

Bayesian model comparison of Higgs couplings

Johannes Bergström*

*Departament d'Estructura i Constituents de la Matèria and Institut de Ciències del Cosmos,
Universitat de Barcelona, Diagonal 647, E-08028 Barcelona, Spain*

Stella Riad†

*Department of Theoretical Physics, School of Engineering Sciences,
KTH Royal Institute of Technology-AlbaNova University Center, Roslagstullsbacken 21,
106 91 Stockholm, Sweden*

(Received 18 December 2014; published 14 April 2015)

We investigate the possibility of contributions from physics beyond the Standard Model (SM) to the Higgs couplings, in the light of the LHC data. The work is performed within an interim framework where the magnitude of the Higgs production and decay rates are rescaled through Higgs coupling scale factors. We perform Bayesian parameter inference on these scale factors, concluding that there is good compatibility with the SM. Furthermore, we carry out a Bayesian model comparison on all models where any combination of scale factors can differ from their SM values and find that typically models with fewer free couplings are strongly favored. We consider the evidence that each coupling individually equals the SM value, making the minimal assumptions on the other couplings. Finally, we make a comparison of the SM against a single “not-SM” model and find that there is moderate to strong evidence for the SM.

DOI: [10.1103/PhysRevD.91.075008](https://doi.org/10.1103/PhysRevD.91.075008)

PACS numbers: 12.60.-i, 12.60.Fr, 14.80.Ec

I. INTRODUCTION

The discovery of a boson with a mass of approximately 125.5 GeV was announced in July 2012 by the ATLAS and CMS experiments at the Large Hadron Collider (LHC) at CERN [1,2]. This discovery is compatible with previous data from proton-antiproton collisions at $\sqrt{s} = 1.96$ TeV at the Tevatron [3]. Using all of the available data, with a total luminosity of 25 fb^{-1} from the proton-proton collisions with energies of $\sqrt{s} = 7$ and 8 TeV runs at the LHC, properties of the Higgs boson properties, such as spin, parity, mass, and the couplings to other Standard Model (SM) particles, has been further investigated [4–7]. So far, however, there are no indications of major deviations from the properties of the SM Higgs boson, and the boson does in fact seem to be a CP -even scalar [8–10]. The discovery of the Higgs boson marks an important milestone in the history of particle physics, especially for our understanding of electroweak symmetry breaking and the generation of particle masses [11–14].

The Higgs boson was discovered through its decays into gauge bosons, i.e., WW , ZZ , and $\gamma\gamma$ pairs. However, using the full set of data from the LHC, there is now also evidence for decays into fermions, $b\bar{b}$ and $\tau^+\tau^-$ [15,16].

Even though the properties of the new boson so far are compatible with those of the SM Higgs, the possibility for new physics in the Higgs sector should be investigated. New physics can manifest itself in different ways, some of which

can be detected since they would give rise to a rescaling of the magnitude, or change in the structure, of the Higgs boson couplings. Hence, a natural step forward in experimental Higgs physics is precision measurements of the Higgs boson couplings to fermions and gauge bosons. From existing data, there are bounds on the couplings of the boson. However, hadron colliders are in general not ideal for Higgs precision measurements, and thus, in order to determine the couplings with significantly greater precision, the upgrade of the LHC to 14 TeV will not suffice, and instead lepton colliders, such as a Higgs factory, are needed [17].

The status of the Higgs couplings as measured by the LHC can be studied by means of so-called *Higgs coupling scale factors*, introduced by the LHC Higgs cross section working group as an interim treatment of the Higgs couplings [18]. Coupling scale factors are introduced in order to rescale the magnitudes of the Higgs production and decay rates, which is especially useful since the experimental data from the collaborations are presented in terms of so-called *global signal strengths*. Since the analysis of the data indicate that a CP -even scalar is preferred to a CP -odd one, we assume a single underlying CP -even scalar boson at a mass of about 125 GeV, and furthermore we assume a simplifying zero-width approximation. This so-called *interim framework* has been used by the ATLAS and CMS collaborations as well as in several phenomenological studies [4,7,19–29]. In this work, we use the software HIGGSIGNALS 1.2.0 for the implementation of the LHC data in the form of a χ^2 -function [30–32].

In the present work, we apply Bayesian inference within the framework of coupling scale factors. We shall use

*bergstrom@ecm.ub.edu
†sriad@kth.se

Bayesian parameter inference to get a rough idea of how the parameters are constrained. However, since the most important question is rather that of which model best describes the data, we will focus on model comparison—in particular of different models in which any combination of couplings differ from their SM values. This framework makes it possible to compare many models to each other at once, but the main advantage is that it is possible to obtain evidence *in favor* of simpler models—in the present case for models where the couplings are given by their SM values.

The paper is organized as follows. In Sec. II, we give an introduction to Higgs physics and the concept of coupling scale factors. In Sec. III, we discuss the Bayesian method used in the present work, especially model comparison in the context of Higgs couplings. In addition, we discuss the models used in the present work as well as the priors used. In Sec. IV, we discuss the results of parameter estimation and the results concerning the different questions addressed using model comparison. Finally, in Sec. V, we give a short summary and give our conclusions.

II. HIGGS PHYSICS

Whether the discovered particle at 125.5 GeV actually is *the* SM Higgs boson, or only a part of some bigger picture, is an important question which needs to be investigated. In general, additional degrees of freedom in the Higgs sector will influence the Higgs couplings to the SM particles as well as the loop-induced production and decay modes.

One common way to investigate the possibility of new physics in the Higgs sector is to study and compare specific renormalizable models for beyond the SM physics, such as two-Higgs-doublet models [33–35], composite Higgs models [36,37], a dilaton model [21], and supersymmetric models [38,39]. However, in these cases, the comparisons are only made between these specific models and the SM, and obviously lack in generality. Another way is to consider the SM extended with effective operators, resulting from new physics above the TeV scale. Since this new physics is heavy by assumption, it will give rise to modifications of the couplings, which are suppressed by the scale of new physics. These modifications are, however, not necessarily small in magnitude if the scale of new physics is low [40,41]. In both of the frameworks discussed above, some or all of the Higgs boson couplings will be altered, both the magnitude and in principle also the tensor structure of the couplings, even though these modifications are often heavily suppressed. Another approach is simply to not consider a physical and realistic model, but instead make a statistical analysis based on the “naive” rescaling of the magnitude of the Higgs couplings. In such a framework, it is only possible to investigate whether there are any significant deviations of the couplings from their SM values, without attempting to interpret the underlying physics. Thus, the relevant result of the analysis is whether

the couplings deviate from their SM value or not, rather than the exact value of the couplings. This treatment with coupling scale factors shall be considered here.

A. Production modes

Four production modes of the Higgs boson in the SM are significant at the LHC. The predominant production mode is the loop-induced gluon fusion $gg \rightarrow H$, with heavy quarks running in a triangular loop, with the main contribution coming from the top quark. Since this process is loop induced, it is of particular interest in searches for new physics. The subdominant processes are vector boson fusion, $qq' \rightarrow qq'H$; associated production with a vector boson, $q\bar{q} \rightarrow WH/ZH$; and the associated production with a top-quark pair, $q\bar{q}/gg \rightarrow t\bar{t}H$. We will use the notation where $l = e, \mu$ and q stands for any quark.

B. Decay modes

The Higgs boson can decay either to a fermion-anti-fermion pair or two gauge bosons. At present, the Higgs boson have been detected in five decay channels at the LHC, namely the $\gamma\gamma$, $ZZ^{(*)}$ (in turn followed by a decay to $4l, 2l2\nu, 2l2q, 2l2\tau$), $WW^{(*)}$ (followed by decays to $l\nu l\nu, l\nu q\bar{q}$), $b\bar{b}$, and $\tau^+\tau^-$, which then decay leptonically and hadronically. Since the Higgs’s coupling to fermions is proportional to the fermion mass, the heaviest fermion mode, which is kinematically accessible, will have the largest partial decay width in the SM. Of the detected decay modes only $H \rightarrow \gamma\gamma$ is loop induced, which is of particular interest for searches beyond the SM. In addition to the observed decay modes, the $H \rightarrow \mu^+\mu^-$ and $H \rightarrow Z\gamma$ channels have been investigated at the LHC. However, the Higgs boson has not been detected in either of them, and there are therefore only (rather loose) upper bounds in these channels at present [42].

C. Definition of coupling scale factors

The LHC Higgs results are commonly presented in terms of global signal strengths, defined as

$$\mu = \frac{\sigma(X) \cdot \text{BR}(H \rightarrow Y)}{\sigma(X)_{\text{SM}} \cdot \text{BR}(H \rightarrow Y)_{\text{SM}}}, \quad (1)$$

where $\sigma(X)$ is the cross section for the production mode X and $\text{BR}(H \rightarrow Y)_{\text{SM}}$ is the branching ratio of the decay mode Y . In the case of a SM process, the value of μ is naturally 1. In the SM, the Higgs boson couples to the other particles with couplings y_i^{SM} , where $i \in \{t, b, \tau, \mu, W, Z\}$. The couplings to the fermions are the Yukawa couplings

$$y_f^{\text{SM}} = \frac{m_f}{v}, \quad (2)$$

where m_f is the mass of the fermion, $f \in \{b, t, \tau, \mu\}$, and v is the Higgs vacuum expectation value. The upper

perturbative limit for these couplings is around 4π . For the gauge couplings, we have

$$y_W^{\text{SM}} = \frac{2m_W^2}{v}, \quad y_Z^{\text{SM}} = \frac{m_Z^2}{v}, \quad (3)$$

where m_W and m_Z are the W and Z masses, respectively. Note that these couplings are dimensionful.

A simple extension of the SM can be made by rescaling the magnitude of the SM decay and production rates, which effectively leads to a rescaling of the Higgs couplings by so-called coupling scale factors, κ_i . For the processes which exist at tree level in the SM, the couplings are rescaled as

$$y_i = \kappa_i \cdot y_i^{\text{SM}}. \quad (4)$$

Naturally, the SM is recovered for $\kappa_i = 1$. In addition, coupling scale factors can be introduced for the loop-induced processes. We introduce κ_g and κ_γ for the $gg \rightarrow H$ and $H \rightarrow \gamma\gamma$, respectively. In principle, a scale factor, $\kappa_{Z\gamma}$, could be introduced for a third loop-induced process $H \rightarrow Z\gamma$. However, since the sensitivity is nowhere close to the region of the SM prediction and since the inference for the other parameters will not be affected, we will not include this as a free parameter. The cross section of the process $ii \rightarrow H \rightarrow ff$ is then given by

$$\begin{aligned} & (\sigma \cdot \text{BR})(ii \rightarrow H \rightarrow ff) \\ &= \sigma_{\text{SM}}(ii \rightarrow H) \cdot \text{BR}_{\text{SM}}(H \rightarrow ff) \cdot \frac{\kappa_i^2 \kappa_f^2}{\kappa_H^2}, \end{aligned} \quad (5)$$

where κ_i and κ_f corresponds to the initial and final states, respectively, and κ_H is the scale factor for the total Higgs decay width.

The coupling scale factors κ_g and κ_γ can be considered either as functions of the other coupling scale factors or free parameters of the fit if new physics is allowed to participate in the loops. In the SM, these scale factors have the values $\kappa_g = \kappa_\gamma = 1$. However, in the case when only the tree-level scale factors are varied, the scale factors of the loop-induced processes will vary depending on the other scale factors. The effects of the rescaled tree-level couplings would have to be cancelled by some new physics, if these parameters were fixed to their SM values. If the scale couplings are free, new physics is allowed to propagate in the loop.

Furthermore, the factor κ_g can be defined in two different ways, either in terms of partial cross sections or decay widths. In the present case, we define the coupling scale factor $\kappa_g(\kappa_t, \kappa_b)$ using the cross sections, since gluon fusion is the more important process. Thus, the scale factor is given by

$$\kappa_g^2(\kappa_b, \kappa_t) = \frac{\kappa_t^2 \cdot \sigma_{ggH}^{tt} + \kappa_b^2 \cdot \sigma_{ggH}^{bb} + \kappa_t \kappa_b \cdot \sigma_{ggH}^{tb}}{\sigma_{ggH}^{tt} + \sigma_{ggH}^{bb} + \sigma_{ggH}^{tb}}. \quad (6)$$

In terms of the other κ 's, κ_γ is given by

$$\kappa_\gamma^2(\kappa_b, \kappa_t, \kappa_\tau, \kappa_W) = \frac{\sum_{ij} \kappa_i \kappa_j \cdot \Gamma_{\gamma\gamma}^{ij}}{\sum_{i,j} \Gamma_{\gamma\gamma}^{ij}}, \quad (7)$$

where $\Gamma_{\gamma\gamma}^{ij}$ are the partial decay widths and the pairs (i, j) are given by $bb, tt, \tau\tau, WW, bt, b\tau, bW, t\tau, tW, \tau W$ [18].

In addition, the total Higgs width scales with a coupling scale factor, which is defined in terms of the other coupling scale factors as

$$\kappa_H^2 = \sum_X \kappa_X^2 \cdot \text{BR}_{\text{SM}}(H \rightarrow X), \quad (8)$$

where the summation runs over all possible decay modes in the SM. This parametrization requires that the resonance width is small, and therefore the zero-width approximation is assumed. In principle, new physics could contribute to the total Higgs width, which occurs if, for instance, the Higgs can decay to dark matter particles. In this case, κ_H should be a free parameter; see for example Ref. [43]. For an extensive description of the concept of coupling scale factors, see Ref. [18].

In the present work, we shall focus on the coupling scale factors in two settings. First, the scale factors corresponding to the SM tree-level couplings (and which are currently constrained by LHC data), i.e., the Higgs couplings to $b\bar{b}$, $t\bar{t}$, $\tau^+\tau^-$, $\mu^+\mu^-$, WW , and ZZ , have the possibility to be varied. In the second case, the loop-induced processes have the potential to be scaled as well, through the variation of κ_g and κ_γ . We shall not consider the total decay width to be a free parameter in the present case. Again, the information on effective scale couplings from LHC data were implemented using the HIGGSIGNALS 1.2.0 software.

Note that the new particle is assumed to “resemble” the SM Higgs boson in a certain way. In principle, however, new physics will not only change the magnitude of the couplings but also their tensor structure. These new couplings usually are referred to as *anomalous couplings*, and the general statistical method of analysis, to be presented in the next chapter, would be applicable in that case as well.

III. STATISTICAL APPROACH

In this work, we will make use of Bayesian probability theory, in which each proposition is associated with a probability or *plausibility*, defined to lie between 0 and 1. This is the only consistent extension of boolean logic incorporating uncertainty [44–46].

In Bayesian inference, the laws of probability are used to infer which underlying hypotheses, assumption, or data model [47] is preferred by some given set of data. Of interest is *Bayes's theorem*, which can be used to reverse the order of the conditioning, denoted by “[|]”,

$$\Pr(A|B) = \frac{\Pr(B|A) \Pr(A)}{\Pr(B)}. \quad (9)$$

Thus, two different hypotheses or models can be compared using the data \mathbf{D} , through calculation of the *posterior odds*, given by

$$\frac{\Pr(M_i|\mathbf{D})}{\Pr(M_j|\mathbf{D})} = \frac{\Pr(\mathbf{D}|M_i) \Pr(M_i)}{\Pr(\mathbf{D}|M_j) \Pr(M_j)}. \quad (10)$$

The *prior odds* $\Pr(M_i)/\Pr(M_j)$ quantify how much more plausible one model is than the other *a priori*. This ratio is typically taken equal to unity, which, however, must be considered more carefully in some cases. The *evidence* $\mathcal{Z}_i = \Pr(\mathbf{D}|M_i)$ is the likelihood of the model, a measure of how well the model describes, or rather predicted, the data. The *Bayes factor* $B_j^i = \mathcal{Z}_i/\mathcal{Z}_j$ is the ratio of the evidences of the two models and quantifies how much better M_i describes the data than M_j .

Given that the model M contains the free parameters Θ , the evidence is given by

$$\begin{aligned} \mathcal{Z} = \Pr(\mathbf{D}|M) &= \int \Pr(\mathbf{D}, \Theta|M) d^N \Theta \\ &= \int \Pr(\mathbf{D}|\Theta, M) \Pr(\Theta|M) d^N \Theta \\ &= \int \mathcal{L}(\Theta) \pi(\Theta) d^N \Theta, \end{aligned} \quad (11)$$

where $\mathcal{L}(\Theta) \equiv \Pr(\mathbf{D}|\Theta, M)$ is the *likelihood function*. The prior probability density of the parameters is given by $\pi(\Theta) \equiv \Pr(\Theta|M)$, and should always be normalized, i.e., it should integrate to unity. The assignment of priors are probably the most discussed and controversial part of Bayesian inference. This is often far from trivial, but nevertheless this assignment is an important, even essential, part of any Bayesian analysis.

The Bayes factors, or rather the posterior odds, are interpreted or “translated” into ordinary language using the so-called *Jeffreys scale*, given in Table I as used in, e.g., Refs. [48,49] (“log” denotes the natural logarithm). Even though the Bayes factor in general will favor the correct

TABLE I. The Jeffreys scale, which is used for interpretation of Bayes factors, odds, and model probabilities. The posterior model probabilities for the preferred model are calculated assuming only two competing hypotheses and equal prior probabilities. Note that log denotes natural logarithm.

$ \log(\text{odds}) $	odds	$\Pr(M_1 \mathbf{D})$	Strength of evidence
< 1.0	$\lesssim 3:1$	$\lesssim 0.75$	Inconclusive
1.0	$\approx 3:1$	≈ 0.75	Weak evidence
2.5	$\approx 12:1$	≈ 0.92	Moderate evidence
5.0	$\approx 150:1$	≈ 0.993	Strong evidence

model once “enough” data have been obtained, the evidence is often highly dependent on the choice of prior.

Under the assumption that a model M is true, complete inference of its parameters is given by the posterior distribution,

$$\Pr(\Theta|\mathbf{D}, M) = \frac{\Pr(\mathbf{D}|\Theta, M) \Pr(\Theta|M)}{\Pr(\mathbf{D}|M)} = \frac{\mathcal{L}(\Theta) \pi(\Theta)}{\mathcal{Z}}. \quad (12)$$

In this case, the evidence is only a normalization factor, since it is independent of the values of the parameters Θ , and it is therefore often disregarded in parameter estimation. However, the actual values of the parameter within a prespecified model are often not of the greatest interest. Instead, the primary question is usually which model, or set of models, is preferred by the data.

After model comparison, there might still be a significant amount of uncertainty regarding which model actually is the best, and this uncertainty should not be ignored when making inference on parameters. Model uncertainty can be taken into account by calculating the model-averaged posterior distribution [50,51]

$$\Pr(\eta|\mathbf{D}) = \sum_i \Pr(\eta|H_i, \mathbf{D}) \Pr(H_i|\mathbf{D}), \quad (13)$$

which is the average of the individual distributions over the full space of the models considered, weighted by the posterior model probabilities. Averaging over models can be done for both prior and posterior distributions; however, the parameters η , which could be derived, obviously need to be well defined in all of the models. The posterior in Eq. (12) is obtained by setting all prior model probabilities, except one, equal to zero. For applications in physics and cosmology, see Refs. [51–53].

The main result of Bayesian parameter inference is the posterior and its marginalized versions (usually in one or two dimensions). Commonly, point estimates such as the posterior mean or median are given together with *credible intervals (regions)*, which are defined as intervals (regions) containing a certain amount of posterior probability. These regions are not unique, without further restrictions, similarly to classical confidence intervals, and in general they do not describe all the information contained in the posterior. We use MULTINEST [54–56] for the evaluation of all evidences and posterior distributions in this work.

A. Model comparison and Higgs couplings

We want to determine whether there is any evidence in the LHC data for deviations from the SM values of the couplings, i.e., if $\kappa_i \neq \kappa_i^{\text{SM}}$, or if $\kappa_i = \kappa_i^{\text{SM}}$ is sufficient to describe the data. In other words, we are interested in *if* there is a deviation from the SM couplings, and not precisely how large it is, given that it is nonzero. For each coupling, this gives two distinct cases, and in order to

differentiate between them, we want to perform Bayesian model comparison. Note that, from a statistical viewpoint, a model with $\kappa_i = \kappa_i^{\text{SM}}$ can also be interpreted as a model where there is some nonzero, but negligible (given current data) deviation from the SM value; see Ref. [50] for further discussion. Beforehand, it is not specified whether the other couplings, i.e., the couplings with indices $j \neq i$, should be fixed to their SM value or not, which gives rise to a complication. In principle, there is an important distinction since, without making the assumption of a particular model, any combination of the couplings can deviate from the SM values.

Thus, we can consider the models H_α , with $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, where each $\alpha_i = 0$ if $\kappa_i = \kappa_i^{\text{SM}}$ and $\alpha_i = 1$ if $\kappa_i \neq \kappa_i^{\text{SM}}$. In total, there are 2^n models, where n is the number of free parameters. In fact, we can consider α as a discrete parameter, for which the posterior odds are given by

$$\frac{\Pr(\alpha|\mathbf{D})}{\Pr(\beta|\mathbf{D})} = \frac{\Pr(\mathbf{D}|\alpha) \Pr(\alpha)}{\Pr(\mathbf{D}|\beta) \Pr(\beta)} = \frac{\mathcal{Z}_\alpha \pi_\alpha}{\mathcal{Z}_\beta \pi_\beta}, \quad (14)$$

where the calculable Bayes factor $\mathcal{B}_\beta^\alpha = \mathcal{Z}_\alpha / \mathcal{Z}_\beta$ quantifies how much better α describes the data than β . The natural baseline model is $\text{SM} = \bar{0} = (0, 0, \dots, 0)$, and all the \mathcal{B}_β^α can be obtained from the Bayes factors with respect to the SM, $\mathcal{B}_{\text{SM}}^\alpha$, as $\mathcal{B}_\beta^\alpha = \mathcal{B}_{\text{SM}}^\alpha / \mathcal{B}_{\text{SM}}^\beta$. If also finite prior probabilities are assigned to the full set of models, finite posteriors $\Pr(\alpha|\mathbf{D})$ can be calculated, even though we will typically refrain from doing this. Calculating the Bayes factor does, however, require assignment of priors on the couplings in all the models, which is nontrivial and will be discussed in detail in Sec. III D.

A different, but equivalent, approach is to instead consider a *single* model with a prior which is a mixture of the continuous prior and a point mass at the SM value,

$$\pi(\kappa_i) = (1 - p_i) f_i(\kappa_i) + p_i \delta(\kappa_i - \kappa_i^{\text{SM}}), \quad (15)$$

for each coupling [57]. Here the continuous part of the prior, given by f_i (which is normalized to unity), corresponds to the prior assuming $\alpha_i = 1$ and is assigned a total probability $1 - p_i$, while the SM value of the coupling is assigned a probability p_i . Note that α is a function of κ and hence that the priors and posteriors of α can be calculated from the distributions obtained using (15). In addition, the Bayes factors (which are independent of the prior on α) can be calculated using (14) by factoring out the prior odds.

B. Inclusion of individual couplings

In the previous section, we discussed the comparison of 2^n models, with different numbers of Higgs scale couplings kept free. However, when n grows in size, comparing this large number of models to each other rapidly becomes less transparent.

One can test if a particular variable should be included by comparing the cases $\kappa_i = \kappa_i^{\text{SM}}$ and $\kappa_i \neq \kappa_i^{\text{SM}}$ and hence calculating the Bayes factors

$$\mathcal{B}_i = \frac{\Pr(\mathbf{D}|\alpha_i = 0)}{\Pr(\mathbf{D}|\alpha_i = 1)}. \quad (16)$$

Again, however, one has to decide what to do with the *other* couplings, i.e., which priors to assign them. Possibilities could be:

- (i) fixed to the SM value (\mathcal{S}),
- (ii) free and different from the SM value (\mathcal{F}), or
- (iii) either of the above, i.e., an average (\mathcal{A}).

The evidences are given by the likelihoods integrated not only over the prior on κ_i but also over the prior on all other couplings. In particular,

$$\Pr(\mathbf{D}|\alpha_i) = \sum_{\alpha_i^*} \Pr(\mathbf{D}|\alpha_i, \alpha_i^*) \pi(\alpha_i^*), \quad (17)$$

which depends on the prior on $\alpha_i^* = (\alpha_1, \dots, \alpha_{i-1}, \alpha_{i+1}, \dots, \alpha_n)$. The evidences in Eq. (17) are simply the evidences discussed in Sec. III A. The three cases then correspond to $p_k = \pi(\alpha_k = 0)$ being equal to either $p_k = 1, 0$ or some intermediate values, most naturally 0.5 (see Ref. [50] for detailed discussion). The results are only expected to be independent of this choice in the case where the constraints on one parameter are independent of the values of the others.

C. Single comparison with SM

In physics, there is often a theoretically *a priori* motivated “baseline” model which all extended models are usually compared to. In the present case, the obvious choice for such a reference model is the SM. Furthermore, the Bayesian model comparison treats all models on equal footing, which enables quantification of how much the SM is favored with respect to extended models. Again, this could be done in the context of specific renormalizable high-energy models, but here we will focus on the effective case only considering the rescaled couplings.

We want to compare the SM with a model “not-SM,” or $\overline{\text{SM}}$. The question is how this model for comparison should be defined. For example, one could compare the SM with a model with only a single coupling free, which is just one of the cases discussed in the previous chapter. However, this is obviously not satisfactory since there are many such models [58], and at the same time, we are completely neglecting models with two or more couplings free [59]. Alternatively, one could compare with the most general model in which all couplings are free. However, the issue is the same, still neglecting the possibility that there could be significant deviations in more than one coupling, but not in all at once. The most general model could be punished for the inclusion of the couplings for which the SM value is

preferred. Therefore, the most appropriate comparison appears to be the one between the SM and a model in which each coupling *either* takes the SM values or differs from it.

Indeed, probability theory again yields

$$\Pr(\mathbf{D}|\overline{\text{SM}}) = \sum_{\alpha} \Pr(\mathbf{D}|\alpha)\pi(\alpha|\overline{\text{SM}}), \quad (18)$$

and all the above cases are just cases for a specific choice of prior $\pi(\alpha|\overline{\text{SM}})$. Due to lack of further information, we take $\pi(\alpha_i|\overline{\text{SM}}) = 1/2$, which means that in the $\overline{\text{SM}}$ model it is equally probable that each coupling deviates (significantly) from the SM value, as it is that there is no (or negligible) deviation. In this case, however, the couplings of the model $\overline{\text{SM}}$ equal the SM couplings with prior probability $1/2^n$. This part of $\overline{\text{SM}}$, i.e., the part that is statistically equivalent to the SM, can of course then just be excluded in the analysis, and this will be our default choice. In principle, however, one could also motivate its inclusion by saying there could still be a deviation from the SM values, but a negligible one. Note that adding any additional couplings, which are unconstrained by data, does not affect the comparison of SM and $\overline{\text{SM}}$.

D. Choice of prior

As discussed in Sec. II, we will consider the two cases: (i) all tree-level couplings are allowed to vary, with the loop-induced couplings calculated assuming no additional contribution from new physics; (ii) all couplings, including the loop-induced ones, are allowed to vary, which implies that new physics is allowed to participate in the loop processes. In the first case, there are six free parameters ($\kappa_t, \kappa_b, \kappa_\tau, \kappa_W, \kappa_Z, \kappa_\mu$), whereas there are eight free parameters in the second case, adding (κ_γ, κ_g). Note that the default “SM” values of these couplings are those calculated in Eqs. (6) and (7) assuming no new particles, which do not necessarily correspond to the actual exact SM value (equal to 1). In addition, the value of the scale factor for the total Higgs width, κ_H , will depend on the other ones according to Eq. (8). We shall, however, not consider this as a free parameter in either case.

To calculate the evidence of the models in which the couplings differ from the SM value, a prior for each coupling κ_i is needed. The assignment of prior is an important task since not only the posteriors within each model depend on it but perhaps more importantly so does the evidence. It is therefore important to take care to include as much known information into the prior without making any assumptions based on the data under consideration.

1. Default: Uniform

A common choice is to take a uniform prior on each of the couplings in order to implement *a priori* “ignorance,” usually unbounded or with “wide enough” limits. However,

such a uniform prior cannot quantify ignorance in a parameter, if not only because a uniform prior in one parameter will not be uniform in a parameter given by a nonlinear transformation of the first one. Second, an unbounded (improper) prior often gives meaningless answers for the evidence, and so do many priors in the limit when their widths go to infinity. However, this does not necessarily imply that the uniform prior as such is useless or in general should be avoided. As any prior, it can be used when it is motivated, and we shall use it in Sec. IV A to get a rough idea of what the parameter constraints on the different couplings are.

2. Couplings: Uniform

In the case where only the tree-level couplings are free, one can consider the actual couplings appearing in the Lagrangian as the free parameters. In the Higgs sector, there are Yukawa couplings for the Higgs coupling to fermions as well as the Higgs coupling to the gauge bosons. In principle, one could argue that *a priori* all couplings should be of order 1. Hence, a roughly uniform prior on each of the couplings, with an effective upper limit of some constant of order 1 would seem appropriate [60]. However, if the measured couplings have a small (absolute) errors compared to 1, this will lead to a very strong “Occam effect” which will strongly disfavor modifications of the couplings and give strong preference to the SM values. This is indeed the case, with the possible exception of the top Yukawa, and the masses of the SM particles differ by many orders of magnitude, a discrepancy commonly known as the “flavor puzzle” [61]. Hence, all models with additional couplings will be severely disfavored if this prior is taken, and so we will not perform a detailed analysis of this case, even though these conclusions are worth bearing in mind.

3. Logarithmic

Dropping the assumption that the couplings should be of order 1, it might seem more appropriate that instead the order of magnitude of the couplings are *a priori* unknown. Thus, the choice is instead a logarithmically uniform prior on y_i between some lower limit and the perturbative upper limit, taken as 4π . The lower limit must be chosen by hand; we will use 10^{-7} as the default choice. However, it turns out that the results are insensitive to changing this lower limit by at least a few orders of magnitude. Furthermore, for simplicity, we will always assume positive couplings. In most cases, the sensitivity to the sign of the couplings is small, with the exception of the case when κ_W and κ_t have different signs, which can enhance the rate of the $H \rightarrow \gamma\gamma$ [4,7]. However, there is no clear sign of this enhancement in the data, which implies that the total mass of the mode in that region will not be much larger than in the region with positive couplings, and therefore the effect on the evidences will be very small.

4. Gaussian

Instead of assuming, as in the previous cases, that the couplings are *a priori* unrelated to the SM ones, one can consider that many SM extensions, such as the ones mentioned in Sec. II, will all lead to modifications of roughly the same size as the SM couplings. Without considering a specific model, we cannot determine the sizes of these contributions nor if they should be positive or negative. Hence, all we can do is say that we *a priori* expect $\langle \kappa_i \rangle = \kappa_i^{\text{SM}}$, and a typical deviation of $\sigma(\kappa_i) = s_i = \mathcal{O}(1)$. Out of all the (prior) distributions on the real numbers with these constraints, there is a unique one which has maximal entropy (or equivalently “minimal information”), namely the Gaussian distribution [44,62]. We will consider values $s_i = s$ in the range 1–4 as appropriate, with a default value of $s = 2$.

Finally, we mention that one in principle could consider the SM augmented with additional higher-dimensional effective operators. These modify the SM couplings by an amount proportional to v^2/Λ^2 , where Λ is the scale of new physics [41,63–66]. These operators could be implemented in a Bayesian analysis such as in Refs. [66,67], but one could also utilize the expected sizes of the additional contributions in the present analysis by using a prior on Λ and from this obtain priors on the κ_i 's. If one expects that Λ could be of any order of magnitude, much of the prior would be piled up close to the SM values, which would imply that it would be possible to obtain significant evidence against the couplings taking those values, but not in favor. However, if the scale of new physics is assumed to be close to the electroweak scale as in Ref. [66], the typical modification would be of order 1, in which case one will get a result similar to the one for the Gaussian prior above.

To summarize, in the present work, we shall consider the following models and priors. In the case with only the tree-level couplings free, we shall make the analysis both using a logarithmic prior, which is placed directly on the actual couplings, and a Gaussian prior, which instead is placed on the coupling scale factors, κ_i . In the second case, where both tree- and loop-level couplings are free, we shall only make an analysis using the Gaussian prior placed on the scale factors. In this case, it should be noted that the expectation value of the now free parameters κ_g and κ_γ are the values given from the other scale factors, i.e., the values given by Eqs. (6) and (7), and not the SM value of these scale factors (which is 1).

IV. RESULTS

A. Default parameter constraints

In this section, we obtain the “default” parameter constraints on the coupling scale factors by calculating the likelihood using HIGGSIGNALS 1.2.0 and imposing a uniform prior on the κ 's with zero as the lower limit and a “large enough” upper limit. Although this prior does not

impose *a priori* ignorance, and it cannot be used for model comparison, the derived parameter constraints will be valid as long as the uniform prior is reasonable in the region of parameter space which is not completely ruled out by the data. A fixed Higgs boson mass of $m_H = 125.5$ GeV was used and will be used throughout this work.

Similar to the model comparison performed later, we first simultaneously estimate only the scale factors present at tree level and then additionally also the loop-induced scale factors. In addition to these two cases, we shall consider the special case where new physics only contribute to the loop-induced processes and thus only the scale factors corresponding to these processes, i.e., κ_g and κ_γ , are free.

In Fig. 1, we present the results in terms of one- and two-dimensional posterior distributions. In the two-dimensional plots, the blue shading denotes the natural logarithm of the posteriors, and the black contours denote the 1σ and 3σ credible regions, [68] while the one-dimensional posteriors are also black in the plots on the diagonal. Superimposed on these, in red, are the 1σ and 3σ contours as well as the one-dimensional posteriors for the case when only the tree-level scale factors are free. As previously discussed, κ_g and κ_γ are given as functions of the free scale factors. Finally, the same quantities are presented in green (in the bottom right) for the case when the tree-level couplings remain fixed at their SM values but new physics is allowed to participate in the loop-induced processes. The SM values are marked with stars and vertical lines.

In all three fits, *all* the SM values are inside (or extremely close to) the 1σ regions, which is in fact rather unlikely. As expected, adding κ_g and κ_γ to the set of free parameters will relax constraints on the six free tree-level couplings. The main effects should be seen in the scale factors corresponding to the particles which give the main contribution to the loop processes. Thus, the largest effect will be for the top quark which gives the absolutely dominating effect to the loop in the gluon fusion process, while a smaller effect should also be seen in the bottom quark coupling. Apart from the top quark, this is also the only particle that participates in both the gluon fusion and $H \rightarrow \gamma\gamma$ processes. The modifications to the other couplings are marginal. In a similar manner, the constraints on the loop-induced couplings are weaker in the eight-dimensional fit than in the two-dimensional one.

Finally, from the plots in Fig. 1 we can conclude that there is quite strong support for the couplings having non-zero values, with the exception of κ_μ , κ_t , and to some extent κ_b , in the eight-parameter fit.

B. Model comparison: All models

Although the previous results were interesting, they were all derived under the assumption that the scale factors actually differed from those of the SM. Following Sec. III, we would instead like to perform model comparison. We will use the priors discussed in Sec. III D and aim to

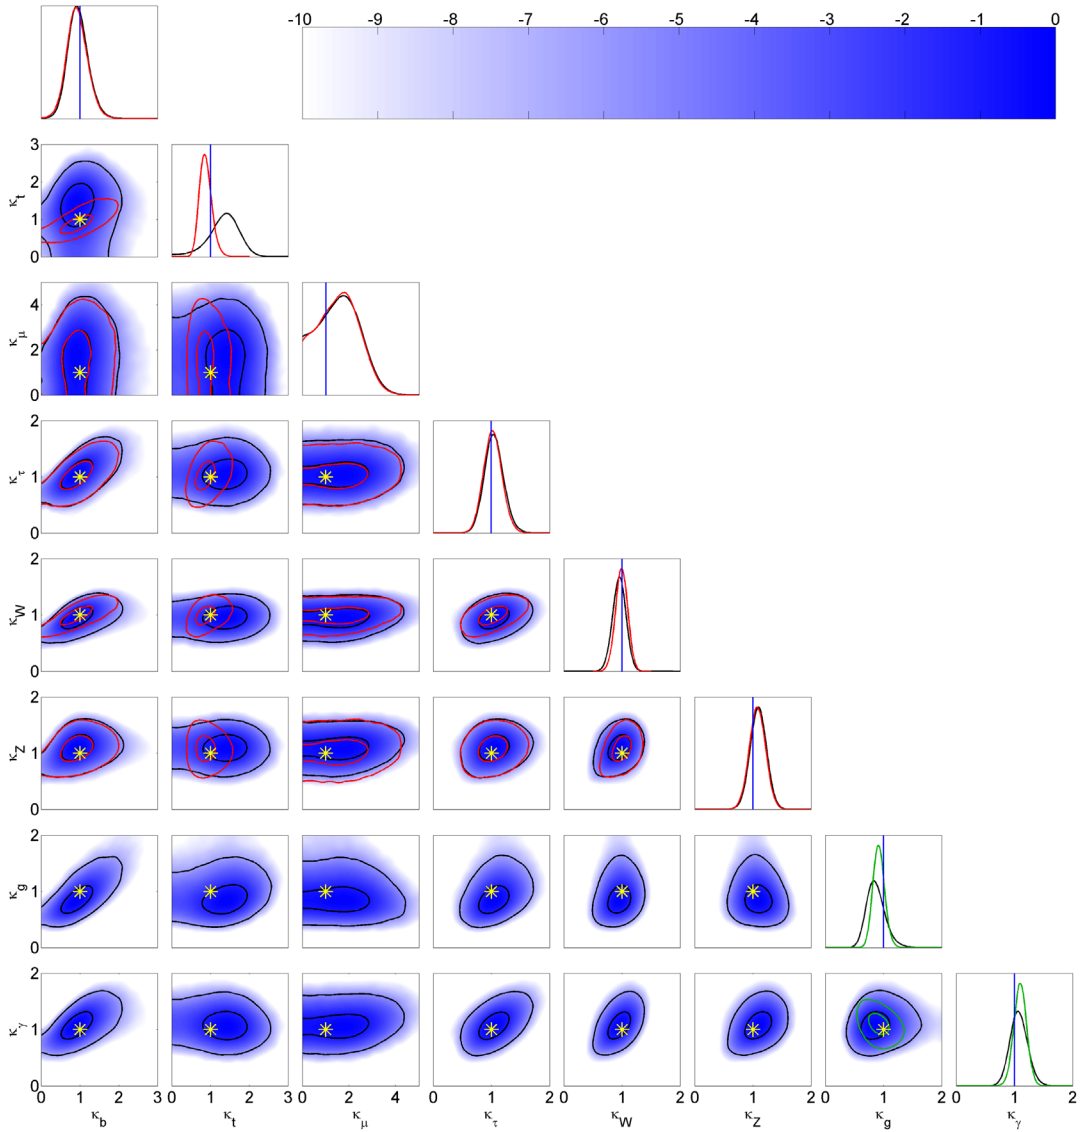


FIG. 1 (color online). Results of (default) parameter estimation. Two-dimensional log-posterior distribution (blue shading), 1σ and 3σ Bayesian credible regions, and one-dimensional posteriors (both black) of the eight-parameter fit. 1σ and 3σ credible regions and one-dimensional posteriors of the tree-level six-parameter fit (red) and the two-parameter fit (green). The SM values of unity are marked with vertical lines and stars, respectively.

evaluate how much the results depend on these different prior choices.

In this section, we follow Sec. III A and compare models with any combination of free parameters. In particular, we use MULTINEST with the priors in (15) and $p_i = p_0$ chosen so that the posterior over the space of models becomes as uniform as possible, and so all values of α will be sampled adequately [69]. There are in total 2^n models, with $n = 6$ when the tree-level couplings are free and $n = 8$ when also loop-processes are included.

In the left panel of Fig. 2, we present the logarithms of Bayes factors for all of the $2^6 = 64$ models, compared to the SM and using the logarithmic prior (on $[10^{-7}, 4\pi]$) for the tree-level couplings, i.e., the Yukawa couplings and

gauge boson couplings. The models are divided into unicolored groups depending on the number of couplings which are free. In the model to the far left in the figure, all couplings are free, the models in the next group have five parameters free, etc., until the model to the far right, which is the SM (and has no visible bar since $\log B = 0$). The blue stars are the values calculated by extrapolating the comparison of the SM with the models with a single coupling free and then assuming that adding an additional parameter has the same effect on $\log B$ regardless of the assumptions on the other parameters. This would be exact if the shape of the likelihood as a function of each parameter did not depend on the values of the other parameters. Although not exact, it seems that treating all of the parameters as

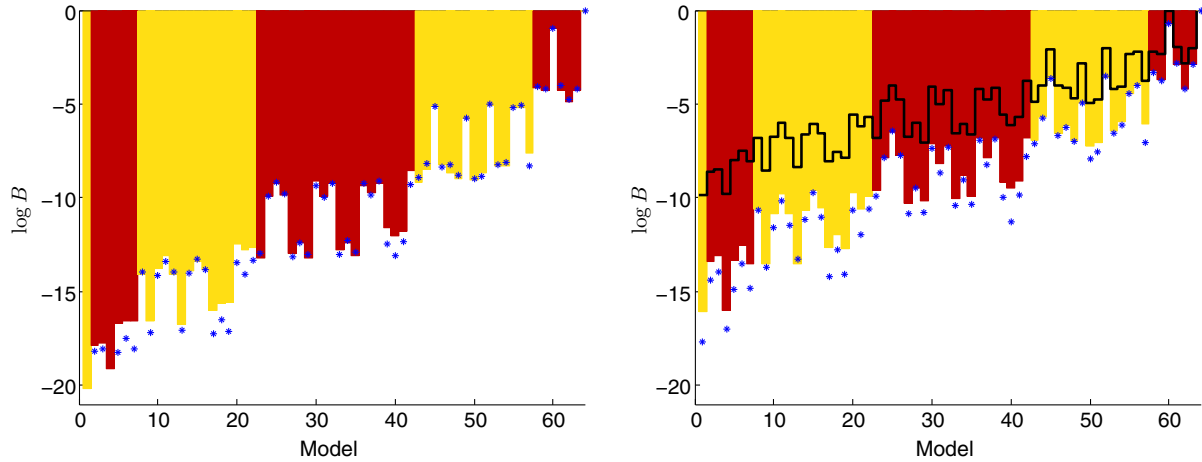


FIG. 2 (color online). Left: logarithms of Bayes factors (with respect to the SM) with the logarithmic prior on tree-level couplings. Each unicolored block of bars have the same number of free parameters, from the right: 0 (the SM), 1, etc., to the most general model with all couplings free to the far left. The blue stars are the values obtained by extrapolating the values of the single-coupling models, assuming independence. Right: same as the left plot but with Gaussian priors. The bars correspond to a standard deviation of $s = 4$, and the solid black line corresponds to $s = 1$.

independent gives a reasonable approximation for the model comparison.

As expected, there is a clear trend. The larger the number of free couplings, the smaller the values of $\log B$, i.e., the stronger the evidence against that model. Hence, the evidence against the model with all couplings free is very strong. Adding any of the parameters makes the model worse with about the same amount, with the exception of κ_μ , which only decreases the evidence with a small amount (roughly one log-unit). Letting κ_W free, corresponding to the most heavily constrained coupling, will have the largest effect on the evidence of the model.

Furthermore, one should remember that the log-odds only equals $\log B$ when the priors are equal. In this case, one might argue that the SM should have a larger prior than any of the other models, perhaps the same as all the other models together, which (assuming that prior is uniformly distributed) would lead to the log-odds being $\log 2^6 \approx 4$ smaller than the $\log B$'s in the plot. Again, we note that the dependence on the prior limits is very weak. For example, decreasing the lower limit to 10^{-15} would lead to a decrease of $\log B$ smaller than 0.7 for the addition of each coupling.

In the right panel of Fig. 2, we present $\log B$ for the same models, but with Gaussian priors on the coupling scale factors. The bars are obtained using a standard deviation of $s = 4$, and the solid black line is obtained using $s = 1$. Naturally, the choice of priors affects the exact values of the evidences, but the general trend is the same in all cases. Adding a parameter with a Gaussian prior is not as influential as adding one with a logarithmic prior, and the difference between the two Gaussian priors is only about one log-unit per parameter.

Next, we consider the case when also the loop-induced couplings are allowed to differ from the SM values, or rather those calculated in Eqs. (6) and (7), giving a total of

$2^8 = 256$ combinations of free couplings. The same Gaussian priors as in the right panel of Fig. 2 have been used, but with the expectation values of κ_g and κ_γ given by Eqs. (6) and (7), since this is the expectation without any contribution from new physics. The trend is similar to the previous case with tree-level couplings in that models with few free couplings are preferred to models with more free couplings. However, when approaching the models with most free parameters, there seems to be a “leveling off” in the sense that adding more parameters is less damaging. This makes sense because, if the parameter constraints deteriorates when more free parameters are added, the evidence will tend to be larger than what would otherwise be expected. Finally, in a similar way to the previous case, one might consider the SM not on equal footing with each of the other models, making the posterior odds smaller than the Bayes factor (now with $\log 2^8 \approx 5.5$ log-units).

C. Inclusion of individual couplings

In the previous section, we studied how all the different combinations of free couplings compared to each other. Although some conclusions could be drawn, the result was not completely transparent. In this section, we instead follow Sec. III B and evaluate the evidence for or against the inclusion of each individual coupling.

In Fig. 4, we give the logarithms of the Bayes factors in Eq. (16), i.e., *against* the inclusion of each of the couplings, both for the case of the six tree-level couplings with logarithmic priors and for the Gaussian priors on the scale factors. Here we use the value $s = 2$ for the standard deviation, although the difference from $s = 4$ and $s = 1$ as used previously is expected to be quite small. As in Eq. (17), the other (nuisance) couplings are either fixed to their SM values (S), allowed to vary with the same priors




FIG. 3 (color online). Logarithms of Bayes factors (with respect to the SM) for models with up to eight free parameters, with a Gaussian prior with standard deviations $s = 4$ (bars) and $s = 1$ (black line).

as the coupling of interest (\mathcal{F}), or averaged over these two cases (\mathcal{A}). However, in Eq. (17), the size of each contribution is proportional to the evidence of that particular model, and since typically the evidences are much larger when the other couplings equal their SM values [70], the average is dominated by these components. Hence, the result for \mathcal{A} equals that of \mathcal{S} to a very good approximation. Note that these Bayes factors are evaluated separately using dedicated MULTINEST runs. Hence, these numbers might differ somewhat from those which can be read from Figs. 2 and 3. The Bayes factors in the table have significantly smaller numerical errors of about 0.1.

Some general conclusions which can be drawn are that the logarithmic prior yields a stronger preference for the SM couplings than the Gaussian (as in previous chapter), and \mathcal{S} stronger than \mathcal{F} (which is reasonable since the constraints are relaxed).

The Higgs decay to $\mu^+\mu^-$ is rather weakly constrained, and the results for this coupling are quantitatively different to the other tree-level couplings. For the logarithmic priors, there is barely weak evidence in favor of the SM, while for the Gaussian case, there is not even that. Moving on to the other tree-level couplings, for the log prior, there is weak to moderate evidence for all the couplings, with \mathcal{F} giving about 1–2 log-units weaker preference than \mathcal{S} and \mathcal{A} . For the Gaussian prior, the evidence is also weak to moderate, but typically weaker than the logarithmic case. For the Gaussian prior for the tree-level couplings, there is no significant difference between the cases where the loop-induced couplings are free or not.

The loop-induced couplings enter only in two cases, both with a Gaussian prior. When the tree-level coupling scale factors are fixed, there is just moderate evidence in favor of




FIG. 4 (color online). Logarithms of Bayes factors against inclusion of couplings for the eight coupling scale factors. Values larger than 0 means the SM value of the coupling is preferred. The other couplings are either fixed to their SM values (\mathcal{S}) or allowed to vary with the same prior as the coupling of interest (\mathcal{F}) or averaged over these two cases (\mathcal{A}). Since typically the evidences are much larger when the other couplings equal their SM values, the average is dominated by these components, and hence \mathcal{A} yields essentially the same result as \mathcal{S} .

the SM values for both κ_g and κ_γ , while in the case when the other couplings are free, this preference essentially disappears completely.

However, as discussed in Sec. III B, making the weakest assumption on the tree-level couplings, Bayesian probability theory tells us that one really ought to use the model-averaged results (the cases with “ \mathcal{S} ” in Fig. 4). Hence, we conclude that the couplings moderately prefer the SM values for $b\bar{b}$, $t\bar{t}$, $W\bar{W}$, $Z\bar{Z}$, and $\tau^+\tau^-$ for both logarithmic and Gaussian priors. For $g\bar{g}$ and $\gamma\gamma$ the preference is barely moderate, and for the coupling to $\mu^+\mu^-$, the evidence is barely weak or none at all.

D. SM vs $\overline{\text{SM}}$

We consider the $\overline{\text{SM}}$ model as discussed in Sec. III C, with the most appropriate assumption is that all the couplings can either take their SM value, or differ from it, with a prior probability of 0.5 for each. The special case where all couplings simultaneously take on their SM values would typically be excluded from, but could also be included in, the $\overline{\text{SM}}$.

In Table II, we present the comparison of $\overline{\text{SM}}$ with the above model for these two cases and for the different continuous priors. When the SM part is excluded, the evidence for the SM is actually just about strong for the logarithmic and Gaussian (with $s = 2$) priors on all couplings and moderate for the case of tree couplings. In the present case, the evidence of the $\overline{\text{SM}}$ is dominated

TABLE II. $\log B$ between SM and $\overline{\text{SM}}$ for three different priors, in the two cases when the values of the couplings is either included in $\overline{\text{SM}}$ or not.

Prior	$\log B$
Log (with SM)	-3.8
Log (no SM)	-4.9
Gauss (tree, with SM)	-3.1
Gauss (tree, no SM)	-3.5
Gauss (all, with SM)	-4.4
Gauss (all, no SM)	-4.8

by the contribution from models with a single coupling free, weighted by their priors within the $\overline{\text{SM}}$.

In the second case, there is also a contribution from the part equivalent with the SM, which can be relatively large (and even dominating in the logarithmic case). Still, the conclusions do not change significantly, although the 1.1 log-units difference for the logarithmic prior takes the evidence for the SM from just about strong to moderate.

V. SUMMARY AND CONCLUSIONS

We have performed a Bayesian analysis of the LHC Higgs data and used an interim framework where the magnitude of the Higgs couplings is rescaled by coupling scale factors, whereas the tensor structure of the couplings is unaltered with respect to the SM. In the present work, we have limited our discussion to the couplings which are constrained by the LHC, in total six tree-level couplings and two loop-induced couplings.

We have performed Bayesian parameter inference on these coupling scale factors in the following three cases: either the tree-level couplings, the loop-level couplings, or both simultaneously, were free. In each case, the SM values

were well within the 1σ -region. However, when all couplings were free, neither κ_t nor κ_μ were well constrained and could in principle be zero.

Since the most important question is rather that of which model best describes the data, we have instead focused on the Bayesian model comparison, considering models with either only the tree-level couplings in the Lagrangian or all couplings, allowed to vary. In the first case, we used both a logarithmic prior, which was imposed directly on the tree-level couplings, and a Gaussian prior, imposed on the coupling scale factors. In the second case, when the loop-induced couplings were also treated as free parameters, the analysis was made with a Gaussian prior imposed on the coupling scale factors. In each case, we performed model comparison between models with one, several, or all of the couplings free. The larger the number of free parameters, the more disfavored the model was.

We have considered a single coupling at a time in the cases where the other couplings could either be fixed to the SM values or allowed to vary with the same prior as the coupling of interest. The favored models are those with the couplings fixed to the SM value, although the evidence is virtually nonexistent for the coupling to $\mu^+\mu^-$. All this was performed with the combinations of free parameters and priors discussed above. Finally, we discussed the definition of the model $\overline{\text{SM}}$ and compared this single model to the SM, finding that the SM is moderately to strongly favored.

ACKNOWLEDGMENTS

J. B. acknowledges partial support from the European Union FP7 ITN INVISIBLES (Marie Curie Actions, Grant No. PITN-GA-2011-289442). S. R. acknowledges support from the Swedish Research Council (Vetenskapsrådet, Contract No. 621-2011-3985).

-
- [1] G. Aad *et al.* (ATLAS Collaboration), *Phys. Lett. B* **716**, 1 (2012).
 - [2] S. Chatrchyan *et al.* (CMS Collaboration), *Phys. Lett. B* **716**, 30 (2012).
 - [3] T. Aaltonen *et al.* (CDF and D0 Collaborations), *Phys. Rev. D* **88**, 052014 (2013).
 - [4] ATLAS collaboration, Tech. Rep. ATLAS-CONF-2014-009.
 - [5] S. Chatrchyan *et al.* (CMS Collaboration), *Phys. Rev. D* **89**, 092007 (2014).
 - [6] S. Chatrchyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **01** (2014) 096.
 - [7] CMS collaboration, Tech. Rep. CMS-PAS-HIG-14-009.
 - [8] G. Aad *et al.* (ATLAS Collaboration), *Phys. Lett. B* **726**, 120 (2013).
 - [9] S. Chatrchyan *et al.* (CMS Collaboration), *Phys. Rev. Lett.* **110**, 081803 (2013).
 - [10] V. Khachatryan *et al.* (CMS Collaboration), arXiv:1411.3441.
 - [11] P. Higgs, *Phys. Lett.* **12**, 132 (1964).
 - [12] P. W. Higgs, *Phys. Rev. Lett.* **13**, 508 (1964).
 - [13] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Phys. Rev. Lett.* **13**, 585 (1964).
 - [14] F. Englert and R. Brout, *Phys. Rev. Lett.* **13**, 321 (1964).
 - [15] S. Chatrchyan *et al.* (CMS Collaboration), *Nat. Phys.* **10**, 557 (2014).
 - [16] ATLAS collaboration, Tech. Rep. ATLAS-CONF-2013-108.
 - [17] M. E. Peskin, arXiv:1312.4974.

- [18] A. David *et al.* (LHC Higgs Cross Section Working Group), [arXiv:1209.0040](#).
- [19] Tech. Rep. ATLAS-CONF-2014-010.
- [20] G. Belanger, B. Dumont, U. Ellwanger, J. F. Gunion, and S. Kraml, *Phys. Rev. D* **88**, 075008 (2013).
- [21] P. P. Giardino, K. Kannike, I. Masina, M. Raidal, and A. Strumia, *J. High Energy Phys.* **05** (2014) 046.
- [22] J. Ellis and T. You, *J. High Energy Phys.* **06** (2013) 103.
- [23] M. Baak, J. Cúth, J. Haller, A. Hoecker, R. Kogler, K. Mönig, M. Schott, and J. Stelzer, *Eur. Phys. J. C* **74**, 3046 (2014).
- [24] C. Englert, A. Freitas, M. M. Mühlleitner, T. Plehn, M. Rauch, M. Spira, and K. Walz, *J. Phys. G* **41**, 113001 (2014).
- [25] P. Bechtle, S. Heinemeyer, O. Stål, T. Stefaniak, and G. Weiglein, *J. High Energy Phys.* **11** (2014) 039.
- [26] P. M. Ferreira, R. Santos, J. F. Gunion, and H. E. Haber, *Phys. Rev. D* **89**, 115003 (2014).
- [27] G. Cacciapaglia, A. Deandrea, G. D. La Rochelle, and J.-B. Flament, *Phys. Rev. Lett.* **113**, 201802 (2014).
- [28] K. Cheung, J. S. Lee, and P.-Y. Tseng, *J. High Energy Phys.* **05** (2013) 134.
- [29] K. Cheung, J. S. Lee, and P.-Y. Tseng, *Phys. Rev. D* **90**, 095009 (2014).
- [30] P. Bechtle, O. Brein, S. Heinemeyer, G. Weiglein, and K. E. Williams, *Comput. Phys. Commun.* **181**, 138 (2010).
- [31] P. Bechtle, S. Heinemeyer, O. Stål, T. Stefaniak, and G. Weiglein, *Eur. Phys. J. C* **74**, 2711 (2014).
- [32] P. Bechtle, O. Brein, S. Heinemeyer, O. Stål, T. Stefaniak, G. Weiglein, and K. E. Williams, *Eur. Phys. J. C* **74**, 2693 (2014).
- [33] S. Chang, S. K. Kang, J.-P. Lee, K. Y. Lee, S. C. Park, and J. Song, *J. High Energy Phys.* **09** (2014) 101.
- [34] A. Celis, V. Ilisie, and A. Pich, *J. High Energy Phys.* **07** (2013) 053.
- [35] A. Celis, V. Ilisie, and A. Pich, *J. High Energy Phys.* **12** (2013) 095.
- [36] M. Montull, F. Riva, E. Salvioni, and R. Torre, *Phys. Rev. D* **88**, 095006 (2013).
- [37] M. Carena, L. Da Rold, and E. Pontón, *J. High Energy Phys.* **06** (2014) 159.
- [38] N. Craig, [arXiv:1309.0528](#).
- [39] K. Blum, R. T. D’Agnolo, and J. Fan, *J. High Energy Phys.* **01** (2013) 057.
- [40] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek, *J. High Energy Phys.* **10** (2010) 085 (2010).
- [41] R. Contino, M. Ghezzi, C. Grojean, M. Mühlleitner, and M. Spira, *J. High Energy Phys.* **07** (2013) 035.
- [42] G. Aad *et al.* (ATLAS Collaboration), *Phys. Lett. B* **738**, 68 (2014).
- [43] M. Pospelov and A. Ritz, *Phys. Rev. D* **84**, 113001 (2011).
- [44] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, England, 2003).
- [45] R. Cox, *Am. J. Phys.* **14**, 1 (1946).
- [46] T. J. Loredo, in *Maximum-Entropy and Bayesian Methods*, edited by T. Fougere (Kluwer Academic Publishers, Dordrecht, 1990), pp. 81–142.
- [47] Typically, “model” will refer to any assumption from which data can be predicted and is not necessarily a full, realistic physical model.
- [48] R. Trotta, *Contemp. Phys.* **49**, 71 (2008).
- [49] M. Hobson *et al.*, *Bayesian Methods in Cosmology* (Cambridge University Press, Cambridge, England, 2010).
- [50] J. A. Hoeting, D. Madigan, A. E. Raftery, and C. T. Volinsky, *Stat. Sci.* **14**, 382 (1999).
- [51] D. Parkinson and A. R. Liddle, *Stat. Anal. Data Min.* **6**, 3 (2013).
- [52] M. Vardanyan, R. Trotta, and J. Silk, *Mon. Not. R. Astron. Soc.* **397**, 431 (2009).
- [53] M. Vardanyan, R. Trotta, and J. Silk, *Mon. Not. R. Astron. Soc.* **413**, L91 (2011).
- [54] F. Feroz and M. Hobson, *Mon. Not. R. Astron. Soc.* **384**, 449 (2008).
- [55] F. Feroz, M. Hobson, and M. Bridges, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [56] F. Feroz, M. Hobson, E. Cameron, and A. Pettitt, [arXiv:1306.2144](#).
- [57] In the general, nonseparable case, all the quantities in the equation can depend on couplings κ_j for $j < i$.
- [58] This could be remedied by comparing with a model in which any of the couplings are free, but only one at a time.
- [59] Again, the constraint $\kappa_i = \kappa_i^{\text{SM}}$ can also be interpreted as holding only to a very good approximation, but not exactly.
- [60] See Ref. [71] for a similar application in the lepton sector of the SM and Ref. [72] for a more general discussion of naturalness in Bayesian inference.
- [61] M. Libanov and F.-S. Ling, [arXiv:1105.6035](#).
- [62] D. S. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial* (Oxford University Press, New York, 2006).
- [63] T. Corbett, O. J. P. Eboli, J. Gonzalez-Fraile, and M. C. Gonzalez-Garcia, *Phys. Rev. D* **87**, 015022 (2013).
- [64] T. Corbett, O. J. P. Eboli, J. Gonzalez-Fraile, and M. C. Gonzalez-Garcia, *Phys. Rev. D* **86**, 075013 (2012).
- [65] S. Fichet, *Nucl. Phys.* **B884**, 379 (2014).
- [66] B. Dumont, S. Fichet, and G. von Gersdorff, *J. High Energy Phys.* **07** (2013) 065.
- [67] J. de Blas *et al.*, [arXiv:1410.4204](#).
- [68] Defining the contours by the usual χ^2 -thresholds on $-2 \log(\mathcal{L}(\theta)/\mathcal{L}_{\max})$, with $\mathcal{L}(\theta)$ the Bayesian marginal likelihood, yields virtually identical contours in all cases.
- [69] We check that the statistical error on each Bayes factor is reasonably small by considering the effective sample size $(\sum w_i)^2 / \sum w_i^2$, with the w_i ’s the weight of each sample belonging to a certain model. Typically, it is of the order of $10^2 - 10^3$, but in a few cases, it is slightly smaller.
- [70] The muon coupling is an exception, but since this is due to a lack of constraints rather than the existence of a tension with the SM value, this has no effect.
- [71] J. Bergström, D. Meloni, and L. Merlo, *Phys. Rev. D* **89**, 093021 (2014).
- [72] S. Fichet, *Phys. Rev. D* **86**, 125029 (2012).