This is the accepted manuscript made available via CHORUS. The article has been published as:

Quantifying concordance in cosmology

Sebastian Seehars, Sebastian Grandis, Adam Amara, and Alexandre Refregier

Phys. Rev. D **93**, 103507 — Published 10 May 2016

DOI: [10.1103/PhysRevD.93.103507](https://doi.org/10.1103/PhysRevD.93.103507)

Quantifying Concordance in Cosmology

Sebastian Seechars,^{1,*} Sebastian Grandis,^{1,2,3,†} Adam Amara,^{1,‡} and Alexandre Refregier^{1,§}

¹*ETH Zurich, Department of Physics, Wolfgang-Pauli-Str. 27, 8093 Zurich, Switzerland*

²*Faculty of Physics, Ludwig-Maximilian University, Scheinerstr. 1, 81679 Munich, Germany*

³*Excellence Cluster Universe, Boltzmannstr. 2, 85748 Garching, Germany*

Quantifying the concordance between different cosmological experiments is important for testing the validity of theoretical models and systematics in the observations. In earlier work, we thus proposed the Surprise, a concordance measure derived from the relative entropy between posterior distributions. We revisit the properties of the Surprise and describe how it provides a general, versatile, and robust measure for the agreement between datasets. We also compare it to other measures of concordance that have been proposed for cosmology. As an application, we extend our earlier analysis and use the Surprise to quantify the agreement between WMAP 9, Planck 13 and Planck 15 constraints on the Λ CDM model. Using a principle component analysis in parameter space, we find that the large Surprise between WMAP 9 and Planck 13 ($S = 17.6$ bits, implying a deviation from consistency at 99.8% confidence) is due to a shift along a direction that is dominated by the amplitude of the power spectrum. The Planck 15 constraints deviate from the Planck 13 results ($S = 56.3$ bits), primarily due to a shift in the same direction. The Surprise between WMAP and Planck consequently disappears when moving to Planck 15 ($S = -5.1$ bits). This means that, unlike Planck 13, Planck 15 is not in tension with WMAP 9. These results illustrate the advantages of the relative entropy and the Surprise for quantifying the disagreement between cosmological experiments and more generally as an information metric for cosmology.

I. INTRODUCTION

The current standard model of cosmology, a flat universe dominated by dark matter and dark energy (Λ CDM), is in good agreement with a wealth of cosmological probes. As we collect and analyze ever more precise cosmological data, an important question is whether or not the different available probes are consistent with each other within a cosmological model.

For this purpose, different measures of agreement between cosmological datasets have been introduced [1–9]. In [10], we proposed a new measure that we called Surprise. It is derived from the relative entropy, or Kullback-Leibler divergence [11], between two posteriors and is a global measure of consistency over the entire cosmological parameter space. In [10], we also applied the Surprise to a historical sequence of CMB experiments. One of the key results was a large Surprise when comparing the constraints from observations of the Wilkinson Microwave Anisotropy Probe (WMAP) [12, 13] to the 2013 release of Planck data [14, 15]. In further work [16], we applied the Surprise measure to a wide range of other cosmological probes and quantified their relative information gains and tensions between them.

In this paper, we revisit the properties of the Surprise and compare it to other proposed measures of agreement between cosmological datasets. We derive expressions for these different measures in the limit of Gaussian likelihoods and a linear model and illustrate the

results using a one-dimensional toy model. As an application of the Surprise, we reexamine the consistency of Planck and WMAP in light of the 2015 data release from Planck [17, 18]. Motivated by the functional form of the Surprise, we analyze how we can identify directions in parameter space that are causing it. We discuss how these results illustrate the versatility of the Surprise as an information metric for concordance in cosmology.

The paper is organized as follows. In section II we review methods that were proposed for assessing the agreement of cosmological probes and compare them to the notion of Surprise. In section III we apply the Surprise to WMAP and Planck constraints. We conclude in section IV.

II. CONCORDANCE OF DATASETS

A number of approaches to quantifying the agreement between constraints from different datasets have been used in cosmology. Basically, each method consists of two steps. First, a measure for the agreement between constraints needs to be devised. Second, the measure needs to be interpreted on a reference scale that indicates the degree of (dis-)agreement between the datasets. In this section, we will revisit such a measure of concordance called Surprise that we introduced in [10], before summarizing the main alternative measures proposed for cosmology.

A. Cosmological parameter estimation

Inference in cosmology is typically based on the likelihood $p(\mathcal{D}|\Theta)$, the expected distribution of the data $\mathcal{D}$

* sebastian.seechars@phys.ethz.ch

† grandis@usm.lmu.de

‡ adam.amara@phys.ethz.ch

§ alexandre.refregier@phys.ethz.ch

given a model for the data that depends on parameters Θ . The constraints on Θ from the data $\mathcal{D}$ can be derived via Bayes' theorem:

$$p(\Theta|\mathcal{D}) = \frac{p(\mathcal{D}|\Theta)p(\Theta)}{p(\mathcal{D})} \quad (1)$$

where $p(\Theta)$ is the prior, $p(\Theta|\mathcal{D})$ is the posterior, and $p(\mathcal{D})$ is the evidence, defined as

$$p(\mathcal{D}) = \int d\Theta p(\mathcal{D}|\Theta)p(\Theta). \quad (2)$$

The evidence is the probability distribution for observing the data $\mathcal{D}$ given the prior knowledge on the parameters Θ .

It is instructive to study the details of the different measures for the agreement between constraints in the limit of a linear model and a Gaussian likelihood and prior. In this case, posteriors and evidences can be evaluated analytically based on a relation between the data covariance and parameter covariance given by the Fisher matrix. While the full multi-dimensional treatment is left for Appendix A, we will illustrate some of the results with a simple, one-dimensional toy model. For this we assume that we collect two independent data points d_1 and d_2 which measure the same parameter θ and follow a Gaussian likelihood:

$$p(d_i|\theta) = \frac{1}{\sqrt{2\pi s_i^2}} \exp\left[-\frac{1}{2}\left(\frac{d_i - \theta}{s_i}\right)^2\right] \quad (3)$$

$$\equiv \mathcal{N}(d_i; \theta, s_i),$$

where s_i is the standard deviation corresponding to the Gaussian error on the d_i measurement. We assume that we start with some Gaussian prior distribution for θ with mean θ_p and variance σ_p^2 . We also assume that $\sigma_p^2 \gg \sigma_i^2$, i.e. that the prior is weak compared to the measurements. In this case, the posterior is again Gaussian and given by

$$p(\theta|d_i) \simeq \mathcal{N}(\theta; \theta_i, \sigma_i), \quad (4)$$

where $\theta_i = d_i$ and $\sigma_i = s_i$ for this simple toy model. Even though posterior and likelihood are identical in this case, we will keep distinguishing between θ_i and d_i in order to make the difference between data and parameter space more apparent.

B. Surprise

The Surprise, introduced in [10], is a measure for consistency between posterior distributions and hence operates in parameter space. We consider the case of two datasets $\mathcal{D}_1$ and $\mathcal{D}_2$ which are used to derive posteriors $p(\Theta|\mathcal{D}_1)$, $p(\Theta|\mathcal{D}_2)$, and eventually even the joint posterior $p(\Theta|\mathcal{D}_1, \mathcal{D}_2)$. The Surprise can be used to analyze either the compatibility of the separately analyzed posteriors $p(\Theta|\mathcal{D}_1)$ and $p(\Theta|\mathcal{D}_2)$, or the posteriors from a Bayesian

update $p(\Theta|\mathcal{D}_1)$ and $p(\Theta|\mathcal{D}_1, \mathcal{D}_2)$. As the approach is equivalent for both cases, we will in the following denote both the individually and jointly analyzed posterior of $\mathcal{D}_2$ as $p(\Theta|\mathcal{D}_2)$ to simplify notation.

To measure the difference between $p(\Theta|\mathcal{D}_1)$ and $p(\Theta|\mathcal{D}_2)$ we use the relative entropy, or Kullback-Leibler divergence [11], $D(p(\Theta|\mathcal{D}_2)||p(\Theta|\mathcal{D}_1))$ defined as follows:

$$D(p(\Theta|\mathcal{D}_2)||p(\Theta|\mathcal{D}_1)) = \int d\Theta p(\Theta|\mathcal{D}_2) \log\left(\frac{p(\Theta|\mathcal{D}_2)}{p(\Theta|\mathcal{D}_1)}\right). \quad (5)$$

Here and in the following, $\log$ denotes the natural logarithm, so this is the definition for the relative entropy in units of *nats*. To convert the relative entropy (and later the Surprise) into *bits*, one can simply divide the values in *nats* by $\log(2)$.

The relative entropy is a convenient measure for comparing posteriors as it is zero if and only if both distributions are identical and is positive otherwise [11]. It is however a directed measure, i.e. it is not invariant under interchanging $p(\Theta|\mathcal{D}_2)$ and $p(\Theta|\mathcal{D}_1)$. The relative entropy is furthermore independent of reparametrizations of the model [11]. As the relative entropy measures both changes in precision and shifts in parameter space when updating or replacing data, it can be interpreted as the information gain when going from $p(\Theta|\mathcal{D}_1)$ to $p(\Theta|\mathcal{D}_2)$. It has been used in this sense to study the information gains from a historical sequence of CMB experiments [10] and from combining observations from a wide range of cosmological probes [16].

To distinguish contributions to D from gains in precision and shifts in parameter space, one can use constraints from one dataset to forecast the information gain that is expected for another dataset assuming that both are well described by the same model. Whenever the observed gains strongly differ from prior expectations, we interpret this as a systematic effect that is not explained by the model or included in the likelihood. From the posterior $p(\Theta|\mathcal{D}_1)$, and for consistent datasets, data $\mathcal{D}_2$ is anticipated to be a realization from

$$p(\mathcal{D}_2|\mathcal{D}_1) = \int d\Theta p(\Theta|\mathcal{D}_1)p(\mathcal{D}_2|\Theta) \quad (6)$$

which is sometimes called the posterior predictive distribution of $\mathcal{D}_1$. On average, the posterior predictive $p(\mathcal{D}_2|\mathcal{D}_1)$ predicts a relative entropy of

$$\langle D \rangle \equiv \int d\mathcal{D}_2 p(\mathcal{D}_2|\mathcal{D}_1) D(p(\Theta|\mathcal{D}_2)||p(\Theta|\mathcal{D}_1)). \quad (7)$$

We can hence compare the observed relative entropy D to the expected relative entropy $\langle D \rangle$. A measure of consistency, the Surprise of the constraints derived from $\mathcal{D}_2$, was consequently defined in [10] as:

$$S \equiv D(p(\Theta|\mathcal{D}_2)||p(\Theta|\mathcal{D}_1)) - \langle D \rangle. \quad (8)$$

By construction, we expect S to scatter around zero. If S is positive, the posteriors are more different than expected a priori. If S is negative, the constraints are more consistent than expected a priori.

The concept of expected relative entropy is also well known from Bayesian experimental design, where the aim is to design an experiment which maximizes $\langle D \rangle$ [19]. Many algorithms have been proposed to perform this optimization in general non-linear cases (see e.g. [20–22]). Huan and Marzouk [22], for example, rewrite the expected relative entropy for a joint analysis of $\mathcal{D}_1$ and $\mathcal{D}_2$ as:

$$\langle D \rangle = \int d\mathcal{D}_2 \int d\Theta (\log p(\mathcal{D}_2|\Theta) - \log p(\mathcal{D}_2|\mathcal{D}_1)) \times p(\mathcal{D}_2|\Theta)p(\Theta|\mathcal{D}_1). \quad (9)$$

Given a sample of size n from $p(\Theta|\mathcal{D}_1)$, $\langle D \rangle$ can then be estimated via

$$\langle D \rangle \approx \frac{1}{n} \sum_{i=1}^n (\log p(\mathcal{D}_2^i|\Theta_i) - \log p(\mathcal{D}_2^i|\mathcal{D}_1)), \quad (10)$$

$$p(\mathcal{D}_2^i|\mathcal{D}_1) \approx \frac{1}{n} \sum_{j=1}^n p(\mathcal{D}_2^i|\Theta_j), \quad (11)$$

where $\mathcal{D}_2^i$ is a sample from $p(\mathcal{D}_2|\Theta_i)$ for each Θ_i of the sample from $p(\Theta|\mathcal{D}_1)$. Together with approaches for calculating the relative entropy between non-Gaussian priors and posteriors (see e.g. [16, 23]), this approach allows for efficient calculation of the Surprise in general, non-Gaussian cases. In order to interpret the observed Surprise, one would however also need to estimate the relative entropy for each data $\mathcal{D}_2^i$ to estimate the expected distribution of the Surprise.

In the standard model of a flat Λ CDM cosmology, constraints on the parameters from the CMB are so stringent that an analysis assuming a linear model for the observables and Gaussianity of prior and likelihood is a good approximation. Any updates from or comparisons between CMB datasets can hence be approximated by this limit. Analytic results for the Surprise in this case have been shown to depend only on the covariances and means of the respective posteriors in [10]. For the Surprise, the result is simply given by

$$S = \frac{1}{2} (\Delta\mu\Sigma_1^{-1}\Delta\mu - \langle \Delta\mu\Sigma_1^{-1}\Delta\mu \rangle), \quad (12)$$

where $\Delta\mu$ is the difference between the means of $p(\Theta|\mathcal{D}_1)$ and $p(\Theta|\mathcal{D}_2)$ and Σ_1 is the covariance of $p(\Theta|\mathcal{D}_1)$. $\langle \Delta\mu\Sigma_1^{-1}\Delta\mu \rangle$ is given by

$$\langle \Delta\mu\Sigma_1^{-1}\Delta\mu \rangle = \text{tr}(\mathbb{1} \pm \Sigma_2\Sigma_1^{-1}), \quad (13)$$

where Σ_2 is the covariance of $p(\Theta|\mathcal{D}_2)$ and the $-$ holds when $\mathcal{D}_2$ was used to update $p(\Theta|\mathcal{D}_1)$ and the $+$ when $p(\Theta|\mathcal{D}_2)$ is derived independently of $\mathcal{D}_1$.

Equation (12) is a quite intuitive measure, as it measures the shift in the means relative to the covariance matrix of the prior information and is hence a natural generalization of the typical one-dimensional notion of $\Delta\mu/\sigma$ tensions. As shown in [10], S follows a generalized

χ^2 distribution and can be rewritten as a weighted sum of χ^2 variables Z_i with one degree of freedom,

$$S = \frac{1}{2} \left(\sum_{i=1}^d \lambda_i Z_i - \text{tr}(\mathbb{1} \pm \Sigma_2\Sigma_1^{-1}) \right) = \frac{1}{2} \sum_{i=1}^d \lambda_i (Z_i - 1) \quad (14)$$

where λ_i are the eigenvalues of $\mathbb{1} \pm \Sigma_2\Sigma_1^{-1}$ and d is the dimensionality of the parameter space.

Given a particular Surprise value S , one can use the cumulative distribution of the generalized χ^2 random variable defined in equation (14) to calculate the probability for measuring a Surprise that deviates from zero by more than S . This probability is the so-called p -value for the hypothesis that the posteriors are derived from datasets which are consistent within the model. If the p -value is small, the datasets are unlikely to be consistent within the model. To compute the p -value for S from mean and covariance of the posteriors, one can use an algorithm by Davies [24], for instance, implemented in the R package **CompQuadForm** by Duchesne and De Micheaux [25].

The one-dimensional toy model of section II A is now a simple limit of equation (12). In the case when d_2 is used to update $p(\theta|d_1)$, S is given by

$$S = \frac{1}{2} \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \left(\frac{(\theta_1 - \theta_2)^2}{\sigma_1^2 + \sigma_2^2} - 1 \right). \quad (15)$$

When d_2 is used to derive the posterior separately from d_1 , the toy model results in:

$$S = \frac{1}{2} \left(\frac{(\theta_1 - \theta_2)^2}{\sigma_1^2} - \frac{\sigma_2^2 + \sigma_1^2}{\sigma_1^2} \right). \quad (16)$$

In this toy model, the Surprise hence quantifies the difference between d_1 and d_2 by comparing their mean constraints θ_1 and θ_2 on the model. The error by which the difference is weighted depends on the scenario: deviations of θ_1 and θ_2 are compared to $\sigma_1^2 + \sigma_2^2$ for separately analyzed posteriors and to σ_1^2 only for joint analyses. The second summand of both expressions ensures that the Surprise is 0 on average.

As already mentioned earlier, the relative entropy, and hence the Surprise, are independent of the parametrization of the model. Indeed, equation (12) suggests a simple reparameterization of the parameter space when comparing constraints. An intuitive choice is a space Ψ in which $p(\Theta|\mathcal{D}_1)$ has parameters that are uncorrelated, have unit variance, and mean zero. In this space, $\tilde{\Sigma}_1 = \mathbb{1}$ by construction and the Surprise is given by

$$S = \frac{1}{2} \left(\tilde{\mu}_2^T \tilde{\mu}_2 - \text{tr}(\mathbb{1} \pm \tilde{\Sigma}_2) \right), \quad (17)$$

where $\sim$ indicates that the moments are now taken in the new space Ψ . Each shift away from zero in Ψ independently contributes to S . As the constraints $p(\Psi|\mathcal{D}_1)$ are very simple, tensions between Gaussian constraints can then be simply visualized by shifts in the one-dimensional marginal distributions of the new parameters Ψ .

C. Other measures

The agreement between datasets within a given model can be estimated from the evidence in data space or the posterior in model space. In this section, we summarize two evidence and one posterior based method that have been applied to cosmology in the past. We illustrate the methods with the one-dimensional toy model introduced in section II A, but the results for the analysis of the general linear Gaussian model in higher dimensions are given in Appendix A.

The most common tool [2–7] for assessing the agreement between cosmological observations was introduced by Marshall *et al.* [2]. It is defined as

$$R \equiv \frac{p(\mathcal{D}_1, \mathcal{D}_2)}{p(\mathcal{D}_1)p(\mathcal{D}_2)}, \quad (18)$$

with the idea that it compares the evidence for $\mathcal{D}_1$ and $\mathcal{D}_2$ when both have to be described by the same parameters of the model (numerator) to the evidence when each dataset is allowed to be described by different parameter values (denominator). The value R is then interpreted on Jeffrey’s scale, indicating consistent datasets when $R > 1$ and inconsistencies otherwise. It can be used to compare multi-dimensional, non-Gaussian constraints and numerical techniques for estimating the Bayesian evidence are well established.

For the one-dimensional toy model introduced in section II A, R evaluates to the following relation [26]:

$$\log R = -\frac{1}{2} \left[\log \left(\frac{s_2^2 + \sigma_1^2}{\sigma_p^2} \right) + \frac{(d_2 - \theta_1)^2}{s_2^2 + \sigma_1^2} - \frac{(d_2 - \theta_p)^2}{\sigma_p^2} \right] \quad (19)$$

The terms in the bottom row of equation (19) compare d_2 to the mean model θ_1 from the posterior of d_1 and to the mean model θ_p of the prior. For consistent datasets, i.e. when d_1 and d_2 are drawn from $p(d_1, d_2)$, we expect that those cancel each other on average. The term in the top row, however, is data independent and affected only by changes in the variance. For this example, $\log R$ is therefore expected to vary about the variance dependent term. Hence, the weaker the prior, the more one tends to underestimate the tensions when interpreting $\log R$ relative to 1. The example shows that interpreting the ratio independent of specific properties of prior and posterior can be non-trivial, as deviations from $R = 1$ can be attributed to inconsistencies between the data and differences between variances at the same time.

To analyze the agreement between constraints on the Hubble constant and the age of the Universe from low-redshift measurements and CMB observations, Verde *et al.* [8] use a modified version of the R measure. The idea is to perform a translation in the parameters Θ of one of the likelihoods $p(\mathcal{D}_i|\Theta)$, with $\bar{p}(\mathcal{D}_i|\Theta)$ being the

shifted distribution. For the translation where the maxima of the two distributions coincide, the evidence is written as

$$\bar{p}(\mathcal{D}_1, \mathcal{D}_2)|_{\max \mathcal{D}_1 = \max \mathcal{D}_2} = \int d\Theta \bar{p}(\mathcal{D}_2|\Theta)p(\Theta|\mathcal{D}_1), \quad (20)$$

where we chose to shift $\bar{p}(\mathcal{D}_2|\Theta)$ in this example. The measure of tension by Verde *et al.* [8] is then defined as

$$T \equiv \frac{\bar{p}(\mathcal{D}_1, \mathcal{D}_2)|_{\max \mathcal{D}_1 = \max \mathcal{D}_2}}{p(\mathcal{D}_1, \mathcal{D}_2)}. \quad (21)$$

The intuition is that $\bar{p}(\mathcal{D}_1, \mathcal{D}_2)|_{\max \mathcal{D}_1 = \max \mathcal{D}_2}$ is the evidence for maximally consistent datasets $\mathcal{D}_1$ and $\mathcal{D}_2$, to which the actual evidence is compared. T is interpreted on a slightly modified Jeffrey’s scale indicating tensions when $\log T > 1$.

Looking at the toy model, we find for $\log T$

$$\log T = \frac{1}{2} \frac{(d_2 - \theta_1)^2}{\sigma_1^2 + \sigma_2^2}. \quad (22)$$

Similarly to R , this measure compares the deviations between d_2 and the mean model for the data from the posterior of d_1 . It however does neither depend on the prior nor show the data independent variance term, implying that $\log T$ is fairly problem independent for this toy model. On average, we expect this expression to be of order 1, or of order d for general, d -dimensional parameter spaces, for consistent data. We hence find that $\log T$ tends to larger values as the size of the parameter space increases. Interpreting it relative to 1 will lead to an overestimation of tensions for parameter spaces with dimensionality greater than 2.

MacCrann *et al.* [9] define a measure of consistency between two datasets in parameter space by looking at the best-fit point Θ_{joint} of the joint constraints and its likelihood $p(\mathcal{D}_i|\Theta_{\text{joint}})$ given either $\mathcal{D}_1$ or $\mathcal{D}_2$. Within a sample $\{\Theta_i\}_{i=1}^N$ from the posterior $p(\Theta|\mathcal{D}_i)$ one can compute the percentile of likelihoods $p(\mathcal{D}_i|\Theta_i)$ that are smaller than $p(\mathcal{D}_i|\Theta_{\text{joint}})$. This percentile is then interpreted as the confidence for the two datasets being consistent. This technique is able to compare multi-dimensional, non-Gaussian constraints. The p -value of the best-fit point can be estimated from Monte Carlo Markov chains (MCMCs) of the individual posteriors.

For the toy model, the best fit point of the joint distribution is simply the mean θ_{joint} of the joint posterior. Without loss of generality, we look at the expected distribution of θ_{joint} as predicted by the posterior $p(\theta|d_1)$ of d_1 and the likelihood $p(d_2|\theta)$ assuming that the model correctly describes both sets of data. We find that, for consistent data, θ_{joint} is expected to be drawn from the following normal distribution:

$$\theta_{\text{joint}} \leftarrow \mathcal{N} \left(\theta; \theta_1, \sigma_1 \sqrt{\frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}} \right). \quad (23)$$

Comparing the distribution for θ_{joint} (23) to the posterior $p(\theta|d_1)$ of d_1 (4), we find that even though the means coincide, the expected standard deviation of θ_{joint} is smaller than σ_1 . We hence expect that, on average, the likelihood $p(d_1|\theta_{\text{joint}})$ of the best fit point is larger than the likelihood of a sample from the posterior $p(\theta|d_1)$. The percentile of $p(d_1|\theta_{\text{joint}})$ will therefore underestimate the tension introduced by the update. Only in the limit where $\sigma_1^2 \gg \sigma_2^2$, i.e. where the constraints from d_1 are weak compared to the joint constraints, is θ_{joint} expected to be distributed approximately according to $p(\theta|d_1)$.

D. Comparison of measures

In this short summary, we have seen that it is possible to define and estimate a variety of measures of agreement between general, non-Gaussian constraints. Interpreting them on a fixed, problem-independent scale can however over- or underestimate the tensions in the data.

While computationally expensive, a potential way out are simulation based approaches, where the measure is calibrated from Monte Carlo realizations of mock data from the posterior predictive of one of the constraints in question. Following Huan and Marzouk [22], such an approach was briefly outlined in IIB for the Surprise. Also Larson *et al.* [27], for example, take such an approach to comparing the constraints from WMAP 9 and Planck 13 and compare the maximum likelihood estimates of each experiment based on a simulation study. The authors create Monte Carlo realizations of CMB maps from an input power spectrum and process these maps into WMAP- and Planck-like observations by adding the corresponding noise features and masks. They then estimate the spectra of the simulated maps and find the best-fit points for each simulated WMAP and Planck map. The sample of differences between the simulated best-fit points is then compared to the observed best-fit difference on a parameter by parameter basis. A further method for dealing with potentially inconsistent datasets, proposed by Hobson *et al.* [1], is to allow for additional weight parameters when combining the likelihoods in the estimation procedure. In this approach, datasets that introduce tensions in the parameter constraints are down-weighted in order to avoid biases from systematics.

Another possibility is to focus on situations where analytic results from a linear Gaussian model are a good approximation. This limit tends to be a good description for flat Λ CDM constraints as soon as CMB data from Planck or WMAP is involved. In these cases, the distributions and expectations for most of the reviewed measures can be derived (see Appendix A). The strength of the Surprise in this scenario is its flexibility: it can be both used to analyze separately derived constraints and constraints from a Bayesian update. Operating in parameter space, it furthermore only depends on mean and covariance of prior and posterior which can be robustly estimated from standard MCMC samples. Finally, the Surprise follows a

relatively simple distribution given by the weighted sum of d χ^2 distributed variables, where d is the dimension of the parameter space. Given again only the covariance of prior and posterior, the p -value of the observed Surprise can be evaluated with standard numerical tools [24, 25]. We provide a Python module for calculating the relative entropy, the Surprise and its p -value in the linear Gaussian case at <https://github.com/seeH/surprise>.

III. APPLICATION TO PLANCK AND WMAP

The first application of the Surprise in [10] was a study of the agreement between constraints on a flat Λ CDM cosmology from a historical sequence of CMB experiments. One of the main results of this study was the detection of a significant Surprise when comparing constraints from WMAP 9 data [12, 13] and the constraints from the 2013 release of Planck data [14, 15].

In this section, we revisit this comparison in light of the 2015 Planck release and perform a more detailed study of the Surprise between WMAP and Planck. We analyze the Surprise between the posteriors derived from the WMAP 9 data on temperature and polarization [12, 13], the Planck 13 temperature data [14, 15] together with low- ℓ WMAP polarization, and finally the Planck 15 data on temperature and polarization [17]. We work in the six-dimensional flat Λ CDM parameter space given by the Hubble constant H_0 , the cold dark matter density today $\Omega_c h^2$, the baryonic matter density $\Omega_b h^2$, the optical depth to reionization τ , and the amplitude A_s and spectral index n_s of the power spectrum of primordial curvature fluctuations. As highlighted in section IIB, we need mean and covariance of these parameters for each of the posteriors. We estimate those moments from Monte Carlo Markov chains that were provided by the Planck team for the 2013 and 2015 Planck releases [18, 28].

Due to cosmic variance correlations between the data from WMAP and Planck, the joint likelihood of both datasets cannot be factorized into the individual likelihoods. For a joint analysis, one would rather have to find the correct joint likelihood function that takes the cosmic variance correlations into account. For simplicity, we hence consider the posteriors from the separately analyzed data sets. When estimating expected relative entropy and Surprise, however, we have to ignore the correlations due to cosmic variance and the large scale WMAP polarization data used in the Planck 2013 analysis to constrain τ . If we were able to take the correlations into account, they would decrease the expected relative entropy and increase the Surprise.

Table I shows the numerical values for the observed relative entropy, expected relative entropy, Surprise, and p -value of the Surprise. As already found in [10], the large Surprise ($S = 17.6$) when going from WMAP 9 to Planck 13 constraints indicates significant tension between those datasets (p -value of 0.002) [29]. In light of the 2015 release of Planck, however, the large Surprise

FIG. 1. Information gains and Surprise values for different combinations of WMAP and Planck results. The dark blue bars show the overall information gain, while the light blue bars show the Surprise. We can see that, while Planck 13 is in strong tension with both WMAP 9 and Planck 15, the constraints from WMAP 9 and Planck 15 agree very well.

TABLE I. Relative entropy and Surprise estimates in bits for comparisons between WMAP and Planck data. The p -value is the prior probability of observing a Surprise that is greater or equal (less or equal) than S if S is greater (smaller) than zero when assuming consistent data.

Data combination			D	$\langle D \rangle$	S	p -value
WMAP 9	→	Planck 13	27.3	9.6	17.6	2×10^{-3}
WMAP 9	→	Planck 15	7.6	12.6	-5.1	7×10^{-2}
Planck 13	→	Planck 15	68.4	12.0	56.3	4×10^{-5}

between WMAP 9 and Planck 13 has vanished, implying that the constraints of WMAP 9 and Planck 15 are in good agreement ($S = -5.1$). The Surprise between the Planck 13 and Planck 15 releases ($S = 56.3$) furthermore shows that the shifts in parameter space are significantly larger than expected a priori (p -value of 4×10^{-5}).

Looking at the marginalized constraints in Figure 2, however, these results are hard to relate to the seemingly mild changes between the Planck 13 and Planck 15 results. To better understand the Surprise between the Planck 13 constraints and the WMAP 9 and Planck 15 results, we argued in section II B that the functional form of the Surprise suggests to standardize the WMAP constraints. Standardization means that we reparametrize our theory such that the WMAP 9 constraints in the new parameters are uncorrelated with mean 0 and standard deviation 1. The advantage of the standardized parameter is that each component independently contributes to the Surprise (see equation (17)), potentially allowing us to identify directions in parameter space that generate the large Surprise.

Standardization is not unique and we choose to work in the eigenbasis of the correlation matrix for numerical stability. The correlation matrix is related to the covariance

matrix via $\text{Corr}(\theta_i, \theta_j) = \text{Cov}(\theta_i, \theta_j) / \sigma(\theta_i)\sigma(\theta_j)$, where $\sigma(\theta_i)$ is the standard deviation of θ_i . Given the eigendecomposition of the correlation matrix Corr of the parameters Θ :

$$\text{Corr} = U^T P U \quad (24)$$

and a diagonal matrix F containing the standard deviations of Θ , the following parameters Ψ follow a standard normal WMAP 9 posterior:

$$\Psi = \left(\sqrt{P}\right)^{-1} U F^{-1} (\Theta - \mu(\Theta)). \quad (25)$$

Here, $\sqrt{P}$ is simply containing the square roots of the eigenvalues P and $\mu(\Theta)$ contains the means of the parameters Θ . The matrix U defines how the new parameters Ψ are related to the original cosmological parameters. It is shown in Figure 3 together with the inverse eigenvalues P^{-1} for each eigenvector. We use the inverse eigenvalues in the following to label the new parameters.

The marginalized posteriors in the new parameter space Ψ are shown in Figure 4. Note that Figure 4 shows the same constraints as Figure 2 in a different parametrization of the Λ CDM model. Note also that the new parametrization is entirely derived from the WMAP 9 constraints without any knowledge of the Planck data. By construction, the WMAP 9 constraints have mean 0, standard deviation 1, and show no correlations between the individual ψ_i components of Ψ . The striking feature of Figure 4 is the deviation of the Planck 13 constraints from the WMAP 9 mean by more than 5σ in the direction of the parameter space with the dominant eigenvalue. The Planck 15 constraints, on the other hand, agree with the WMAP 9 constraints in this direction to great accuracy, thereby explaining the vanishing Surprise.

Looking at the eigenvector that corresponds to the largest eigenvalue (see first column of Figure 3), we see that the direction that is responsible for the tension between Planck 13 and the WMAP 9 and Planck 15 constraints is dominated by A_s with additional contributions from τ and $\Omega_c h^2$. Larson *et al.* [27] come to a very similar conclusion based on the simulation-based approach to comparing WMAP and Planck data that was outlined in section II D. They point out that there is a mismatch in amplitude between Planck 13 data and WMAP 9 for small ℓ which dominate the constraints on the cosmological parameters and also conclude that the shifts in the $\Omega_c h^2$ - A_s direction are larger than expected when taking cosmic variance correlations into account. A_s and τ dominantly affect the overall amplitude of the temperature power spectrum. The Planck team indeed changed their calibration scheme from the 2013 to the 2015 release [30, 31], and the strong change in the Surprise is most likely a direct consequence of this change in the data.

It is not easy to interpret the actual magnitude of the negative Surprise between WMAP 9 and Planck 15. In

FIG. 2. Marginalized posteriors of WMAP 9 (black), Planck 13 (blue), and Planck 15 (green) constraints in the parameters of a flat Λ CDM model. The contours cover 68% and 95% of the overall posterior volume and the dotted lines show the means.

FIG. 3. Matrix plot of U (defined in equation (24)), showing the eigenvectors (columns) of the correlation matrix of the WMAP 9 posterior in the cosmological parameter space together with the inverse eigenvalues given by the diagonal of P (column labels). We plot the absolute value of U to show the strength with which each cosmological parameter contributes to the eigenvectors that define the new parameter space (see color scale). For the first new parameter (labeled by 120.6), for example, we can see that it gets the strongest contribution from A_s , but furthermore points into the τ and $\Omega_c h^2$ direction.

principle, a negative Surprise indicates that the agreement between the constraints is better than expected from statistical fluctuations. The p -value of 0.07 for the Surprise being -5.1 or smaller is estimated assuming independence of Planck and WMAP power spectra. How-

ever, the measurements are correlated due to cosmic variance as they both measure the same CMB. We would therefore expect the measurements to be more consistent than predicted from the WMAP 9 posterior. The confidence of the over-consistency between WMAP 9 and Planck 15 is small anyways and taking the correlations into account would most likely diminish it.

IV. CONCLUSIONS

In this work, we have revisited the Surprise [10], a measure for the agreement between cosmological datasets within a cosmological model. The Surprise is based on the relative entropy between two posteriors and is a global measure of consistency over the entire parameter space. In the past, methods based on evidence ratios [2–8] or the likelihood of the joint best-fit to two datasets within the posteriors of the individual constraints [9] have been used to compare different sets of cosmological data. For general non-Gaussian distributions, these measures can be more straightforward to calculate than the Surprise. However, we conclude from a Gaussian toy model analysis that interpreting those measures on a fixed, problem independent scale can be misleading. To make precise statements on the compatibility of datasets for arbitrary constraints, a calibration step based on Monte Carlo simulations would be needed.

Alternatively, one can focus on situations where this calibration can be done analytically. CMB observations, for example, are constraining flat Λ CDM to a precision where posteriors are well approximated by Gaussian distributions. Whenever updating or comparing CMB con-

FIG. 4. Marginalized posteriors of WMAP 9 (black), Planck 13 (blue), and Planck 15 (green) constraints in the parametrization Ψ given by equation (25). The contours cover 68% and 95% of the overall posterior volume and the dotted lines show the means. The labels show the inverse eigenvalues of the correlation matrix given by the diagonal of P defined in equation (24). The relation between Ψ and the cosmological parameters is shown in Figure 3. By construction, the WMAP 9 constraints on Ψ are uncorrelated, have mean 0 and standard deviation 1. Each deviation from 0 in the means of the Planck constraints independently contributes to the Surprise between WMAP and Planck.

straints on Λ CDM, analytical results for Gaussian distributions can hence be used to simplify the interpretation. The Surprise is particularly useful in this limit. Being derived from a general framework that measures both tensions and gains in precision between posterior distributions in the same units of bits, it can handle Bayesian updates as well as independently analyzed data. It furthermore depends only on the mean and covariance of the posteriors which can be robustly estimated from standard MCMC samples. For consistent data, the Surprise is expected to follow a relatively simple generalized χ^2 distribution. The p -value for observing a given Surprise, assuming that the data is consistent within the model, can be estimated from the posterior moments. For convenience, we provide a `Python` module for calculating the relative entropy, the Surprise and its p -value in the linear Gaussian case at <https://github.com/seeh/surprise>.

The Surprise was first applied to a historical sequence of CMB constraints in [10], where a significant disagreement between WMAP 9 and the constraints from the 2013 release of Planck data was detected ($S = 17.6$ bits, p -value of 0.002). When taking the 2015 release of Planck data into account, we find that Planck 2015 deviates strongly from the Planck 2013 results ($S = 56.3$ bits), but is in good agreement with the WMAP 9 constraints as indicated by a negative Surprise ($S = -5.1$ bits). By analyzing the posteriors in the principal components of the WMAP 9 constraints, we were able to detect the di-

rection in parameter space that is the primary cause for the Surprise in the update from WMAP 9 to Planck 13. It is mainly composed by A_s , τ , and $\Omega_c h^2$ and shifts along this direction lead to a change in the amplitude of the predicted temperature power spectrum. The inconsistency between Planck 13 and WMAP 9 as detected by the Surprise is hence most likely the consequence of a systematic in the calibration of the Planck data which was resolved in the 2015 release of the Planck team [27, 30, 31].

Our study shows that the Surprise is a reliable measure of agreement between cosmological constraints. While there exist other measures for the agreement between datasets that can be easier to calculate for general distributions, the strength of the Surprise is its interpretability for well constrained models. As more cosmological probes are able to put tight constraints on Λ CDM parameters, estimating the Surprise between the constraints can help to detect systematic issues in data or model in the future.

ACKNOWLEDGMENTS

We acknowledge use of the Planck Legacy Archive. Planck (<http://www.esa.int/Planck>) is an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. This work was in part supported by the Swiss National Science Foundation (grant number 200021_143906).

-
- [1] M. P. Hobson, S. L. Bridle, and O. Lahav, *Monthly Notices of the Royal Astronomical Society* **335**, 377 (2002), astro-ph/0203259.
 - [2] P. Marshall, N. Rajguru, and A. Slosar, *Physical Review D* **73**, 067302 (2006).
 - [3] L. Amendola, V. Marra, and M. Quartin, *Monthly Notices of the Royal Astronomical Society* **430**, 1867 (2013).
 - [4] C. Heneka, V. Marra, and L. Amendola, *Monthly Notices of the Royal Astronomical Society* **439**, 1855 (2014), 1310.8435.
 - [5] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, *Physical Review D* **90**, 063501 (2014), 1405.7272.
 - [6] N. V. Karpenka, F. Feroz, and M. P. Hobson, *Monthly Notices of the Royal Astronomical Society* **449**, 2405 (2015), 1407.5496.
 - [7] M. Raveri, *ArXiv e-prints* (2015), arXiv:1510.00688.
 - [8] L. Verde, P. Protopapas, and R. Jimenez, *Physics of the Dark Universe* **2**, 166 (2013).
 - [9] N. MacCrann, J. Zuntz, S. Bridle, B. Jain, and M. R. Becker, *Monthly Notices of the Royal Astronomical Society* **451**, 2877 (2015).
 - [10] S. Seehars, A. Amara, A. Refregier, A. Paranjape, and J. Akeret, *Physical Review D* **90**, 023533 (2014).
 - [11] S. Kullback and R. A. Leibler, *The Annals of Mathematical Statistics* **22**, 79 (1951).
 - [12] C. L. Bennett, D. Larson, J. L. Weiland, and N. Jarosik, *The Astrophysical Journal Supplement Series* **208**, 54 (2013).
 - [13] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright, *The Astrophysical Journal Supplement Series* **208**, 19 (2013).
 - [14] Planck Collaboration, *Astronomy and Astrophysics* **571**, A1 (2014).
 - [15] Planck Collaboration, *Astronomy and Astrophysics* **571**, A15 (2014).
 - [16] S. Grandis, S. Seehars, A. Refregier, A. Amara, and A. Nicola, *ArXiv e-prints* (2015), arXiv:1510.06422.
 - [17] Planck Collaboration, *ArXiv e-prints* (2015), arXiv:1502.01582v1.
 - [18] Planck Collaboration, *ArXiv e-prints* (2015), arXiv:1502.01589v2.
 - [19] D. V. Lindley, *The Annals of Mathematical Statistics* **27**, 986 (1956).
 - [20] C. C. Drovandi, J. M. McGree, and A. N. Pettitt, *Computational Statistics and Data Analysis* **57**, 320 (2013).
 - [21] Q. Long, M. Scavino, R. Tempone, and S. Wang, *Computer Methods in Applied Mechanics and Engineering* **259**, 24 (2013).
 - [22] X. Huan and Y. M. Marzouk, *Journal of Computational Physics* **232**, 288 (2013).
 - [23] J. Skilling, in *American Institute of Physics Conference Series*, Vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint (2004) pp. 395–405.
 - [24] R. B. Davies, *Journal of the Royal Statistical Society Series C Applied Statistics*, 323 (1980).
 - [25] P. Duchesne and P. L. De Micheaux, *Computational Statistics and Data Analysis* **54**, 858 (2010).
 - [26] The asymmetry between d_1 and d_2 in this expression is an

artifact of the limit of a wide prior. As the expressions are much simpler in this limit we decided to anyways show these results. The general expression in Appendix A does not show this asymmetry and all conclusions are not affected by the limit.

- [27] D. Larson, J. L. Weiland, G. Hinshaw, and C. L. Bennett, *The Astrophysical Journal Letters* **801**, 1 (2015).
- [28] Planck Collaboration, *Astronomy and Astrophysics* **571**, A16 (2014).
- [29] Note that the numbers of the comparison between WMAP 9 and Planck 13 do not perfectly agree with the numbers quoted in [10] because the official Planck chains use only low- ℓ WMAP polarization while [10] used the Planck temperature data together with the full WMAP polarization data.
- [30] Planck Collaboration, *ArXiv e-prints* (2015), arXiv:1502.01587.
- [31] Planck Collaboration, *ArXiv e-prints* (2015), arXiv:1507.02704.

Appendix A: Linear Gaussian model

Formally, we consider two experiments that are described by data $\mathcal{D}_i$ and likelihood $p(\mathcal{D}_i|\Theta)$ where Θ are the parameters of a model for the data and $i = 1, 2$. Using Bayes' theorem, both experiments can put individual constraints on the parameters of the model via

$$p_i \equiv p(\Theta|\mathcal{D}_i) = \frac{p(\mathcal{D}_i|\Theta)p(\Theta)}{p(\mathcal{D}_i)}, \quad (\text{A1})$$

where we call p , p_i , and $p(\mathcal{D}_i)$ prior, posterior, and evidence, respectively. The evidence is the prior probability for observing data $\mathcal{D}_i$ given by:

$$p(\mathcal{D}_i) = \int d\Theta p(\mathcal{D}_i|\Theta)p(\Theta). \quad (\text{A2})$$

Both sets of data can also be used to put joint constraints on the model parameters. This either requires knowledge of the joint likelihood $\mathcal{L}(\Theta; \mathcal{D}_1, \mathcal{D}_2) = p(\mathcal{D}_1, \mathcal{D}_2|\Theta)$ or independence of the measurement in the sense that the joint likelihood can be factorized $\mathcal{L}(\Theta; \mathcal{D}_1, \mathcal{D}_2) = \mathcal{L}_1(\Theta; \mathcal{D}_1)\mathcal{L}_2(\Theta; \mathcal{D}_2)$. If independence is assumed, the joint posterior is given by:

$$p_{\text{joint}}(\Theta|\mathcal{D}_1, \mathcal{D}_2) = \frac{\mathcal{L}_1(\Theta; \mathcal{D}_1)\mathcal{L}_2(\Theta; \mathcal{D}_2)p(\Theta)}{p(\mathcal{D}_1, \mathcal{D}_2)}. \quad (\text{A3})$$

It is worth noting that this is formally equivalent to using the constraints from one set of data as a prior for the analysis of the other. For the evidence, a similar statement holds:

$$\begin{aligned} p(\mathcal{D}_1, \mathcal{D}_2) &= \int d\Theta p(\mathcal{D}_1|\Theta)p(\mathcal{D}_2|\Theta)p(\Theta) \\ &= p(\mathcal{D}_1) \int d\Theta p(\mathcal{D}_2|\Theta)p(\Theta|\mathcal{D}_1) \\ &\equiv p(\mathcal{D}_1)p(\mathcal{D}_2|\mathcal{D}_1). \end{aligned} \quad (\text{A4})$$

For the case of a Gaussian prior with mean Θ_p and covariance Σ_p , a linear model for the data $F_i(\Theta) = F_i^0 + M_i\Theta$, and a Gaussian likelihood with covariance C_i , the posterior is Gaussian with moments [10]:

$$\Sigma_i = (\Sigma_p^{-1} + M_i^T C_i^{-1} M_i)^{-1}, \quad (\text{A5})$$

$$\Theta_i = \Theta_p + \Sigma_i M_i^T C_i^{-1} (\mathcal{D}_i - F(\Theta_p)). \quad (\text{A6})$$

Also $p(\mathcal{D}_i)$ is Gaussian (in the data) with mean $F(\Theta_p)$ and covariance $C_i + M_i \Sigma_p M_i^T$ and we use the following notation for this:

$$p(\mathcal{D}_i) = \mathcal{N}(\mathcal{D}_i; F(\Theta_p), C_i + M_i \Sigma_p M_i^T) \quad (\text{A7})$$

We now have all the ingredients to study the various techniques outlined in section II C in more detail.

We start with the R measure by Marshall *et al.* [2], defined as

$$R = \frac{p(\mathcal{D}_1, \mathcal{D}_2)}{p(\mathcal{D}_1)p(\mathcal{D}_2)}. \quad (\text{A8})$$

Using equation (A4), we can rewrite R as

$$R = \frac{p(\mathcal{D}_2|\mathcal{D}_1)}{p(\mathcal{D}_2)}, \quad (\text{A9})$$

i.e. as the ratio between the evidence for $\mathcal{D}_2$ given the posterior from $\mathcal{D}_1$ to the evidence for $\mathcal{D}_2$ from the prior. In the linear Gaussian model, we know both distributions:

$$p(\mathcal{D}_2|\mathcal{D}_1) = \mathcal{N}(\mathcal{D}_2; F_2(\Theta_1), C_2 + M_2 \Sigma_1 M_2^T) \quad (\text{A10})$$

$$p(\mathcal{D}_2) = \mathcal{N}(\mathcal{D}_2; F_2(\Theta_p), C_2 + M_2 \Sigma_p M_2^T) \quad (\text{A11})$$

Using this, we can rewrite the logarithm of R as

$$\begin{aligned} \log R = & -\frac{1}{2} \left[\log \frac{\det(C_2 + M_2 \Sigma_1 M_2^T)}{\det(C_2 + M_2 \Sigma_p M_2^T)} \right. \\ & + (\mathcal{D}_2 - F_2(\Theta_1))^T (C_2 + M_2 \Sigma_1 M_2^T)^{-1} (\mathcal{D}_2 - F_2(\Theta_1)) \\ & \left. - (\mathcal{D}_2 - F_2(\Theta_p))^T (C_2 + M_2 \Sigma_p M_2^T)^{-1} (\mathcal{D}_2 - F_2(\Theta_p)) \right]. \end{aligned} \quad (\text{A12})$$

Equation (A12) hence depends on the difference between the covariances of $p(\mathcal{D}_2|\mathcal{D}_1)$ and $p(\mathcal{D}_2)$ as well as on the difference between the model of prior and posterior means Θ_q and Θ_1 and the data $\mathcal{D}_2$. Note in particular that $\log R$ depends on the prior moments Θ_p and Σ_p because of $p(\mathcal{D}_2)$. This is a well known problem of the evidence when used with a wide prior that is supposed to be uninformative. On average, i.e. when calculating the expected value of (A12) under $p(\mathcal{D}_1, \mathcal{D}_2)$, all the data dependent terms cancel each other and we find

$$\begin{aligned} \langle \log R \rangle & \equiv \int d\mathcal{D}_1 d\mathcal{D}_2 p(\mathcal{D}_1, \mathcal{D}_2) \log R \\ & = -\frac{1}{2} \log \frac{\det(C_2 + M_2 \Sigma_1 M_2^T)}{\det(C_2 + M_2 \Sigma_p M_2^T)} \end{aligned} \quad (\text{A13})$$

The second evidence-based measure by Verde *et al.* [8] is defined as:

$$T = \frac{\bar{p}(\mathcal{D}_1, \mathcal{D}_2)|_{\max \mathcal{D}_1 = \max \mathcal{D}_2}}{p(\mathcal{D}_1, \mathcal{D}_2)}, \quad (\text{A14})$$

where $\bar{p}(\mathcal{D}_1, \mathcal{D}_2)|_{\max \mathcal{D}_1 = \max \mathcal{D}_2}$ is defined as the evidence when shifting Θ in likelihood $p(\mathcal{D}_2|\Theta)$ such that the maximum of $p(\Theta|\mathcal{D}_1)$ is at the same position in parameter space as the maximum of $p(\Theta|\mathcal{D}_2)$ (see section II C for more details).

In the linear Gaussian model, shifting the parameter in the likelihood $p(\mathcal{D}_2|\Theta)$ is equivalent to changing the linear model for $\mathcal{D}_2$ to $F'_2(\Theta) = F_2^0 + M_2(\Theta + \Theta')$. We have to choose Θ' such that the mean of the new posterior Θ'_2 is equal to Θ_1 :

$$\Theta'_2 = \Theta_1 + \Sigma_2 M_2^T C_2^{-1} (\mathcal{D}_2 - F_2(\Theta_1)) \stackrel{!}{=} \Theta_1. \quad (\text{A15})$$

Solving for Θ' and plugging it into (A14), we find that T evaluates to:

$$\log T = \frac{1}{2} (\mathcal{D}_2 - F_2(\Theta_1))^T K (\mathcal{D}_2 - F_2(\Theta_1)), \quad (\text{A16})$$

where

$$K = C_2^{-1} M_2 \Sigma_2 (\Sigma_1 - \Sigma_2)^{-1} \Sigma_2 M_2^T C_2^{-1}. \quad (\text{A17})$$

The measure T hence measures the difference between the data $\mathcal{D}_2$ and the prediction for the data from the posterior of $\mathcal{D}_1$. Averaging (A16) over $p(\mathcal{D}_1, \mathcal{D}_2)$, we find that this measure is expected to vary around the following value for consistent datasets:

$$\langle \log T \rangle = \frac{d}{2}, \quad (\text{A18})$$

where d is the dimensionality of the parameter space Θ .

Finally, we take a look at the measure by MacCrann *et al.* [9] which is the p-value of the best-fit point of the joint analysis on the posteriors of the individual constraints. Without loss of generality, we consider the distribution of the best-fit point of the joint distribution as compared to the distribution of the posterior $p(\Theta|\mathcal{D}_1)$. In the linear Gaussian model the best fit point is simply given by the mean of p_{joint} :

$$\Theta_{\text{joint}} = \Theta_1 + \Sigma_{\text{joint}} M_2^T C_2^{-1} (\mathcal{D}_2 - F_2(\Theta_1)) \quad (\text{A19})$$

with $\Sigma_{\text{joint}} = (\Sigma_1^{-1} + M_2^T C_2^{-1} M_2)^{-1}$ and Θ_1 and Σ_1 being mean and covariance of $p(\Theta|\mathcal{D}_1)$. Under $p(\mathcal{D}_2|\mathcal{D}_1)$, Θ_{joint} is expected to be normally distributed with mean Θ_1 and covariance $\Sigma_1 - \Sigma_{\text{joint}}$:

$$\Theta_{\text{joint}} \leftarrow \mathcal{N}(\Theta; \Theta_1, \Sigma_1 - \Sigma_{\text{joint}}). \quad (\text{A20})$$

Using the p -value of Θ_{joint} under p_1 is hence a good approximation for the significance of the shift only if $\Sigma_{\text{joint}} \ll \Sigma_1$, i.e. when the joint posterior is much tighter than the individual constraints, and is otherwise underestimating it.