

A FREQUENCY-CALIBRATED BAYESIAN SEARCH FOR NEW PARTICLES

BY SHIRIN GOLCHI¹ AND RICHARD LOCKHART¹

Simon Fraser University

The statistical procedure used in the search for new particles is investigated in this paper. The discovery of the Higgs particles is used to lay out the problem and the existing procedures. A Bayesian hierarchical model is proposed to address inference about the parameters of interest while incorporating uncertainty about the nuisance parameters into the model. In addition to inference, a decision making procedure is proposed. A loss function is introduced that mimics the important features of a discovery problem. Given the importance of controlling the “false discovery” and “missed detection” error rates in discovering new phenomena, the proposed procedure is calibrated to control for these error rates.

1. Introduction. The standard model (SM) of particle physics is a theory that describes the dynamics of subatomic particles. The Higgs particle is an essential component of the SM; its existence explains why other elementary particles are massive [Englert and Brout (1964), Guralnik, Hagen and Kibble (1964), Higgs (1964), Guralnik (2009)]. The existence of the Higgs boson was confirmed by experiments run at the Large Hadron Collider (LHC) at the European organization for nuclear research, known as CERN, in 2012 [CMS Collaboration (2012), ATLAS Collaboration (2012)]. LHC is a high energy collider specifically designed and constructed to detect the Higgs particle. Two beams of protons circulating at very high energies in the LHC collide inside two detectors (ATLAS and CMS). Collisions between a proton in one beam and a proton in the other beam result in generation of new particles, possibly including the Higgs boson; each such collision is an event. Some of the particles generated can be tracked and measured in the detectors. However, the Higgs particle, if generated, decays extremely quickly into other known SM particles and cannot be detected directly. Instead, the existence of the Higgs particle must be inferred by looking for those combinations of detectable particles that are predicted by the SM.

Once a Higgs particle has been created by one of several “production mechanisms” in a proton–proton collision, there are several different processes, called “decay modes”, through which the particle may decay. The decay process can

Received May 2016; revised January 2018.

¹Supported by Natural Sciences and Engineering Research Council of Canada—NSERC.

Key words and phrases. Bayes rule, decision set, Higgs boson, linear loss function, sequential Monte Carlo.

be reconstructed based on the detected collision byproducts. Events with reconstructed processes that match one of the possible Higgs decay modes and pass other selection criteria (called cuts) are recorded as “Higgs candidates”. The “invariant mass” of the unobserved particle is computed from masses, energies and momenta of the decay products such that it reflects the conservation of mass-energy. A histogram of the estimator of the mass is then created for each decay mode (some analyses including $H \rightarrow 2$ photons, i.e., Higgs decays to two photons, use unbinned likelihood fits instead of histograms). However, there are other processes, not involving the Higgs boson, that can result in the generation of Higgs event byproducts which also pass the cuts; these are called background events. Thus the histogram created is either a mixture of background events and events in which a Higgs particle was created or just a histogram of background events if the Higgs particle does not exist.

Luckily, the SM predicts, as a function of the mass of the Higgs particle and the energy of the colliding beams, the expected rate at which events generate Higgs particles (a quantity proportional to the so-called Higgs cross section). It also predicts, as a function of the same parameters, the probability that the particle will decay by a given decay mode and produce byproducts which pass the cuts. Effectively then, the SM predicts that if the Higgs particle exists and has mass m_H then there will be a bump on the histogram of invariant masses whose size and shape are completely predicted by unknown mass, m_H , and other measured quantities such as beam energy. The statistical problem is to determine whether or not such a bump exists and if so at what mass it is centered.

The search procedure has two stages. The first stage is discovery, or search, and the second is called exclusion. In the first stage, for each mass in a range of possible masses the null hypothesis that there is no Higgs particle is tested against the alternative that it exists and has this specific mass. A p -value, called a local p -value, is computed for each mass; this local p -value is then plotted against mass, typically on a log-scale, and the smallest value of this local value is reported as a quantity of significance [see Figure 7 and the discussion in Section 7 in [ATLAS Collaboration \(2012\)](#) and Figure 15 and related discussion in [CMS Collaboration \(2012\)](#)]. At the same time a global p -value is computed by approximating the distribution of the smallest local p -value. A claim of “discovery” should require this global p -value to be less than $1 - \Phi(5)$ (a 5σ effect— Φ is the standard normal cumulative density function). The second stage is referred to as exclusion where for each mass in the range under consideration the null hypothesis that the particle exists, at this mass, with the predicted signal strength, is tested against a background only hypothesis.

In the searches for the Higgs particle described in [CMS Collaboration \(2012\)](#) and [ATLAS Collaboration \(2012\)](#), in the discovery phase, theoretical predictions are used to define the signal function for different production mechanisms and analysis categories. An overall signal strength parameter in the form of a unitless scaling factor, generally denoted by μ , is used in the model. The standard model predicts that $\mu = 1$ and the null hypothesis that there is no Higgs is represented

by $\mu = 0$. The procedures described in [CMS Collaboration \(2012\)](#) and [ATLAS Collaboration \(2012\)](#) treat the alternative to $\mu = 0$ as $\mu > 0$. Treating μ as a free parameter. It is worth noting that the signal strength parameter μ is the focus of the local likelihood ratio tests since the mass is fixed for each test and the background model parameters and any other unknown parameters of the signal function are treated as nuisance parameters. Consequently, simultaneous two dimensional confidence intervals for the mass of the Higgs particle and the signal strength parameter are reported for different decay modes.

In order to fix some notation we now give some further details of the two stages for the current procedure. Some key weaknesses of the existing method are highlighted in the process.

Search: For each mass value, $m \in (m_0, m_n)$, in the search window, a likelihood ratio test [see [Cowan et al. \(2011\)](#) for details] is performed for the hypotheses

$$(1.1) \quad H_0 : \mu = 0 \quad \text{vs.} \quad H_A : \mu > 0.$$

Note that the overall signal strength, $\mu = 1$, is not used under the alternative. For each $m \in (m_0, m_n)$, a p -value is obtained. The smallest of these local p -values is described as the *local significance* and the mass at which this minimum is achieved is reported as an initial estimate of the mass of the Higgs particle. We note that this is not the final estimate of the mass; a further detailed analysis of the decay products in the relevant events is used to provide an estimate with uncertainties. The plot of local p -values does not itself provide an interval estimate for the mass of the detected particle.

This local significance corresponds to a testing procedure which rejects the null hypothesis if any of a family of local test statistics, $v(m)$ (e.g., log likelihood ratios) indexed by the unknown mass m is larger than a predefined level κ . [Gross and Vitells \(2010\)](#) proposed a method for estimating the “global p -value” of this testing procedure as the null probability that the maximum of the local test statistics is greater than the observed maximum. Their method is based on the results of [Davies \(1987\)](#) and gives

$$(1.2) \quad \begin{aligned} P_G &= P_{H_0} \left(\max_m v(m) > \kappa \right) \\ &\approx P_{H_0} (\chi_d^2 > \kappa) + E_{H_0} (N(\kappa)), \end{aligned}$$

where κ in this case is the observed maximum statistic, χ_d^2 , denoting a chi-squared distribution with d degrees of freedom, is the null distribution of $v(m)$ and $E_{H_0} (N(\kappa))$ is the expected number of upcrossings of the level κ by the process $v(m)$. Using the method of [Gross and Vitells \(2010\)](#), one can estimate the (global) type I error rate associated with the discovery procedure that is based on

the local tests,

$$\begin{aligned}
 \alpha_G &= P_{H_0} \left(\max_m v(m) > \kappa_{\alpha_0} \right) \\
 (1.3) \quad &\approx P_{H_0} (\chi_d^2 > \kappa_{\alpha_0}) + E_{H_0} (N(\kappa_{\alpha_0})) \\
 &= \alpha_0 + E_{H_0} (N(\kappa_{\alpha_0})),
 \end{aligned}$$

where κ_{α_0} is the χ_d^2 quantile corresponding to α_0 . Since $E_{H_0} (N(\kappa)) \geq 0$, the actual global type I error rate is larger than the controlled local type I error rates. The size of the difference depends on the specific statistical model and on the search range. For this very reason, discovery is announced if the global p -value shows a 5σ effect.

Exclusion: In the second stage, further investigation is done to exclude regions of m which are unlikely values for the mass of the Higgs boson. The theoretical signal strength is tested at significance level $\alpha_2 = 0.05$ at the exclusion step [Cowan et al. (2011)],

$$(1.4) \quad H_0 : \mu = 1 \quad \text{vs.} \quad H_A : \mu < 1.$$

Any mass whose corresponding theoretical signal strength is rejected is excluded from the range of possible masses for the Higgs particle.

Some implications of the existing procedure are as follows; the two stages serve different purposes: the search stage is sensitive to a variety of signals rather than only that predicted by the SM, potentially searching for any Higgs-like phenomenon. The exclusion stage on the other hand is designed to exclude mass values at which a signal exactly matching the predicted signal is not observed. Different significance levels are used for the two steps which implies that much more caution is taken for testing for signal associated with a new phenomenon than for testing if an observed signal matches the predicted signal.

Some care is needed in interpreting the local and global significance levels particularly when both are presented in the same analysis. Use of a global 5σ rule for declaring “discovery” controls the rate at which incorrect discoveries are announced provided that the rule is applied without any prior screening. To make this clear imagine the following reporting procedure: if the local significance level exceeds 5σ we compute and report the global p -value defined above. If the local significance level does not reach 5σ we neither compute nor report the global p -value. Then the interpretation of any global significance level exceeding 5σ is clear because whenever the global significance level exceeds 5σ so does the local significance level. Thus the global p -value will be computed in every case where it would exceed 5σ and in those cases bears the same interpretation as any small p -value. But when the global p -value is 4.5σ , say, we might or might not report it depending on the local significance level. Thus we cannot easily compute the probability that a reported global significance level will exceed 4.5σ .

It will be seen that in [ATLAS Collaboration (2012)] the word “discovery” is used once in the paper (in the conclusion section); the global significance for a search over the range 110 to 600 GeV is 5.1σ . The word “discovery” is used in [CMS Collaboration (2012)] only in the introduction to describe the goal because the global significance, even over the narrower range of 115 to 130 GeV is only 4.6σ . (We note that both paper titles begin by “Observation of a new . . .” and that except for the crucial word the discussions are quite similar.)

Both papers compute global p -values over more than one mass range, often increasing the global significance by searching over a smaller range. These smaller ranges are not without motivation; they often correspond to searching over mass ranges not excluded in earlier searches.

Data analysis does not end with discovery and exclusion. Having declared a discovery, for instance, confidence intervals for the mass and joint confidence sets for the signal strength parameter μ will be computed. The model will be assessed by checking to see if the signal strength is close to 1 and by checking to see if the μ 's for different production mechanisms appear to be the same. In this paper we focus on the discovery/exclusion part of the problem in a partially Bayesian framework and provide Bayesian tools for this post-discovery analysis.

In the following we consider two strategies. First we propose a purely Bayesian approach where we treat the discovery of the Higgs particle as an inference problem. A Bayesian hierarchical model is defined that captures the main features of the particle discovery problem. Treating the mass of the Higgs boson as the main parameter of interest we estimate the background, signal strength parameter and the hyper-parameters in a fully Bayesian framework where uncertainties about the estimates are expressed by the posterior variance. A joint credible set for the mass and signal strength is proposed as the basis of discovery.

Our second strategy is decision theoretic. Given that a larger penalty is associated with false discovery than with missed detection, we propose a hybrid Bayes-frequentist decision making procedure that allows for controlling frequency error rates at desired levels while using a Bayesian test statistic such as the posterior odds. The proposed method can be perceived as a unified search procedure that takes advantage of the flexibility of the Bayesian framework while adhering to frequency theory requirements.

Our procedure is similar to that of Feldman and Cousins (1998), a generalization of which [Cowan et al. (2011)] is used for construction of two dimensional confidence intervals for mass and signal strength. Feldman and Cousins (1998) discuss the coverage issues of confidence intervals that are built based on the two step procedure without proper conditioning and propose a method for constructing confidence intervals based on ranking the ratio of the likelihood under null and alternative models. Our method is similar to this approach in that the posterior odds are used for constructing the decision set. The final decision set is interpreted as a confidence interval with different confidence levels for values of the parameter (mass) that represent the null and the alternative hypotheses.

Note that systematic errors are not considered in the present work. However, such errors are naturally included in a Bayesian model as priors. More specifically, some systematic errors are difficult to deal with in a non-Bayesian framework. For instance, in computing the signal shape and cross section as a function of beam luminosity it is necessary to carry out quantum mechanical calculations by expansions where the result is an approximation error that typically can only be quantified subjectively. Incorporating that subjectivity explicitly in a prior offers the opportunity to examine the sensitivity of the conclusions to prior uncertainty.

The rest of the paper is organized as follows. In Section 2, we introduce a Bayesian hierarchical model and describe inference as a Bayesian search for the Higgs particle. The inference results are reported for simulated data provided by the CMS group. Section 3 is dedicated to the Bayesian decision making procedure that, while combining the discovery and exclusion steps, is calibrated to obtain desired frequency-theory error-rates. Approximation and sampling techniques are introduced for efficient implementation of calibration. Section 4 follows with concluding remarks and discussion of extensions and future work. In addition, a supplementary document is provided that includes the proposed approach applied to a simple signal detection problem [Golchi and Lockhart (2018)].

2. Bayesian inference. In this section, we propose a Bayesian hierarchical model that captures the features of the problem of discovery of the Higgs particle. The goal is to make inference about two parameters of interest, the mass of the Higgs particle and the signal strength parameter, while adequately handling the nuisance parameters that specify the background. The proposed model is fit to simulated data provided to us by the CMS group and the results are presented.

2.1. Model. Suppose that the data, that is, the invariant masses recorded by the detector, are realizations of a Poisson process whose intensity function is given by the sum of a background intensity function $\Lambda(m)$ and a signal intensity function $s_{m_H}(m)$. The shape of the signal function is known and its location is determined by the unknown parameter, $m_H \in \mathcal{M}$, where $\mathcal{M} = \{\emptyset\} \cup (m_0, m_n)$ [$(m_0, m_n) \subset \mathcal{R}^+ - \{0\}$, that is, m_0 and m_n are strictly positive]. The parameter, m_H , is the unknown mass of the Higgs particle where $m_H \in (m_0, m_n)$ means that the Higgs boson has a mass in the search window, (m_0, m_n) , while $m_H = \emptyset$ refers to the case that the particle does not exist, at least not with a mass in (m_0, m_n) . Note that m_H is the parameter of interest and in the Bayesian framework is treated as a random variable; it should not be confused with m that is used to denote an arbitrary but fixed mass in \mathcal{M} .

We quantify uncertainty in the background by considering the logarithm of the intensity function to be a realization of a Gaussian process,

$$(2.1) \quad \log \Lambda_{\beta, \eta, \sigma^2}(m) \sim \mathcal{GP}(\xi_{\beta, \sigma^2}(m), \rho_{\eta, \sigma^2}(m, m')), \quad m \in (m_0, m_n),$$

with covariance function,

$$(2.2) \quad \rho_{\eta, \sigma^2}(m, m') = \sigma^2 \exp(-\eta(m - m')^2),$$

where σ^2 is the variance parameter and η is the correlation parameter that controls the smoothness of the background function.

We parametrize the mean function, $\xi_{\beta}(m)$, such that the expectation of the background, $E(\Lambda)$, has one of the typical parametric forms currently used to model the background function, for example a fourth order Bernstein polynomial [ATLAS Collaboration (2012)]. Therefore,

$$(2.3) \quad \xi_{\beta, \sigma^2}(m) = \log\left(\sum_{i=0}^4 \beta_i h_i(z(m))\right) - \frac{\sigma^2}{2},$$

where $z : \mathcal{M} \rightarrow (0, 1)$ is an affine transformation of mass onto the unit interval and the basis functions are given by

$$(2.4) \quad h_i(z) = \binom{4}{i} z^i (1 - z)^{4-i}, \quad z \in (0, 1).$$

However, we note that the choice of Bernstein polynomials is commonly made since they are constrained to be positive. Since under our model the log-Gaussian prior guarantees the positivity of background, an unconstrained parametric form may be used for the mean function.

The notation $\Lambda_{\beta, \eta, \sigma^2}$ is used to show the dependence of the background function on the hyper-parameters. For the sake of brevity the subscript is dropped from here on.

We choose the signal function as a Gaussian probability density function with the location parameter m_H [in the current practice [ATLAS Collaboration (2012), CMS Collaboration (2012)] a slightly more complex signal shape called the ‘‘crystal ball function’’ is used]. Thus, the signal function is given by

$$(2.5) \quad s_{m_H}(m) = c_{m_H} \phi\left(\frac{m - m_H}{\varepsilon}\right) \quad \text{for } m_H \in (m_0, m_n),$$

$$(2.6) \quad s_{\emptyset}(m) = 0,$$

where c_{m_H} is a scaling constant (analogous to the cross section), and ϕ is the standard normal probability density function. The standard deviation, ε , controls the spread of the signal function.

The use of finely binned data is common in the physics literature since the size of the data collected is often large. The likelihood function for the parameters is given by

$$(2.7) \quad \pi(\mathbf{y} | m_H, \mu, \Lambda) = \prod_{i=1}^n \frac{\exp(-\Gamma_i) \Gamma_i^{y_i}}{y_i!},$$

where

$$(2.8) \quad \Gamma_i = \int_{m_{i-1}}^{m_i} [\Lambda(m) + \mu s_{m_H}(m)] dm.$$

Here μ is the signal strength, that is, a unitless parameter that allows for credibility of signal sizes that do not match the one predicted by the theory (equivalent to $\mu = 1$). The grid $\mathbf{m} = (m_0, m_1, \dots, m_n)$ is the vector of bin boundaries over the search window. In practice, bin-wise expectations γ_i are distributed according to a multivariate normal distribution whose covariance matrix is obtained by applying (2.2) to the bin centers. This simplification is made to avoid integration of the log-Gaussian process at every evaluation of the likelihood for more efficient computations.

The posterior distribution of the model parameters $\boldsymbol{\theta} = (m_H, \mu, \Lambda, \boldsymbol{\beta}, \eta, \sigma^2)$ given the data \mathbf{y} can be written as

$$(2.9) \quad \pi(\boldsymbol{\theta} | \mathbf{y}) = \frac{\pi(\boldsymbol{\theta})\pi(\mathbf{y} | \boldsymbol{\theta})}{\int \pi(\boldsymbol{\theta})\pi(\mathbf{y} | \boldsymbol{\theta}) d\boldsymbol{\theta}},$$

where

$$(2.10) \quad \pi(\boldsymbol{\theta}) = \pi(m_H)\pi(\mu)\pi(\Lambda | \boldsymbol{\beta}, \eta, \sigma^2)\pi(\boldsymbol{\beta})\pi(\eta)\pi(\sigma^2).$$

We describe the prior distribution on \mathcal{M} by specifying a density with respect to the measure ν which puts a point mass on \emptyset and normalized Lebesgue measure on (m_0, m_n) . That is,

$$(2.11) \quad \nu(A) = \mathbb{1}(\emptyset \in A) + \int_{A \cap (m_0, m_n)} \frac{dm}{m_n - m_0},$$

where $\mathbb{1}(\cdot)$ is an indicator function. The prior on m_H is specified as a density, $\pi(m_H)$, with respect to ν given by

$$(2.12) \quad \pi(m_H) = 0.5.$$

This prior implies that a priori the two events, ‘‘Higgs particle does not exist’’ and ‘‘Higgs particle exists with mass $m_H \in (m_0, m_n)$ ’’ are equally likely. Then given that the Higgs particle exists, all the mass values in the search window are also equally likely. For clarity the prior probability of a subset A of \mathcal{M} is

$$(2.13) \quad \begin{aligned} \Pi(A) &\equiv \int \mathbb{1}(m \in A)\pi(m)\nu(dm) \\ &= \frac{1}{2}\mathbb{1}(\emptyset \in A) + \frac{1}{2} \int_{A \cap (m_0, m_n)} \frac{dm}{m_n - m_0}. \end{aligned}$$

The choice of a particular prior for a real physical parameter like the Higgs mass will inevitably be controversial. Our frequency adjusted decision theoretic approach described in Section 3 will remove much of the impact of the prior; the point is discussed further there.

The signal strength parameter, μ , is assigned a log-normal prior with mean 1 and standard deviation of 0.6 which allows signals of about two tenth to about three times the size of the theoretical signal to remain credible under the prior. Note that the signal strength is meant for searching alternative theories to the SM. Therefore, the prior distribution needs to be specified in accordance with credible alternative signal sizes.

The background prior, $\pi(\Lambda \mid \boldsymbol{\beta}, \eta, \sigma^2)$ is given by (2.1) and the mean hyperparameters, $\boldsymbol{\beta}$, have diffuse normal priors. The correlation parameter, η , is assigned an inverse-Gamma prior with shape and scale parameters equal to one. This prior is considered weakly informative given the affine transformation of m onto $(0, 1)$ explained above. The prior over the variance parameter σ^2 is $\pi(\sigma^2) \propto \frac{1}{\sigma^2}$ which is recommended in the literature as a reference prior [Paulo (2005)].

2.2. Computational issues. The mixture prior on m_H results in difficulties in Markov chain Monte Carlo (MCMC) sampling from (2.9). More specifically, moving efficiently between \emptyset and (m_0, m_n) in the parameter space is not trivial which leads to poor mixing of the Markov chain. To overcome computational difficulties we use sequential Monte Carlo (SMC) [Doucet, de Freitas and Gordon (2001), Del Moral, Doucet and Jasra (2006)]. The SMC samplers are a family of algorithms that take advantage of a sequence of distributions that bridge between a distribution that is straightforward to sample from (for example the prior) and the target distribution. Samples are filtered through the sequence of densities using importance sampling to obtain a sample from the target distribution.

A common approach for defining the sequence of densities between the initial and target distributions is by inducing the likelihood in the model sequentially. Let the filtering sequence of distributions be denoted by

$$(2.14) \quad \pi_0, \pi_1, \dots, \pi_T.$$

Using an annealing schedule $\{\tau_t, t = 0, \dots, T\}$, the t th distribution in the sequence is defined as a power posterior,

$$(2.15) \quad \pi_t = \pi(\boldsymbol{\theta})[\pi(\mathbf{y} \mid \boldsymbol{\theta})]^{\tau_t},$$

where

$$(2.16) \quad 0 = \tau_0 < \tau_1 < \dots < \tau_T = 1.$$

The SMC sampler comprises iterative steps of weighting and sampling. While the samples are moved toward the target distribution through re-sampling with weights calculated according to the current temperature, they are also moved toward higher probability regions under each distribution, using MCMC steps. These steps are explained in Algorithm 1.

The form of the incremental weights \tilde{w}_i depends on the choice of transition kernel K_t in SMC. For a variety of choices for the forward kernel and importance

Algorithm 1 Sequential Monte Carlo**Input:** A temperature schedule $\{\tau_t, t = 0, \dots, T\}$ A MCMC transition kernel K_t 1: Generate an initial sample $\theta_0^{1:N} \sim \pi_0$;2: $W_1^{1:N} \leftarrow \frac{1}{N}$;3: **for** $t := 1, \dots, T - 1$ **do**

- $W_t^i \leftarrow W_{t-1}^i \frac{\tilde{w}_t^i}{\sum \tilde{w}_t^i}$ where $\tilde{w}_t^i = \pi(\mathbf{y} | \theta_t^{(i)})^{\tau_t - \tau_{t-1}}$, $i = 1, \dots, N$;

- Re-sample the particles $\theta_t^{1:N}$ with importance weights $W_t^{1:N}$;

- $W_t^{1:N} \leftarrow \frac{1}{N}$;

- Sample $\theta_{t+1}^{1:N} \sim K_t$;

4: **end for****Return:** Particles $\theta_T^{1:N}$.

weights in SMC see [Del Moral, Doucet and Jasra \(2006\)](#). In [Algorithm 1](#) K_t is chosen as an MCMC transition kernel that results in the simplified form of the incremental weights as the proportion of two consecutive densities evaluated for the current sample.

The proposal distributions used for the the MCMC step are as follows. For the background at time t a log-Gaussian distribution is used that has a mean equal to the background at time $t - 1$ and covariance matrix following the prior covariance structure but scaled according to the posterior variances at time t . The proposal distribution for m is determined by estimating the marginal distribution of m given \mathbf{y} at time $t - 1$ as

$$(2.17) \quad \hat{\pi}_t(m | \mathbf{y}) = \begin{cases} \frac{\sum_{i=1}^N \mathbf{1}(m_t^i = \emptyset)}{N}, & m = \emptyset, \\ \frac{\sum_{i=1}^N \mathbf{1}(m_t^i \neq \emptyset)}{N} \zeta(m), & m \neq \emptyset, \end{cases}$$

where $\zeta(m)$ is a kernel density estimate of the distribution of $m \neq \emptyset$ at time t . The rationale for this proposal distribution is that for efficient Metropolis–Hastings steps, the proposal distribution needs to be chosen close to the target distribution. Construction of the proposal based on π_{t-1} is based on the assumption that π_{t-1} is close to π_t . A normal proposal is used for μ and the background mean hyper-parameters β and the covariance hyper-parameters are generated from chi-squared proposal distributions.

The temperature schedule, $\{\tau_t\}_{t=1}^T$, is specified adaptively using the approach proposed by [Jasra et al. \(2011\)](#): at each step t , τ_t is determined such that the effective sample size (ESS) is equal to a pre-specified value such as half the sample

size. This is achieved by solving the following equation numerically for τ_t ,

$$(2.18) \quad \text{ESS} = \frac{(\sum_{n=1}^N w_n^t(\tau_t))^2}{\sum_{n=1}^N (w_n^t(\tau_t))^2}.$$

The sampling is stopped when $\tau = 1$ satisfies the above condition.

Note that in simpler models, such as the one introduced for the on/off problem in [Golchi and Lockhart \(2018\)](#), the above computational undertakings would not be necessary. The posterior in [Golchi and Lockhart \(2018\)](#) is obtained by more direct numerical calculations.

2.3. Data and inference results. Confidentiality rules do not allow access to the real data for nonmembers of the Higgs research groups. In this section we fit the proposed model to simulated data provided to us by Matthew Kenzie of the CMS group and described in [CMS Collaboration \(2014\)](#) and [CMS Collaboration \(2013\)](#). The simulation procedure is very complex because it must model not only the predicted behavior of the Higgs boson but also the behavior of the extremely complex CMS detector. Analysis of such simulated data was an essential step in developing the analytic techniques to be used for the real experiment. The simulated data available to us represent the diphoton decay mode invariant mass spectrum (in the range $100 < m_{\gamma\gamma} < 180$ GeV) at centre of mass energy $\sqrt{s} = 8$ TeV. For each of the decay modes there are different Higgs signatures referred to as analysis categories. For the diphoton decay mode there are nine analysis categories. We had access to a list of the invariant mass of each data event together with the corresponding analysis category.

There are several ‘‘production mechanisms’’ through which a Higgs particle can be generated; five such mechanisms were considered for our data. Each such production mechanism leads to a specific predicted signal function. While the production mechanism is not identified in the data the signal function is propagated through the analysis separately for each of the SM Higgs production mechanisms at the LHC and each analysis category. The shape of the signal function in each analysis category, for each of the production modes and at three hypothesized Higgs masses (120, 125 and 130 GeV) has been provided to us in the form of a histogram; the entry in a single bin of such a histogram is the expected number of Higgs events produced by a specific production mechanism in a specific analysis category in the mass range for that bin if the Higgs has the particular hypothesized mass. A handful of the production mechanism, analysis category combinations produce so few expected outcomes that we were not provided the corresponding signal histograms; in the end we have histograms for 41 of the 45 combinations.

We fit the signal function, (2.6), to these histograms and estimate the signal strength, c_{m_H} , and signal width, ε . We then extrapolate from the three signal masses we were given to obtain the signal function corresponding to other masses. We pool the data for all nine analysis categories and bin the data according to the




FIG. 1. *Simulated data representing the invariant masses of events.*

signal histograms with 322 bins. Figure 1 shows the histogram of the data. The reason for pooling the analysis categories is that some of these categories have very few data points and that the computational burden of a full analysis exceeded our capabilities. We also use the sum of the signal functions over the production modes and analysis categories as a single signal function for each mass.

Note that, in principle, our method can handle the complications produced by having several combinations of production mechanism/decay mode/analysis category. However, simulated data is provided to us for only one decay mode and in practice, low event counts in some of the analysis categories can be a source of problems in making inference separately in these categories.

The results are summarized in Figures 2(a), 2(b) and 3. Figure 2(a) shows the marginal posterior density estimates of mass over five selected steps of the sequential algorithm. The colour of the curves gets darker as the posterior gets closer to




FIG. 2. *The results of the analysis of the simulated data representing (a) the posterior density of the mass evolving through the sequential sampler as the likelihood is induced into the power posterior sequentially; the estimated mass of the Higgs particle ($\hat{m}_H \approx 126$) is specified by the vertical red line, and (b) the joint posterior density contours of mass and signal strength; the red cross shows the position of mass and μ pair that maximizes the posterior.*




FIG. 3. The results of the analysis of the simulated data representing the evolution of the 95% credible intervals for the background sample paths from the prior (light blue bands) to the posterior (dark blue bands). The red dots are the observations; the posterior background mean and the background posterior mean plus the estimated signal are shown by the solid dark blue line and the dashed line, respectively.

the target posterior. The red vertical line shows the maximum a posteriori mass value ($\hat{m}_H \approx 126$). A 95% credible interval for the mass of the Higgs particle based on the analysis of the simulated data is (124.58, 127.30).

Figure 2(b) presents the joint posterior of mass and the signal strength parameter. The contours are 2-d kernel density estimates of the joint posterior and the red cross shows the (m_H, μ) pair that maximizes the posterior.

In Figure 3 the blue bands show the 95% credible intervals for the background sample paths; the wide bands show the prior uncertainty for the background function. The background uncertainty decreases at each step of SMC as the data outweighs the prior. The background posterior mean and the background posterior mean plus the estimated signal are shown as a solid dark blue line and a dashed line, respectively.

3. Decision making. In this section we consider the problem from a decision theoretic point of view. We define a linear loss function and derive the Bayes rule that can be used as an alternative to the current discovery/exclusion method for reporting one or more possible mass values for the Higgs particle. The Bayes procedure is calibrated to satisfy specified frequency theory error rates.

3.1. *Structure.* The required ingredients of a decision theory problem are a model with the corresponding parameter space, a decision space which is a set of possible actions to take, and a loss function [Berger (1980)].

The model was introduced in Section 2.1. However, the procedure we now suggest could be used regardless of the specific details of the model. We define the decision space as the set of all possible subsets, $S \subset \mathcal{M}$ where \mathcal{M} was defined in Section 2.1 as the union of the interval (m_0, m_n) and the single point \emptyset that represents the case that the Higgs particle does not exist. The interpretation of S is that, $m \in S$ if, having observed the data, we wish to retain m as a possible value of the true mass. For instance, if $\emptyset \in S$, the results suggest that it is possible that the Higgs particle does not exist (at least not with a mass in the search window).

The next step is to define a loss function that reflects two goals. First we would like to include the correct parameter value in the decision set. Therefore we charge a penalty if the correct value is excluded from the decision set. Second, we would like to exclude from S any incorrect parameter value. So we charge a penalty for including any parameter value that is not the true value. For the time being, suppose that the parameter space, \mathcal{M} , and the decision set, S , are discrete. Let $\ell_i(m_j)$ and $\ell_e(m_j)$ denote the loss values that respectively correspond to the case where m_j is not the mass of the Higgs particle ($m_j \neq m_H$) but included in the decision set and the case where m_j is the true mass of the Higgs particle ($m_j = m_H$) but excluded from the decision set. We refer to $\ell_i(m_j)$ and $\ell_e(m_j)$ as inclusion and exclusion losses respectively. Allowing ℓ_i and ℓ_e to depend on m_j permits us later on to adjust these functions to give desired error rates.

The following linear loss function accounts for all the possible decision scenarios with the corresponding losses,

$$\begin{aligned}
 L_D(m_H, S) = & \sum_{m_j \in \mathcal{M}} \ell_i(m_j) \mathbb{1}(m_j \in S) \mathbb{1}(m_j \neq m_H) \\
 (3.1) \quad & + \ell_e(m_H) \mathbb{1}(m_H \notin S),
 \end{aligned}$$

where the subscript D in L_D shows the momentary discreteness assumption. However, we wish to treat the mass as a continuous variable over the interval of interest. Moreover, the case $m_H = \emptyset$ needs to be displayed explicitly in the formulas since it is treated differently than the rest of the parameter space in terms of error rates. Therefore, we begin by rewriting the loss function as

$$\begin{aligned}
 L_D(m_H, S) = & \left[\sum_{m_j \in S \cap (m_0, m_n)} \ell_i(m_j) \mathbb{1}(m_j \neq m_H) \right. \\
 (3.2) \quad & \left. + \ell_i^\emptyset \mathbb{1}(\emptyset \in S) + \ell_e(m_H) \mathbb{1}(m_H \notin S) \right] \mathbb{1}(m_H \in (m_0, m_n)) \\
 & + \left[\sum_{m_j \in S \cap (m_0, m_n)} \ell_i(m_j) \mathbb{1}(m_j \neq m_H) + \ell_e^\emptyset \mathbb{1}(\emptyset \notin S) \right] \mathbb{1}(m_H = \emptyset).
 \end{aligned}$$

To pass to the continuous case we now replace the sum over $S \cap (m_0, m_n)$ by an integral with respect to the measure ν defined in (2.11). To do so we replace $\mathbb{1}(m_j \neq m_H)$ by $\mathbb{1}(m_H \notin (m_{j-1}, m_j))$ and write

$$(3.3) \quad \begin{aligned} & \sum_{m_j \in S \cap (m_0, m_n)} \ell_i(m_j) \mathbb{1}(m_H \notin (m_{j-1}, m_j)) \\ &= \sum_{m_j \in S \cap (m_0, m_n)} \frac{\ell_i(m_j)(m_n - m_0)}{m_j - m_{j-1}} \frac{\delta m_j}{m_n - m_0}, \end{aligned}$$

where $\delta m_j = m_j - m_{j-1}$. Then we take a limit as the largest δm_i converges to 0 and require the rescaled loss

$$(3.4) \quad \frac{\ell_i(m_j)(m_n - m_0)}{m_j - m_{j-1}}$$

to have a limit as $n \rightarrow \infty$ (where n represents the number of masses) and $m_j \rightarrow m$ which we still denote by $\ell_i(m)$; this abuse of notation should cause no confusion. The new function ℓ_i is a “loss density” with respect to normalized Lebesgue measure on (m_0, m_n) ; we also denote $\ell_i(\emptyset) = \ell_i^\emptyset$ for compact notation. Values of the function ℓ_i have units of “loss” as do the quantities ℓ_e^\emptyset and $\ell_e(m_H)$. Note that the exclusion loss is only defined for $m_i = m_H$. The result is the loss function

$$(3.5) \quad \begin{aligned} L(m_H, S) &= \int_S \ell_i(m) \mathbb{1}(m \neq m_H) \nu(dm) \\ &\quad + \ell_e(m_H) \mathbb{1}(m_H \notin S) \mathbb{1}(m_H \in (m_0, m_n)) \\ &\quad + \ell_e^\emptyset \mathbb{1}(\emptyset \notin S) \mathbb{1}(m_H = \emptyset) \\ &= \left[\int_{S \cap (m_0, m_n)} \ell_i(m) \mathbb{1}(m \neq m_H) \frac{dm}{m_n - m_0} \right. \\ &\quad \left. + \ell_i^\emptyset \mathbb{1}(\emptyset \in S) + \ell_e(m_H) \mathbb{1}(m_H \notin S) \right] \mathbb{1}(m_H \in (m_0, m_n)) \\ &\quad + \left[\int_{S \cap (m_0, m_n)} \ell_i(m) \mathbb{1}(m \neq m_H) \frac{dm}{m_n - m_0} \right. \\ &\quad \left. + \ell_e^\emptyset \mathbb{1}(\emptyset \notin S) \right] \mathbb{1}(m_H = \emptyset). \end{aligned}$$

The indicator $\mathbb{1}(m \neq m_H)$ can be dropped from the first integral without changing its value to give the simplified form

$$(3.6) \quad \begin{aligned} L(m_H, S) &= \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} \\ &\quad + [\ell_i^\emptyset \mathbb{1}(\emptyset \in S) + \ell_e(m_H) \mathbb{1}(m_H \notin S)] \mathbb{1}(m_H \in (m_0, m_n)) \\ &\quad + \ell_e^\emptyset \mathbb{1}(\emptyset \notin S) \mathbb{1}(m_H = \emptyset), \end{aligned}$$

where the term $\int_{S \cap (m_0, m_n)} \ell_i(m) d\frac{m}{m_n - m_0}$ is the loss due to including incorrect mass values in S .

By averaging the loss function (3.6) with respect to the marginal posterior $\pi(m | \mathbf{y})$ the posterior expected loss or the Bayes risk is obtained as follows.

$$\begin{aligned}
 r_{\pi(m|\mathbf{y})}(S) &= E_{\pi(m|\mathbf{y})}[L(m, S)] \\
 &= \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} \\
 &\quad + \int_{S^c \cap (m_0, m_n)} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} \\
 &\quad + \ell_e^\emptyset \mathbb{1}_{S^c}(\emptyset) \pi(\emptyset | \mathbf{y}) + \ell_i^\emptyset \mathbb{1}_S(\emptyset) (1 - \pi(\emptyset | \mathbf{y})).
 \end{aligned}
 \tag{3.7}$$

In (3.7) the terms in (3.6) were rearranged to make it easier to identify the Bayes procedure. The Bayes rule is obtained by minimizing the Bayes risk with respect to S .

THEOREM 3.1. *The Bayes rule, that is, the decision rule that minimizes $r_{\pi(m|\mathbf{y})}(S)$, is given by*

$$S = \begin{cases} \left\{ m \in (m_0, m_n) : \frac{\ell_i(m)}{\ell_e(m)} < \pi(m | \mathbf{y}) \right\}, \\ \frac{\ell_i^\emptyset}{\ell_e^\emptyset} \geq \frac{\pi(\emptyset | \mathbf{y})}{1 - \pi(\emptyset | \mathbf{y})}, \\ \left\{ m \in (m_0, m_n) : \frac{\ell_i(m)}{\ell_e(m)} < \pi(m | \mathbf{y}) \right\} \cup \{\emptyset\}, \\ \frac{\ell_i^\emptyset}{\ell_e^\emptyset} < \frac{\pi(\emptyset | \mathbf{y})}{1 - \pi(\emptyset | \mathbf{y})}. \end{cases}
 \tag{3.8}$$

A proof is provided in Appendix A.

By the above theorem, the optimal Bayes decision set is obtained by including all the values of $m \in (m_0, m_n)$ whose posterior density is greater than the ratio of inclusion loss to exclusion loss. The \emptyset is included if the corresponding posterior odds ratio is greater than its inclusion to exclusion loss ratio. The loss ratios can be specified such that the Bayes decision set has certain frequency coverage. In the following we explain the calibration process.

3.2. Calibration and error rate estimation. As mentioned before, the proposed procedure is calibrated to give desired frequency theory properties such as error rates. The loss ratios $\ell_i(m)/\ell_e(m)$ can be adjusted to satisfy the type I error rates required in particle physics applications. The same effect could be achieved in principle by keeping the loss ratio fixed and adjusting the prior. However, fixing the prior makes the computations below more straightforward. The global Type I

error rate and false exclusion rates are controlled by solving the following equations for $\ell_i(\emptyset)/\ell_e(\emptyset)$ and $\ell_i(m)/\ell_e(m)$, respectively.

$$(3.9) \quad P(\emptyset \notin S \mid m_H = \emptyset) = P\left(\frac{\pi(\emptyset \mid \mathbf{y})}{1 - \pi(\emptyset \mid \mathbf{y})} \leq \frac{\ell_i^\emptyset}{\ell_e^\emptyset} \mid m_H = \emptyset\right) = \alpha_1$$

and

$$(3.10) \quad \begin{aligned} &P(m \notin S \mid m_H = m, \mu = 1) \\ &= P\left(\pi(m \mid \mathbf{y}) < \frac{\ell_i(m)}{\ell_e(m)} \mid m_H = m, \mu = 1\right) \\ &= \alpha_2, \end{aligned}$$

for $m \in (m_0, m_n)$ and where $P(A)$ is the probability of event A .

The Associate Editor has observed that these calibrations are not strictly frequentist in nature because of the way we are handling the background Λ . A strict frequentist procedure would regard the function Λ as a parameter. Thus such a frequentist procedure would have

$$(3.11) \quad P(\emptyset \notin \hat{S} \mid m_H = \emptyset, \Lambda) = \alpha_1,$$

$$(3.12) \quad P(m \notin \hat{S} \mid m_H = m, \Lambda, \mu = 1) = \alpha_2,$$

where \hat{S} is used instead of S to indicate the dependence on data. Such calibration is likely impossible if we insist on exact calibration in small samples. Traditional statistical procedures such as likelihood ratio tests are calibrated by letting the critical value for the test statistic depend on the nuisance parameters. In other words, equation (3.11) is required to hold only at the estimated value $\hat{\Lambda}$ of Λ ; approximate calibration is then achieved by parametric bootstrapping. The procedure we describe below is parallel but averages over those Λ which remain credible after seeing the data, that is, according to the posterior distribution of Λ ; details are given in the next two subsections. Our procedure and parametric bootstrapping are both properly calibrated in large samples.

In (3.10) and (3.12) exclusion is done for the Standard Model, that is, only at $\mu = 1$. In (3.12) parametric bootstrapping can be used with Λ replaced by an estimate. As for (3.9), standard large sample results predict that this will achieve calibration at the true but unknown background.

Solving equations (3.9) and (3.10) for the loss ratios $\ell_i^\emptyset/\ell_e^\emptyset$ and $\ell_i(m)/\ell_e(m)$ requires obtaining the $\alpha_1 100\%$ and $\alpha_2 100\%$ quantiles of the distributions of the posterior odds $\pi(\emptyset \mid \mathbf{y})/(1 - \pi(\emptyset \mid \mathbf{y}))$, and the posterior probability density, $\pi(m \mid \mathbf{y})$ under the null hypotheses for the two tests, that is, $m_H = \emptyset$ and $m_H = m$, respectively.

Unfortunately, under most realistic models the distribution of the posterior functionals cannot be obtained in closed form. In Johnson (2013) the uniformly most

powerful Bayesian test (UMPBT) for one-parameter exponential family is developed based on the same idea, that is, maximizing the probability that the Bayes factor is smaller than a certain threshold under the null model. Johnson (2013) briefly visits the Higgs problem and reports the size of a Bayes factor equivalent to the local significance level of $\alpha_0 = 3 \times 10^{-7}$. However, to be able to obtain the UMPBT, a normal model is used.

The results in Johnson (2013) cannot be used under our model. Therefore, we need to estimate percentiles of the distribution of the posterior using Monte Carlo. However, this requires intense computation since for each generated data set at each iteration of the Monte Carlo we need to run the SMC algorithm to estimate the posterior. This Monte Carlo within Monte Carlo scheme is computationally costly on its own, while satisfying the small significance level in the physics application requires a large number of iterations to estimate precise tail quantiles adding to the computational intensity.

Note that calibration to produce desired frequency error rates offsets the effect of the prior in decision making. Since the posterior odds can be written as a product of the prior odds and marginal likelihood ratio, the loss ratio we obtain and use in decision making, responds to the choice of the prior. Nevertheless the prior must be expected to influence the inferences and certainly plays a role in the computational costs of the procedure. In particular the prior of Section 2 is computationally expensive and we have not implemented our calibration in this context as a result. Instead we replaced the conditional prior on μ , given $m_H = m$ for a mass $m \in (m_0, m_n)$ by a point mass at $\mu = 1$.

This choice was not popular with reviewers of this paper. Here is some discussion of the issue from a frequency theory perspective. For a frequentist the performance of a procedure \hat{S} is evaluated by using the frequentist risk. This is the function of the unknown parameters m_H, μ and Λ given by

$$(3.13) \quad R(m_H, \mu, \Lambda; \hat{S}) = E\{L(m_H, \hat{S}) \mid m_H, \mu, \Lambda\}.$$

This risk can be written as a sum of two components, R_\emptyset and R_E , given by

$$(3.14) \quad \begin{aligned} R_\emptyset(m_H, \mu, \Lambda) &= \{\ell_i^\emptyset \mathbb{1}(m_H \neq \emptyset) + \ell_e^\emptyset \mathbb{1}(m_H = \emptyset)\} P(\emptyset \in \hat{S} \mid m_H, \mu, \Lambda) \end{aligned}$$

and

$$(3.15) \quad \begin{aligned} R_E(m_H, \mu, \Lambda) &= \int_{m_0}^{m_n} \ell_i(m) P(m \in \hat{S} \mid m_H, \mu, \Lambda) \frac{dm}{m_n - m_0} \\ &\quad + \ell_e(m_H) \mathbb{1}(m_H \neq \emptyset) P(m_H \in \hat{S} \mid m_H, \mu, \Lambda). \end{aligned}$$

The component R_\emptyset depends only on the random set $\hat{S} \cap \{\emptyset\}$ while the second depends only on $\hat{S} \cap (m_0, m_n)$. Any such random subset of $\{\emptyset\}$ has a natural interpretation as a hypothesis test which rejects the null hypothesis of no Higgs if and

only if the random subset is empty. The level of this test is

$$(3.16) \quad P(\emptyset \notin \hat{S} \mid m_H = \emptyset, \mu, \Lambda),$$

and the power at $m \neq \emptyset, \mu$ and Λ is

$$(3.17) \quad \text{Pow}(m, \mu, \Lambda) = P(\emptyset \notin S \mid m_H = m, \mu, \Lambda).$$

Since the frequentist risk is additive our procedure minimizes

$$(3.18) \quad R_\emptyset(m_H, \mu, \Lambda; \hat{S})$$

calibrated as at (3.9), that is, it minimizes the (Bayesian) expected risk averaged over Λ subject to our constraint. It can also be thought of as incorporating a hypothesis test of $H_0 : m_H = \emptyset$. The level of the test is

$$(3.19) \quad \alpha_1 = P(\emptyset \notin \hat{S} \mid m_H = \emptyset, \Lambda),$$

and the power is the function

$$(3.20) \quad \text{Pow}(m, \mu, \Lambda) = P(\emptyset \notin \hat{S} \mid m_H = m, \mu, \Lambda),$$

for $m \in (m_0, m_n)$. Using a prior on μ which puts a point mass on $\mu = 1$ will maximize $\text{Pow}(m, 1, \Lambda)$ (averaged over Λ). The prior of Section 2 will maximize $\text{Pow}(m, \mu, \Lambda)$ averaged over μ and Λ . Our reviewers argue that we should not limit our attention to $\mu = 1$ but we observe that this amounts to trading sensitivity at the theory being considered, the Standard Model, against sensitivity at theories which make similar predictions to those of the Standard Model but with $\mu \neq 1$.

The issue is important so we take the time to consider some elementary analogues. In a problem testing a one-sided hypothesis about a normal mean the most powerful level α test does not depend on the particular alternative; there is a uniformly most powerful test. After calibration any Bayesian procedure analogous to ours will then simply give back this uniformly most powerful test. In more complex problems the most powerful test, produced by the Neyman–Pearson Lemma, will depend on the alternative. For each given prior on the alternative we would get a different calibrated Bayesian-frequentist test which maximized the power, averaged with respect to that prior. This sort of problem is our situation. The procedure we present below makes the choice to use the highly informative point mass prior on μ largely for computational reasons. To highlight the computational intensity for the case that μ is allowed to vary, consider that the SMC algorithm when run in serial to generate 1000 samples takes about two hours. Calibration requires simulating the posterior multiple times to obtain the distribution of the test statistic. Even 100 simulations would take about eight days which needs to be repeated for each pair of mass and μ to obtain the exclusion threshold. [The time complexity is $O(N \times M \times G)$, where N is the number of posterior samples, M is the number of Monte Carlo simulations and G is the number of (m, μ) pairs.] Of course, by parallelizing over a large number of processors, full calibration can

be feasible. However, such computational undertaking is beyond the scope of the present work. These computational considerations also make it very challenging to carry out a full power study comparing a variety of priors on the signal strength parameter.

To address calibration with affordable computation, we combine importance sampling and approximation techniques: we replace the SMC algorithm with a Laplace approximation in each Monte Carlo algorithm and use importance sampling to reduce the number of iterations required to obtain tail probability estimates for the Bayesian statistic’s distribution for a fixed level of precision. Again, for the on/off problem [Golchi and Lockhart (2018)] approximation is not required as the posterior can be obtained by simpler numerical calculations.

3.2.1. *Approximation.* In the following, we explain the Laplace approximation to the marginal posterior distribution of the mass of the Higgs particle. This approximation is used as a fast alternative to sampling the posterior distribution to speed up the calibration Monte Carlo. The hyper-parameters, β , η and σ^2 are held fixed at their maximum a posteriori estimates in the calibration. Consider reparametrizing the model in terms of $\Psi = \log \Lambda$. The approximation method, inspired by Rue, Martino and Chopin (2009), is based on a Gaussian approximation to the conditional distribution, $\tilde{\pi}(\Psi | m, \mathbf{y})$, that is,

$$(3.21) \quad \tilde{\pi}(m | \mathbf{y}) = \frac{\pi(m, \Psi | \mathbf{y})}{\tilde{\pi}(\Psi | m, \mathbf{y})} \Big|_{\Psi=\Psi^*}.$$

The Gaussian approximation, $\tilde{\pi}(\Psi | \mathbf{y}, m)$, is obtained by numerically approximating the mode and curvature of $\pi(\Psi | \mathbf{y}, m)$;

$$(3.22) \quad \pi(\Psi | \mathbf{y}, m) \propto \exp \left\{ -\frac{1}{2}(\Psi - \xi)^T \Sigma^{-1}(\Psi - \xi) + \log \pi(\mathbf{y} | \Psi, m) \right\},$$

where ξ and Σ are the mean vector and covariance matrix, respectively. Consider the Taylor expansion of the n components of the log likelihood around the initial values Ψ_0 ,

$$(3.23) \quad \begin{aligned} \log \pi(\mathbf{y} | \Psi, m) &= \sum_{i=1}^n g_i(\Psi_i) \\ &\approx \sum_{i=1}^n \left[g_i(\Psi_{0i}) + g'_i(\Psi_{0i})(\Psi_i - \Psi_{0i}) + \frac{g''_i(\Psi_{0i})}{2}(\Psi_i - \Psi_{0i})^2 \right] \\ &= \sum_{i=1}^n \left[a_i(\Psi_{0i}) + b_i(\Psi_{0i})\Psi_i - \frac{1}{2}c_i(\Psi_{0i})\Psi_i^2 \right], \end{aligned}$$

where

$$(3.24) \quad g_i(\Psi) = \log(\pi(y_i | \Psi, m)),$$

$$(3.25) \quad a_i(\Psi_{0i}) = g_i(\Psi_{0i}) - g'_i(\Psi_{0i})\Psi_{0i} + \frac{g''_i(\Psi_{0i})}{2}\Psi_{0i}^2,$$

$$(3.26) \quad b_i(\Psi_{0i}) = g'_i(\Psi_{0i}) - \Psi_{0i}g''_i(\Psi_{0i}),$$

$$(3.27) \quad c_i(\Psi_{0i}) = -\frac{g''_i(\Psi_{0i})}{2}.$$

The above expressions are given explicitly in Appendix B. Note that g_i and its derivatives and consequently (3.24)–(3.27) depend on m but the dependence is not explicitly expressed for the sake of conciseness. Therefore, we have

$$(3.28) \quad \begin{aligned} \tilde{\pi}(\Psi | \mathbf{y}, m) &\propto \exp\left\{-\frac{1}{2}(\Psi - \boldsymbol{\xi})^T \Sigma^{-1}(\Psi - \boldsymbol{\xi})\right. \\ &\quad \left. + \sum_{i=1}^n \left[a_i(\Psi_{0i}) + b_i(\Psi_{0i})\Psi_i - \frac{1}{2}c_i(\Psi_{0i})\Psi_i^2 \right]\right\} \\ &\propto \exp\left\{-\frac{1}{2}\Psi^T (\Sigma^{-1} + \text{diag}(\mathbf{c}_0))\Psi + (\Sigma^{-1}\boldsymbol{\xi} + \mathbf{b}_0)^T \Psi\right\}, \end{aligned}$$

where $\mathbf{b}_0 = (b_1(\Psi_{01}), \dots, b_n(\Psi_{0n}))^T$ and $\mathbf{c}_0 = (c_1(\Psi_{01}), \dots, c_n(\Psi_{0n}))^T$. The mean (mode) of the approximate Gaussian distribution, $\tilde{\pi}(\Psi | \mathbf{y}, m)$, is obtained by repeatedly solving $(\Sigma^{-1} + \text{diag}(\mathbf{c}_t))\Psi_{t+1} = (\Sigma^{-1}\boldsymbol{\xi} + \mathbf{b}_t)$ for Ψ_{t+1} until convergence, where \mathbf{b}_t and \mathbf{c}_t are updated matrices at iteration t . The approximate covariance matrix of Ψ is $\Sigma^{-1} + \text{diag}(\mathbf{c})$, where \mathbf{c} is the convergent value of the sequence \mathbf{c}_t . Therefore, the approximate marginal distribution can be obtained up to a normalizing constant as follows,

$$(3.29) \quad \begin{aligned} \tilde{\pi}(m | \mathbf{y}) &\propto \pi(m) \int \tilde{\pi}(\Psi | m, \mathbf{y}) d\Psi \\ &\propto \pi(m) |\Sigma^{-1} + \text{diag}(\mathbf{c})|^{-\frac{1}{2}}. \end{aligned}$$

3.2.2. Importance sampling for estimating error rates. As mentioned before, to evaluate the error rates associated with the Bayesian testing procedure, tail probabilities of the posterior functionals need to be estimated. Accurate Monte Carlo estimates for probabilities of rare events are only obtained with large Monte Carlo samples (in the order of 10^7 and larger in this application). In this section we introduce an importance Monte Carlo algorithm that is used to obtain tail probability estimates with lower variances.

We focus on the global type I error rate of the Bayesian testing procedure, that is,

$$(3.30) \quad \alpha_1 = P\left(\frac{\pi(\emptyset | \mathbf{y})}{1 - \pi(\emptyset | \mathbf{y})} < \frac{\ell_i^\emptyset}{\ell_e^\emptyset} \mid m_H = \emptyset\right) = P(\pi(\emptyset | \mathbf{y}) < q_\emptyset \mid m_H = \emptyset),$$

where $q_\emptyset = \ell_i^\emptyset / (\ell_e^\emptyset + \ell_i^\emptyset)$. While in calibrating the Bayes procedure the goal is to estimate q_\emptyset to satisfy a determined α_1 , suppose, for the time being, that α_1 is to be estimated for a given q_\emptyset .

To estimate α_1 using basic Monte Carlo, data, \mathbf{y}_i , is generated in each iteration under the null hypothesis, $m_H = \emptyset$, and $\pi(\emptyset | \mathbf{y}_i)$ is obtained. The Monte Carlo estimate of α_1 based on a (large) sample of N posterior values is given by

$$(3.31) \quad \hat{\alpha}_1 = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(\pi(\emptyset | \mathbf{y}_i) \in (0, q_\emptyset)).$$

However, under the null hypothesis, the event that $\pi(\emptyset | \mathbf{y}_i)$ falls below q_0 is rare and an unaffordably large N is required to obtain a nonzero estimate for α_1 .

Importance sampling is a popular method for simulating rare events [Rubino and Tuffin (2009)]. The idea is to generate samples from an importance distribution under which the event of interest is more likely to occur and weight the samples according to the original distribution of interest. To use importance Monte Carlo, here, we seek an importance distribution under which small values of $\pi(\emptyset | \mathbf{y})$ are more likely to occur.

Let us remind ourselves of the model under the null and alternative hypotheses,

$$(3.32) \quad H_0 : \text{The Higgs particle does not exist, that is, } m_H = \emptyset,$$

$$(3.33) \quad H_A : \text{The Higgs particle exists with a mass } m_H \in (m_0, m_n).$$

Clearly we expect the event $\pi(\emptyset | \mathbf{y}) < q_\emptyset$ to occur with high probability under the alternative. Therefore we can use the model under H_A as the importance distribution. The importance weights are then given by

$$(3.34) \quad \begin{aligned} W_i &= \frac{\pi(\mathbf{y} | H_0)}{\pi(\mathbf{y} | H_A)} \\ &= \frac{\pi(\mathbf{y} | m = \emptyset)}{\int_{m_0}^{m_n} \pi(\mathbf{y} | m) dm} \\ &= \frac{\pi(\emptyset | \mathbf{y})\pi(\mathbf{y})/\pi(\emptyset)}{\int_{m_0}^{m_n} [\pi(m | \mathbf{y})\pi(\mathbf{y})/\pi_A(m)] dm} \\ &= \frac{\pi(\emptyset | \mathbf{y})}{\pi(\emptyset) \int_{m_0}^{m_n} [\pi(m | \mathbf{y})/\pi_A(m)] dm}, \end{aligned}$$

where $\pi(\emptyset)$ and $\pi_A(m)$ are the priors over the mass under H_0 and H_A , respectively. The importance Monte Carlo estimate of α_1 based on a sample generated under the alternative model is given by

$$(3.35) \quad \tilde{\alpha}_1 = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(\pi(\emptyset | \mathbf{y}_i) \in (0, q_\emptyset)) W_i.$$

Algorithm 2 Importance Monte Carlo calibration algorithm

Input: Predetermined significance level, α_1 .

- 1: **for** $i := 0, 1, \dots, N$ **do**
 - a- Generate $m_i \sim \pi_A(m)$;
 - b- Generate $\Lambda_i \sim \pi(\Lambda \mid \mathbf{y}_{\text{obs}})$ (sample a realization from the posterior sample generated by Algorithm 1);
 - c- Generate data, $\mathbf{y}_i \sim \pi(\mathbf{y} \mid \Lambda_i, m_i)$;
 - d- Obtain $\tilde{\pi}(m \mid \mathbf{y}_i)$ using (3.29);
 - e- Obtain W_i using (3.34).
- 2: **end for**
- 3: Solve (3.35) to obtain q_0 .

Return: Discovery threshold q_0 .

For calibration, however, (3.35) is solved for q_\emptyset with a given significance level, α_1 . Algorithm 2 outlines the calibration steps. As mentioned earlier, the marginal posterior of the mass is obtained by integrating the background over its posterior distribution, that is, the procedure is calibrated for the most likely realizations of the background function. This is done by repeatedly sampling from the posterior sample of the background generated by Algorithm 1 in step 1-b of Algorithm 2.

3.3. *Calibrated discovery threshold and the decision set.* We run Algorithm 2 for the model introduced in Section 2.1 with $\mu = 1$ to obtain the discovery threshold. The lower (3×10^{-7}) 100% quantile of the null distribution of the posterior probability, $\pi(\emptyset \mid \mathbf{y})$ is $q_0 = 0.0066$. The discovery decision based on the simulated data, \mathbf{y}_{sim} is made by comparing the estimated null posterior probability to the obtained threshold. The SMC-based point estimate of $\pi(\emptyset \mid \mathbf{y}_{\text{sim}})$ is zero, that is, no samples with $m = \emptyset$ remain in the final posterior sample. Therefore, we conclude that the simulated data contains adequate evidence of the existence of the Higgs particle and we do not include $m = \emptyset$ in the decision set.

Having concluded that $\emptyset \notin S$, to obtain the final decision set we only need to obtain the exclusion thresholds, $q(m)$, for $m \in (100, 180)$, such that,

$$(3.36) \quad P(\pi(m \mid \mathbf{y}) < q(m) \mid m_H = m) = \alpha_2,$$

where $\alpha_2 = 0.05$ is a common choice. Obtaining the exclusion thresholds is analogous to that of discovery threshold given in Section 3.2. Since the exclusion controlled error rates are not as small as the discovery thresholds even basic Monte Carlo can provide accurate estimates. However, the approximation in (3.21) is pointwise and the computational burden increases with the size of the mass grid. Therefore, the time complexity of the exclusion step is $O(N^2 \times M)$ for a mass grid of size N and M Monte Carlo iterations. To reduce the computational cost we use a coarse discretization of the mass spectrum and use a kernel smooth of the




FIG. 4. Exclusion threshold curve (red line) over the marginal posterior density estimate of mass together with the decision interval (the interval within two gray vertical lines). In (b) the graph is zoomed in for better visualization, the points are the decision thresholds obtained for the coarse grid over mass.

exclusion thresholds for this coarse grid. Figure 4 shows the exclusion threshold together with the histogram of the posterior sample of mass in the search window. The boundaries of the decision set are determined by the mass values where the estimated posterior density is higher than the exclusion threshold. These boundaries are specified in Figure 4 by the grey vertical lines. The decision set obtained by the exclusion threshold cuts is $\hat{S} = (122.40, 129.57)$ which is more conservative than the 95% Bayesian credible set reported in Section 2.3, that is, $(124.58, 127.30)$.

4. Conclusion and discussion. In this article a Bayesian hierarchical model is introduced to make inference about the mass and signal strength of the Higgs particle as well as estimating the background function and its hyper-parameters in a fully Bayesian framework. A sequential Monte Carlo algorithm is used to sample the posterior distributions of the model parameters. The model is fit to data produced by computer models that simulate the behaviours of the detectors. The analysis results resemble the analysis of the real Higgs data reported in physics literature—the posterior distribution of the mass is peaked around the reported mass of the Higgs particle, that is, $\hat{m}_H \approx 126$ [ATLAS Collaboration (2012), CMS Collaboration (2012)]. The 95% credible Bayesian interval is reported as the set of credible values of the Higgs boson based on this analysis.

In addition, we have proposed a frequency-calibrated Bayesian procedure that can be used as an alternative to the detection/exclusion method used in the search for the Higgs particle. In a decision theoretic framework, we define a linear loss function that summarizes the possible outcomes of searching for a new particle with the associated losses. The Bayes rule is obtained by minimizing the expected loss and is the basis of decision making.

Our procedure is calibrated to give required frequency theory error rates. A calibration algorithm is proposed in which the posterior distributions are obtained

by a fast Laplace approximation instead of SMC, thereby making the calibration step computationally feasible. An importance Monte Carlo for simulation of rare events is proposed that allows for calibration of the Bayes procedure to meet the small type I error rate typically used in discovery of new phenomena. We calibrate our procedure according to the error rate requirements in particle physics. A decision set is reported that excludes the null hypothesis (the hypothesis of “no Higgs”) and contains a range of masses that remain plausible according to discovery and exclusion error rates of 3×10^{-7} and 0.05 respectively.

As mentioned earlier, the nature of our procedure resembles that of [Feldman and Cousins \(1998\)](#) for constructing confidence intervals that discuss the coverage issues of confidence intervals built based on the two step procedure without proper conditioning. However, we use the posterior odds ratios instead of likelihood ratios and report confidence sets with different levels of confidence for parameter values that represent the null hypothesis than those under the alternative hypothesis.

We conclude with a number of issues deserving acknowledgment and/or further work: There are a number of issues surrounding the inclusion of μ as an unknown parameter to be estimated. In Section 2 we have described a fully Bayesian analysis including μ . There will inevitably be controversy surrounding any choice of prior; we believe that the prior need to captures the physicists’ consensus view of likely signal strength parameter values. Prior specification should, of course, involve a sensitivity analysis to study the robustness of the results to prior choice. Such sensitivity analysis is out of scope of the present paper but can be the focus of further investigations.

Another issue with μ is the possibility of including it in our decision theoretic approach. The decision theoretic procedure of Section 3 may be extended to consider the parameter space

$$(4.1) \quad \mathcal{M}_2 := \{\emptyset\} \cup \{(m, \mu) \mid m_0 < m < m_n, \mu > 0\}$$

to give a set $S_2 \subset \mathcal{M}_2$ which would be a variable level confidence set for the pair (m_H, μ) . The losses would be extended to such pairs and we could try to calibrate more error rates than just those at (3.9) and (3.10). The resulting calibrations would be computationally much more demanding and would often yield a complex set of pairs (m_H, μ) for μ near 0. Preliminary testing has shown that calibration will require a lot more computation and likely some care in implementation; as a result we have not yet pursued this problem.

The set S of Section 3 combines discovery and exclusion in the sense that a value of m is excluded if it is not in S . The new set S_2 can be interpreted to do exclusion if we declare m to be excluded provided $(m, 1) \notin S_2$. The key is that we are excluding the standard model value of μ .

The Bayes procedure will be quite sensitive to the prior for μ ; the decision theoretic procedure should be less so because calibration of the error rates reduces the impact of the prior. It does not eliminate the effect since the prior determines

an ordering on the possible data sets which is used to determine which data sets correspond to a set S_2 including a given pair (m, μ) .

We remark that a likelihood ratio statistic is a maximum over mass values of a family of test statistics depending on a mass m . This maximization produces the well known irregularities in the large sample theory of the test. Our proposal effectively uses the Bayes factor as a test statistic; this statistic replaces maximization with averaging, a much more regular process. We expect this regularization will make our calibration procedures work better with our test statistic than they might with a likelihood ratio.

A power study of the procedure in Section 3, comparing the effects of priors should be carried out. Such a study would require significant computational resources.

After the set S has been found further data analysis will be needed. If $\emptyset \notin S$, a discovery will be declared and it will be necessary to continue the analysis. One needs to:

- Give a suitable confidence set or credible region for m_H ;
- Give a suitable confidence set or credible region for the pair (m_H, μ) ;
- Assess the assertion that $\mu = 1$ —If $\mu \neq 1$ then we might be discovering a particle which is Higgs-like but not the specific object predicted by the standard model; alternatively we might be discovering that systematic errors have not been adequately handled;
- Assess the assertion that μ does not depend on the production mechanism of the particle in question.

Our procedure and the standard procedure could, in principle, have discovered more than one peak had these been present, in the sense that our confidence set might find two or more disjoint intervals in (m_0, m_n) .

Interpretation of such an occurrence is properly part of the post discovery analysis described above if we do not want that interpretation to contribute, potentially, to error rates in the formal discovery declaration step.

Finally, systematic errors were ignored in our discussion. Those errors are naturally included in a Bayesian model as priors. For systematic errors that result from measurement errors for parameters in the calculations leading to the mass histogram, one might imagine non-Bayesian (frequentist) approaches. But there are systematic errors which are much harder to deal with in frequentist terms. For instance, in computing the signal shape and cross section as a function of beam luminosity it is necessary to carry out quantum mechanical calculations by expansions. These expansions are carried out at some order described by phrases such as next to leading order (NLO) or next to next to leading order (NNLO). The result is an approximation error which typically can only be quantified subjectively. Incorporating that subjectivity explicitly in a prior offers the opportunity to examine the sensitivity of the conclusions to the specification of this prior uncertainty.

An adaptation of the proposed approach is presented for a signal detection problem, referred to as the on/off problem, in Golchi and Lockhart (2018). As mentioned before, due to simplicity of the problem and model, the proposed approach can be implemented with less computational cost in this framework.

APPENDIX A: PROOF OF THEOREM 3.1

PROOF. Consider the problem of deciding whether or not mass value $m^* \in \mathcal{M}$ should be included in the decision set. Following the Bayes rule, m^* is included if the change in risk associated with inclusion of m^* is negative. To avoid measure theoretic complications we restrict the decision space to finite unions of open intervals and $\{\emptyset\}$. First, consider the case that $m^* \in (m_0, m_n)$. Since the prior distribution of m on the interval (m_0, m_n) is absolutely continuous, the addition of any point like m^* (as well as any other zero measure set) to the decision set leaves the risk unchanged. Let S be the decision set before including m^* . We consider the change in the risk as a result of adding the interval $(m^* - \delta, m^* + \delta)$ for very small δ to S . Let $S_0 = S \cup (m^* - \delta, m^* + \delta)$. We include m^* in the final decision set if and only if

$$(A.1) \quad r(S_0) - r(S) < 0,$$

where

$$(A.2) \quad \begin{aligned} & r(S_0) - r(S) \\ &= \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} + \int_{m^* - \delta}^{m^* + \delta} \ell_i(m) \frac{dm}{m_n - m_0} \\ &+ \int_{S^c \cap (m_0, m_n)} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} \\ &- \int_{m^* - \delta}^{m^* + \delta} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} - \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} \\ &- \int_{S^c \cap (m_0, m_n)} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} \\ &= \{\ell_i(m^*) - \ell_e(m^*) \pi(m^* | \mathbf{y})\} 2\delta + o(\delta). \end{aligned}$$

The above expression is negative for all sufficiently small $\delta > 0$ if and only if

$$(A.3) \quad \ell_i(m^*) - \ell_e(m^*) \pi(m^* | \mathbf{y}) < 0,$$

or

$$(A.4) \quad \frac{\ell_i(m^*)}{\ell_e(m^*)} < \pi(m^* | \mathbf{y}).$$

For the case that $m^* = \emptyset$, let $S_0 = S \cup \{\emptyset\}$,

$$\begin{aligned}
 & r(S_0) - r(S) \\
 &= \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} + \ell_i^\emptyset (1 - \pi(\emptyset | \mathbf{y})) \\
 &+ \int_{S^c \cap (m_0, m_n)} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} \\
 \text{(A.5)} \quad & - \ell_e^\emptyset \pi(\emptyset | \mathbf{y}) - \int_{S \cap (m_0, m_n)} \ell_i(m) \frac{dm}{m_n - m_0} \\
 & - \int_{S^c \cap (m_0, m_n)} \ell_e(m) \pi(m | \mathbf{y}) \frac{dm}{m_n - m_0} \\
 &= \ell_i^\emptyset (1 - \pi(\emptyset | \mathbf{y})) - \ell_e^\emptyset \pi(\emptyset | \mathbf{y}).
 \end{aligned}$$

Therefore, 0 is added to S if

$$\text{(A.6)} \quad \ell_i^\emptyset (1 - \pi(\emptyset | \mathbf{y})) - \ell_e^\emptyset \pi(\emptyset | \mathbf{y}) < 0,$$

$$\text{(A.7)} \quad \frac{\ell_i^\emptyset}{\ell_e^\emptyset} < \frac{\pi(\emptyset | \mathbf{y})}{1 - \pi(\emptyset | \mathbf{y})}.$$

□

APPENDIX B: LIKELIHOOD EXPANSION TERMS

The explicit expressions for (3.24)–(3.27) are given below:

$$\text{(B.1)} \quad g_i(\Psi) = -\exp(\Psi) - s_i + y_i \log(\exp(\Psi) + s_i) - \log(y_i!),$$

where

$$\text{(B.2)} \quad s_i = \int_{m_{i-1}}^{m_i} s_{m_H}(m) dm,$$

$$\begin{aligned}
 \text{(B.3)} \quad a_i(\Psi) &= g_i(\Psi) - g_i'(\Psi)\Psi + \frac{g_i''(\Psi)}{2}\Psi^2 \\
 &= -\exp(\Psi) - s_i + y_i \log(\exp(\Psi) + s_i) - \log(y_i!) \\
 &+ \left(\exp(\Psi) - y_i \frac{\exp(\Psi)}{\exp(\Psi) + s_i} \right) \Psi \\
 &+ \left(-\exp(\Psi) + y_i \frac{\exp(\Psi)s_i}{(\exp(\Psi) + s_i)^2} \right) \Psi^2,
 \end{aligned}$$

$$\begin{aligned}
 \text{(B.4)} \quad b_i(\Psi) &= g_i'(\Psi)\Psi + \Psi g_i''(\Psi) \\
 &= -\exp(\Psi)(1 - \Psi) + y_i \frac{\exp(\Psi)}{\exp(\Psi) + s_i} \left(1 - \frac{\Psi s_i}{\exp(\Psi) + s_i} \right),
 \end{aligned}$$

and

$$(B.5) \quad c_i(\Psi) = -\frac{g_i''(\Psi)}{2} = \frac{1}{2} \exp(\Psi) \left(1 - y_i \frac{\exp(\Psi) s_i}{(\exp(\Psi) + s_i)^2} \right).$$

Data and code. The data and code required for reproducing the results in the paper are provided at <https://github.com/sgolchi/BPD>.

Acknowledgments. We are grateful to Louis Lyons and Matthew William Kenzie of the CMS group who provided to us the simulated data as well as insight on the problem throughout the course of this work. We also greatly appreciate the very substantial time and effort spent by the Associate Editor and the reviewers to help improve the manuscript.

SUPPLEMENTARY MATERIAL

Supplement A: Frequency-calibrated Bayesian analysis of the on/off problem (DOI: [10.1214/18-AOAS1138SUPP](https://doi.org/10.1214/18-AOAS1138SUPP); .pdf). In this supplementary file, an adaptation of the proposed approach is described for a simple signal detection problem referred to as the on/off problem that helps better understanding of the methods.

REFERENCES

- ATLAS COLLABORATION (2012). Observation of a new particle in the search for the standard model higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B* **716** 1–29.
- BERGER, J. O. (1980). *Statistical Decision Theory: Foundations, Concepts, and Methods*. Springer, New York. [MR0580664](#)
- CMS COLLABORATION (2012). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B* **716** 30–61.
- CMS COLLABORATION (2013). Updated measurements of the Higgs boson at 125 GeV in the two photon decay channel. Technical Report CMS-PAS-HIG-13-001, CERN, Geneva.
- CMS COLLABORATION (2014). Observation of the diphoton decay of the Higgs boson and measurement of its properties. *Eur. Phys. J. C* **74** 3076.
- COWAN, G., CRANMER, K., GROSS, E. and VITELLS, O. (2011). Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C* **71** 1554–1573.
- DAVIES, R. B. (1987). Hypothesis testing when a nuisance parameter is present only under the alternative. *Biometrika* **74** 33–43. [MR0885917](#)
- DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 411–436. [MR2278333](#)
- DOUCET, A., DE FREITAS, N. and GORDON, N., eds. (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York. [MR1847783](#)
- ENGLERT, F. and BROUT, R. (1964). Broken symmetry and the mass of gauge vector mesons. *Phys. Rev. Lett.* **13** 321–323. [MR0174314](#)
- FELDMAN, G. J. and COUSINS, R. D. (1998). Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D* **57** 3873–3889.

- GOLCHI, S. and LOCKHART, R. (2018). Supplement to “A Frequency-calibrated Bayesian search for new particles.” DOI:10.1214/18-AOAS1138SUPP.
- GROSS, E. and VITELLS, O. (2010). Trial factors for the look elsewhere effect in high energy physics. *Eur. Phys. J. C* **70** 525–562.
- GURALNIK, G. S. (2009). The history of the Guralnik, Hagen and Kibble development of the theory of spontaneous symmetry breaking and gauge particles. *Internat. J. Modern Phys. A* **24** 2601–2627.
- GURALNIK, G. S., HAGEN, C. R. and KIBBLE, T. W. B. (1964). Global conservation laws and massless particles. *Phys. Rev. Lett.* **13** 585–587.
- HIGGS, P. W. (1964). Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.* **13** 508–509. [MR0175554](#)
- JASRA, A., STEPHENS, D. A., DOUCET, A. and TSAGARIS, T. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scand. J. Stat.* **38** 1–22. [MR2760137](#)
- JOHNSON, V. E. (2013). Uniformly most powerful Bayesian tests. *Ann. Statist.* **41** 1716–1741. [MR3127847](#)
- PAULO, R. (2005). Default priors for Gaussian processes. *Ann. Statist.* **33** 556–582. [MR2163152](#)
- RUBINO, G. and TUFFIN, B., eds. (2009). *Rare Event Simulation Using Monte Carlo Methods*. Wiley, Chichester. [MR2742378](#)
- RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 319–392. [MR2649602](#)

DEPARTMENT OF STATISTICS AND ACTUARIAL SCIENCE
SIMON FRASER UNIVERSITY
8888 UNIVERSITY DR
BURNABY BC V5A 1S6
CANADA
E-MAIL: sgolchi@sgolchi.ca
lockhart@sfu.ca