}) + 2 \ln \mathcal{L}_J(\theta_C^{\text{ML}}). \quad (\text{D1})$$

Note that in the limit of uncorrelated data this reduces to $Q_{\text{DML}}^C = -2 \ln \mathcal{L}_J(\theta_J^{\text{ML}}) + 2 \ln \mathcal{L}_1(\theta_1^{\text{ML}}) + 2 \ln \mathcal{L}_2(\theta_2^{\text{ML}})$ which is similar to the expressions used in [20].

In the GLM, it can be shown that, up to constants which are not important for our purpose, we get the following quadratic form in data space:

$$\begin{aligned} Q_{\text{DML}}^C &= X^T [(\mathbb{I} - \mathbb{P}_J)^T \Sigma_J^{-1} (\mathbb{I} - \mathbb{P}_J) \\ &\quad - (\mathbb{I} - \mathbb{P}_C)^T \Sigma_J^{-1} (\mathbb{I} - \mathbb{P}_C)] X \\ &\equiv X^T A_{\text{DML}}^C X, \end{aligned} \quad (\text{D2})$$

where we have used the joint projector $\mathbb{P}_J = M_J C_J M_J^T \Sigma_J^{-1}$ and the projector under parameter duplication written as $\mathbb{P}_C = M_C C_C M_C^T \Sigma_J^{-1}$. Now, we can rewrite the joint projector in the following way:

$$\mathbb{P}_J = M_C D_C (D_C^T C_C^{-1} D_C)^{-1} D_C^T M_C^T \Sigma_J^{-1}, \quad (\text{D3})$$

while the projector in the case of parameter duplication can be expressed as

$$\mathbb{P}_C = M_C (M_C^T \Sigma_J^{-1} M_C)^{-1} M_C^T \Sigma_J^{-1}. \quad (\text{D4})$$

Then, using the above expressions it is straightforward to show that the joint set of parameters is a subset of the duplicate set, since $\mathbb{P}_J \mathbb{P}_C = \mathbb{P}_C \mathbb{P}_J = \mathbb{P}_J$. Therefore, we can use theorem (5.2.5) in [47] to show that, at the ML level,

$$\begin{aligned} Q_{\text{DML}}^C &\sim \chi^2(\text{rank}(\mathbb{I} - \mathbb{P}_J) - \text{rank}(\mathbb{I} - \mathbb{P}_C)) \\ &= \chi^2(N_C - N_J), \end{aligned} \quad (\text{D5})$$

where N_C and N_J are the number of parameter duplicates and the number of joint parameters respectively. Note that, in the limit of uncorrelated data sets $N_C = N_1 + N_2$, where N_1 and N_2 are the number of relevant parameters for the first and second data sets respectively.

In contrast with the case of data split the exact statistics of the parameter split DML estimator is a chi-square, which also means that Q_{DML}^C is optimal.

The exact statistics of Q_{DML}^C can also be obtained by explicitly computing the eigenvalues of $A_{\text{DML}}^C S_J = \mathbb{P}_C^T - \mathbb{P}_J^T$, where $S_J = \Sigma_J + M_J C_{\Pi} M_J^T$. Notice that $M_J C_{\Pi} M_J^T$ represents

a prior that is fully correlated between the split parameters, whereas our parameter split analysis assumes separate priors that are uncorrelated. This is necessary since otherwise the split parameters would be expected to vary according to \mathcal{C}_Π leading to a counterfactually large expected improvement from fitting them separately. Conversely, the split parameter technique cannot employ fully correlated priors because no matter how weak such a prior is, it would force the split parameter posterior means to the same values (see Appendix B).

We can now turn to the distribution of Q_{DMAP} , with parameter copies. This can be written as

$$\begin{aligned} Q_{\text{DMAP}}^C &\equiv -2 \ln \mathcal{L}_J(\theta_J^p) + 2 \ln \mathcal{L}_J(\theta_C^p) \\ &= Q_{\text{DML}}^C + X_J^T [\tilde{M}_J^T \mathcal{C}_\Pi^{-1} \mathcal{C}_{pJ} \mathcal{C}_J^{-1} \mathcal{C}_{pJ} \mathcal{C}_\Pi^{-1} \tilde{M}_J \\ &\quad - \tilde{M}_C^T \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} \mathcal{C}_C^{-1} \mathcal{C}_{pC} \mathcal{C}_{\Pi C}^{-1} \tilde{M}_C] X_J \\ &\equiv X_J^T A_{\text{DMAP}}^C X_J, \end{aligned} \quad (\text{D6})$$

where we have used that for Gaussian priors the likelihood at the point of maximum posterior is given by $-2 \ln \mathcal{L}(\theta_{pJ}) = X^T [(\mathbb{I} - \mathbb{P}_J)^T \Sigma_J^{-1} (\mathbb{I} - \mathbb{P}_J) + \tilde{M}_J^T \mathcal{C}_\Pi^{-1} \mathcal{C}_{pJ} \mathcal{C}_J^{-1} \times \mathcal{C}_{pJ} \mathcal{C}_\Pi^{-1} \tilde{M}_J] X$ for the joint, and similarly for the parameter copy case. In the above, the copy prior covariance is defined as $\mathcal{C}_{\Pi C} = \text{diag}(\mathcal{C}_\Pi, \mathcal{C}_\Pi)$.

To calculate the exact distribution of Q_{DMAP}^C , we follow the same procedure as in the case of data splits in Appendix C. Therefore, we start with the computation of the matrix $A_{\text{DMAP}}^C \mathcal{S}_J$ whose spectrum completely specifies the distribution of Q_{DMAP}^C as a sum of independent Gamma distributed variables. It can be shown that this matrix reduces to

$$\begin{aligned} A_{\text{DMAP}}^C \mathcal{S}_J &= \mathbb{P}_C^T - \mathbb{P}_J^T + \tilde{M}_J^T \mathcal{C}_\Pi^{-1} \mathcal{C}_{pJ} M_J^T \\ &\quad - \tilde{M}_C^T \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} \mathcal{C}_C^{-1} \mathcal{C}_{pC} (\mathcal{C}_{\Pi C}^{-1} \mathcal{C}_C + D_C D_C^T) M_C^T. \end{aligned} \quad (\text{D7})$$

It can then be also shown that the nonzero eigenvalues of Eq. (D7) are also the eigenvalues of the matrix:

$$\begin{aligned} \mathcal{A}_{\text{DMAP}}^C &= \mathbb{I}_{2N} + D_C \mathcal{C}_\Pi^{-1} \mathcal{C}_{pJ} D_C^T - \mathcal{C}_{pC}^{-1} D_C \mathcal{C}_{pJ} D_C^T \\ &\quad - \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} D_C D_C^T + \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} \mathbb{J}_{2N}, \end{aligned} \quad (\text{D8})$$

where we have defined, for convenience, the exchange matrix $\mathbb{J}_{2N} \equiv D_C D_C^T - \mathbb{I}_{2N}$ that exchanges the off diagonal blocks with the diagonal ones. Note that the above expression is written in terms of quantities that can be obtained from MCMC samples of the posterior of both the parameter copy and joint chains.

Either one can use Eq. (D8) to compute the exact distribution or one can approximate it by a chi-squared distribution matching the mean of the exact distribution as a first order Patnaiks' approximation [48]. The mean of the




FIG. 11. Spectrum of the exact distribution of Q_{DMAP}^C . Different lines correspond to different SN split, as shown in the legend. Modes that have eigenvalues below 0.05 are not shown.

exact distribution and the number of degrees of freedom of the chi-squared approximation is given by

$$\begin{aligned} \text{tr}[\mathcal{A}_{\text{DMAP}}^C] &= N + \text{tr}[\mathcal{C}_\Pi^{-1} \mathcal{C}_{pJ}] - \text{tr}[\mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC}] \\ &\quad - \text{tr}[\mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC} (\mathbb{I}_{2N} - \mathcal{C}_{\Pi C}^{-1} \mathcal{C}_{pC}) \mathbb{J}_{2N}] \\ &= N_{\text{eff}}^C - N_{\text{eff}}^J + \text{tr}[\mathcal{C}_\Pi^{-1} (\mathcal{C}_{p1C} + \mathcal{C}_{p2C} - \mathcal{C}_\Pi) \\ &\quad \times \mathcal{C}_\Pi^{-1} (\mathcal{C}_{p12C} + \mathcal{C}_{p21C})]. \end{aligned} \quad (\text{D9})$$

We can furthermore calculate the variance of the distribution as it is proportional to the trace of the matrix $(\mathcal{A}_{\text{DMAP}}^C)^2$. This, however, does not significantly simplify and in practical applications it is significantly easier to compute the variance numerically.

All the results in this Appendix agree, in the uncorrelated limit, with the results in [20].

We now compute the two exact distributions for our SN example considering only Q_{DMAP}^C as Q_{DML}^C is chi-square distributed. The eigenvalues of both the Q_{DMAP}^C matrix, as in Eq. (D8), are shown in Fig. 11.

As we can see, since the prior is not informative for the $z_{\text{cut}} = 0.3$ case, the eigenvalues result in a chi-squared distribution with 2 degrees of freedom. On the other hand, the presence of a mildly informative prior for the $z_{\text{cut}} = 0.7$ case makes one zero eigenvalue for DML different from zero and slightly negative. We, however notice that, since the distribution for Q_{DMAP}^C is exact for directions that are either fully data and fully prior constrained, contrarily to what happens in the data split case, negative eigenvalues that arise from our mild inconsistency in accounting for the priors on split parameters are usually a small correction.

It is possible to use in practice these eigenvalues to check whether there is a difference in the statistical significance of the Q_{DMAP}^C exact and approximate distribution. We find

that, in this case, the misestimate of statistical significance is subpercent.

APPENDIX E: ARBITRARILY SPLIT PARAMETERS

In this Appendix we generalize the discussion of the parameter split estimators to the case where we consider parameters that are multiply split or not split at all.

Quantities associated with the split and unsplit part of the parameter space will be denoted by the subscripts ‘‘C’’ and ‘‘U’’ respectively. We denote the unsplit posterior parameters with θ_U^p and the n posterior parameter copies with $\theta_{SC}^p = (\theta_{1C}^p, \theta_{2C}^p, \dots, \theta_{nC}^p)^T$. Therefore, the full posterior parameter vector can be written as $\theta_C^p = (\theta_{SC}^p, \theta_U^p)^T$. In what follows, the total number of copy parameters will be nN_C , where N_C is the number of split parameters, and the number of unsplit parameters will be N_U ; therefore, $N_p = nN_C + N_U$ is the total number of parameters in the final parameter vector.

Note also that the joint analysis deals with the original $N_J = N_C + N_U$ parameters in total. The joint parameter vector will be denoted as $\theta_J^p = (\theta_{CJ}^p, \theta_{UJ}^p)^T$, where the two parts θ_{CJ} and θ_{UJ} correspond to the parameter subspaces that are split and unsplit in the parameter split methodology, respectively.

The design matrix D_C relates the joint quantities with the copy ones. Constructing this appropriately is then enough to generalize our analysis as described in the previous sections. Let D_{SC} be the $nN_C \times N_C$ dimensional design matrix related to the part of parameter space that is being copied n times; thus, $D_{SC}^T = (\mathbb{I}_C, \dots, \mathbb{I}_C)$ with n instances of the identity matrix \mathbb{I}_C of dimensions $N_C \times N_C$. Let also \mathbb{I}_U be the $N_U \times N_U$ identity matrix related to the N_U unsplit parameters. Then, the design matrix takes the form

$$D_C = \begin{pmatrix} D_{SC} & \mathbb{O} \\ \mathbb{O} & \mathbb{I}_U \end{pmatrix}, \quad (\text{E1})$$

where \mathbb{O} is the vector with the appropriate number of zeros in each case. Thus, the full design matrix has dimensions of $(nN_C + N_U) \times (N_C + N_U)$.

We first consider parameter shifts of the form $\theta_{iC}^p - \theta_{jC}^p$ between the i th and j th copies, where $i, j = 1, 2, \dots, n$ run over all the n parameter copies. Then, we can express the general form of the covariance between the two parameter differences $\theta_{iC}^p - \theta_{jC}^p$ and $\theta_{kC}^p - \theta_{lC}^p$, with $k, l = 1, 2, \dots, n$, as

$$\langle (\theta_{iC}^p - \theta_{jC}^p)(\theta_{kC}^p - \theta_{lC}^p)^T \rangle = C_{pikC} + C_{pjlc} - C_{pilC} - C_{pjkc}. \quad (\text{E2})$$

These matrices then construct the covariance that is associated with the split copy part of the parameter space. Since there is no shift in the unsplit parameters, the

parameter differences and covariances associated with the unsplit part of the parameter space is zero.

We now turn to the discussion of update parameter differences. In this case, we consider differences between the posterior parameters from a joint analysis, namely θ_J^p , and the copy parameter vector $(\theta_{iC}^p, \theta_U^p)^T$ which includes the unsplit copy parameters as well as the i th copy parameter set. We thus form the parameter differences in update form as

$$\begin{aligned} \Delta\theta_C^U &\equiv (\theta_{iC}^p, \theta_U^p)^T - \theta_J^p \\ &= (\theta_{iC}^p - \theta_{CJ}^p, \theta_U^p - \theta_{UJ}^p)^T. \end{aligned} \quad (\text{E3})$$

Note that θ_{UJ}^p are generally different from θ_U^p , and that the unsplit parameters can be correlated with the split parameters. We can then explicitly calculate the parts of the covariance between such update parameter differences.

We begin by considering the covariance of the split parameter differences, which results in

$$\langle (\theta_{iC}^p - \theta_{CJ}^p)(\theta_{jC}^p - \theta_{CJ}^p)^T \rangle = C_{pijC} - C_{pCJ}, \quad (\text{E4})$$

where $C_{pCJ} = \langle (\theta_{CJ}^p)(\theta_{CJ}^p)^T \rangle - \langle (\theta_{CJ}^p) \rangle \langle (\theta_{CJ}^p)^T \rangle$ is the covariance of the parameters in the split part of the joint set. We have used the fact that $\langle (\theta_{iC}^p)(\theta_{CJ}^p)^T \rangle - \langle (\theta_{iC}^p) \rangle \langle (\theta_{CJ}^p)^T \rangle = C_{pCJ}$.

Similarly to the above, we can calculate the covariance of the unsplit parameter differences as

$$\langle (\theta_U^p - \theta_{UJ}^p)(\theta_U^p - \theta_{UJ}^p)^T \rangle = C_{pU} - C_{pUJ}, \quad (\text{E5})$$

where the covariance matrices C_{pU} and C_{pUJ} correspond to the unsplit part of the copy and joint parameter sets respectively in the same manner as for the split parameters above. Finally, we can calculate the covariance between split and unsplit parameter differences, which yields

$$\langle (\theta_{iC}^p - \theta_{CJ}^p)(\theta_U^p - \theta_{UJ}^p)^T \rangle = C_{piU} - C_{pCUJ}. \quad (\text{E6})$$

In the above we have defined the covariance C_{piU} between the copy i and unsplit copy parameters and C_{pCUJ} between the split and unsplit joint parameters, again as above.

We can now comment on the relation between the Q_{DM}^C and Q_{UDM}^C estimators and their statistical significance. As in the case of two parameter copies without unsplit parameters, which is discussed in Sec. VII A, their significance is the same for the maximum likelihood parameters, since

$$\begin{pmatrix} \theta_{iC}^{\text{ML}} \\ \theta_U^{\text{ML}} \end{pmatrix} - \theta_J^{\text{ML}} = C_J D_C^T C_C^{-1} \begin{pmatrix} D_{SC} \theta_{iC}^{\text{ML}} - \theta_{SC}^{\text{ML}} \\ \mathbf{O} \end{pmatrix},$$

where $\theta_{SC}^{\text{ML}} = (\theta_{1C}^{\text{ML}}, \theta_{2C}^{\text{ML}}, \dots, \theta_{nC}^{\text{ML}})^T$ and the zero vector \mathbf{O} has length N_U .

At the maximum posterior level the two statistics can differ, however, since the update parameter shifts contain only one copy of the prior in the joint but the prior is applied once to each set in the split analysis. As we did in Sec. VII A here we can also define the joint parameter estimate $\tilde{\theta}_J^p = (\tilde{\theta}_{CJ}^p, \theta_{UJ}^p)^T$ that counts the prior n times and has covariance $\tilde{C}_{pJ}^{-1} = C_J^{-1} + D_C^T C_{\Pi C}^{-1} D_C = D_C^T C_{pC}^{-1} D_C$. Then, the update parameter shifts would be defined as

$$\begin{pmatrix} \theta_{iC}^p \\ \theta_U^p \end{pmatrix} - \tilde{\theta}_J^p = \tilde{C}_{pJ} D_C^T C_{pC}^{-1} \begin{pmatrix} D_{SC} \theta_{iC}^p - \theta_{SC}^p \\ \mathbf{0} \end{pmatrix}, \quad (\text{E7})$$

where $\mathbf{0}$ is the zero vector of length N_U . Therefore, the statistical significance of the update differences $(\theta_{iC}^p, \theta_U^p)^T - \tilde{\theta}_J^p$ is the same as that of the parameter shifts $\theta_C^p - D_{SC} \theta_{iC}^p$ since they are related by a linear and invertible transformation. We can then always use $\tilde{\theta}_J^p$ to rewrite the update parameter difference as

$$\begin{pmatrix} \theta_{iC}^p \\ \theta_U^p \end{pmatrix} - \theta_J^p = \left[\begin{pmatrix} \theta_{iC}^p \\ \theta_U^p \end{pmatrix} - \tilde{\theta}_J^p \right] + (\tilde{\theta}_J^p - \theta_J^p). \quad (\text{E8})$$

To complete the generalization in the case of n parameter copies with unsplit parameters, we now discuss how the

statistics of goodness-of-fit loss both at the ML level, through $Q_{\text{DML}}^C \equiv X_J^T A_{\text{DML}}^C X_J$, and at the level of MAP, through $Q_{\text{DMAP}}^C \equiv X_J^T A_{\text{DMAP}}^C X_J$, can be computed. To do so we can directly follow the discussion in Appendix D to construct the matrices $A_{\text{DML}}^C \mathcal{S}_J$ and $A_{\text{DMAP}}^C \mathcal{S}_J$, respectively for ML and MAP, where we define the covariance matrix $\mathcal{S}_J \equiv \Sigma_J + M_J C_{\Pi} M_J^T$ for Gaussian priors.

Doing so it can be shown that the expressions for both ML and MAP goodness-of-fit loss statistics remain invariant compared to the results in Appendix D, provided that one uses the design matrix in Eq. (E1). Therefore, it is still true that the joint projector is a subset of the copy one, and thus Q_{DML}^C is chi-squared distributed with $\text{rank}(\mathbb{I} - \mathbb{P}_J) - \text{rank}(\mathbb{I} - \mathbb{P}_C)$ degrees of freedom.

Furthermore, Q_{DMAP}^C can be approximated by a chi-square distribution by matching moments of the approximate and exact distributions; the mean will then be given by the equivalent of Eq. (D9) if we define $N_{\text{eff}}^C = nN_C + N_U - \text{tr}[C_{\Pi C}^{-1} C_{pC}]$.

At last we highlight that, with the MCMC chain of multiple parameter copies we can easily construct the distribution of parameter differences and proceed with the statistical significance calculation as in Sec. VII B to compute the overall statistical significance of multiple parameter shifts.

-
- [1] L. Verde, T. Treu, and A. G. Riess, *Nat. Astron.* **3**, 891 (2019).
 - [2] P. Marshall, N. Rajguru, and A. Slosar, *Phys. Rev. D* **73**, 067302 (2006).
 - [3] F. Feroz, B. C. Allanach, M. Hobson, S. S. AbdusSalam, R. Trotta, and A. M. Weber, *J. High Energy Phys.* **10** (2008) 064.
 - [4] M. C. March, R. Trotta, L. Amendola, and D. Huterer, *Mon. Not. R. Astron. Soc.* **415**, 143 (2011).
 - [5] L. Amendola, V. Marra, and M. Quartin, *Mon. Not. R. Astron. Soc.* **430**, 1867 (2013).
 - [6] L. Verde, P. Protopapas, and R. Jimenez, *Phys. Dark Universe* **2**, 166 (2013).
 - [7] C. L. Bennett, D. Larson, J. L. Weiland, and G. Hinshaw, *Astrophys. J.* **794**, 135 (2014).
 - [8] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, *Phys. Rev. D* **90**, 063501 (2014).
 - [9] N. V. Karpenka, F. Feroz, and M. P. Hobson, *Mon. Not. R. Astron. Soc.* **449**, 2405 (2015).
 - [10] D. Larson, J. L. Weiland, G. Hinshaw, and C. L. Bennett, *Astrophys. J.* **801**, 9 (2015).
 - [11] G. E. Addison, Y. Huang, D. J. Watts, C. L. Bennett, M. Halpern, G. Hinshaw, and J. L. Weiland, *Astrophys. J.* **818**, 132 (2016).
 - [12] M. Raveri, *Phys. Rev. D* **93**, 043522 (2016).
 - [13] S. Seehars, A. Amara, A. Refregier, A. Paranjape, and J. Akeret, *Phys. Rev. D* **90**, 023533 (2014).
 - [14] S. Seehars, S. Grandis, A. Amara, and A. Refregier, *Phys. Rev. D* **93**, 103507 (2016).
 - [15] S. Grandis, D. Rapetti, A. Saro, J. J. Mohr, and J. P. Dietrich, *Mon. Not. R. Astron. Soc.* **463**, 1416 (2016).
 - [16] G. E. Addison, D. J. Watts, C. L. Bennett, M. Halpern, G. Hinshaw, and J. L. Weiland, *Astrophys. J.* **853**, 119 (2018).
 - [17] A. Nicola, A. Amara, and A. Refregier, *J. Cosmol. Astropart. Phys.* **01** (2019) 011.
 - [18] J. L. Weiland, K. Osumi, G. E. Addison, C. L. Bennett, D. J. Watts, M. Halpern, and G. Hinshaw, *Astrophys. J.* **863**, 161 (2018).
 - [19] Y. Huang, G. E. Addison, J. L. Weiland, and C. L. Bennett, *Astrophys. J.* **869**, 38 (2018).
 - [20] M. Raveri and W. Hu, *Phys. Rev. D* **99**, 043506 (2019).
 - [21] P. Motloch and W. Hu, *Phys. Rev. D* **97**, 103536 (2018).
 - [22] P. Motloch and W. Hu, *Phys. Rev. D* **99**, 023506 (2019).
 - [23] S. Adhikari and D. Huterer, *J. Cosmol. Astropart. Phys.* **01** (2019) 036.
 - [24] M. Kerscher and J. Weller, *SciPost Phys. Lect. Notes* **9**, 26 (2019)..
 - [25] W. Handley and P. Lemos, *Phys. Rev. D* **100**, 043504 (2019).
 - [26] Y. Huang, G. Addison, and C. Bennett, *Astrophys. J.* **882**, 124 (2019).

- [27] W. Lin and M. Ishak, [arXiv:1909.10991](https://arxiv.org/abs/1909.10991).
- [28] J. Zhang, L. Hui, and A. Stebbins, *Astrophys. J.* **635**, 806 (2005).
- [29] M. Chu and L. Knox, *Astrophys. J.* **620**, 1 (2005).
- [30] S. Wang, L. Hui, M. May, and Z. Haiman, *Phys. Rev. D* **76**, 063503 (2007).
- [31] A. Abate and O. Lahav, *Mon. Not. R. Astron. Soc.* **389**, L47 (2008).
- [32] E. J. Ruiz and D. Huterer, *Phys. Rev. D* **91**, 063009 (2015).
- [33] J. L. Bernal, L. Verde, and A. J. Cuesta, *J. Cosmol. Astropart. Phys.* **02** (2016) 059.
- [34] F. Köhlinger, B. Joachimi, M. Asgari, M. Viola, S. Joudaki, and T. Tröster, *Mon. Not. R. Astron. Soc.* **484**, 3126 (2019).
- [35] P. Lemos, F. Köhlinger, W. Handley, B. Joachimi, L. Whiteway, and O. Lahav, [arXiv:1910.07820](https://arxiv.org/abs/1910.07820).
- [36] D. M. Scolnic *et al.*, *Astrophys. J.* **859**, 101 (2018).
- [37] A. G. Riess, S. Casertano, W. Yuan, L. M. Macri, and D. Scolnic, *Astrophys. J.* **876**, 85 (2019).
- [38] J. Guy, P. Astier, S. Nobili, N. Regnault, and R. Pain (SNLS Collaboration), *Astron. Astrophys.* **443**, 781 (2005).
- [39] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
- [40] A. Lewis and S. Bridle, *Phys. Rev. D* **66**, 103511 (2002).
- [41] A. Lewis, [arXiv:1910.13970](https://arxiv.org/abs/1910.13970).
- [42] H. Hotelling, *Biometrika* **28**, 321 (1936).
- [43] M. Wand and M. Jones, *Kernel Smoothing*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability (Taylor & Francis, London, 1994).
- [44] K. Aylor, M. Joy, L. Knox, M. Millea, S. Raghunathan, and W. L. K. Wu, *Astrophys. J.* **874**, 4 (2019).
- [45] M. Raveri, *Phys. Rev. D* **101**, 083524 (2020).
- [46] L. Knox and M. Millea, *Phys. Rev. D* **101**, 043533 (2020).
- [47] A. M. Mathai and S. B. Provost, *Quadratic Forms in Random Variables* (Taylor & Francis, London, 1992).
- [48] P. B. Patnaik, *Biometrika* **37**, 78 (1950).