Hunting down the best model of inflation with Bayesian evidence

Jérôme Martin, Christophe Ringeval, Roberto Trotta

PHYSICAL REVIEW D **83,** 063524 (2011)

Hunting down the best model of inflation with Bayesian evidence

Jérôme Martin*
Institut d'Astrophysique de Paris, UMR 7095-CNRS, Université Pierre et Marie Curie, 98bis boulevard Arago, 75014 Paris, France

Christophe Ringeval†
Institute of Mathematics and Physics, Centre for Cosmology, Particle Physics and Phenomenology, Louvain University, 2 Chemin du Cyclotron, 1348 Louvain-la-Neuve, Belgium

Roberto Trotta‡
Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK
(Received 22 October 2010; published 25 March 2011)

We present the first calculation of the Bayesian evidence for different prototypical single field inflationary scenarios, including representative classes of small field and large field models. This approach allows us to compare inflationary models in a well-defined statistical way and to determine the current "best model of inflation." The calculation is performed numerically by interfacing the inflationary code FIELDINF with MULTINEST. We find that small field models are currently preferred, while large field models having a self-interacting potential of power $p > 4$ are strongly disfavored. The class of small field models as a whole has posterior odds of approximately 3:1 when compared with the large field class. The methodology and results presented in this article are an additional step toward the construction of a full numerical pipeline to constrain the physics of the early Universe with astrophysical observations. More accurate data (such as the Planck data) and the techniques introduced here should allow us to identify conclusively the best inflationary model.

DOI: 10.1103/PhysRevD.83.063524 PACS numbers: 98.80.Cq, 98.70.Vc

I. INTRODUCTION

The theory of inflation represents a cornerstone of the standard model of modern cosmology [1–4] (for a review see, e.g., Refs. [5–9]). By definition, it is a phase of accelerated expansion which is supposed to take place in the very early universe, somewhere between the electroweak to the grand unified theory energy scales, i.e., between $\sim 10^3$ GeV and $\sim 10^{15}$ GeV [10]. Inflation allows us to understand several puzzles which plagued the preinflationary standard model and that could not be understood otherwise. Without inflation, the standard model of cosmology would remain incomplete and highly unsatisfactory. The most spectacular achievement of inflation is that, combined with quantum mechanics, it provides a convincing mechanism for the origin of the cosmological fluctuations [the seeds of galaxies and of Cosmic Microwave Background (CMB) anisotropies] [11–15] and it predicts that their spectrum should be almost scale invariant (i.e. equal power on all spatial scales) [6,16,17] which is fully consistent with the observations [18]. This part of the scenario is particularly remarkable since it combines general relativity and quantum mechanics.

However, the physical nature of the inflaton (the field driving inflation) and its relation with the standard model of particle physics and its extensions remain elusive. Moreover the shape of its potential is not known except, of course, that it must be sufficiently flat. This is not so surprising since, as mentioned above, the inflationary mechanism is supposed to take place at energy scales larger than typically ~ 1 TeV, in a regime where particle physics is not known and has not been tested at accelerators. Another crucial aspect of the inflationary scenario is how it ends and how it is connected to the subsequent hot big-bang phase. It is believed that, after the slow-roll period, the field reaches the bottom of its potential, oscillates, and decays into radiation [19–22]. In this way, inflation is smoothly connected to the radiation-dominated epoch. However, the energy density at which the radiation-dominated era starts is not accurately known, although some new constraints on the reheating have recently been obtained in Refs. [23–25].

Despite the fact that it has become a cornerstone of modern cosmology, inflation is not as observationally constrained as the other components of the standard model. To improve on this situation, full numerical approaches can be put in place in order to use, in an optimal way, the astrophysical data now at our disposal [26–32]. This should allow investigations on the "fine structure" of the inflationary scenario. This program is particularly timely since new high-accuracy astrophysical observations, such as the European Space Agency Planck data [33], among others, will be released soon. They will provide an unprecedented window of opportunity to learn about inflation.

*jmartin@iap.fr
†christophe.ringeval@uclouvain.be
‡r.trotta@imperial.ac.uk

In this article, we are concerned with the question of how to evaluate the performance of a given inflationary model to explain the data as compared with others. This problem can be dealt within Bayesian inference [34] (see, e.g., Ref. [35] for an application to inflationary model comparison). In fact, Bayesian statistics can be used at two levels. The first level is to determine which model parameter values are favored by the data within a given inflationary model, and this for all models. To this end, one needs to compute the model's predictions for the relevant observables, such as the CMB, the galaxy power spectra, etc., and then use the experimental data to extract the posterior probability distributions of the model parameters given the data and the theoretical priors. The second level is to use Bayesian inference for model comparison. At this level, one has to calculate, for each model, the global likelihood (also known as the evidence, or model likelihood) which is obtained by integrating the usual likelihood over all of the model parameters' values, weighted by their prior probability distribution. The resulting quantity can be used to compute the posterior probability of the model, given the available data, thus updating our prior belief in each of the inflationary models in light of the observations. The Bayesian approach to model comparison has the advantage of automatically incorporating a quantitative notion of "Occam's razor," i.e., more complex inflationary models are assigned a larger posterior probability only if their complexity is effectively required to explain the data.

On the practical side, these two levels in Bayesian inference can be implemented by adopting appropriate numerical algorithms to integrate the power spectrum for a given inflationary model. This has been routinely available for several years now and, in this paper, we use the public code FIELDINF [18,36,37]. This inflationary code is then coupled with a CMB perturbation code, such as CAMB [38], and then linked with an appropriate algorithm capable of delivering both the posterior distributions for each model's parameters as well as the Bayesian evidence of each model. The evidence is computed using the publicly available MULTINEST code [39–41], which implements the nested sampling algorithm, employed as an add-on sampler to COSMOMC [42].

On the theoretical side, one has to choose classes of scenarios that are representative of the inflationary landscape and that one wishes to analyze. In this article, we focus on large and small field models for reasons specified in the following. The reheating stage is described via the reheating parameter as introduced in Refs. [18,23]. Moreover, since the choice of priors is always relevant in problems of model comparison, we have paid particular attention to their physical motivation and we carefully investigate this question both for the parameters describing the inflationary potential and for the reheating.

This article is organized as follows. In the next section, Sec. II, we present the models studied, paying special attention to the reheating part and the so-called reheating parameter. In Sec. III, we recall the definition of the Bayesian evidence, describing in detail how the priors on the free parameters characterizing each scenario are chosen. We also explain how its calculation is implemented numerically. Finally, in Sec. IV, we present our results and discuss their physical implications. Readers already familiar with the inflationary models, techniques and methods can directly jump to Sec. III B. Perhaps the most important outcome of our article is that it sketches a general method which allows us to quantify and determine the "best" model of inflation (within the list of models considered here).

II. INFLATIONARY COSMOLOGICAL PERTURBATIONS

In this section, after having briefly recalled how the theory of cosmological perturbations of quantum-mechanical origin allows us to derive the inflationary predictions, we present the scenarios studied here, discuss our choice of parametrization and motivate it based on physical considerations.

A. Choosing the inflationary potential

In order to compare inflation with various astrophysical observables, one must first determine the power spectrum of the density perturbations defined by the following expression:

$$\mathcal{P}_\zeta(k) \equiv \frac{k^3}{2\pi^2}|\zeta_k|^2, \tag{1}$$

where ζ_k is the comoving curvature perturbation in Fourier space and is a conserved quantity on super-Hubble length scales [6–9,43].

This power spectrum depends on the shape of the inflaton's potential, and thus, on its free parameters, which have to be specified. It is common to describe the landscape of possible single field inflationary models with three different archetypal classes: large field models, small field models, and hybrid inflation. This simple approach is based on the following considerations. Any inflaton potential $V(\phi)$ can always be Taylor expanded as

$$V(\phi) = V_0 \pm \alpha\left(\frac{\phi}{M_{\rm Pl}}\right)^2 + \ldots. \tag{2}$$

According to the value of the coefficients of the expansion, one obtains different classes of models. If the constant term V_0 vanishes, then one obtains a large field model [4,44]. Instead of restricting ourselves to a massive scenario, a simple generalization is to consider an arbitrary power index p, not necessarily fixed to $p = 2$ [45]. If the constant term is not zero, then one obtains a small field model [2,3] (with a negative second term) or an effective hybrid model [46,47] (with a positive second term). Again, instead of

considering only a quadratic term, it is more generic to let the power index unspecified. This leads to the three classes mentioned before.

An important question is whether the other terms of the Taylor expansion are under control. This has led to a debate on the question of whether vacuum expectation values of ϕ larger than the Planck mass are meaningful or not [5,48,49]. In the simple approach used here, we do not take part in this discussion and consider sub- as well as super-Planckian vacuum expectation values. Moreover, hybrid inflation is an intrinsic multiple field scenario (with the above potential, inflation could not actually stop) which cannot always be described by a single field approach [50,51]. Indeed, in a multiple field model, the presence of entropy perturbations can cause the evolution of $\zeta_{\boldsymbol{k}}$ on large scales and this effect can modify the power spectrum during the preheating stage. Since this type of effect is model dependent, it must be studied for each scenario and, for this reason, it is wiser, in a first step, to focus on simpler models. For this reason, we will consider in the following only the large and small field scenarios having, respectively, the following potentials:

$$V(\phi) = M^4\left(\frac{\phi}{M_{\rm Pl}}\right)^p \quad \text{(large field)}, \tag{3}$$

and

$$V(\phi) = M^4\left[1 - \left(\frac{\phi}{\mu}\right)^p\right] \quad \text{(small field)}. \tag{4}$$

Of course, this has to be considered as a first step toward a more complete scan of the inflationary landscape. The large field model is characterized by two parameters, the energy scale M and the power index p. The small field potential is characterized by three parameters, M, μ, and p. We come back to the issue of the prior distributions to assign to each parameter in Sec. III B.

B. Describing the reheating

In order to compare an inflationary model with observations, we also need to take into account the reheating stage which takes place after the end of inflation and before the onset of the radiation-dominated era. This is compulsory since one needs to know the actual value of a physical wave number during inflation from its observed value today. For instance, the amplitude of the power spectrum P_* is measured at a given wave number, typically $k_*/a_0 = 0.05\ \mathrm{Mpc}^{-1}$, where a_0 denotes the present-day scale factor. During inflation, the corresponding physical wave number is stretched back to

$$\frac{k_*}{a} = \frac{k_*}{a_0}(1 + z_{\rm end})\mathrm{e}^{N_{\rm end}-N}, \tag{5}$$

where $z_{\rm end}$ is the redshift at which inflation ended, $N_{\rm end}$ the total number of e-folds during inflation and $N \equiv \ln a$ the number of e-folds at the time considered during inflation. The quantity k_*/a is uncertain precisely due to the existence of the reheating. Assuming instantaneous transitions between inflation, reheating, radiation, and matter era, one can simplify

$$1 + z_{\rm end} = (1 + z_{\rm eq})\left(\frac{\rho_{\rm reh}}{\rho_{\rm eq}}\right)^{1/4}\frac{a_{\rm reh}}{a_{\rm end}}, \tag{6}$$

where "reh" and "eq," respectively, stand for the end of reheating and the equality between the energy density of radiation and matter. The so-called reheating parameter $R_{\rm rad}$ [18,23] describes the evolution of the Universe during the reheating stage and is defined by

$$R_{\rm rad} \equiv \frac{a_{\rm end}}{a_{\rm reh}}\left(\frac{\rho_{\rm end}}{\rho_{\rm reh}}\right)^{1/4}, \tag{7}$$

such that Eq. (6) becomes

$$1 + z_{\rm end} = \frac{1}{R_{\rm rad}}\left(\frac{\rho_{\rm end}}{\rho_{\rm r_0}}\right)^{1/4}, \tag{8}$$

where $\rho_{\rm r_0}$ is the energy density of radiation today.[1] As a result, $R_{\rm rad}$ encodes all of our ignorance on how the reheating influences the observable inflationary power spectra. In fact, it is for inflation what the optical depth τ is for CMB observations. The latter encodes how much reionization of the Universe affects the measured CMB anisotropies (independently of the details of the reionization history, at least at first order) while $R_{\rm rad}$ plays a similar role for the reheating. As it should be clear from Eq. (7), $R_{\rm rad}$ quantifies the deviation from a reheating era which would be radiationlike.

In fact, as discussed in Ref. [23], Eq. (7) can be recast into various equivalent forms. In terms of the number of e-folds during reheating $\Delta N = N_{\rm reh} - N_{\rm end} = \ln(a_{\rm reh}/a_{\rm end})$, one has

$$\ln R_{\rm rad} = \frac{\Delta N}{4}(-1 + 3\bar{w}_{\rm reh}), \tag{9}$$

where $\bar{w}_{\rm reh}$ stands for the mean equation of state parameter

$$\bar{w}_{\rm reh} \equiv \frac{1}{\Delta N}\int_{N_{\rm end}}^{N_{\rm reh}}\frac{P(n)}{\rho(n)}\mathrm{d}n. \tag{10}$$

Here $P(n)$ and $\rho(n)$ are the instantaneous total pressure and energy density of the Universe during reheating. This description is completely general since no assumption about the physical properties of the effective fluid dominating the matter content of the Universe during reheating has been made. One can also express ΔN in terms of $\bar{w}_{\rm reh}$ such that

$$\ln R_{\rm rad} = \frac{1 - 3\bar{w}_{\rm reh}}{12(1 + \bar{w}_{\rm reh})}\ln\left(\frac{\rho_{\rm reh}}{\rho_{\rm end}}\right). \tag{11}$$

[1]The density parameter of radiation today is $\Omega_{\rm r_0} \simeq 2.471 \times 10^{-5}h^2$.

As expected, one can verify explicitly that $R_{\rm rad} = 1$ if $\bar{w}_{\rm reh} = 1/3$.

III. BAYESIAN MODEL COMPARISON

In this section, we briefly review Bayesian model comparison, which we adopt to compare the performance of our inflationary models (for further details, see, e.g., [34]). As a preliminary remark, we notice that if one seeks to determine the most economical description of the inflationary potential in light of the available data, Bayesian model comparison is well suited, in that classical statistics only allows us to reject hypotheses, not to confirm them (see also Ref. [52] for alternative model selection criteria). Therefore, while some simpler models might become ruled out in a classical sense (i.e., their parameter space can become completely constrained by the data, until no viable region remains), classical statistics does not allow one to rank the remaining models in any way. Bayesian model comparison, with its natural inclusion of the Occam's razor effect, is therefore the only available tool to quantify in a self-consistent way our preference for a specific model.

A. The Bayesian evidence

Bayesian model comparison aims at computing the posterior probability of a model in view of the available data. The fundamental idea behind the procedure is that "economic" models that fit the data well are rewarded for their predictivity, while models with a large number of free parameters that turn out not to be required by the data are penalized for the wasted parameter space. Therefore, in a Bayesian sense, the best model is the one that achieves the best compromise between quality of fit and simplicity. One of the attractive features of Bayesian model comparison is that it automatically embodies a quantitative version of Occam's razor, i.e., the principle of simplicity.

Here and in the following, by "model" we denote a choice of inflationary potential, together with a specification of its free parameters, Θ_j, and of their prior probability distribution, $p(\Theta_j|\mathcal{M}_j)$. The specification of the prior is fundamental for model comparison, as the prior shape and range influence the Occam's razor effect. From Bayes' theorem, the posterior probability of model $\mathcal{M}_j$ given the data d, $p(\mathcal{M}_j|d)$, is related to the Bayesian evidence (or model likelihood) $p(d|\mathcal{M}_j)$ by

$$p(\mathcal{M}_j|d) = \frac{p(d|\mathcal{M}_j)p(\mathcal{M}_j)}{p(d)}, \tag{12}$$

where $p(\mathcal{M}_j)$ is the prior belief in model $\mathcal{M}_j$. In Eq. (12), $p(d) = \sum_i p(d|\mathcal{M}_i)p(\mathcal{M}_i)$ is a normalization constant (where the sum runs over all available known models $\mathcal{M}_i$, $i = 1, \ldots, N$) and

$$p(d|\mathcal{M}_j) = \int \mathrm{d}\Theta_j p(d|\Theta_j, \mathcal{M}_j)p(\Theta_j|\mathcal{M}_j) \tag{13}$$

is the Bayesian evidence, where $p(d|\Theta_j, \mathcal{M}_j)$ is the likelihood. The Bayesian evidence is thus the average likelihood under the prior, and is the central quantity for Bayesian model comparison.

Given two competing models, $\mathcal{M}_0$ and $\mathcal{M}_1$, the posterior odds among them are given by

$$\frac{p(\mathcal{M}_0|d)}{p(\mathcal{M}_1|d)} = B_{01}\frac{p(\mathcal{M}_0)}{p(\mathcal{M}_1)}, \tag{14}$$

where we have introduced the factor B_{01} as defined as the ratio of the models' evidences

$$B_{01} \equiv \frac{p(d|\mathcal{M}_0)}{p(d|\mathcal{M}_1)}. \tag{15}$$

The Bayes factor thus updates our relative state of belief in two models from the prior odds to the posterior odds. Large values of B_{01} denote a preference for $\mathcal{M}_0$, and small values of B_{01} denote a preference for $\mathcal{M}_1$. The "Jeffreys' scale" (Table I) gives an empirical prescription for translating the values of B_{01} into strengths of belief.

Given two or more models, specified in terms of their parametrization and priors on the parameters, it is straightforward (although sometimes computationally challenging) to compute the Bayes factor. Depending on the problem at hand, semianalytical [54,55] and numerical [39,40,56–58] techniques are available. In the usual case where the prior over models is taken to be noncommittal [i.e., $p(\mathcal{M}_j) = 1/N$], the model with the largest Bayes factor ought to be preferred. Thus, the computation of B_{01} allows us to select one (or a few) promising model(s) from a set of known models. This framework has recently been extended to evaluate the probability that the set of known models is incomplete, see Ref. [59].

Finally, we can also summarize our findings in terms of posterior probability for the entire class of models being considered here, large field or small field. From Bayes' theorem, the posterior probability for, e.g., the small field class (SF) is given by

$$p(\mathrm{SF}|d) = \sum_{i=1}^{n_{\rm SF}} \frac{p(d|\mathrm{SF}_i)p(\mathrm{SF}_i)}{p(d)}, \tag{16}$$

TABLE I. Empirical scale for evaluating the strength of evidence when comparing two models, $\mathcal{M}_0$ versus $\mathcal{M}_1$ (so-called Jeffreys' scale, here slightly modified following the prescriptions given in [34,53]). The right-most column gives our convention for denoting the different levels of evidence above these thresholds.

$\lvert \ln B_{01} \rvert$	Odds	Strength of evidence
<1.0	$\lesssim 3{:}1$	Inconclusive
1.0	$\sim 3{:}1$	Weak evidence
2.5	$\sim 12{:}1$	Moderate evidence
5.0	$\sim 150{:}1$	Strong evidence

where

$$p(d) = \sum_{i=1}^{n_{\rm SF}} p(d|{\rm SF}_i)p({\rm SF}_i) + \sum_{j=1}^{n_{\rm LF}} p(d|{\rm LF}_j)p({\rm LF}_j), \quad (17)$$

and $n_{\rm SF} = 3$ is the number of small field models considered in the class, while $n_{\rm LF} = 6$ is the number of large field models, as explained in the next section. Regarding the choice of priors for the models, in view of comparing the viability of large field and small field inflation, it is natural to divide equally the prior probability between the two classes, and then further subdivide it equally among the models in each class, so that $p({\rm SF}_j) = 1/(2n_{\rm SF})$ and $p({\rm LF}_j) = 1/(2n_{\rm LF})$. For reasons that shall become clear below, it will be convenient to consider the Bayes factor between the various models and the large field model with $p = 2$ ($\rm LF_2$), and it is therefore useful to divide both the numerator and the denominator of Eq. (16) by the evidence of $\rm LF_2$, obtaining

$$p({\rm SF}|d) = \frac{\sum_i^{n_{\rm SF}} B_{i*}p({\rm SF}_i)}{\sum_i^{n_{\rm SF}} B_{i*}p({\rm SF}_i) + \sum_j^{n_{\rm LF}} B_{j*}p({\rm LF}_j)} \quad (18)$$

$$= \frac{\langle B_{i*}\rangle_{\rm SF}}{\langle B_{i*}\rangle_{\rm SF} + \langle B_{i*}\rangle_{\rm LF}} \quad (19)$$

$$= \left(1 + \frac{\langle B_{i*}\rangle_{\rm LF}}{\langle B_{i*}\rangle_{\rm SF}}\right)^{-1}, \quad (20)$$

where we have defined

$$\langle B_{i*}\rangle_{\rm SF} \equiv \frac{1}{n_{\rm SF}} \sum_{i=1}^{n_{\rm SF}} B_{i*}, \quad (21)$$

$$\langle B_{i*}\rangle_{\rm LF} \equiv \frac{1}{n_{\rm LF}} \sum_{i=1}^{n_{\rm LF}} B_{i*}, \quad (22)$$

and in the above B_{i*} denotes the Bayes factor between model i and the $\rm LF_2$ model.

It is also instructive to consider the Bayesian complexity associated with each model, defined as [60]

$$\mathcal{C}_{\rm b} = -2[\mathcal{D}_{\rm KL}(P, \pi) - \widehat{\mathcal{D}_{\rm KL}}], \quad (23)$$

where, here, π denotes the prior distribution and $\mathcal{D}_{\rm KL}(P, \pi)$ is the Kullback-Leiber divergence between the posterior P and the prior, π, namely,

$$\mathcal{D}_{\rm KL}(P, \pi) \equiv \int P(\theta|d) \log\frac{P(\theta|d)}{\pi(\theta)} {\rm d}\theta. \quad (24)$$

In Eq. (23), $\widehat{\mathcal{D}_{\rm KL}}$ denotes a point estimate for the KL divergence. It has been shown in [34,61] that the Bayesian complexity measures the number of model parameters that the data can constrain. Evaluated together with the evidence, the complexity helps to assess whether the parametrization of a model is excessive for the constraining power of the available data (for details, see [61]). The complexity can be expressed as

$$\mathcal{D}_{\rm KL} = \langle\chi^2\rangle - \hat{\chi}^2, \quad (25)$$

where $\chi^2 \equiv -2\ln\mathcal{L}$ and the expectation value is taken with respect to the posterior. The second term, $\hat{\chi}^2$ is a plug-in estimate that can be taken to be, for example, the best-fit χ^2 value or the value of the χ^2 at the posterior mean. Here we adopt the best-fit value, following [61].

As mentioned above, the evidence is computed using the publicly available MULTINEST code [39–41], which implements the nested sampling algorithm. The gist of nested sampling is that the multidimensional evidence integral of Eq. (13) is recast into a one-dimensional integral. This is accomplished by defining the prior volume x as ${\rm d}x \equiv p(\Theta_j|\mathcal{M}_j){\rm d}\Theta_j$ so that

$$x(\lambda) = \int_{\mathcal{L}(\Theta_j)>\lambda} p(\Theta_j|\mathcal{M}_j){\rm d}\Theta_j, \quad (26)$$

where the integral is over the parameter space enclosed by the iso-likelihood contour $\mathcal{L}(\Theta_j) = \lambda$. So $x(\lambda)$ gives the volume of parameter space above a certain level λ of the likelihood (for a specific model $\mathcal{M}_j$). Then the Bayesian evidence, Eq. (13), can be written as

$$p(d|\mathcal{M}_j) = \int_0^1 \mathcal{L}(x){\rm d}x, \quad (27)$$

where $\mathcal{L}(x)$ is the inverse of Eq. (26). Samples from $\mathcal{L}(x)$ can be obtained by drawing uniformly samples from the likelihood volume within the iso-contour surface defined by λ. The standard deviation on the value of the log evidence can be estimated as $(H/n_{\rm live})^{1/2}$, where H is the negative relative entropy and $n_{\rm live}$ is the number of live points adopted, which in our case is $n_{\rm live} = 1000$ (see Ref. [39] for details). We have checked that our evidence values are robust (within error bars) if one increases $n_{\rm live}$ to 5000. The posterior distributions have also been cross-checked with standard Metropolis-Hastings Markov Chain Monte Carlo.

B. Choice of priors

Since our aim is to evaluate the evidence of large and small field models, it is absolutely crucial to choose well-motivated priors for the parameters describing the potential. In order to see why it is so, it is instructive to consider the evidence of a simple, one-parameter toy case, where there is only one single parameter θ, whose prior density under model $\mathcal{M}$ is given by $p(\theta|\mathcal{M})$. We shall further assume that the likelihood is much more sharply peaked than the prior (i.e., the quantity θ has been well measured), so that $p(\theta) \approx$ const in the range $\delta\theta$ where the likelihood $\mathcal{L}(\theta)$ is appreciably different from zero. Then the evidence of model $\mathcal{M}$, Eq. (13), is approximately equal to

$$p(d|\mathcal{M}) \approx \mathcal{L}(\theta_{\rm ML})\delta\theta p(\theta_{\rm ML}|\mathcal{M}), \quad (28)$$

where $\theta_{\rm ML}$ is the value that maximizes the likelihood function. Since the prior must be normalized, $p(\theta_{\rm ML}|\mathcal{M}) \approx 1/\Sigma$, where Σ is the characteristic width of the prior. Therefore, one finds that $p(d|\mathcal{M}) \propto \Sigma^{-1}$, i.e., the evidence scales inversely proportionally to the width of the prior. The term $\delta\theta/\Sigma$ is the so-called ''Occam's factor,'' which penalizes models with a large ''wasted'' parameter space under the prior, i.e., models for which the characteristic width of the likelihood is much smaller than that of the prior, $\delta\theta/\Sigma \ll 1$. Hence the *a priori* plausible range of parameter values determines the strength of the Occam's penalty term, and for this reason it has to be carefully chosen on the basis of physical considerations.[2]

Going back to the potentials (3) and (4), we notice that the parameter M, common to both classes of models, is *a priori* unknown, and is observationally determined by the overall normalization of the power spectrum, P_*. Since the scale of M is unknown, it is appropriate to adopt a prior flat on $\ln M$, to reflect the fact that we are giving equal *a priori* probability to all orders of magnitude within some suitably chosen lower and upper limits. A flat prior on $\ln M$ is equivalent to a flat prior on $\ln P_*$, and therefore in our numerical sampling we swap $\ln M$ for $\ln P_*$ as a fundamental parameter. Since the overall power spectrum normalization is common to all models, the precise range of values under the prior for $\ln P_*$ becomes irrelevant (as long as the range is sufficiently wide to encompass the support of the likelihood), as all models share the same Occam's razor penalty from this common parameter. In practice, we chose $\ln P_* \in [2.7 \times 10^{-10}, 4.0 \times 10^{-10}]$, but because of the above argument the Bayes factor between our models would remain unchanged even if this range was arbitrarily enlarged.

For large field models, we chose to adopt a flat prior in the range $0.2 < p < 5$. The lower limit is arbitrarily chosen to encompass all proposed large field potentials having a fractional power [62,63]. In principle, one could imagine an arbitrarily small p (which would suggest the use of a Jeffreys' prior, instead) but, up to now, there is no theoretical motivation to do so. On the other hand, there is no strong theoretical reason not to consider a model with, say, $p = 7$. However, we know that the data already strongly disfavor models with $p > 5$ (as a matter of fact, even models with $p > 3$ are disfavored [23]) and therefore one expects that the evidence of models with $p > 5$ (fixed) would be strongly disfavored. Furthermore, if one wanted to enlarge the prior range to $p > 5$ it would be easy to rescale the evidence to account for the enlarged parameter space, since the likelihood is close to 0 for $p > 5$. This would lead to a larger Occam's penalty and thus to a lower evidence, see Eq. (28).

For small field models, we have chosen a flat prior $p \in [2.4, 10]$ as our representative class since $p = 2$ is a very special case. As discussed in Ref. [18], approaching the value $p = 2$ is numerically tricky and we have chosen the lower bound as the closest, but different, possible value of $p > 2$. Models with $p < 2$ [64,65] might, in principle, be included but would constitute another class of models since this would require to cross the $p = 2$ barrier. Moreover, models with negative p correspond to very different physical regimes. For instance, the model with $p = -4$ is nothing but the Coulomb potential of brane inflation and was analyzed in detail in [66]. For the reasons detailled in Sec. I, and at this stage of the analysis, we do not include those cases. The upper bound for p has been chosen typically an order of magnitude higher. Theoretically, as already mentioned above, small field models are archetypal of inflationary potentials which can be Taylor expanded in the (small) field values, in units of a given vacuum expectation value μ. As a result, too large values of p would appear quite unnatural. Concerning μ, its scale is *a priori* unknown and, therefore, we have chosen a flat prior on $\log(\mu/M_{\rm Pl})$ in the range $[-1, 2]$. On one hand, if one has a theoretical prejudice of viewing the small field models as representative of Taylor expanded potential (as was done above), and, in particular, in the supersymmetric framework, one would expect $\mu < M_{\rm Pl}$ to keep the supergravity corrections under control. On the other hand, other theoretical approaches do not forbid super-Planckian vacuum expectation values [67] since one can always consider that this potential is obtained, not from a Taylor expansion but exactly from a more fundamental theory. The corrections would therefore not be controlled by the ratio of the vacuum expectation value to the Planck mass but by the ratio of the energy density to the Planck density. Hence our prior range is chosen in such a way as to extend above the Planck mass. Concerning the boundary values, in the limit $\mu/M_{\rm Pl} \gg 1$, one can show that the two first slow-roll parameters, and hence all observable predictions, do not longer depend on both μ and p. As a result, it is straightforward to show from Eq. (28) that the corresponding Bayes factors would be unchanged for larger values of μ. In the limit $\mu/M_{\rm Pl} \ll 1$, the first slow-roll parameter becomes tiny and the second one becomes μ independent such that, again, the observable predictions, and thus the likelihood and the evidence, are no longer sensitive to μ.

Finally, in addition to the two broad classes of large field and small field models, we have introduced in our model space finer subdivisions leading to more specific model classes. Motivated by the above prior discussion, it is natural to further distinguish between small field models allowing super-Planckian expectation values (i.e., with $\log\mu/M_{\rm Pl} > 0$) from the ones that do not

[2]Notice that parameters which are unconstrained by the data are not penalized by the Occam's factor, i.e., if the likelihood's width is similar to the prior range, then $\delta\theta/\Sigma \sim 1$ and the Occam's factor effect vanishes.

TABLE II. Inflationary models considered in this analysis and priors on their parameters. All priors are taken to be uniform (i.e. flat) in the variable and range specified, see the text for a detailed justification. In the last row, n is the number of free parameters related to the inflationary sector.

Parameter	Small field models, Eq. (4)			Large field models, Eq. (3)					
	SFI_s	SFI_l	SFI_f	LFI_p	$\mathrm{LFI}_{2/3}$	LFI_1	LFI_2	LFI_3	LFI_4
Normalization, $\ln P_*$	$[2.7\times10^{-10}, 4.0\times10^{-10}]$			$[2.7\times10^{-10}, 4.0\times10^{-10}]$					
Exponent, p	[2.4, 10]			[0.2, 5]	2/3	1	2	3	4
Vacuum expectation, $\log(\mu/M_{\mathrm{Pl}})$	$[-1, 0]$	[0, 2]	$[-1, 2]$	Not applicable					
Reheating, $\ln R$	$[-46, 15]$			$[-46, 15]$					
n	4	4	4	3	2	2	2	2	2

($\log\mu/M_{\mathrm{Pl}} < 0$). In the large field class, we have also singled out some models having a peculiar interest such as the genuine chaotic massive inflation model ($p = 2$), monodromy inflation ($p = 2/3$), linear inflation ($p = 1$), and the self-interacting potential ($p = 3$ or $p = 4$). Of course, one must restrict oneself to the positive part of the potential when necessary. Therefore we consider a total of 9 classes of models.

Having parametrized the evolution of the Universe during the reheating in the previous section, one must now discuss the choice of the prior on the reheating parameter. As shown in Ref. [18], instead of working with R_{rad} introduced in Eq. (7), it is more convenient to work with the rescaled reheating parameter R defined by

$$R \equiv R_{\mathrm{rad}} \frac{\rho_{\mathrm{end}}^{1/4}}{M_{\mathrm{Pl}}}. \tag{29}$$

As we recap in the appendix, R_{rad} exhibits trivial correlations with the normalization of the power spectrum P_* which can be easily removed by considering R instead. Notice that once the inflationary model is specified, ρ_{end} is known and Eq. (29) is nothing but a rescaling. Clearly, the order of magnitude of the different physical quantities appearing in Eqs. (7) and (29) is unknown and this suggests that we choose a flat prior on $\ln R$. The next step is to determine the prior boundaries. In fact, using the expression of R_{rad} given before, one also has

$$\ln R = \frac{1-3\bar{w}_{\mathrm{reh}}}{12(1+\bar{w}_{\mathrm{reh}})} \ln\left(\frac{\rho_{\mathrm{reh}}}{M_{\mathrm{Pl}}^4}\right) + \frac{1+3\bar{w}_{\mathrm{reh}}}{6(1+\bar{w}_{\mathrm{reh}})} \ln\left(\frac{\rho_{\mathrm{end}}}{M_{\mathrm{Pl}}^4}\right). \tag{30}$$

Positivity energy conditions in general relativity impose that $\bar{w}_{\mathrm{reh}}$ cannot exceed unity and we want to separate inflation from reheating such that $\bar{w}_{\mathrm{reh}}$ cannot be less than $-1/3$. Moreover, $\rho_{\mathrm{nuc}} < \rho_{\mathrm{reh}} < \rho_{\mathrm{end}}$, where ρ_{nuc} is the energy density at big-bang nucleosynthesis, which we take to be $\rho_{\mathrm{nuc}}^{1/4} = 10$ MeV, this implies that $-46 < \ln R < 15 + (1/3)\ln(\rho_{\mathrm{end}}/M_{\mathrm{Pl}}^4)$. Since there is no preferred value for $\ln R$, we initially take the maximal possible theoretically allowed range $[-46, 15]$. However, for each given model parameter values, we then reject all $\ln R$ values not satisfying the consistency bound $\ln R < 15 + (1/3) \times \ln(\rho_{\mathrm{end}}/M_{\mathrm{Pl}}^4)$. Finally, notice that this description of reheating via the $\ln R$ parameter and its prior range is common to all models.

To conclude the discussion on priors, we have chosen flat priors on the standard cosmological parameters centered around their currently measured values, i.e., for the density parameter of baryons $\Omega_{\mathrm{b}}h^2$, of dark matter $\Omega_{\mathrm{dm}}h^2$, the angular size of the sound horizon at last scattering θ and the optical depth τ. We also marginalize over the amplitude of the unresolved SZ signal with a flat prior in the range $A_{\mathrm{SZ}} \in [0, 2]$. These prior choices do not impact on our evidence result for the inflationary models as all models share the same standard cosmological parameters and their respective priors. We, moreover, assume throughout a flat universe as predicted by cosmic inflation.

The models we consider and the priors on the relevant inflationary parameters are summarized in Table II.

IV. RESULTS AND DISCUSSION

In this section, we present our model comparison results for the classes of models described above. Concerning the data, we have used the seven years Wilkinson Microwave Anisotropies Probe (WMAP7) data [68–70] complemented with the Hubble Space Telescope constraints on the Hubble constant today, $H_0 = 74.2 \pm 3.6$ km/s/Mpc [71]. Our findings are summarized in Fig. 1, where we show the Bayes factors for each model, computed with respect to the large field model with $p = 2$.

Within the class of large field models, we can see that models with $p \geq 3$ are disfavored, at the "weak evidence level" for $p = 3$ and at the "strong evidence" level for $p = 4$. Clearly, one can conclude that models with even larger (and fixed) values of p would be even more strongly disfavored, so that they can be effectively ruled out. We have chosen the large field $p = 2$ model as our "reference model" (the one with respect to which the Bayes factors are computed) because it plays the role of a watershed point: large field models with shallower potentials are preferred by the Bayesian evidence, with $p = 1$ and $p = 2/3$ gathering slightly more than "weak evidence" in their favor. However, the evidence is not strong enough to allow one to conclude a definite preference for these models. The more generic large field model with $p \in [0.2, 5]$ is also weakly preferred over LF_2, and this despite

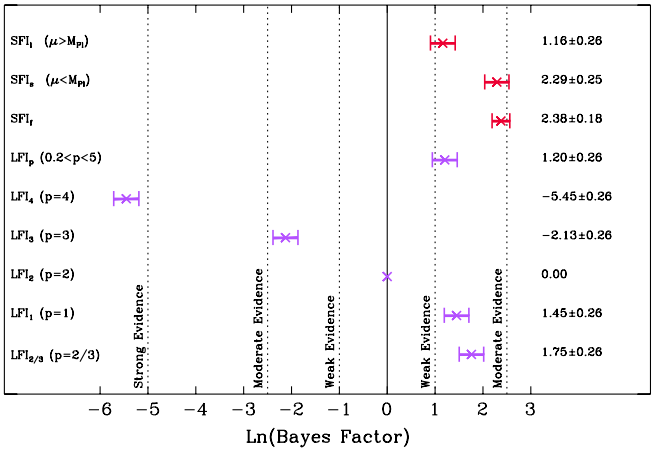

FIG. 1 (color online). Results for the Bayes factor between different inflationary models considered in the analysis. The names of the models are specified on the left of the figure. The Bayes factor are computed taking massive large field model as the reference model, and the results are given in the column on the right of the plot. The dotted vertical lines indicate the thresholds of weak, moderate, and strong evidence, as per Table I.

the extra parameter of the former, which incurs an Occam's razor penalty. As expected, the performance of this model, as measured by the evidence, falls in between the steep potentials ($p > 2$) and the shallower ones ($p < 2$).

Moving on to small field models, we remark that their overall performance is superior to our reference large field model LF_2, but quite comparable to the shallower large field models, despite the fact that small field models have one or even two parameters more than large field models. The very best models of inflation are small field. Within the error bars, the evidence cannot distinguish between a model with an upper cutoff at M_{Pl} and one that allows μ to go above the Planck mass. Models with purely super-Planckian expectation values are only very slightly disfavored, by about 1 unit in the log evidence. Therefore, we can conclude that the data are currently not sufficient to distinguish between the two scenarios.

Further insight in the model comparison outcome can be garnered by investigating simultaneously the Bayesian complexity and the evidence (or the Bayes factor) of the models considered here (see Ref. [61] for further details about the interpretation of the complexity). The Bayesian complexity, Eq. (25), has been computed for each model from a pure Markov Chain Monte Carlo run whose convergence has been monitored by using the R statistics implemented in COSMOMC [42]. The chains have been stopped as soon as the estimated errors were below 3%, which corresponds to a total number of samples ranging from 5×10^4 to 4×10^5 depending on the underlying inflationary model. The variance of our complexity estimate is obtained from the variance of four subchains of equal length randomly selected from the post burn-in samples. Both quantities are displayed in Fig. 2, where the horizontal axis gives the value of the number of input parameters for each model (both inflationary and cosmological) minus the Bayesian complexity, which we denote by the symbol $\Delta\mathcal{C}_{\mathrm{b}}$. A value of $\Delta\mathcal{C}_{\mathrm{b}}$ close to zero means that the model parameters are well constrained by the data, while $\Delta\mathcal{C}_{\mathrm{b}} > 0$ gives an estimate of the effective number of parameters remaining unconstrained by the data.

The value of $\Delta\mathcal{C}_{\mathrm{b}}$ for the large field models with $p > 2$ is generally smaller, and reaches $\Delta\mathcal{C}_{\mathrm{b}} \approx 0$ for $p = 4$, the model with the lowest evidence. This is a consequence of the tension between these models and the data, which leads to the reheating parameter becoming more and more constrained as p increases: for $p = 4$, we find a 2σ lower limit $\ln R > -2.1$, thus leading to an increase in the value of the complexity by about 1 unit. Since the models with $p = 3$ and $p = 4$ have the smallest Bayes factor while exhibiting values of $\Delta\mathcal{C}_{\mathrm{b}}$ close to 0 (meaning that all of their free parameters are well constrained), we can conclude that those models are genuinely disfavored by the data. On the other hand, for the models having a similar Bayes factor, Fig. 2 shows that the larger number of free parameters in the small field models corresponds to an increase in the number of unconstrained parameters $\Delta\mathcal{C}_{\mathrm{b}}$ with respect to its value for the large field models with $p \leq 1$. This indicates that the extra inflationary parameters in the small field class are not being constrained by the data. Therefore, we are led to conclude that while a slight preference for small field models is beginning to accumulate, it is too early to be able to conclusively favor small field models over large field ones. It is expected that Planck data will be able to conclusively pass judgement on this issue.

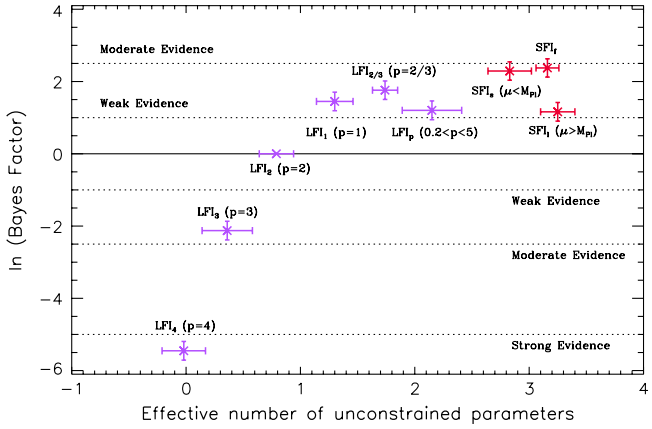

FIG. 2 (color online). Bayes factor versus the effective number of unconstrained parameters ($\Delta\mathcal{C}_{\mathrm{b}}$) for all large and small field models. The steeper LFI models are genuinely disfavored by the data, as all of their free parameters are well constrained. Small field models being favored by the evidence still have unconstrained parameters, and therefore it is too early to conclusively rule out shallower ($p < 2$) large field models, despite the fact that they exhibit a slightly smaller Bayes factor.

A consistent picture emerges when one considers the Bayesian complexity of the two models with the largest number of parameters in each class, namely SFI_f and LFI_p, with 4 and 3 inflationary parameters, respectively. For both cases, we find a similar complexity, $\mathcal{C}_\mathrm{b} \simeq 5.9$, which suggests that current data can constrain up to approximately 2 inflationary parameters. This is because our models have all $N = 5$ noninflationary parameters in common, including the SZ amplitude, and 4 of them are well constrained and contribute approximately 4 units to the Bayesian complexity. This leads to the conclusion that WMAP7 data are still insufficiently powerful to fully constrain the whole inflationary sector as parametrized in this work (see also Refs. [72,73]).

We can also evaluate the posterior probability for the entire class of small field scenarios. From Eq. (18) we find

$$p(\mathrm{SF}|d) \simeq 0.77 \pm 0.03, \tag{31}$$

and, therefore, $p(\mathrm{LF}|d) \simeq 0.23 \pm 0.03$. Therefore, the probability of the small field scenario has risen from 50% in the prior to 77% in the posterior. This represents posterior odds of $\sim$3:1 in favor of small field inflation, as compared with large field inflation. Although, as explained above, this shift in the odds is by no means conclusive, it does represent an indication that large field inflation is getting increasingly under pressure from the data [74].

Finally, it is important to assess the robustness of our results with respect to reasonable changes in our choice of models' priors. Our choice to divide the prior probability equally between the LF class and the SF class reflects the desire to compare both classes of models on an equal footing *a priori*. Another natural choice for the models' prior would be to split the prior mass equally among models, i.e., to assign $p(\mathrm{SF}_i) = p(\mathrm{LF}_i) = 1/(n_\mathrm{SF} + n_\mathrm{LF})$. This choice would, however, result in prior odds of 2:1 in favor of the LF class, which seems contrived, given that it arises solely from the fact that we have double as many LF models as SF models. Even with this (unfair to the SF class) prior choice, the posterior probability for SF would be $p(\mathrm{SF}|d) \simeq 0.6$ [up from an initial prior probability $p(\mathrm{SF}) = 1/3$], so our result of a (slight) preference for SF models stands.

Finally, we notice that our result is robust with respect to the inclusion of further models under either the SF or LF class, provided such models are disfavored by the data (as they would be, e.g., for $p > 5$ in the LF class). Inclusion of such highly disfavored models would result in their Bayes factors with respect to LF_2 being close to 0, hence the average values defined in Eqs. (21) and (22) would simply be rescaled by the new (larger) number of models in each class. However, the posterior probability of SF models only depends on the ratio of the average Bayes factors [see Eq. (18)], hence such rescaling factors would largely cancel out (for a detailed discussion of this rearrangement of prior probability in a similar context, see Ref. [59]). This holds true provided the overall number of models in each class is not widely different. We do not have any reason to believe that this should be the case. However, if one of the model classes truly had a much larger number of potential models in it, one would have to carefully reconsider the choice of giving both classes equal *a priori* mass: after all, a class of models with a smaller number of physically distinct possibilities in it is *a priori* more predictive than a class with a large number of possible distinct models.

V. CONCLUSION

To summarize, this article presented the first calculation of the Bayesian evidence for different classes of inflationary scenarios, explaining from first principles how physically meaningful priors could be derived for the fundamental parameters of the models. Among the models studied here, small field models appear to be favored, albeit still in a fairly mild way. This result must be viewed as a first step toward a more exhaustive exploration of the inflationary landscape. With the techniques introduced here and the high-accuracy CMB data soon available, we have paved the way to the identification of the best inflationary scenario.

ACKNOWLEDGMENTS

We would like to thank Patrick Peter and Jean-Philippe Uzan for useful discussions. R. T. would like to thank the Office of the Mayor of the City of Paris for partial support and the Institut d'Astrophysique de Paris (IAP) for hospitality. This work is partially supported by the Belgian Federal Office for Science, Technical and Cultural Affairs, under the Inter-University Attraction Pole Grant No. P6/11.

APPENDIX: OPTIMAL REHEATING PARAMETER

As discussed in Ref. [23], the reheating parameter R_rad can also be explicitly related to observable quantities and to the number of e-folds N_* at which a given scale k_* leaves the Hubble radius, i.e. when $k_* = a(N_*)H(N_*)$ during inflation

$$\ln R_\mathrm{rad} \simeq (N_\mathrm{end} - N_*) + N_0 - \frac{1}{4}\ln(8\pi^2 P_*) + \frac{1}{4}\ln\left(\frac{72}{r}\frac{V_\mathrm{end}}{V_*}\right), \tag{A1}$$

where V_* stands for the potential evaluated at the e-fold N_*, i.e., when $\phi = \phi(N_*)$. The quantity r is the primordial tensor-to-scalar ratio and

$$N_0 \equiv \ln\left(\frac{k_*/a_0}{\rho_{\mathrm{r}_0}^{1/4}}\right). \tag{A2}$$

Using the Friedmann-Lemaître equations together with Eq. (29), the rescaled reheating parameter now reads

$$\ln R \simeq (N_{\rm end} - N_*) + N_0 + \frac{1}{2}\ln\left(\frac{9}{2}\frac{V_{\rm end}}{V_*}\right), \tag{A3}$$

which clearly no longer depends explicitly on P_*. There is thus a great advantage of sampling the model parameters on $\ln R$ rather than $\ln R_{\rm rad}$ to prevent the unwanted degeneracies appearing in Eq. (A1). This was the approach adopted recently in Ref. [23] where, in most of this work, the parameter $\ln R$ was used. Additional constraints can be further obtained if one introduces extra assumptions on the equation of state parameter. This question was also studied in Ref. [23], but only for specific cases (unlike the description of that work given in Ref. [75]). Let us notice that in Ref. [75], the reheating phase was marginalized over by using a prior on $\Delta N_* = N_{\rm end} - N_*$ for a particular value of $k_* = 0.05\ {\rm Mpc}^{-1}$. From the above formula, it is clear that one can always trade the parameter $\ln R$ with ΔN_* provided some values for r and P_* are assumed, although it might seem awkward to introduce a scale dependent prior for a background quantity (the prior changes if one chooses another value for k_*). A drawback of this approach is that this introduces correlations between the parametrization of reheating and the normalization of the power spectrum when one has to determine the prior range. Reference [75] fixes $20 < \Delta N_* < \Delta N_*^{\rm end}$, where $\Delta N_*^{\rm end}$ corresponds to the value of ΔN_* when $\rho_{\rm end} = \rho_{\rm reh}$. It is easy to see that this choice excludes from the prior models that are physically legitimate. For instance, a small field model with $p = 3$, $\bar{w}_{\rm reh} = -0.3$, and $\mu = 0.01 M_{\rm Pl}$ is such that $17.2 < \Delta N_* < 46.0$ (for this model, the lower bound is always smaller than 20 for $\mu \in [0.01, 10]$). Of course, this also depends on the choice of $\rho_{\rm nuc}$ which is not given very precisely. If one takes $\rho_{\rm nuc} \simeq (100\ {\rm MeV})^4$, the lower bounds become $\simeq 19.3$ but if one chooses the extreme value $\rho_{\rm nuc} \simeq (1\ {\rm MeV})^4$, it is $\simeq 15.1$. When constraining model parameters, these considerations does not really matter since, as shown in Ref. [23], these model parameters are in fact disfavored by the data. But, clearly, this will affect the calculation of the evidence. The above considerations show that priors should always be chosen and justified from physical considerations.

Finally, in this article we have restricted ourselves to a standard post-reheating thermal history. But the approach used here can in fact be straightforwardly generalized to a nonstandard thermal history before big-bang nucleosynthesis. For instance, if one assumes that, inserted into the radiation-dominated era, there is actually a phase of evolution dominated by a fluid X, the equation of state parameter of which is given by $w_{\rm X}$, one could define a new parameter $R_{\rm X}$ by

$$\ln R_{\rm X} \equiv \frac{1 - 3\bar{w}_{\rm X}}{12(1 + \bar{w}_{\rm X})}\ln\left(\frac{\rho_{\rm X}^{\rm end}}{\rho_{\rm X}^{\rm start}}\right), \tag{A4}$$

where $\rho_{\rm X}^{\rm start}$ is the energy density at the beginning of the epoch dominated by the fluid X and $\rho_{\rm X}^{\rm end}$ the energy density at the end. Then, nothing is changed in the above description except that the parameter $R_{\rm rad}$ should now be replaced with $R_{\rm rad}R_{\rm X}$. Moreover, if after the X-dominated period, there is another, say, Y-dominated period, then one can define the parameter $R_{\rm Y}$ and $R_{\rm rad}$ should now be replaced with $R_{\rm rad}R_{\rm X}R_{\rm Y}$. Obviously, this works for an arbitrary number of new epochs. This also means that these nonstandard thermal histories are not really observable (unless one has a definite model for the reheating) since the new parameters $R_{\rm X}$ and $R_{\rm Y}$ are in fact completely degenerated with $R_{\rm rad}$.

[1] A. H. Guth, Phys. Rev. D **23**, 347 (1981).
[2] A. D. Linde, Phys. Lett. **108B**, 389 (1982).
[3] A. Albrecht and P. J. Steinhardt, Phys. Rev. Lett. **48**, 1220 (1982).
[4] A. D. Linde, Phys. Lett. **129B**, 177 (1983).
[5] A. D. Linde, Lect. Notes Phys. **738**, 1 (2008).
[6] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger, Phys. Rep. **215**, 203 (1992).
[7] J. Martin, Braz. J. Phys. **34**, 1307 (2004).
[8] J. Martin, Lect. Notes Phys. **669**, 199 (2005).
[9] J. Martin, Lect. Notes Phys. **738**, 193 (2008).
[10] Z.-K. Guo, D. J. Schwarz, and Y.-Z. Zhang, arXiv:1008.5258.
[11] V. F. Mukhanov and G. V. Chibisov, JETP Lett. **33**, 532 (1981).
[12] S. W. Hawking, Phys. Lett. **115B**, 295 (1982).
[13] A. A. Starobinsky, Phys. Lett. **117B**, 175 (1982).
[14] A. H. Guth and S. Y. Pi, Phys. Rev. Lett. **49**, 1110 (1982).
[15] J. M. Bardeen, P. J. Steinhardt, and M. S. Turner, Phys. Rev. D **28**, 679 (1983).
[16] E. D. Stewart and D. H. Lyth, Phys. Lett. B **302**, 171 (1993).
[17] A. R. Liddle, P. Parsons, and J. D. Barrow, Phys. Rev. D **50**, 7222 (1994).
[18] J. Martin and C. Ringeval, J. Cosmol. Astropart. Phys. 08 (2006) 009.
[19] M. S. Turner, Phys. Rev. D **28**, 1243 (1983).
[20] L. Kofman, A. D. Linde, and A. A. Starobinsky, Phys. Rev. D **56**, 3258 (1997).
[21] B. A. Bassett, S. Tsujikawa, and D. Wands, Rev. Mod. Phys. **78**, 537 (2006).
[22] A. Mazumdar and J. Rocher, Phys. Rep. **497**, 85 (2011).
[23] J. Martin and C. Ringeval, Phys. Rev. D **82**, 023511 (2010).
[24] K. Nakayama, S. Saito, Y. Suwa, and J. Yokoyama, J. Cosmol. Astropart. Phys. 06 (2008) 020.

[25] S. Kuroyanagi, C. Gordon, J. Silk, and N. Sugiyama, Phys. Rev. D **81**, 083524 (2010).
[26] C. Ringeval, Lect. Notes Phys. **738**, 243 (2008).
[27] D. S. Salopek, J. R. Bond, and J. M. Bardeen, Phys. Rev. D **40**, 1753 (1989).
[28] I. J. Grivell and A. R. Liddle, Phys. Rev. D **61**, 081301 (2000).
[29] S. M. Leach and A. R. Liddle, Phys. Rev. D **63**, 043508 (2001).
[30] J. A. Adams, B. Cresswell, and R. Easther, Phys. Rev. D **64**, 123514 (2001).
[31] A. Makarov, Phys. Rev. D **72**, 083517 (2005).
[32] S. Bird, H. V. Peiris, and D. Baumann, Phys. Rev. D **80**, 023534 (2009).
[33] J.-M. Lamarre *et al.*, New Astron. Rev. **47**, 1017 (2003) .
[34] R. Trotta, Contemp. Phys. **49**, 71 (2008).
[35] G. Ballesteros, J. A. Casas, J. R. Espinosa, R. Ruiz de Austri, and R. Trotta, J. Cosmol. Astropart. Phys. 03 (2008) 018.
[36] C. Ringeval, P. Brax, C. v. de Bruck, and A.-C. Davis, Phys. Rev. D **73**, 064035 (2006).
[37] C. Ringeval, http://theory.physics.unige.ch/~ringeval/fieldinf.html.
[38] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).
[39] F. Feroz and M. P. Hobson, Mon. Not. R. Astron. Soc., **384**, 449 (2008).
[40] F. Feroz, M. P. Hobson, and M. Bridges, Mon. Not. R. Astron. Soc. **398**, 1601 (2009).
[41] R. Trotta, F. Feroz, M. P. Hobson, L. Roszkowski, and R. Ruiz de Austri, J. High Energy Phys. 12 (2008) 024.
[42] A. Lewis and S. Bridle, Phys. Rev. D **66**, 103511 (2002).
[43] D. J. Schwarz, C. A. Terrero-Escalante, and A. A. Garcia, Phys. Lett. B **517**, 243 (2001).
[44] A. D. Linde, JETP Lett. **38**, 176 (1983).
[45] E. Silverstein and A. Westphal, Phys. Rev. D **78**, 106003 (2008).
[46] A. D. Linde, Phys. Rev. D **49**, 748 (1994).
[47] E. J. Copeland, A. R. Liddle, D. H. Lyth, E. D. Stewart, and D. Wands, Phys. Rev. D **49**, 6410 (1994).
[48] D. H. Lyth and A. Riotto, Phys. Rep. **314**, 1 (1999).
[49] A. D. Linde, Contemp. Phys. **5**, 1 (2005).
[50] S. Clesse, C. Ringeval, and J. Rocher, Phys. Rev. D **80**, 123534 (2009).
[51] S. Clesse, Phys. Rev. D **83**, 063518 (2011).
[52] A. R. Liddle, Mon. Not. R. Astron. Soc. **377**, L74 (2007).
[53] C. Gordon and R. Trotta, Mon. Not. R. Astron. Soc. **382**, 1859 (2007).
[54] R. Trotta, Mon. Not. R. Astron. Soc. **378**, 72 (2007).
[55] A. F. Heavens, T. D. Kitching, and L. Verde, Mon. Not. R. Astron. Soc. **380**, 1029 (2007).
[56] P. Mukherjee, D. Parkinson, and A. R. Liddle, Astrophys. J. **638**, L51 (2006).
[57] M. Kilbinger *et al.*, arXiv:0912.1614 [Mon. Not. Roy. Astron. Soc. (to be published)].
[58] P. Serra, A. Heavens, and A. Melchiorri, Mon. Not. R. Astron. Soc. **379**, 169 (2007).
[59] M. C. March, G. D. Starkman, R. Trotta, and P. M. Vaudrevange, Mon. Not. R. Astron. Soc. **410**, 2488 (2011).
[60] S. Kullback and R. Leibler, Ann. Math. Stat. **22**, 79 (1951).
[61] M. Kunz, R. Trotta, and D. Parkinson, Phys. Rev. D **74**, 023503 (2006).
[62] L. McAllister, E. Silverstein, and A. Westphal, Phys. Rev. D **82**, 046003 (2010).
[63] R. Kallosh and A. Linde, J. Cosmol. Astropart. Phys. 11 (2010) 011.
[64] L. Alabidi and D. H. Lyth, J. Cosmol. Astropart. Phys. 05 (2006) 016.
[65] L. Alabidi and J. E. Lidsey, Phys. Rev. D **78**, 103519 (2008).
[66] L. Lorenz, J. Martin, and C. Ringeval, J. Cosmol. Astropart. Phys. 04 (2008) 001.
[67] P. Brax and J. Martin, Phys. Rev. D **72**, 023518 (2005).
[68] E. Komatsu *et al.*, Astrophys. J. Suppl. Ser. **192**, 18 (2011).
[69] D. Larson *et al.*, Astrophys. J. Suppl. Ser. **192**, 16 (2011).
[70] N. Jarosik *et al.*, Astrophys. J. Suppl. Ser. **192**, 14 (2011).
[71] A. G. Riess *et al.*, Astrophys. J. **699**, 539 (2009).
[72] M. Kawasaki and T. Sekiguchi, J. Cosmol. Astropart. Phys. 02 (2010) 013.
[73] D. Parkinson and A. R. Liddle, Phys. Rev. D **82**, 103533 (2010).
[74] D. Boyanovsky, C. Destri, H. J. de Vega, and N. G. Sanchez, Int. J. Mod. Phys. A **24**, 3669 (2009).
[75] M. J. Mortonson, H. V. Peiris, and R. Easther, Phys. Rev. D **83**, 043505 (2011).