Quantified naturalness from Bayesian statistics

S. Fichet[1,2]

[1]*International Institute of Physics, UFRN, Avenida Odilon Gomes de Lima, 1722—Capim Macio, Natal-RN 59078-400, Brazil*
[2]*Laboratoire de Physique Subatomique et de Cosmologie, UJF Grenoble 1, CNRS/IN2P3, 53 Avenue des Martyrs,*
F-38026 Grenoble, France

We present a formulation of naturalness made in the framework of Bayesian statistics, which unravels the conceptual problems related to previous approaches. Among other things, the relative interpretation of the measure of naturalness turns out to be unambiguously established by Jeffreys' scale. Also, the usual sensitivity formulation (the so-called Barbieri-Giudice measure) appears to be embedded in our formulation under an extended form. We derive the general sensitivity formula applicable to an arbitrary number of observables. Several consequences and developments are further discussed. As a final illustration, we work out the map of combined fine-tuning associated with the gauge hierarchy problem and neutralino dark matter in a classic supersymmetric model.

I. INTRODUCTION

The notions of naturalness and fine-tuning are a center of interest in some domains of theoretical physics, like the theoretical side of particle physics and cosmology. Loosely speaking, these notions refer to the propensity of a model to reproduce experimental observations. When they are employed, their effect is to modify our degree of belief in the model examined. Indeed, intuitively, our degree of belief in a model follows its propensity to fulfill experimental constraints. For instance, when the parameters of a model must be adjusted very precisely to satisfy a constraint, the model is said to suffer from a lack of naturalness, or to have a fine-tuning problem, and this consideration typically decreases our degree of belief in the model.

But these considerations, even if they are taken to be intuitive by a certain fraction of people, remain fully subjective and unquantified. To be more precise and eventually extract some objective information from these intuitive observations, two things are necessary. First, it is necessary to define a consistent measure of naturalness. Second, it is also necessary to have a rule telling how the measure of naturalness should be mapped to our degree of belief. This second point is important, because the measure of naturalness would not be usable if the subjectivity were not under control.

Several naturalness issues appear in particle physics—in particular, with the gauge hierarchy problem, the strong *CP* puzzle, and the flavor puzzle—as well as in cosmology, with the cosmological constant problem, cosmological coincidence, and the flatness problem, the latter of these being resolved by the inflationary theory. We refer to the Appendix for a short reminder of those different issues.

To our knowledge, it is in the context of the gauge hierarchy problem that a measure of naturalness—i.e., of fine-tuning—was first built. Indeed, supersymmetric (SUSY)

models solve the gauge hierarchy problem up to a certain degree, leaving a so-called little hierarchy problem.[1] In the seminal Refs. [1,2], the amount of fine-tuning is defined as the sensitivity of the electroweak scale (characterized by the *Z*-boson mass) with respect to the model parameters. An *ad hoc* formula quantifying the fine-tuning is then derived:

$$\max_i \left| \frac{\partial \log m_Z^2}{\partial \log \theta_i} \right|. \tag{1}$$

In this context, the formula gives a measure of the amount of cancellation between the SUSY parameters, which are typically $O(\text{TeV})$, necessary to reproduce the *Z*-boson mass, 1 order of magnitude below.

This sensitivity measure (often called the Barbieri-Giudice measure) is largely exploited in the SUSY literature. We refer, for example, to Refs. [3–6] for recent work making use of it. However, this formulation has also been criticized, either for its limitations, or at the conceptual level. At the conceptual level, maybe the most straightforward remark is that there is no rule connecting the sensitivity measure to our degree of belief. The interpretation of the numbers provided by Eq. (1) is therefore fully subjective.

Several attempts in the literature have already been made to produce alternative definitions, in particular in Refs. [7,8]. Among other things, Ref. [7] introduces the key notion of the probability distribution of parameters, while Ref. [8] also introduces the important notion of the volume of parameter space. Other propositions have also been discussed in Refs. [9–12]. However, even though all of these alternative propositions are well motivated and

---

[1]The little hierarchy problem comes from the tension remaining between the electroweak and TeV scales, and is in fact an issue common to a lot of (all?) models of physics beyond the Standard Model.

contain interesting ingredients, it is unfortunately still possible to find conceptual problems and criticisms. Overall, one may find that all those measures of fine-tuning are a bit *ad hoc* or lack for a robust framework.

It is with the will of offering a solid framework to the notion of naturalness that we present an approach based on Bayesian statistics. Among other things, the link between the naturalness measure and the degree of belief will be established from Jeffreys' scale. Also, our approach turns out to contain the sensitivity measure in a generalized form. Once embedded in this framework, the usual limitations and problems of the sensitivity measure vanish. An attempt in this direction has been made in Ref. [5]; it is, to our knowledge, the only paper containing this idea.

The article is organized as follows: Naturalness problems, the sensitivity formulation, and its conceptual flaws are reviewed in a generic way in Sec. II. Section III is devoted to the basics of Bayesian model comparison relevant for our purpose, such that our presentation is self-contained from the point of view of Bayesian statistics. We then expose the Bayesian approach to naturalness and its implications in Sec. IV. Finally, Sec. V is devoted to some application of the results, focusing mainly on the gauge hierarchy problem in supersymmetric models.

II. FINE-TUNING, PUZZLES, AND SENSITIVITY

In this section, we discuss in a generic way naturalness problems and the sensitivity formulation. The presentation is aimingly transverse, and applicable to any naturalness problem. Along these lines, some of the statements might appear weak or lacking in solid definitions. These inconsistencies will be highlighted in the last paragraph. The critical point of view will be adopted only in this last part. The rest of the section is supporting the sensitivity formulation.

Throughout this section, we will consider a dimensionless quantity δ defined in a given model $\mathcal{M}$ with parameters θ_i. We will assume that this δ is subject to experimental constraints (or any other piece of information exterior to the model). In all generality, one can say that a naturalness problem appears when δ is constrained to values that it is not expected to take. In particular, it can be of an order different from $O(1)$ when it was expected to be $O(1)$, or it can be of $O(1)$ when it was not expected to be $O(1)$. For instance, the gauge hierarchy, cosmological constant, and strong CP problems are represented in the first category, with δ being m_Z^2/M_{Pl}^2, $\rho_\Lambda/M_{\mathrm{Pl}}^4$, and $\theta/2\pi$, respectively. The Universe flatness problem and cosmological coincidence fall into the second category, with δ being ρ/ρ_c and ρ_Λ/ρ_M, respectively. However, this splitting into two categories is in fact artificial. Indeed, one has always the freedom to transform one into the other by redefining $\delta \to 1/(\delta - 1)$. Therefore, in this section, whatever the naturalness problem is, we will always choose to define δ as a number unnaturally smaller than 1.

In all generality, δ is a function of the model parameters, $\delta(\theta_i)$. As we are concerned with the values that δ can potentially take, its dependence with respect to the parameters is crucial. In the limiting case where δ does not depend at all on the parameters, it is completely determined by the model $\mathcal{M}$. In that case, $\mathcal{M}$ is totally predictive, or in other words, totally natural. In the opposite direction, the more δ depends on θ_i, the more one has to adjust precisely these parameters to satisfy the experimental constraint—or, in other words, the more $\mathcal{M}$ is fine-tuned.

It is then tempting to define a measure of naturalness by making use of the derivative of δ with respect to the parameters. An appropriate quantity has to involve the logarithm of δ to measure a relative variation, and the logarithm of θ_i to maintain independence with respect to the choice of units, provided that θ_i is dimensionful. Therefore, this measure has to be based on the quantity $|\partial \log\delta/\partial \log\theta_i|$. Two such measures have been provided (see Refs. [1,2,9,13]) in the context of the gauge hierarchy problem and its supersymmetric solution:

$$c_a = \max_i \left| \frac{\partial \log\delta}{\partial \log\theta_i} \right| \quad \text{or} \quad c_b = \sqrt{\sum_i \left(\frac{\partial \log\delta}{\partial \log\theta_i} \right)^2}. \quad (2)$$

Whatever the exact definition is, we will denote this kind of quantity as c. With this definition, a fully predictive model has $c = 0$, and a model requiring infinite fine-tuning has $c \to \infty$.

However, in between these two extreme cases, one can also identify a particular threshold, when $c = 1$. This is the particular case where δ is directly an input parameter of $\mathcal{M}$. The strong CP puzzle and the flavor puzzle in the Standard Model, as well as the flatness of the Universe and cosmic coincidence in the Standard Cosmological Model are all examples of such a case. Depending on the context and on the opinions, this situation is sometimes considered as being a ''puzzle'' and not a ''problem.'' However, this kind of consideration is subjective. It depends ultimately on whether the scientist wishes to find a model more natural than the one with $c = 1$, or if he is satisfied with that one. In any case, from the strict point of view of sensitivity, $c = 1$ appears well as a limit between predictivity and fine-tuning.

Let us propose two toy examples to illustrate the c measure. To explain the smallness of δ, one often has to invoke ''special cancellations'' between the θ_i's. It is, for example, the case in the gauge hierarchy problem, where cancellations between $O(M_{\mathrm{Pl}}^2)$ quantum contributions need to occur to obtain m_Z^2; or the cosmological constant problem, in which cancellations between $O(M_{\mathrm{Pl}}^4)$ quantum contributions have to occur to make Λ vanish. Let us sketch this by $\delta \propto 1 - \theta$, where θ is a parameter expected to be $O(1)$. To produce $\delta \ll 1$, θ has to be tuned to be close to 1. The c measure is then $|\partial \log\delta/\partial \log\theta| = \theta/\delta$. c is proportional to θ, which is $O(1)$, and to $1/\delta$, which grows with the precision of cancellation required. This quantity is

thus well measuring the amount of cancellation necessary to get $\delta \ll 1$.

The second toy example is the situation where $\delta \propto e^{-\theta}$, with θ still an $O(1)$ parameter. This case of "exponential suppression" appears, for example, in the Randall-Sundrum setup for solving the gauge hierarchy problem, in inflationary theories to explain why $\rho/\rho_c - 1$ is so small, and also in the dimensional transmutation arising when an asymptotically free theory becomes confining in the infrared. In that case, the c measure is $|\partial \log\delta/\partial \log\theta| = \theta$. It does not depend on δ but only on the $O(1)$ parameter. Comparing $c = \theta/\delta$ and $c = \theta$, one can see that the "exponential" model is more natural by a factor of δ with respect to the "cancellation" model.

This way of formulating a measure of naturalness using sensitivity seems well justified, even if rather *ad hoc*. However, taking a closer look, one can identify several conceptual flaws, more or less linked together, some of them being already obvious in what we have written above.

Firstly, the notion of "expectation" for the value of a quantity, used throughout the section, is not rigorously defined. Even if one tries to express things differently, at some point this notion appears and requires a precise definition. Secondly, the notion of parameter space does not appear in this formulation. This is a bit worrying, as we are concerned with all potential values that δ could take. These two remarks are particularly suggestive of the Bayesian approach, which will be presented in the following sections. But the third, worst issue is the following: there is no rule telling us how to interpret the sensitivity measure in terms of a degree of belief. This holds for the absolute interpretation of c, and also at the level of the relative interpretation, when one compares two different values of c. For example, we said above that c in the exponential toy model is enhanced by a factor of δ with respect to the cancellation model. But does it really mean that our relative degree of belief between the two models should be given by the value δ? Or maybe δ^2, or $\sqrt{\delta}$? Finally, we can notice the freedom of redefinition of δ. For instance, if one redefines $\delta \rightarrow \delta^{100}$, c is scaled by a factor of 100. Given the absence of a rule to interpret c, this fact does not constitute a problem in itself. Instead, it can be taken as a constraint of consistency. That is, it would be good if the interpretation of c varied consistently with a redefinition of δ, so that the conclusions would remain unchanged.

III. BAYESIAN MODEL COMPARISON

The aspects of Bayesian statistics relevant for our purpose are briefly reviewed in this section. For any additional details, we refer the reader to the comprehensive review in Ref. [14] and references therein, and the textbook in Ref. [15].

Within the framework of Bayesian statistics, the notion of probability is defined as a measure of the degree of belief about a proposition. On the other hand, one also knows that whatever the definition of probability p is, the axioms of probability theory entail Bayes's law:

$$p(A|B) = p(B|A)\frac{p(A)}{p(B)}, \qquad (3)$$

which, with any additional true information I, takes the form

$$p(A|B, I) = p(B|A, I)\frac{p(A|I)}{p(B|I)}. \qquad (4)$$

This well-known result gains a crucial meaning when applied to probability as a degree of belief. Indeed, replacing A by any hypothesis H, and B by the known information available (called d for "data"), the previous equality becomes

$$p(H|d, I) = p(d|H, I)\frac{p(H|I)}{p(d|I)}. \qquad (5)$$

In Eq. (5), $p(H|I)$ is the probability (i.e., the degree of belief) given to the hypothesis without taking the data into account, which is called the prior probability, or just the "prior." $p(H|d, I)$ is the probability of the hypothesis once the data is taken into account, called the posterior probability. One thus sees that the Bayes formula, applied to a piece of information d and a hypothesis H, tells how our degree of belief in H should be updated in the light of d. It is the remaining term, $p(d|H, I)/p(d|I)$, which performs this action. $p(d|H, I)$ is the probability of obtaining the data, assuming that the hypothesis is true. But taken as a function of H, this quantity is not a probability anymore, and is called a likelihood function, denoted as $\mathcal{L}(H)$. It has to be normalized by the constant $p(d|I)$, which is called the Bayesian evidence. The Bayesian evidence is the sum over all possible realizations of H:

$$p(d|I) = \sum_{H} p(d|H, I)p(H|I). \qquad (6)$$

Two main applications follow from Eq. (5): parameter inference and model comparison. We will be interested in the latter for our purpose. For model comparison, it is the Bayesian evidence [Eq. (6)] which will play the main role.

Let us consider Eq. (5), where the hypothesis H is "model $\mathcal{M}$ is true," and there is no additional proposition I. The equation becomes

$$p(\mathcal{M}|d) = p(d|\mathcal{M})\frac{p(\mathcal{M})}{p(d)}. \qquad (7)$$

Applying it to two models (which can be the same model with two different priors), $\mathcal{M}_0$ and $\mathcal{M}_1$, and eliminating the unknown constant $p(d)$, one obtains the equation

$$\frac{p(\mathcal{M}_0|d)}{p(\mathcal{M}_1|d)} = \frac{p(d|\mathcal{M}_0)}{p(d|\mathcal{M}_1)}\frac{p(\mathcal{M}_0)}{p(\mathcal{M}_1)}. \qquad (8)$$

The quantity $\frac{p(\mathcal{M}_0)}{p(\mathcal{M}_1)}$ is called the prior odds, while $\frac{p(\mathcal{M}_0|d)}{p(\mathcal{M}_1|d)}$ is called the posterior odds. The crucial quantity is the ratio of the Bayesian evidences $\frac{p(d|\mathcal{M}_0)}{p(d|\mathcal{M}_1)}$, denoted as B_{01}, and called the Bayes factor.

The Bayes factor tells us how the relative degree of belief between two models is updated given information d. A Bayes factor larger (smaller) than 1 will favor $\mathcal{M}_0$ ($\mathcal{M}_1$). Bayes factors are usually interpreted with respect to Jeffreys' scale [16], given in Table I. This scale is empirically calibrated, with thresholds at values of the odds 3:1, 12:1, and 150:1, representing weak, moderate, and strong evidence in favor of $\mathcal{M}_0$, respectively. It can also be convenient to consider the logarithm $\log B_{01}$.

Note that for a model with continuous parameters, the Bayesian evidence [Eq. (6)] takes an integral form:

$$p(d|I) = \int_{\mathcal{D}} p(d|\theta, \mathcal{M})p(\theta|\mathcal{M}). \tag{9}$$

It is then the average of the likelihood function over the parameter space $\mathcal{D}$, weighted by the prior density of the parameters within the model $p(\theta|\mathcal{M})$.

Bayesian model comparison tells us how the odds between two models should be modified by taking into account an external piece of information d. It formalizes two competing effects: quality of fit and predictivity. The first of these is the usual measure of deviation between data and prediction, given by the likelihood function. The second is a principle of economy, i.e., a formalization of Occam's razor. It will enter in the form of the notion of volume in the parameter space. Roughly speaking, provided that the volume of parameter space allowed by the likelihood is smaller than the one allowed by priors (i.e., that data is informative), the Bayes factor will favor the model with the smaller prior volume. Or, in other words, it favors the model which is the more predictive with respect to data. This notion of volume is closely related to Fisher information, which, in this context, is a measure of the intrinsic amount of information that the likelihood function and priors contain [17]. For our purposes, we will consider the "observed" Fisher information, defined as $I\{f\}(x) = |-\frac{\partial^2 \log f}{\partial x^i x^j}|$. For example, Fisher information of a normal density with variance σ^2 is $1/\sigma^2$, and Fisher information of a uniform density over the volume V is $1/V^2$. In the present work, it will mainly be this second aspect of predictivity that will matter.

TABLE I. The empirical Jeffreys' scale calibrating the odds between model $\mathcal{M}_0$ and model $\mathcal{M}_1$.

$|\log B_{01}|$	Odds	Probability	Strength of evidence
<1.0	$\lesssim$ 3:1	<0.750	Inconclusive
1.0	$\sim$3:1	0.750	Weak evidence
2.5	$\sim$12:1	0.923	Moderate evidence
5.0	$\sim$150:1	0.993	Strong evidence

To end this section, let us discuss the prior density $p(\theta|\mathcal{M})$. The choices of both the functional form and the range of the prior density are critical. The range, conservatively, should be taken to be as wide as possible. It can be crucial to have ranges which are intrinsically bounded, such that prior volumes remain finite. On the other hand, the functional form of the density is often chosen to be the least informative one possible; i.e., the most objective. Several approaches based on Fisher information (Jeffreys prior) or the Kullback-Leibler divergence (reference priors, see, e.g., Ref. [18]) have been elaborated to construct such priors.

In this work, we will make use of the principle of indifference, which is an approach to minimize the amount of subjective information about a problem. This principle states that our *a priori* degree of belief about a problem should be invariant under transformations considered as irrelevant for the problem [19,20]. Applied to continuous variables, this condition constrains the objective densities and can sometimes fully determine them. For example, a change in coordinates $x' = x + a$ should not influence our *a priori* degree of belief on x. This transformation is thus considered irrelevant. This imposes the condition $p(x + a) = p(x)$, which constrains p to be the uniform density. Another important example is the one of a dimensionful quantity, μ. The principle of indifference states that our *a priori* degree of belief $p(\mu)$ should not depend on the choice of units, such that $\mu' = a\mu$ has the same prior as μ. This translates as the condition $p(\mu) = ap(a\mu)$, which sets $p(\mu) \propto \mu^{-1}$, called a logarithmic prior, since $\mu^{-1}d\mu = d\log\mu$. As a lot of our observables and parameters are dimensionful, this logarithmic prior will be omnipresent.

IV. NATURALNESS IN A BAYESIAN FRAMEWORK

We present in this section the Bayesian approach to the notion of naturalness. First, let us set up the notations. We consider a model $\mathcal{M}$, with a set of dimensionful parameters $\theta = (\theta_1, \ldots, \theta_n)$, spanning the parameter space $\mathcal{D}$ of dimension n. We consider a set of m dimensionful observables $\mathcal{O}(\theta) = (\mathcal{O}_1, \ldots, \mathcal{O}_m)$ (with $m \leq n$) predicted by this model, taking the measured value $\mathcal{O}_{ex}$ on the subset of the parameter space $\mathcal{D}_{ex}$ of dimension n-m. Data other than the $\mathcal{O}$ measurement are collectively called d, and the likelihood function $p(\mathcal{O} = \mathcal{O}_{ex}|\theta, \mathcal{M})$ is denoted as $\mathcal{L}_{\mathcal{O}}(\theta)$. An amount of precision Σ is associated with the measurement of $\mathcal{O}$. It can be, for instance, the covariance matrix of a multivariate normal law, or it can be more generally given by the Fisher information of the likelihood $\mathcal{L}_{\mathcal{O}}$ (as a function of $\mathcal{O}$), $I\{\mathcal{L}_{\mathcal{O}}\}(\mathcal{O}) = \Sigma^{-1}$.

Calling $\mathcal{O}$ an "observable" is somewhat misleading. In fact, it just has to be a quantity constrained by experimental data (or any other exterior piece of information). Note that compared to the δ defined in Sec. II, we let the $\mathcal{O}$ be dimensionful. We emphasize that we restrict ourselves to dimensionful observables and parameters for the sake of

simplicity. The consequence of this choice is to make logarithmic priors appear everywhere. However, the whole approach is general to any prior. The generalization does not present difficulty, and it will be explained in the last subsection. Finally, the reason for the restriction $m \leq n$ is that $m > n$ is similar to $m = n$ from the naturalness point of view. This point will be discussed afterwards, in the last subsection.

A. Probability formulation

Loosely speaking, naturalness is the propensity of a given model to reproduce the experimental observation. Using the notations we adopted, the usual translation of this idea is

(i) "Sensitivity of $\mathcal{O}$ with respect to θ, in the vicinity of a point θ_{ex} belonging to $\mathcal{D}_{ex}$."

This leads to the c measure already presented above,

$$c_a = \max_i \left| \frac{\partial \log \mathcal{O}}{\partial \log \theta_i} \right|_{\theta = \theta_{ex}} \tag{10}$$

or

$$c_b = \sqrt{\sum_i \left(\frac{\partial \log \mathcal{O}}{\partial \log \theta_i} \right)^2} \Bigg|_{\theta = \theta_{ex}}. \tag{11}$$

However, an alternative formulation for naturalness, arguably as intuitive as the first one, is

(ii) "Probability of having $\mathcal{O} = \mathcal{O}_{ex}$ in the model."

It is this second formulation which will be our starting point. As it involves a notion of probability, it necessarily has a Bayesian character. This formulation is translated as the probability

$$p(\mathcal{O} = \mathcal{O}_{ex} | \mathcal{M}, d). \tag{12}$$

But we are interested in this quantity as a function of the hypothesis $(\mathcal{M}, d)$. Taken as such, it is not a probability, but instead a Bayesian evidence, as defined in Eq. (6). Due to the absence of normalization, this evidence alone is not usable. Instead, it has to appear inside a Bayes factor. As a measure of naturalness, we therefore have to consider a Bayes factor which compares our hypothesis $(\mathcal{M}, d)$ to another hypothesis $(\mathcal{M}', d')$:

$$B = \frac{p(\mathcal{O} = \mathcal{O}_{ex} | \mathcal{M}, d)}{p(\mathcal{O} = \mathcal{O}_{ex} | \mathcal{M}', d')}. \tag{13}$$

This well-defined quantity plays the main role in our approach.

Two comments are in order. First, it is clear that such a measure of naturalness has a relative character. In this framework, comparing the naturalness of two models $\mathcal{M}$ and $\mathcal{M}'$ is certainly possible, but an absolute statement about the naturalness of $\mathcal{M}$ has to be done with care. To do so, $\mathcal{M}'$ would have to be defined such that it constitutes an absolute reference. How this can be realized will be discussed further in the section. Second, we emphasize that

the distinction between the model $\mathcal{M}$ and the data d is artificial. Indeed, d could as well be considered as a part of $\mathcal{M}$. It just depends on how $\mathcal{M}$ is defined. It is convenient to keep this separation explicit for the discussion, and to stress that d and d' need not be identical.

Let us now specify the different options available for $(\mathcal{M}', d')$. If one takes the two pieces of data to be identical, $d = d'$, and applies B to two realistic models $\mathcal{M}$ and $\mathcal{M}'$, it provides a measure of the relative naturalness between these models. In particular, it makes sense to apply B to the same model with two different prior densities. For instance, one can compare the naturalness of two different regions of the parameter space $\mathcal{D}_{ex}$. One can even choose punctual priors; that is, priors that select a single point belonging to $\mathcal{D}_{ex}$. In that case, B measures the relative naturalness between two points of $\mathcal{D}_{ex}$. This brings us back to a local measure of naturalness, just like the c measure. Note that the selection of a single point of the parameter space also happens if pieces of data $d^{(\prime)}$ are sufficiently constraining. A necessary condition for that is to have at least as many observables in $d^{(\prime)}$ as parameters in $\mathcal{M}^{(\prime)}$.

Now, let the two models be identical, $\mathcal{M} = \mathcal{M}'$, and let the pieces of data be different, $d \neq d'$. This time, B indicates the change in naturalness induced by going from data d to data d'. Following the literature, this kind of quantity may be dubbed as the "naturalness price" or "fine-tuning price" associated with the change of data. A recent work along this line is Ref. [21].

Finally, how can $\mathcal{M}'$ be defined such that it constitutes an absolute naturalness reference? In Sec. II, we already identified the two limiting cases of total predictivity and infinite fine-tuning. We also identified a threshold in between, when observables $\mathcal{O}_{1...m}$ are input parameters. To define an absolute reference for naturalness, this threshold as well as the limit of total predictivity may be employed. This suggests two ways of defining a reference model: (1) One can consider that $\mathcal{M}'$ is an ideal, fully natural model satisfying $\mathcal{O} = \mathcal{O}_{ex}$ everywhere in its parameter space. We will denote this ideal model as $\mathcal{X}$. (2) Or, $\mathcal{M}'$ may be a hypothetical "puzzle" model, in which the $\mathcal{O}_{1...m}$ would be directly input parameters. This model will be denoted as $\mathcal{P}$. We call this a puzzle model because, from the point of view of sensitivity, it stands at the threshold between predictivity and fine-tuning.

These different possibilities for $\mathcal{M}'$ and their implications will be discussed later in the section. What we have obtained up to now is a well-defined measure of naturalness, under the form of a Bayes factor. Unlike in other approaches, a mapping (Jeffreys' scale) between this measure and our degree of belief exists. The measure is therefore usable, and different applications are possible depending on what one defines as being $(\mathcal{M}', d')$. Starting from now, we will continue the development to show that this probability formulation, instead of being an alternative to the sensitivity formulation, actually embeds it.

B. Apparition of a sensitivity measure

From this point until the end of the section, it is assumed that the measurement of $\mathcal{O}$ is sufficiently precise, such that one can consider the Laplace approximation of the likelihood function. That is, the log-likelihood can be expanded around a maximum $\theta_{\max} \in \mathcal{D}_{\mathrm{ex}}$ as

$$\log \mathcal{L}_\mathcal{O}(\theta) \simeq \log \mathcal{L}_{\max} + \frac{\partial^2 \log \mathcal{L}_\mathcal{O}}{\partial \theta^i \partial \theta^j} \bigg|_{\theta_{\max}}$$
$$\times \frac{(\theta^i - \theta^i_{\max})(\theta^j - \theta^j_{\max})}{2!}, \qquad (14)$$

in which the first-order derivatives of $\mathcal{L}_\mathcal{O}$ vanish, since $\theta_{\max}$ is an extremum. This expansion corresponds to approximating the likelihood as a (multivariate) normal law. Let us rewrite the right-hand term by introducing the Jacobian matrix of the observables $J_{\mathcal{O}ij} = (\frac{\partial \mathcal{O}_i}{\partial \theta_j})$,

$$\frac{\partial^2 \log \mathcal{L}_\mathcal{O}}{\partial \mathcal{O}^i \partial \mathcal{O}^j} J_{\mathcal{O}ik} J_{\mathcal{O}jl} \bigg|_{\theta_{\max}} \frac{(\theta^k - \theta^k_{\max})(\theta^l - \theta^l_{\max})}{2!}. \quad (15)$$

One recognizes in that expression the quantity $\partial^2 \log \mathcal{L}_\mathcal{O}/\partial \mathcal{O}^i \partial \mathcal{O}^j$, which up to a minus sign is the observed Fisher information associated with the $\mathcal{O}$ measurement, $I\{\mathcal{L}_\mathcal{O}\} \times (\mathcal{O})$. We rescale $\mathcal{O}$ by $\mathcal{O}_{\mathrm{ex}}$ to make appear a dimensionless Jacobian and a dimensionless Fisher information associated with $\mathcal{O}/\mathcal{O}_{\mathrm{ex}}$, such that Eq. (15) becomes

$$\frac{\partial^2 \log \mathcal{L}_\mathcal{O}}{\partial \log \mathcal{O}^i \partial \log \mathcal{O}^j} J_{\log \mathcal{O}ik} J_{\log \mathcal{O}jl}|_{\theta_{\max}} \frac{(\theta^k - \theta^k_{\max})(\theta^l - \theta^l_{\max})}{2!}. \tag{16}$$

The dimensionless Fisher information $-\partial^2 \log \mathcal{L}_\mathcal{O}/\partial \log \mathcal{O}^i \partial \log \mathcal{O}^j$ describes the amount of relative uncertainty associated with $\mathcal{O}$. We will denote it as Σ^{-1} from now on.

Given this expansion of $\mathcal{L}_\mathcal{O}$, we can reconsider our central quantity, the Bayesian evidence $p(\mathcal{O} = \mathcal{O}_{\mathrm{ex}}|\mathcal{M}, d)$. This evidence can be written as a continuous sum over all the values of the parameters:

$$p(\mathcal{O} = \mathcal{O}_{\mathrm{ex}}|\mathcal{M}, d) = \int_\mathcal{D} \mathcal{L}_\mathcal{O}(\theta) p(\theta) d\theta. \tag{17}$$

It is the average of $\mathcal{L}_\mathcal{O}(\theta)$ weighted by the prior density of the model parameters $p(\theta) = p(\theta|\mathcal{M}, d)$. We will denote the Fisher information associated with this prior density as $I\{p(\theta)\} = |V|^{-1}$, and designate $|V|^{1/2}$ as the "prior volume."

Provided that the likelihood is informative with respect to the prior, Eq. (17) takes the form

$$p(\mathcal{O} = \mathcal{O}_{\mathrm{ex}}|\mathcal{M}, d) = \mathcal{L}_{\max} \frac{|\Sigma|^{1/2}}{|V|^{1/2}} \int_{\mathcal{D}_{\mathrm{ex}}} \frac{1}{C} d\sigma(\theta). \quad (18)$$

Here, $d\sigma(\theta)$ is the induced integration measure on the manifold $\mathcal{D}_{\mathrm{ex}}$, and C is the Jacobian factor,

$$C = |\det(J_{\log \mathcal{O}} J^t_{\log \mathcal{O}})|^{1/2}. \tag{19}$$

From the point of view of Fisher information, C measures how much information about the parameters θ is contained in $\mathcal{O}/\mathcal{O}_{\mathrm{ex}}$ regardless of the uncertainty Σ. The interesting fact is that C is a generalized version of the sensitivity measure c, such that Eq. (18) makes the link between the two formulations of naturalness.

Some remarks are in order. Firstly, Eq. (18) holds in the limit where $C|V|^{1/2} \gg |\Sigma|^{1/2}$ over $\mathcal{D}_{\mathrm{ex}}$. We will designate the likelihood as informative when this condition is fulfilled. When the condition is not satisfied, the overlap between the likelihood and the prior has to be taken into account properly, and the Bayesian evidence tends to $\mathcal{L}_{\max}$. Secondly, we emphasize that $J_{\log \mathcal{O}} J^t_{\log \mathcal{O}}$ is indeed an $m \times m$ matrix. Its size does not depend on the number of parameters, but on the number of observables. Thirdly, by choosing a punctual prior, or if the other data d are sufficiently constraining, $\mathcal{D}_{\mathrm{ex}}$ is reduced to a single point θ_0, and the integral $\int_{\mathcal{D}_{ex}} C^{-1} d\sigma(\theta)$ is reduced to $C^{-1}|_{\theta_0}$. In such a situation, we get closer from the sensitivity definition, which is a local measure.

We also emphasize that the derivatives which appear within C depend on the choice of prior. Indeed, in all generality, these derivatives are performed with respect to the "prior repartition function" $G(\theta)$, defined such that $p(\theta)d\theta = dG(\theta)$. Let us illustrate this fact by considering a single observable and a flat prior on all parameters, restricted to the volume $[a_1, b_1] \times \cdots \times [a_n, b_n]$. The prior volume is $|V|^{1/2} = (b_1 - a_1) \times \cdots \times (b_n - a_n)$, and the Jacobian factor is $C = \sqrt{\sum_i |\partial \log \mathcal{O}/\partial \theta_i|^2}$. Instead, if one chooses a logarithmic prior for all parameters $p(\theta_i) \propto \theta_i^{-1}$, the prior volume becomes $|V|^{1/2} = (\log b_1 - \log a_1) \times \cdots \times (\log b_n - \log a_n)$, and the Jacobian factor becomes $C = \sqrt{\sum_i |\partial \log \mathcal{O}/\partial \log \theta_i|^2}$. The derivatives in C are then made with respect to $\log \theta$ instead of θ. In the case of dimensionful parameters, the choice of the logarithmic prior has a particular meaning, because it is the more objective prior.

Through these several remarks, we can finally state that $\int_{\mathcal{D}_{ex}} C^{-1} d\sigma(\theta)$ reduces to the c measure of the sensitivity formulation, provided that one considers a single observable, a single point in $\mathcal{D}_{\mathrm{ex}}$, and that one gives logarithmic priors to the parameters. It is more precisely the expression c_b, used in Ref. [13], which is exactly reproduced. The measure c_a is an approximation of c_b when one of the components of the gradient dominates.

Interestingly, the average $\int_{\mathcal{D}_{ex}} C^{-1} d\sigma(\theta)$ has been proposed in Ref. [7], in an attempt to normalize the c measure. In our approach this quantity arises naturally, and we also see that in itself it does not help to interpret the c measure. On the other hand, the use of the volume of parameter space has been proposed in Ref. [8], in an attempt to build an alternative measure. These different ideas, somewhat

intuitive as such, become rigorously usable once they appear together through Eq. (18). Our approach also justifies Bayesian studies which introduce C^{-1} as a "naturalness prior." This is not new; it was already explained in Ref. [5]. However, we add that the prior of the parameters has to correspond to the derivatives made in C to keep the approach consistent.

The factor C is a generalization of the c measure. Among other things, it tells us the information content of several possibly correlated observables. Let us consider the case of two observables. In this case, C is nothing but the norm of the wedge product of the gradients,

$$C = \|\nabla \log \mathcal{O}_1 \wedge \nabla \log \mathcal{O}_2\|, \qquad (20)$$

which is also

$$C = (\|\nabla \log \mathcal{O}_1\|^2 \|\nabla \log \mathcal{O}_2\|^2 - (\nabla \log \mathcal{O}_1 . \nabla \log \mathcal{O}_2)^2)^{1/2}. \qquad (21)$$

It is instructive to discuss the behavior of C depending on the correlation between the two observables which is induced by the model. One can rewrite C as

$$C = C_1 C_2 \sqrt{1 - \rho^2}, \qquad (22)$$

where C_1 and C_2 are the one-dimensional sensitivities and ρ is the correlation in the model, defined by

$$\rho = \frac{|\nabla \log \mathcal{O}_1 . \nabla \log \mathcal{O}_2|}{\|\nabla \log \mathcal{O}_1\| \|\nabla \log \mathcal{O}_2\|}. \qquad (23)$$

If the observables are independently predicted, the two gradients are orthogonal, and thus $\rho = 0$. Equation (21) reduces in that case to the product of the one-dimensional C measures. On the contrary, if the observables are correlated within the model, one has $\rho > 0$, and C decreases. In the Bayesian point of view, this should be interpreted as the fact that it is more economical for a model to predict correlated observables than independent observables. One may be worried that Eq. (21) tends to zero in the limit of total correlation, when the two observables are linearly dependent. However, the formula does not apply in that limit. Indeed, recall that the condition for having informative data is $C|V|^{1/2} \gg |\Sigma|^{1/2}$. It translates here as an upper bound on the correlation ρ,

$$\rho \ll \sqrt{1 - \frac{|\Sigma|}{|V|(C_1 C_2)^2}}. \qquad (24)$$

For instance, for a pair of Gaussian measurements, one has $|\Sigma|^{1/2} = \sigma_1 \sigma_2 \sqrt{1 - \rho_{\text{exp}}^2}$. When Eq. (24) is not satisfied, the correlation is too large, such that the two observables are not separately informative. That is, instead of two constraints, the model effectively feels a single constraint $\hat{\mathcal{O}} \equiv \mathcal{O}_1 \sim \propto \mathcal{O}_2$. Instead of Eq. (21), it is then a one-dimensional sensitivity associated with $\hat{\mathcal{O}}$ which appears in the Bayesian evidence [Eq. (18)]. This discussion

illustrates also the fact that C taken outside of the formula Eq. (18) can induce misunderstandings, and has to be interpreted with care.

At this point, puzzling observations can be made about priors and the meaning of logarithms which appear everywhere. What is, after all, the reason for having $\log \mathcal{O}$ in C? Is it for the sake of making C dimensionless? Or is it for the sake of measuring a relative variation, as we naively stated in Sec. II? Here we assumed that $\mathcal{O}$ is dimensionful. By doing so, we avoided this discussion in a first time, since rescaling $\mathcal{O}$ by $\mathcal{O}_{\text{ex}}$ makes C both dimensionless and measuring a relative variation. Also, although the log in $\partial \log \mathcal{O}$ is suggestive of an objective prior, this remains just a way of writing $\partial \mathcal{O}/\mathcal{O}_{\text{ex}}$. And, anyway, speaking about a prior for $\mathcal{O}$ actually does not makes sense for the moment, as $\mathcal{O}$ is determined by the parameters. These observations will be resolved when examining the Bayes factor $B_{\mathcal{MP}}$ in the next subsection.

To summarize, we find that the sensitivity formulation of naturalness turns out to be embedded in the probability formulation. The c measure, Eq. (11), turns out to be a particular case of the factor C, arising in the Bayesian evidence [Eq. (18)]. The only assumption made to obtain this result is the Laplace approximation, i.e., that the likelihood function can be reduced to a normal law. We will now examine the different Bayes factors B that can be constructed, and the role taken by the C measure.

C. The different versions of B

1. $B_{\mathcal{MX}}$

As a warmup, let us examine the Bayes factor comparing a model $\mathcal{M}$ to the fully natural model $\mathcal{X}$:

$$B_{\mathcal{MX}} = \frac{p(\mathcal{O} = \mathcal{O}_{\text{ex}} | \mathcal{M})}{p(\mathcal{O} = \mathcal{O}_{\text{ex}} | \mathcal{X})}. \qquad (25)$$

By definition, $\mathcal{X}$ satisfies $\mathcal{O} = \mathcal{O}_{\text{ex}}$ in all its parameter space. The evidence of this ideal model is thus $p(\mathcal{O} = \mathcal{O}_{\text{ex}} | \mathcal{X}) = \mathcal{L}_{\text{max}}$. Recall that $\mathcal{L}_{\text{max}}$ is an overall normalization constant, which will be canceled once we consider the ratio of evidences. Assuming that the piece of data $\mathcal{O} = \mathcal{O}_{\text{ex}}$ is informative for $\mathcal{M}$, $B_{\mathcal{MX}}$ takes the form

$$B_{\mathcal{MX}} = |\Sigma|^{1/2}/|V|^{1/2} \int_{\mathcal{D}_{\text{ex}}} C^{-1} d\sigma(\theta). \qquad (26)$$

Clearly, since $\mathcal{M}$ cannot be more natural than $\mathcal{X}$, $B_{\mathcal{MX}}$ cannot be larger than 1. At most, it can tend to 1, if $\mathcal{M}$ tends to be an ideal model like $\mathcal{X}$. Equation (26) is not valid in this limit, as it implies that the data are not informative for $\mathcal{M}$ anymore. Let us interpret what Eq. (26) is telling us as a Bayes factor. One sees that $B_{\mathcal{MX}}$ decreases with $|\Sigma|$. This is because when the constraint $\mathcal{O} = \mathcal{O}_{\text{ex}}$ becomes more precise, $\mathcal{M}$ is penalized, but not $\mathcal{X}$. Also, $B_{\mathcal{MX}}$ decreases as $|V|$ increases, because it penalizes the waste of parameter space of $\mathcal{M}$ excluded

by $\mathcal{O} = \mathcal{O}_{\text{ex}}$. Finally, $B_{\mathcal{M}\mathcal{X}}$ also decreases with the sensitivity C. C measures the amount of information that $\mathcal{O}$ carries about the parameters of $\mathcal{M}$. The larger C is, the more $\mathcal{O}$ contains information, and the stronger the constraint $\mathcal{O} = \mathcal{O}_{\text{ex}}$ is for $\mathcal{M}$, regardless of the experimental uncertainty.

$B_{\mathcal{M}\mathcal{X}}$ is certainly useful to understand the content of the Bayesian evidence. On the other hand, it is *a priori* not very useful in a concrete application, as it will just tell us that the model under consideration is worse than the ideal model. Does it provide a good basis for giving an absolute interpretation to C? The interpretation of C would inevitably depend on $|\Sigma|$. It would be necessary that $|\Sigma|$ be intrinsically bounded from below, independently of the details of the experimental observations. This can actually happen, for instance, in quantum theories when observables do not commute. But as far as we know, nothing of this kind happens in a domain of physics having naturalness issues.

2. $B_{\mathcal{M}\mathcal{P}}$

Consider now the Bayes factor comparing the model $\mathcal{M}$ to a model $\mathcal{P}$ in which the observables $\mathcal{O}_{1\ldots m}$ (or any linear combination of them) are directly input parameters, such that $C = 1$. It is defined by

$$B_{\mathcal{M}\mathcal{P}} = \frac{p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{M})}{p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{P})}. \qquad (27)$$

We recall that the model $\mathcal{P}$ is dubbed a "puzzle" model, since from the point of view of sensitivity, it is at the limit between predictivity and fine-tuning. In this case, the prior density for $\mathcal{O}$ will enter the game, as $\mathcal{O}$ itself is an input parameter of $\mathcal{P}$.

The Bayesian evidence of $\mathcal{P}$ is

$$p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{P}) = \mathcal{L}_{\text{max}} \frac{|\Sigma|^{1/2}}{|V_{\mathcal{O}}|^{1/2}}. \qquad (28)$$

$|V_{\mathcal{O}}|^{1/2}$ is the prior volume associated with the parameter $\mathcal{O}$. In this expression, one can introduce the ratio $\mathcal{O}/\mathcal{O}_{\text{ex}}$, as we did for the Bayesian evidence of $\mathcal{M}$, $p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{M})$, given in Eq. (18). By doing so, Σ has the same relative uncertainty $\Sigma = -\partial^2 \log \mathcal{L}_{\mathcal{O}}/\partial \log \mathcal{O}^i \partial \log \mathcal{O}^j$ as the one which appears in $p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{M})$. The two $|\Sigma|^{1/2}$ thus cancel in $B_{\mathcal{M}\mathcal{P}}$, such that

$$B_{\mathcal{M}\mathcal{P}} = \frac{|V_{\mathcal{O}}|^{1/2}}{|V|^{1/2}} \int_{\mathcal{D}_{\text{ex}}} \frac{1}{C} d\sigma(\theta). \qquad (29)$$

With this choice, $C \propto \partial \log \mathcal{O}/\partial \ldots$, and the prior volume $V_{\mathcal{O}}$ is dimensionless. $V_{\mathcal{O}}$ is, however, not determined. To do so, the prior of $\mathcal{O}$ would need to be specified.

To go further and specify a particular prior for $V_{\mathcal{O}}$, it is necessary to impose a condition referring to some principle. Interestingly, there are two different principles, leading to two different conditions, which lead to the same result.

Firstly, one can invoke the principle of indifference, introduced in Sec. III. Applied to a dimensionful quantity, it states that our *a priori* degree of belief should not depend on the unit scale. This is translated as the invariance of $p(\mathcal{O}|\mathcal{P})$ under the transformation $\mathcal{O} \to \mathcal{O} \times b$, which imposes the logarithmic prior $p(\mathcal{O}|\mathcal{P}) \propto \mathcal{O}^{-1}$.

But there is a second principle which gives the same result. In this section, we restricted our discussion to a dimensionful $\mathcal{O}$. However, there is no specification made about the actual dimension of $\mathcal{O}$. It seems legitimate to require that the whole approach lead to the same outcome whatever the dimension of $\mathcal{O}$ is. Said differently, we require that the measure of naturalness not depend on a redefinition of $\mathcal{O}$ changing its dimension. We will designate this property as the "consistency" of the naturalness measure. It is translated as the invariance of $B_{\mathcal{M}\mathcal{P}}$ under the transformation $\mathcal{O} \to \mathcal{O}^a$. The consequence of imposing this condition is once again that $p(\mathcal{O}|\mathcal{P})$ is the logarithmic prior $p(\mathcal{O}|\mathcal{P}) \propto \mathcal{O}^{-1}$.

We thus find that the logarithmic prior is independently motivated by the principle of indifference and by the consistency of the measure. This consistency condition is a kind of principle of indifference, applied to a Bayes factor instead of a probability. Depending on the point of view adopted, one can claim either that the consistency of the measure leads automatically to an objective prior, or that the principle of indifference leads automatically to a consistent measure. In any case, $B_{\mathcal{M}\mathcal{P}}$ is finally invariant under the transformation $\mathcal{O} \to b \times \mathcal{O}^a$.

Provided that prior volumes of both $\mathcal{M}$ and $\mathcal{P}$ can be bounded, $B_{\mathcal{M}\mathcal{P}}$ provides a kind of absolute scale to C. As expected, the puzzle model $\mathcal{P}$ plays the role of a reference in terms of sensitivity, to which $\mathcal{M}$ can be compared. Once the volumes are determined, C is directly related to Jeffreys' scale. The interpretation of C in terms of degree of belief does not depend on the definition of $\mathcal{O}$. Indeed, any redefinition of $\mathcal{O}$ is accompanied by a change in $|V_{\mathcal{O}}|^{1/2} = \int d \log \mathcal{O}$, such that the interpretation of C always remains the same. For instance, in the gauge hierarchy problem, it does not matter anymore whether we take $\mathcal{O} = m_Z$ or $\mathcal{O} = m_Z^2$, because this is compensated for in the prior volume of m_Z, $\int d \log m_Z \to \int d \log m_Z^2$. This consistency property resolves one of the issues raised in Sec. II.

3. Relative naturalness

Finally, let us compare two hypotheses, $(\mathcal{M}_0, d_0)$ and $(\mathcal{M}_1, d_1)$. We make the assumption that the piece of data $\mathcal{O} = \mathcal{O}_{\text{ex}}$ is informative for both models. The Bayes factor

$$B_{01} = \frac{p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{M}_0, d_0)}{p(\mathcal{O} = \mathcal{O}_{\text{ex}}|\mathcal{M}_1, d_1)}, \qquad (30)$$

takes the form

$$B_{01} = \frac{|V_1|^{1/2}}{|V_0|^{1/2}} \int_{\mathcal{D}_{0\mathrm{ex}}} C_0^{-1} d\sigma(\theta) \left(\int_{\mathcal{D}_{1\mathrm{ex}}} C_1^{-1} d\sigma(\theta) \right)^{-1}.$$

$$(31)$$

Here, one can see clearly that this naturalness measure puts in balance both the prior volumes and the sensitivities integrated over the parameter space. $\mathcal{M}_0$ and $\mathcal{M}_1$ can be two different models, or the same model with different priors, or associated with different data d_0 and d_1. When the two models are the same, one has $C_0 = C_1$, but the two domains of integration are still different. If the two hypotheses differ only by the data $d_0 \neq d_1$, B gives the "naturalness price" between the two pieces of data. For example, d_0 could be the pre-LHC constraint on SUSY particle masses, and d_1 the constraint once LHC measurements are taken into account.

Within the same model $\mathcal{M}$, one can make the choice of punctual priors, which select two different points θ_0, θ_1 of $\mathcal{D}_{\mathrm{ex}}$. In that case, the prior volumes cancel, leaving only the Bayes factor

$$B_{01} = \frac{C_1}{C_0}. \qquad (32)$$

The Bayes factor, then, is simply reduced to the comparison of the sensitivities. This quantity shows clearly that the relative sensitivity within a model has to be interpreted on the basis of Jeffreys' scale. It is also true for the usual c measure, which is a particular case of C. This finishes resolving all the issues raised in Sec. II.

D. Generalization and comments

To sum up, the Bayes factor $B_{\mathcal{MP}}$ provides a handle on the absolute interpretation of C and sets the functional form of all quantities, once either the consistency or the indifference principle is required. The relative versions of B contain various possibilities of application, some of them being reminiscent of previous work done in the literature. We will now discuss generalization, absolute naturalness, and some implications of this approach.

1. The general case

First of all, let us explain why we keep a number m of observables $\mathcal{O}_{1\ldots m}$ smaller than or equal to the number of parameters n. If one has n observables nonproportional to each other, the model is fully constrained; i.e., $\mathcal{D}_{\mathrm{ex}}$ has dimension zero. The likelihood function gets in that case one or several maxima, with an uncertainty Σ associated with each of them. If one adds a new constraint, the effect will be to increase the precision, i.e., to reduce Σ, and possibly to decrease the maximum of the likelihood, if this new constraint is not in agreement with the n previous. But from the point of view of the C measure, this new constraint is necessarily reduced around the maximum to a linear combination of the n others. For that reason, the

contribution of this new constraint vanishes in the determinant contained in the Jacobian factor C, and thus cannot influence the sensitivity. Therefore, for the purpose of the naturalness study, it is sufficient to keep $m \leq n$.

The second assumption we made in the beginning of Sec. IV was that our observables and parameters were dimensionful. We found that either applying the indifference principle or requiring consistency of the naturalness measure leads to the invariance of $p(\mathcal{O})$ under $\log \mathcal{O} \to \log \mathcal{O} + b$ and of $B_{\mathcal{MP}}$ under $\log \mathcal{O} \to a \times \log \mathcal{O} + b$, where a is an $m \times m$ matrix and b is an m-vector. These conditions imply the use of the logarithmic prior, such that $C = |\det(J_{\log \mathcal{O}} J_{\log \mathcal{O}}^t)|^{1/2}$, where $J_{\log \mathcal{O}} = \partial \log \mathcal{O}_i / \partial \log \theta_j$ and $V_{\mathcal{O}} = \int d^m \log \mathcal{O}_i$, $V = \int d^n \log \theta_j$. All these properties are the consequences of considering that a transformation law, a change in unit scale, is irrelevant for our degree of belief about the problem. Let us now go to the general case by considering arbitrary, possibly dimensionless observables and parameters. All the results can be generalized, provided the existence of an irrelevant transformation. Let us assume that the transformation $G(\mathcal{O}) \to G(\mathcal{O}) + b$ and the transformation $H(\theta) \to H(\theta) + c$ do not modify our degree of belief. Then, the naturalness measure is invariant under $G(\mathcal{O}) \to a \times G(\mathcal{O}) + b$; the sensitivity takes the form $C = |\det(J_{G(\mathcal{O})} J_{G(\mathcal{O})}^t)|^{1/2}$, where $J_{G(\mathcal{O})} = \partial G(\mathcal{O})_i / \partial H(\theta)_j$; and the prior volumes are $V_{\mathcal{O}} = \int d^m G(\mathcal{O}_i)$, $V = \int d^n H(\theta_j)$.

What we stated above is based on the existence of a continuous irrelevant transformation. However, other kinds of conditions, possibly less obvious, can also be found. For instance, when a theory is isomorphic to itself under a duality transformation, it is possible to find the objective priors of parameters transforming nontrivially under the duality.

Finally, it is important to recall that results obtained in Bayesian statistics depend to some extent on the parametrization of the problem. The choice of parametrization is somehow intricate with the choice of prior for the parameters. The indifference principle (Sec. III) plays a crucial role with respect to this issue. It allows us to minimize the amount of information contained in the priors—or, said differently, it helps us to find a preferred, objective parametrization. For example, it happens that a dimensionless parameter, whose objective prior is unknown, can be seen as a ratio of two dimensionful parameters. This is, for instance, the case of $\tan\beta \equiv v_u/v_d$ in the minimal supersymmetric standard model (MSSM). Given that the objective prior of dimensionful parameters is known, this provides the (nontrivial) objective prior of the dimensionless parameter. Or, equivalently, one can choose these dimensionful parameters as input (we refer to Ref. [22] for an application to the MSSM).

By construction, our Bayesian approach to naturalness inherits all of these features. However, an extra subtlety is that there is freedom of parameterization on both the

parameters θ_i and the observable $\mathcal{O}$. Without referring to the indifference principle, there is no means to favor a particular parametrization, and the naturalness measure is dependent on this parametrization. Once applied, the indifference principle provides the objective priors for both the θ_i and $\mathcal{O}$ (as discussed in the analysis of $B_{\mathcal{MP}}$). As a result, we end up with a unique naturalness measure, depending only on the transformation properties associated with the indifference principle. The general result is given above in this section. Speaking more formally, the indifference principle defines an equivalence class among the parametrizations, and the naturalness measure turns out to be an invariant of this equivalence class. In the usual example of dimensionful parameters, the transformation defining the equivalence class is $\mathcal{O} \rightarrow b \times \mathcal{O}^a$. This is nothing but the 2D set of quantities of arbitrary dimensions. As a consequence, whatever the dimension of $\mathcal{O}$—e.g., $\mathcal{O} \equiv m_Z, m_Z^2$, or $3 \times m_Z^{100}$—the naturalness measure remains the same.

2. About absolute naturalness

The puzzle model $\mathcal{P}$ provides a reference in terms of naturalness. It is the only sensible reference we are able to find. But how can it be defined in practice? Let us try to do this for the gauge hierarchy problem.

The Bayes factor associated with this problem is

$$B_{\mathcal{MP}} = \frac{p(m_Z = m_{Z\text{ex}}|\mathcal{M}, d)}{p(m_Z = m_{Z\text{ex}}|\mathcal{P}, d')}. \qquad (33)$$

The pieces of data d and d' have to be identical, as our goal is not to compare different data. Which information is contained in d? By construction, in our approach, all experimental information available is split into two categories. There is the one that contributes to indicate what the electroweak scale is, which is called $m_Z = m_{Z\text{ex}}$, and the one that doesn't, which is called d. With only the knowledge of d, one would know, for example, the strength of gravity and gauge interactions and the fermion and hadron masses, but not the electroweak boson masses or the Fermi constant. Such a situation is, of course, impossible to imagine in practice, but here we are simply splitting a set of existing information, regardless of the way it was obtained.

Now, what should $\mathcal{P}$ be? It is a model which predicts data d and has m_Z both as an input and an output. We can imagine that it is a kind of quantum field theory in which the weak scale does not receive any quadratic corrections, for some unknown reason. What should be the prior volume of m_Z? We know from both the indifference principle and the consistency of the measure that m_Z should have a logarithmic prior. The bounds of this density remain to be found. Given that d contains the knowledge of gravity, $\mathcal{P}$ has a cutoff at the Planck mass, so $m_Z \leq M_{\text{Pl}}$. On the other hand, as d contains the quark masses, m_Z is bounded from below due to the unitarity of quark scattering by weak currents (see Ref. [23]), which implies roughly $m_Z \gtrsim 10$ GeV. The prior volume V_{m_Z} in the model $\mathcal{P}$ is therefore $V_{m_Z} = \log(M_{\text{Pl}}/10 \text{ GeV}) \approx 40.0$. This completes the definition of $\mathcal{P}$. Using Laplace approximation, the Bayes factor is

$$B_{\mathcal{MP}} = \frac{|V_{m_Z}|^{1/2}}{|V|^{1/2}} \int_{\mathcal{D}_{\text{ex}}} \frac{1}{C_{m_Z}} d\sigma(\theta), \qquad (34)$$

where $C_{m_Z} \propto \partial \log m_Z/\partial \ldots$, and $|V|$ is the prior volume of $\mathcal{M}$. With this equation, for any choice of $\mathcal{M}$, $\int_{\mathcal{D}_{\text{ex}}} C_{m_Z}^{-1} d\sigma(\theta)$ is equal to Jeffreys' scale up to a known constant. Therefore, we get the absolute interpretation of the sensitivity measure.

One may or not be satisfied with this approach. In any case, it illustrates that it is not so obvious to define $\mathcal{P}$ in practice. This, however, does not take away the general results obtained by studying $B_{\mathcal{MP}}$.

3. Second-order fine-tuning

When considering a naturalness map, the following interrogation often appears. The interest of a naturalness map is to select regions of the parameter space which have a relatively low fine-tuning. But suppose that a very tiny region of the parameter space has a very small C, while C is sensibly larger around, in at least one direction. Selecting this tiny region and discarding the zone around it would be itself an action of fine-tuning. We will designate such an issue as a "second-order fine-tuning." How is this taken into account in our framework?

It is easy to guess that there is a relation to the choice of punctual priors, which select single points of $\mathcal{D}_{\text{ex}}$. Indeed, if $1/C$ were integrated around the tiny zone with small C, the particularity of that zone would disappear. Formally, the action of selecting regions with low fine-tuning corresponds to imposing a prior such that $C \leq C_{\text{min}} + \Delta C$, where ΔC is a level of tolerance, and C_{min} is a minimal value. One can construct a Bayes factor comparing two regions $\mathcal{D}_{0\,\text{ex}}, \mathcal{D}_{1\,\text{ex}}$ of the parameter space, containing the minima $C_{0\,\text{min}}, C_{1\,\text{min}}$, respectively, and with the requirement of an upper bound on C. Several versions can be built, depending, for example, on whether C_{min} is considered a common value, or if $C_{\text{min}} = C_{1,0\,\text{min}}$, respectively. These versions correspond to different reasonings. Provided that C^{-1} can be approximated over the domain considered, such Bayes factors can be computed analytically.

For example, let us consider the Bayes factor comparing two regions $\mathcal{D}_{0\,\text{ex}}, \mathcal{D}_{1\,\text{ex}}$ of the parameter space, containing minima $C_{0\,\text{min}}, C_{1\,\text{min}}$ which are not on the boundaries. We impose the condition $C \leq C_{\text{min}} + \Delta C$, where C_{min} is a common value. It can be $\min(C_{0,\text{min}}, C_{1,\text{min}})$ or a smaller value. It does not matter, since it will not appear in the final result. These two domains are denoted as $\mathcal{D}'_{0\,\text{ex}}, \mathcal{D}'_{1\,\text{ex}}$. The Bayes factor reads

$$B_{01} = \int_{\mathcal{D}'_{0\,\text{ex}}} C^{-1} d\sigma(\theta) \left(\int_{\mathcal{D}'_{1\,\text{ex}}} C^{-1} d\sigma(\theta) \right)^{-1}. \qquad (35)$$

When ΔC is not too large, one can take the Laplace approximation of C around $C_{0\min}$ and $C_{1\min}$. In that limit, the integration can be done, and one obtains the Hessian of $\log C$

$$H = \det(\nabla_i \nabla_j \log C)|_{C_{\min}}. \tag{36}$$

As by assumption both boundaries $\partial \mathcal{D}'_{i\,\text{ex}}$ are inside the corresponding boundaries $\partial \mathcal{D}_{i\,\text{ex}}$, the Bayes factor reduces to

$$B_{01} = \frac{C_{1\min}}{C_{0\min}} \frac{H_1^{1/2}}{H_0^{1/2}}. \tag{37}$$

We can see that a new factor $H_1^{1/2}/H_0^{1/2}$ appears in addition to $C_{1\min}/C_{0\min}$. This is the quantity which renders account for the second-order fine-tuning.

More generally, for Bayes factors where the minima are on the boundaries, there will be contributions of the form

$$B_{01} \propto \prod_i \left| \frac{\partial C}{\partial \theta_i} \right|_{C_{1\min}} \left| \frac{\partial \theta_i}{\partial C} \right|_{C_{0\min}} \tag{38}$$

coming in. Terms in Eqs. (37) and (38) provide a comparison of the steepness of C around the two minima. This is properly quantified and interpreted in terms of naturalness with the Bayesian approach.

4. The top Yukawa in the gauge hierarchy problem

A recurring question about the gauge hierarchy problem is whether or not the top quark Yukawa coupling y_t should be considered an input parameter, such that the derivative $\partial m_Z / \partial y_t$ appears in the C_{m_Z} measure. On one hand, one can think of y_t as a simple constant, and not a parameter. In that case, it should not appear in C_{m_Z}. On the other hand, one can think of it as an input parameter, fixed by the experiment. In that case, y_t must appear in the C_{m_Z} measure. So what is the right point of view? Surprisingly, it is the first proposition which makes sense. To understand this, we have to examine more carefully the second proposition.

Indeed, the choice of considering y_t as an input parameter or a constant is just a matter of viewpoint, and should not modify the information content of our study. This implies that if y_t is taken as an input parameter, one has to add to the set of experimental constraints the top quark mass measurement, $m_t = m_{t\,\text{ex}}$. But the observables m_Z and m_t are not independent in the model. Therefore, to study naturalness of the gauge hierarchy problem, they need to be simultaneously taken into account. It is thus the combined sensitivity C_{m_Z,m_t} which has to be used when y_t is seen as an input parameter.

At this point, it is instructive to wonder what is the common fine-tuning associated with a generic observable $\mathcal{O}$ and an observable Θ which is directly an input parameter. The set of input parameters is denoted as $p_i = (\theta_j, \Theta)$. We assume a logarithmic prior for all quantities for concreteness. If $\mathcal{O}$ and Θ are independent in the model, the common sensitivity

$$C_{\mathcal{O},\Theta} = \left\| \frac{\partial \log \mathcal{O}}{\partial \log p_i} \wedge \frac{\partial \Theta}{\partial \log p_i} \right\| \tag{39}$$

factorizes and reduces to

$$C_{\mathcal{O}} = \left\| \frac{\partial \log \mathcal{O}}{\partial \log \theta_i} \right\|, \tag{40}$$

given that $C_\Theta = 1$. But what happens when the two observables are correlated? It turns out that the answer is the same. Whatever the "puzzle" observable Θ is, the sensitivity always reduces to $C_{\mathcal{O},\Theta} = C_{\mathcal{O}}$. We emphasize that, although all priors are chosen to be logarithmic there, these kinds of results hold whatever the priors are.

Let us come back to the top Yukawa and the C_{m_Z,m_t} measure. The previous remark does not apply directly, because the observable is not y_t, but rather the top mass $m_t = y_t \times v$. Thus, y_t does not play the same role as Θ. However, the outcome will in fact be the same. We denote the set of input parameters as $p_i = (\theta_j, y_t)$. The objective prior of m_t is logarithmic, and it implies that the prior of y_t is also logarithmic. We also assume for simplicity that the priors of the θ_j are logarithmic. The sensitivity is then

$$\left\| C_{m_Z,m_t} = \frac{\partial \log m_Z}{\partial \log p_i} \wedge \frac{\partial \log(v y_t)}{\partial \log p_i} \right\|. \tag{41}$$

As m_Z is directly related to v, the gradients $\partial \log m_Z / \partial \log p_i$ and $\partial \log v / \partial \log p_i$ are collinear. The v contribution therefore vanishes in the sensitivity measure. The remaining part contains y_t, which plays the same role as Θ in the previous paragraph. As a consequence, the sensitivity reduces to C_{m_Z}, without the parameter y_t:

$$C_{m_Z,m_t} = \left\| \frac{\partial \log m_Z}{\partial \log \theta_i} \right\|. \tag{42}$$

This result holds, whatever the priors are. Thus, to reply to the initial question, the second proposition in fact gives the same result as the first proposition, after a careful examination: y_t should not appear in C_{m_Z}.

5. Consequences of LHC searches

The existence of a scalar resonance whose properties are roughly compatible with that of a Higgs boson has been established beyond reasonable doubt at the LHC [24,25]. The mass of this new state is a stringent constraint on many models of new physics. Some of them are almost excluded, barring some very specific choices of parameters. As a result, this constitutes a new naturalness problem associated with the Higgs mass constraint. It would therefore be particularly appropriate to study the fine-tuning related to the Higgs, either inside the parameter space of a single model, or comparing two different models. The naturalness measure to use for such study is

$$B_{01} = \frac{p(m_h = m_{h\,\text{ex}}, m_Z = m_{Z\,\text{ex}})|\mathcal{M}_0)}{p(m_h = m_{h\,\text{ex}}, m_Z = m_{Z\,\text{ex}})|\mathcal{M}_1)}. \quad (43)$$

We emphasize once again that the two observables m_h and m_Z, both independently responsible for some amount of fine-tuning, should not be treated separately, because their predictions are correlated in the models.

On the other hand, searches at the LHC and other experiments have not, up to now, shown conclusive evidence of existence of Beyond Standard Model physics. As the idea of new physics (NP) at the multi-TeV scale is in part motivated by the gauge hierarchy problem, one can wonder to what extent NP models are more natural than the Standard Model, given the increasing exclusion limits. Let us answer this question in a very simplified way. For concreteness, we assume the SM to be valid up to the Planck scale. As the origin of the gauge hierarchy problem is an issue of cancellations between square mass parameters, we will consider a one-parameter "model" embedding this property. That is, we just define the EW scale as given by $m_Z^2 = M_{\text{Pl}}^2(1 - \delta)$. We also consider a BSM model suppressing the quadratic corrections to the EW scale at a scale $\tilde{M} < M_{\text{Pl}}$. The EW scale is thus given by $m_Z^2 = \tilde{M}^2(1 - \delta)$ in this model. Picking similar priors for the δ in each hypothesis, the naturalness measure turns out to be

$$B_{\text{NP,SM}} \approx \frac{M_{\text{Pl}}^2}{\tilde{M}^2}. \quad (44)$$

We can see that, unless $\tilde{M}$ is close to M_{Pl}, this ratio indicates an extremely strong fine-tuning of the SM, far beyond the typical value of 150 indicated in Jeffreys' scale. If, for instance $\tilde{M} \approx 100$ TeV, one gets $B_{\text{NP,SM}} \approx 10^{26}$. We emphasize that, for the sake of comparing the SM to a NP model improving substantially the gauge hierarchy problem, there is no need to set up a more evolved analysis. This estimation embeds the large leading contribution, which flushes away any other subleading effects.

V. THE GAUGE HIERARCHY PROBLEM AND NEUTRALINO DARK MATTER IN THE cMSSM

In this section, we apply our results to a concrete problem. We choose to study the naturalness of a classic supersymmetric model, the constrained MSSM (cMSSM), taking into account both the gauge hierarchy problem and the fine-tuning of neutralino dark matter.

Supersymmetry solves the gauge hierarchy problem by embedding the Standard Model fields into supermultiplets, which do not generate quadratic corrections to the Higgs mass. The simplest realistic model this one can build is called the Minimal Supersymmetric Standard Model (MSSM). But the superparticles which accompany the SM particles in the supermultiplets are experimentally constrained to be heavier than their standard partners. This implies that supersymmetry has to be broken. However, with broken SUSY, the gauge hierarchy problem

is not completely solved. Instead, it remains in the form of a certain amount of special cancellations between the SUSY parameters, of typical scale M_{SUSY}, necessary to reproduce the Z boson mass. M_{SUSY} is constrained through both direct and indirect observations, and the LHC experiments are currently improving these direct limits (see, e.g., summary plots of Atlas [26] and CMS [27]). Roughly speaking, M_{SUSY} is at least O (TeV), 1 order of magnitude above the Z mass $m_Z \approx 91$ GeV.

One of the simplest and most widely studied versions of the MSSM with broken SUSY is the constrained MSSM (cMSSM). The parameters of that model are a common gaugino mass $m_{1/2}$, a common scalar mass m_0^2, a common scalar trilinear coupling $A_0 \equiv a_{ij}/y_{ij}$, the ratio of the two Higgs vacuum expectation values $\tan\beta = \langle H_u \rangle / \langle H_d \rangle$, and the sign of the SUSY Higgs mass term, sign(μ) (see, e.g., Ref. [28] for an introduction to SUSY models). However, this is a setup which already takes into account $m_Z = m_{Z\text{ex}}$. As we are interested in the fine-tuning induced by $m_Z = m_{Z\,\text{ex}}$, this constraint must not be incorporated in the model in the first place. Therefore, a new input parameter has to be introduced. It is, in fact, interesting to trade $\tan\beta$ for the dimensionful parameters μ and B_μ of the Higgs sector. Indeed, whereas it is not obvious to find an objective density for $\tan\beta$, the objective densities of μ and B_μ are clearly logarithmic. In practice, it is the former parametrization of the model which is used. In that case, the objective prior of $\tan\beta$ has to be inferred from the priors of μ and B_μ. This remark was made in Ref. [22], where the resulting density is called the "REWSB prior."

Also, the MSSM has another celebrated feature. Its mass spectrum contains the neutralino, a fermion charged only under the weak force, and which is a mixture of neutral Higgsinos and gauginos. If the lightest neutralino $\tilde{\chi}_1^0$ is the lightest particle of the SUSY spectrum, and if a remnant of the $U(1)_R$ symmetry of the SUSY algebra is still present, it cannot decay directly into SM particles and is therefore stable. Such a particle is a good dark matter candidate. Under the assumptions that the Cosmological Standard Model is valid in the early Universe, and that the neutralinos were at the thermal equilibrium for some period, today's neutralino density can be precisely predicted using the Boltzmann equation. This density is the relic remaining after thermal freeze-out, when the neutralino annihilation rate vanishes due to the expansion of the Universe. The relic density predicted strongly depends on the masses and compositions of all particles of the spectrum.

The latest release of the dark matter relic density measured by WMAP7 is $\Omega h_{\text{ex}}^2 = 0.1126 \pm 0.0036$ [29]. One can know where this constraint is satisfied in the plan $(m_{1/2}, m_0)$ by looking at the lines in the plots of Fig. 1. This figure will be described in detail below. Experimental uncertainty, by construction, does not appear in the plots. Lines are set wide only to ease the reading of the color code. Typically in the cMSSM, the dark matter relic density predicted is a bit large compared to observations.
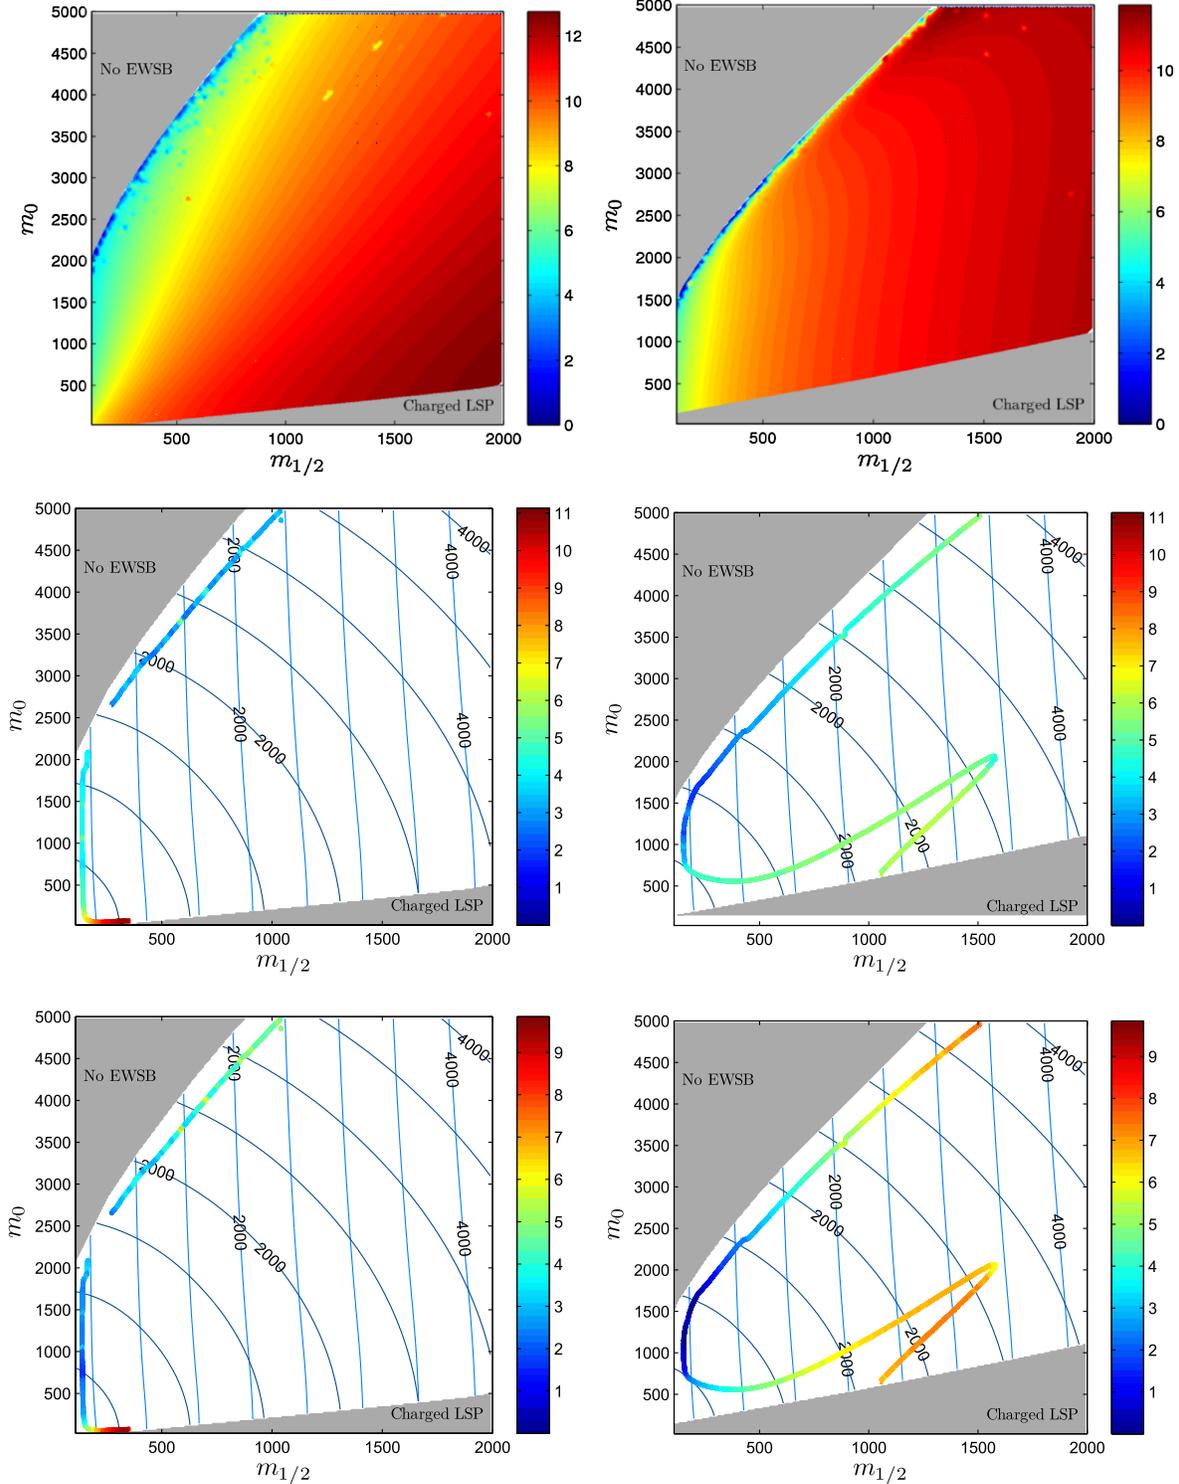
FIG. 1 (color online).   Quantified naturalness in the cMSSM. All of these plots are for $A_0 = 0$, sign$(\mu) = 1$, and $m_t = 172.4$ GeV, with $\tan\beta = 10$ and 50 for the left and right panels, respectively. $m_{1/2}$ and m_0 are given in GeV units. *Top row:* Maps of the logarithm of the electroweak fine-tuning measure C_{m_Z}, normalized to the point of minimal fine-tuning. *Center row:* Maps of the logarithm of the dark matter fine-tuning measure C_Ω. The measure on both plots has the same normalization. *Bottom row:* Maps of the logarithm of the combined electroweak and dark matter fine-tuning measure $C_{m_Z,\Omega}$. The measure on both plots has the same normalization. Blue (gray) and dark blue (dark gray) isolines show the mass of the gluino and the lightest squark with steps of 500 GeV. Following Jeffreys' scale, the relative degree of belief between two points 1 and 2 is given by $|\log(C_2/C_1)|$, such that the threshold values 1, 2.5 and 5 correspond to weak, moderate, and strong evidence for point 1, respectively.

To reproduce this experimental constraint, it is necessary that one or several processes of neutralino annihilation be particularly efficient [30]. There are at least four such processes in the cMSSM.

Firstly, the neutralinos can annihilate through the exchange of a scalar. But this mechanism is only efficient for light sparticles, which are more and more excluded by the LHC. Secondly, the annihilation through the exchange of a Higgsino or $SU(2)$ gaugino is efficient if the mixing with those states is large enough. This happens near the "No EWSB" zone. Thirdly, a coannihilation with a slepton may dominate if it is close to the neutralino mass, and if both particles are not too heavy. This happens near the "Charged LSP" zone. Finally, the exchange of a CP-odd Higgs A^0 is enhanced near the resonance pole, when $m_{A^0} \approx 2m_{\tilde{\chi}_1^0}$. This happens at large $\tan\beta$ and is dubbed "A-pole funnel."

But relying on the efficiency of such processes to obtain the correct value for Ωh^2 requires a rather precise adjustment of parameters. It is therefore an act of fine-tuning. So if one wants to explain dark matter by the neutralino, one ends up with two naturalness problems, one induced by the piece of information $m_Z = m_{Z\,\text{ex}}$ and the other due to $\Omega = \Omega_{\text{ex}}$. To study fine-tuning in the cMSSM, it is therefore the common sensitivity $C_{m_Z, \Omega h^2}$ which must be used. From the sensitivity point of view, one has to consider the set of fundamental parameters $p_i = (m_{1/2}, m_0^2, A_0, \mu, B_\mu)$. All of those parameters are defined at the GUT scale. They all have a logarithmic prior as objective density. On the other hand, although $\Omega = \rho_{\text{CDM}}/\rho_c$ is a density rescaled to be made dimensionless, ρ_{CDM} is dimensionful, so it necessitates a logarithmic prior as well. The common sensitivity measure is therefore

$$C_{m_Z, \Omega} = \left\| \frac{\partial \log m_Z}{\partial \log p_i} \wedge \frac{\partial \log \Omega h^2}{\partial \log p_i} \right\|. \tag{45}$$

We assume that the experimental uncertainties are sufficiently small, such that Eq. (45) holds for all the parameter space. This sensitivity will be denoted as C from now on.

In the MSSM, the top quark mass is given by $m_t = y_t v \sin\beta$. Thus, rigorously, y_t should not be taken as a constant, since what we explained in Subsection IV D about the top Yukawa does not hold here due to the presence of $\sin\beta$. To stay exact, it would be necessary to consider the sensitivity associated with the three observables, C_{m_Z, Ω, m_t}. However, the correction induced from adding the observable m_t is small, because in the $\sin\beta$ contribution the derivative $\partial \log \sin\beta / \partial \log y_t$ is dominant over the other derivatives. Therefore we choose to work only with the observables m_Z, Ωh^2, and keep y_t as a constant.

We have evaluated the dark matter relic density, the sensitivity C, and the SUSY spectrum over slices of the parameter space of the cMSSM. Our analysis was realized using a modified version of the spectrum calculator SOFTSUSY [31] interfaced with MICROMEGAS 2.4 [32] to compute the dark matter relic density. In spectrum calculators, these are not μ and B_μ, which are input parameters,

but m_Z and $\tan\beta$. This is already taken into account in SOFTSUSY to compute the m_Z derivatives, but it has to be carefully considered when implementing the Ωh^2 derivatives. The results obtained are presented in Fig. 1.

Figure 1 shows slices of the parameter space with $A_0 = 0$, $\text{sign}(\mu) = 1$, and $m_t = 172.4$ GeV, for $\tan\beta = 10$ and 50. The low-mass region of the parameter space is increasingly excluded by the LHC bounds on sparticle masses (see Refs. [26,27]). Instead of showing a particular limit, we prefer to plot the gluino and lightest squark masses, and leave the choice to the reader to apply his preferred bound. All the plots display the logarithm of the sensitivity. Following our results in Sec. IV, the difference between two points 1 and 2 is given by $|\log C_2 - \log C_1| = \Delta \log C$, which has to be interpreted on the basis of Jeffreys' scale. That is, $\Delta \log C = 1$, 2.5, 5 correspond to weak, moderate, and strong evidence in favor of point 1, respectively. The statements about naturalness that we will make when discussing the plots are based on this scale.

Plots in the upper line show maps of electroweak fine-tuning. The fact that $\log C_{m_Z}$ drops down near the "No EWSB" zone is due to a feature of the MSSM renormalization group equations, known as the mechanism of "focus point" [33]. In short, the low scale value of the Higgs soft mass $m_{H_u}^2$ becomes generically small in this region, such that the cancellations required to reproduce the Z mass are less important. This feature also implies a large Higgsino fraction for the neutralino, so that the experimental value of Ωh^2 can be reproduced. In this zone of the parameter space, the predictions of m_Z and Ωh^2 are therefore particularly correlated by the model.

Plots in the center line show the dark matter fine-tuning. In the $\tan\beta = 10$ slice, one can see that the coannihilation region has a very strong fine-tuning compared to the focus point region. Formally, the coannihilation region continues all along the "charged LSP zone" with an increasing fine-tuning, but points are so fine-tuned that the numerical analysis does not render them. In the $\tan\beta = 50$ slice, one can see that the A-pole funnel and the focus point region have sensibly the same naturalness. Relative to the $\tan\beta = 10$ focus point, these regions have a weak to moderate fine-tuning. On the other hand, they are strongly more natural than the $\tan\beta = 10$ coannihilation region. At $\tan\beta = 50$, some very fine-tuned coannihilations can also occur on the border, but are not shown on the plot. Dark matter fine-tuning has been previously investigated in the literature; see, e.g., Refs. [3,34], with slightly different definitions for C_Ω.

Finally, plots shown in the lower line are for the combined electroweak and dark matter fine-tuning measure $\log C_{m_Z, \Omega}$. Compared to the dark matter fine-tuning alone, here the C measure increases with $M_{1/2}$ due to the gauge hierarchy problem. In the $\tan\beta = 10$ slice, the coannihilation region is still strongly fine-tuned with respect to the focus point region. In the $\tan\beta = 50$ slice, the fine-tuning of the focus point region increases by $\Delta \log C \approx 5$ between

$m_{1/2} = 500$ GeV and 1500 GeV. At low $m_{1/2}$, this region is the most favored. At high $m_{1/2}$, it is moderately fine-tuned compared to the focus point region at $\tan\beta = 10$. The A-pole funnel and the focus point region at high $M_{1/2}$ are only moderately preferred to the $\tan\beta = 10$ coannihilation region.

VI. CONCLUSION

The degree of naturalness is often intuitively defined as a sensitivity, although this approach suffers from several conceptual flaws. We propose a different definition to formalize naturalness, working in the framework provided by Bayesian statistics. This approach is self-consistent, and interestingly turns out to embed the usual sensitivity definition in a generalized form.

So our approach is not an alternative. It appears that the sensitivity is actually a piece, intuitively guessed, of a larger setting. It is not consistent when taken alone, but the flaws find an explanation once the embedding in the Bayesian framework is done. Somehow, the essential missing piece was the notion of prior volume, which is also intuitive on its own. In this paper, we work out the consistent framework bringing together these notions.

The naturalness measure which appears in this framework is a Bayes factor. The link between the naturalness measure and our degree of belief, which was missing so far, is therefore automatically provided by Jeffreys' scale. The generalized sensitivity which emerges takes into account the fine-tuning of an arbitrary number of correlated observables. We discussed in details the two observable cases.

By studying the Bayes factor involving a "puzzle" model, we found that either the principle of indifference or consistency of the measure are setting the functional form of all quantities. For the sensitivity, it entails that these are the objective prior repartition functions, both for parameters and observables, which appear in the derivatives. As the puzzle model is a reference in terms of sensitivity, this Bayes factor gives a handle on the absolute interpretation of C.

The Bayesian approach resolves without ambiguity the question of whether or not the top Yukawa should enter into the gauge hierarchy measure C_{m_Z}. Also, it accounts for the "second-order fine-tuning," which is induced when it is necessary to adjust precisely a parameter to select a zone with small fine-tuning in the parameter space. Consequences of recent LHC searches are also discussed.

We present a simple illustration of our results by examining the naturalness of a supersymmetric model, the cMSSM. The sensitivity formulas associated with the electroweak scale and dark matter relic density, taken separately or together, are well defined, and differ from some work in the literature. By using Jeffreys' scale, we make statements about the naturalness of the different dark matter annihilation regions. Roughly speaking, the focus point region is the winner of the naturalness comparisons,

while the coannihilation region comes last with a strong evidence gap.

ACKNOWLEDGMENTS

APPENDIX: NATURALNESS PROBLEMS IN PARTICLE PHYSICS AND COSMOLOGY

In this Appendix, we recall some of the main naturalness problems. These are the most commonly discussed, but the list is not intended to be exhaustive.

  (i) Gauge hierarchy problem [28]: The electroweak scale, often represented by the Z-boson mass m_Z, is $O(100$ GeV$)$. On the other hand, the Planck scale $M_{\text{Pl}} = \sqrt{\hbar c/8\pi G_N} = 2.4 \times 10^{18}$ GeV, sets the scale at which the theory of quantum gravity appears. Why is the electroweak scale so small compared to M_{Pl}, when it should receive $O(M_{\text{Pl}}^2)$ quantum contributions?

  (ii) Strong CP puzzle [35]: From neutron electric dipole measurement, one deduces that the θ angle, contributing to the QCD Lagrangian $\mathcal{L}_{\text{QCD}} \supset -1/4g^2 G^{\mu\nu}G_{\mu\nu} + \theta/16\pi^2 G^{\mu\nu}\tilde{G}_{\mu\nu}$, is very small, $\theta < 10^{-12}$, while it could take values in $[-\pi, \pi]$. Then why is it so close to zero?

  (iii) Flavor puzzle [36]: Ratios of successive SM fermion mass eigenvalues, as well as CKM angles, are all of roughly the same order. Why do they follow this particular structure?

  (iv) Cosmological constant problem [37]: The cosmological constant Λ, which appears in Einstein's equations $\mathcal{R}_{\mu\nu} - \frac{1}{2}g_{\mu\nu}\mathcal{R} = 8\pi G_N T_{\mu\nu} + \Lambda g_{\mu\nu}$, is estimated to be of $O(10^{-47})$ GeV^4 by fitting the Standard Cosmological Model to CMB, large-scale structures, and supernovae data. Within quantum field theories, it should receive $O(M_{\text{Pl}}^4)$ contributions, [or $O(M_{\text{SUSY}}^4)$ if there is SUSY]. Then why is it so small?

  (v) Flatness problem [38]: In the Standard Cosmological Model, the curvature of the Universe is given by $1/R^2 = H^2(\rho/\rho_c - 1)$, where H is the Hubble constant ρ is the total energy density contained in the Universe, and $\rho_c = 3H^2/8\pi G_N$. $\rho/\rho_c - 1$ is estimated to be less than 0.01, and $O(10^{-61})$ at the Planck era. Why did the Universe have such a small curvature?

  (vi) Cosmic coincidence [39]: Why are the densities of matter and vaccum energy of the same order of magnitude, i.e., $\rho_M \sim \rho_\Lambda$? And why now?

[1] J. R. Ellis, K. Enqvist, D. V. Nanopoulos, and F. Zwirner, Mod. Phys. Lett. A **01**, 57 (1986).

[2] R. Barbieri and G. F. Giudice, Nucl. Phys. **B306**, 63 (1988).

[3] S. Cassel, D. M. Ghilencea, and G. G. Ross, Nucl. Phys. **B835**, 110 (2010).

[4] A. Strumia, J. High Energy Phys. 04 (2011) 073.

[5] M. E. Cabrera, J. A. Casas, and R. R. de Austri, J. High Energy Phys. 03 (2009) 075.

[6] B. C. Allanach, Phys. Lett. B **635**, 123 (2006).

[7] G. W. Anderson and D. J. Castano, Phys. Lett. B **347**, 300 (1995).

[8] P. Athron and D. J. Miller, Phys. Rev. D **76**, 075010 (2007).

[9] P. Ciafaloni and A. Strumia, Nucl. Phys. **B494**, 41 (1997).

[10] K. L. Chan, U. Chattopadhyay, and P. Nath, Phys. Rev. D **58**, 096004 (1998).

[11] R. Barbieri and A. Strumia, Phys. Lett. B **433**, 63 (1998).

[12] L. Giusti, A. Romanino, and A. Strumia, Nucl. Phys. **B550**, 3 (1999).

[13] J. A. Casas, J. R. Espinosa, and I. Hidalgo, J. High Energy Phys. 11 (2004) 057.

[14] R. Trotta, Contemp. Phys. **49**, 71 (2008).

[15] G. D'Agostini, *Bayesian Reasoning in Data Analysis: A Critical Introduction* (World Scientific, Hackensack, NJ, 2003).

[16] H. Jeffreys, *Theory of Probability* (Oxford University Press, Oxford, England, 1961), 3rd ed. (republished 1998).

[17] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, Springer Texts in Statistics (Springer, New York, 1998).

[18] J. O. Berger, J. M. Bernardo, and D. Sun, Ann. Stat. **37**, 905 (2009).

[19] J. Press, *Subjective and Objective Bayesian Statistics: Principles, Models, and Applications*, Wiley Series in Probability and Statistics (Wiley, New York, 2002), 2nd ed.

[20] E. T. Jaynes, IEEE Trans. Syst. Sci. Cybern. **4**, 227 (1968).

[21] C. Balazs, A. Buckley, D. Carter, B. Farmer, and M. White, arXiv:1205.1568.

[22] B. C. Allanach, K. Cranmer, C. G. Lester, and A. M. Weber, J. High Energy Phys. 08 (2007) 023.

[23] M. S. Chanowitz, M. A. Furman, and I. Hinchliffe, Phys. Lett. **78B**, 285 (1978).

[24] S. Chatrchyan *et al.* (CMS Collaboration), Phys. Lett. B **716**, 30 (2012).

[25] G. Aad *et al.* (ATLAS Collaboration), Phys. Lett. B **716**, 1 (2012).

[26] ATLAS Experiment, public results, https://twiki.cern.ch/twiki/bin/view/AtlasPublic.

[27] CMS Experiment, public results, https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResults.

[28] S. P. Martin, arXiv:hep-ph/9709356.

[29] E. Komatsu *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **192**, 18 (2011).

[30] E. A. Baltz, M. Battaglia, M. E. Peskin, and T. Wizansky, Phys. Rev. D **74**, 103521 (2006).

[31] B. C. Allanach, Comput. Phys. Commun. **143**, 305 (2002).

[32] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov, Comput. Phys. Commun. **174**, 577 (2006); **149**, 103 (2002).

[33] J. L. Feng, K. T. Matchev, and T. Moroi, Phys. Rev. Lett. **84**, 2322 (2000).J. L. Feng, K. T. Matchev, and F. Wilczek, Phys. Lett. B **482**, 388 (2000).

[34] J. R. Ellis, S. F. King, and J. P. Roberts, J. High Energy Phys. 04 (2008) 099.

[35] J. E. Kim and G. Carosi, Rev. Mod. Phys. **82**, 557 (2010).

[36] K. Nakamura *et al.* (Particle Data Group), J. Phys. G **37**, 075021 (2010).

[37] R. Bousso, Gen. Relativ. Gravit. **40**, 607 (2008); J. Polchinski, arXiv:hep-th/0603249.

[38] D. Baumann, arXiv:0907.5424.

[39] N. Arkani-Hamed, L. J. Hall, C. F. Kolda, and H. Murayama, Phys. Rev. Lett. **85**, 4434 (2000).