Metropolis-Hastings via Classification

Tetsuya Kaji^{*} and Veronika Ročková[†]

April 8, 2021

Abstract

This paper develops a Bayesian computational platform at the interface between posterior sampling and optimization in models whose marginal likelihoods are difficult to evaluate. Inspired by adversarial optimization, namely Generative Adversarial Networks (GAN) [32], we reframe the likelihood function estimation problem as a classification problem. Pitting a Generator, who simulates fake data, against a Classifier, who tries to distinguish them from the real data, one obtains likelihood (ratio) estimators which can be plugged into the Metropolis-Hastings algorithm. The resulting Markov chains generate, at a steady state, samples from an approximate posterior whose asymptotic properties we characterize. Drawing upon connections with empirical Bayes and Bayesian mis-specification, we quantify the convergence rate in terms of the contraction speed of the actual posterior and the convergence rate of the Classifier. Asymptotic normality results are also provided which justify the inferential potential of our approach. We illustrate the usefulness of our approach on examples which have posed a challenge for existing Bayesian likelihood-free approaches.

Keywords: Approximate Bayesian Computation, Classification, Generative Adversarial Networks, Likelihood-free Inference, Metropolis-Hastings Algorithm, Markov Chain Monte Carlo.

^{*}Assistant Professor in Econometrics and Statistics; Liew Family Junior Faculty Fellow and Richard N. Rosett Faculty Fellow at the Booth School of Business, University of Chicago

[†]Associate Professor in Econometrics and Statistics and James S. Kemper Faculty Scholar at the Booth School of Business, University of Chicago. The author gratefully acknowledges the support from the James S. Kemper Faculty Fund at the Booth School of Business and the National Science Foundation (Grant No. NSF DMS-1944740).

1 Introduction

Many contemporary statistical applications require inference for models which are easy to simulate from but whose likelihoods are impossible to evaluate. This includes implicit (simulator-based) models [17], defined through an underlying generating mechanism, or models prescribed through intractable likelihood functions.

Statistical inference for intractable models has traditionally relied on some form of likelihood approximation (see [33] for a recent excellent survey). For example, [17] propose kernel log-likelihood estimates obtained from simulated realizations of an implicit model. Approximate Bayesian Computation (ABC) [7, 56, 65] is another simulation-based approach which obviates the need for likelihood evaluations by (1) generating fake data \widetilde{X}_{θ} for parameter values θ sampled from a prior, and (2) weeding out those pairs $(\widetilde{X}_{\theta}, \theta)$ for which \widetilde{X}_{θ} has low fidelity to observed data. The discrepancy between observed and fake data is evaluated by first reducing the two datasets to a vector of summary statistics and then measuring the distance between them. Both the distance function and the summary statistics are critical for inferential success. While eliciting suitable summary statistics often requires expert knowledge, automated approaches have emerged [9, 11, 34]. Notably, [24] proposed a semi-automated approach that approximates the posterior mean (a summary statistic that guarantees first-order accuracy) using a linear model regressing parameter samples onto simulated data. Subsequently, [38] elaborated on this strategy using deep neural networks which are expected to yield better approximations to the posterior mean. Beyond subtleties associated with summary statistics elicitation, ABC has to be deployed with caution for Bayesian model choice [44, 59]. Synthetic likelihood (SL) [55, 72] is another approach for carrying out inference in intractable models by constructing a proxy Gaussian likelihood for a vector of summary statistics. Implicit in the success of both ABC and SL is the assumption that the generating process can produce simulated summary statistics that adequately represent the observed ones. If this compatibility is not satisfied (e.g. in misspecified models), both SL [25] and ABC [27] can provide unreliable estimates. Besides SL, a wide range parametric surrogate likelihood models have been suggested including normalising flows, Gaussian processes or neural networks [10, 21, 33, 51]. Avoiding the need for summary statistics, [34] proposed to use discriminability of the observed and simulated data as a discrepancy measure in ABC. Their accepting/rejecting mechanism separates samples based on a discriminator's ability to tell the real and fake data apart. Similarly as their work, our paper is motivated by the observation that distinguishing two data sets is usually easier if they were simulated with very different parameter values. However, instead of deploying this strategy inside ABC, we embed it directly inside the Metropolis-Hastings algorithm using likelihood approximations obtained from classification.

The Metropolis-Hastings (MH) method generates ergodic Markov chains through an accept-reject mechanism which depends in part on likelihood ratios comparing proposed candidate moves and current states. For many latent variable models, the marginal likelihood is *not* available in closed form, making direct application of MH impossible (see [16] for examples). The pseudo-marginal likelihood method [3] offers a remedy by replacing likelihood evaluations with their (unbiased) estimates. Many variants of this approach have been proposed including the inexact MCWM method (described in [48] and [3]) and its elaborations that correct for bias [49], reduce the variance of the likelihood ratio estimator [16] or make sure that the resulting chain produces samples from the actual (not only approximate) posterior [6]. The idea of using likelihood approximations within MH dates back to at least [49] and has been implemented in a series of works (see e.g [50] and [6] and references therein). Our approach is fundamentally different from typical pseudo-marginal MH algorithms since it *does not* require a hierarchical model where likelihood estimates are obtained through simulation from conditionals of latent data. Our method can be thus applied in a wide range of generative models (where forward simulation is possible) and other scenarios (such as diffusion processes [36]) where PM methods would be cumbersome or time-consuming to implement (as will be seen later in our examples).

Inspired by adversarial optimization, namely the Generative Adversarial Networks (GAN) [32], we reframe the likelihood (ratio) estimation problem as a classification problem using the 'likelihood-ratio trick' [14, 21, 66]. Similarly as with GANs, we pit two agents (a Generator and a Classifier) against one another. Assessing the similitude between the fake data, outputted by the Generator, and observed data, the Classifier provides likelihood

estimators which can be deployed inside MH. The resulting algorithm provides samples from an approximate posterior. GANs have been successful in learning distributions over complex objects, such as images, and have been coupled with MH in [68] to sample from the likelihood. Their method is an elaboration of the discriminator rejection sampling [5] which uses an importance sampling post-processing step using information from a trained discriminator. These two strategies are very different from our proposal here which is concerned with posterior inference about model parameters in a likelihood-free environment.

Our contributions are both methodological and theoretical. We develop a personification of Metropolis-Hastings (MH) algorithm for intractable likelihoods based on Classification, further referred to as MHC. We consider two variants: (1) a fixed generator design which may yield biased samples, and (2) a random generator design which may yield unbiased samples with increased variance. We then describe how and when the two can be combined in order to provide posterior samples with an asymptotically correct location and spread. Our theoretical analysis consists of new convergence rate results for a posterior residual (an approximation error) associated with the Classifier. These rates are then shown to affect the rate of convergence of the stationary distribution, in a similar way as the ABC tolerance level affects the convergence rate of ABC posteriors [26]. Theoretical characterizations of related pseudo-marginal (PM) methods have been, so far, limited to convergence properties of the Markov chain such as mixing rates [3, 16]. Here, we provide a rigorous asymptotic study of the stationary distribution including convergence rates (drawing upon connections to empirical Bayes and Bayesian misspecification), asymptotic normality results and, in addition, polynomial mixing time characterizations of the Markov chain.

To illustrate that our MHC procedure can be deployed in situations when sampling from conditionals of latent data (required for PM) is not practical or feasible, we consider two examples. The first one entails discretizations of continuous-time processes which are popular in finance [13, 36]. The second one is a population-evolution generative model where PM is not straightforward and where ABC methods need strong informative priors and high-quality summaries. In both examples, we demonstrate that MHC offers a reliable practical inferential alternative which is straightforward to implement. We also show very good performance on a Bayesian model choice example (where ABC falls short) and on the famous Ricker model (Section 7.7 in the Appendix) [57] analyzed by multiple authors [24, 33, 72].

The paper is structured as follows. Section 2 and 3 introduce the classification-based likelihood ratio estimator and the MHC sampling algorithm. Section 4 then describes the asymptotic properties of the stationary distribution. Section 5 shows demonstrations on simulated data and, finally, Section 6 wraps up with a discussion.

Notation The following notation will be used throughout the manuscript. We employ the operator notation for expectation, e.g., $P_0 f = \int f dP_0$ and $\mathbb{P}_m^{\theta} f = \frac{1}{m} \sum_{i=1}^m f(X_i^{\theta})$. The ε -bracketing number $N_{[]}(\varepsilon, \mathcal{F}, d)$ of a set \mathcal{F} with respect to a premetric d is the minimal number of ε -brackets in d needed to cover \mathcal{F} .¹ The δ -bracketing entropy integral of \mathcal{F} with respect to d is $J_{[]}(\delta, \mathcal{F}, d) := \int_0^{\delta} \sqrt{1 + \log N_{[]}(\varepsilon, \mathcal{F}, d)} d\varepsilon$. We denote the usual Hellinger semi-metric for independent observations as $d_n^2(\theta, \theta') = \frac{1}{n} \sum_{i=1}^n \int (\sqrt{p_{\theta,i}} - \sqrt{p_{\theta',i}})^2 d\mu_i$. Next, $K(p_{\theta_0}^{(n)}, p_{\theta}^{(n)}) = \sum_{i=1}^n K(p_{\theta_0,i}, p_{\theta,i})$ denotes the Kullback-Leibler divergence between product measures and $V_2(f, g) = \int f |\log(f/g)|^2 d\mu$. Define $\langle a, b \rangle = \sum_{i=1}^d a_i b_i$ for $a, b \in \mathbb{R}^d$.

2 Likelihood Estimation with a Classifier

Our framework consists of a series of i.i.d. observations $\{X_i\}_{i=1}^n \in \mathcal{X}$ realized from a probability measure P_{θ_0} indexed by a parameter $\theta_0 \in \Theta$ which is endowed with a prior $\Pi_n(\cdot)$. We assume that P_{θ} , for each $\theta \in \Theta$, admits a density p_{θ} . Our objective is to draw observations from the posterior density given $X^{(n)} = (X_1, \ldots, X_n)'$ defined through

$$\pi_n(\theta \mid X^{(n)}) = \frac{p_\theta^{(n)}(X^{(n)})\pi(\theta)}{\int_{\Theta} p_\vartheta^{(n)}(X^{(n)}) \ d\Pi(\vartheta)},\tag{2.1}$$

where $p_{\theta}^{(n)} = \prod_{i=1}^{n} p_{\theta}(X_i)$. Our focus is on situations where the likelihood $p_{\theta}^{(n)}$ is too costly to evaluate but can be readily sampled from.

We develop a Bayesian computational platform at the interface between sampling and optimization inspired by Generative Adversarial Networks (GAN) [32]. The premise of

¹A premetric on \mathcal{F} is a function $d: \mathcal{F} \times \mathcal{F} \to \mathbb{R}$ such that d(f, f) = 0 and $d(f, g) = d(g, f) \ge 0$.

GAN's is to discover rich distributions over complex objects arising in artificial intelligence applications through simulation. The learning procedure consists of two entities pitted against one another. A Generator aims to deceive an Adversary by simulating samples that resemble the observed data while, at the same time, the Adversary learns to tell the fake and real data apart. This process iterates until the generated data are indistinguishable by the Adversary. While GAN's have found their usefulness in simulating from distributions over images, here we forge new connections to Bayesian posterior simulation.

Similarly as with GAN's, we assume a Generator transforming a set of latent variables $\widetilde{X} \in \widetilde{\mathcal{X}}$ to collect samples from P_{θ} through a known deterministic mapping $T_{\theta} : \widetilde{\mathcal{X}} \to \mathcal{X}$, i.e. $T_{\theta}(\widetilde{X}) \sim P_{\theta}$ for $\widetilde{X} \sim \widetilde{P}$ for some distribution \widetilde{P} on $\widetilde{\mathcal{X}}$. This implies that we can draw a single set of m observations $\widetilde{X}^{(m)}$ and then filter them through T_{θ} to obtain a sample $\widetilde{X}_{\theta}^{(m)} = T_{\theta}(\widetilde{X}^{(m)})$ from P_{θ} for any $\theta \in \Theta$. Being able to easily draw samples from the model suggests the intriguing possibility of learning density ratios 'by-comparison' [47]. Indeed, the fact that density ratios can be computed by building a classifier that compares two data sets [14, 21] has lead to an emergence of a rich ecosystem of algorithms for model-free inference [12, 35, 51, 53, 66]. Many of these machine learning procedures are based on variants of the 'likelihood ratio trick' (LRT) which builds a surrogate classification model for the likelihood ratio. We embed the LRT within a classical Bayesian sampling algorithm and furnish our procedure with rigorous frequentist-Bayesian inferential theory.

Our approach relies on the simple fact that a cross-entropy classifier (used by the Adversary in the GAN framework) can be deployed to obtain an estimator of the likelihood ratio [32]. Recall that the classification problem with the cross-entropy loss is defined through

$$\max_{D \in \mathcal{D}} \left[\frac{1}{n} \sum_{i=1}^{n} \log D(X_i) + \frac{1}{m} \sum_{i=1}^{m} \log(1 - D(X_i^{\theta})) \right],$$
(2.2)

where \mathcal{D} is a set of measurable classification functions $D : \mathcal{X} \to (0, 1)$ (1 for 'real' and 0 for 'fake' data) and where $X_i^{\theta} = T_{\theta}(\widetilde{X}_i)$ for $\widetilde{X}_i \sim \widetilde{P}$ for $i = 1, \ldots, m$ are the 'fake' data outputted by the Generator. If an oracle were to furnish the true model p_{θ_0} , it is known that the population solution to (2.2) is the 'Bayes classifier' (see [32, Proposition 1])

$$D_{\theta}(X) := \frac{p_{\theta_0}(X)}{p_{\theta_0}(X) + p_{\theta}(X)} \quad \text{for } X \in \mathcal{X}.$$

$$(2.3)$$

Reorganizing the terms in (2.3), the likelihood can be written (see e.g. [66]) in terms of the discriminator function as

$$p_{\theta}^{(n)}(X^{(n)}) = p_{\theta_0}^{(n)}(X^{(n)}) \exp\left(\sum_{i=1}^n \log \frac{1 - D_{\theta}(X_i)}{D_{\theta}(X_i)}\right).$$
(2.4)

The oracle discriminator $D_{\theta}(\cdot)$ depends on p_{θ_0} but can be estimated by simulation. Indeed, one can deploy the Generator to simulate the fake data $\widetilde{X}_{\theta}^{(m)} = T_{\theta}(\widetilde{X}^{(m)})$ and train a Classifier to distinguish them from $X^{(n)}$. The Classifier outputs an estimator $\hat{D}_{n,m}^{\theta}$, for which we will see examples, and which can be plugged into (2.4) to obtain the following likelihood estimator

$$\hat{p}_{\theta}^{(n)}(X^{(n)}) = p_{\theta_0}^{(n)}(X^{(n)}) \exp\left(\sum_{i=1}^n \log \frac{1 - \hat{D}_{n,m}^{\theta}(X_i)}{\hat{D}_{n,m}^{\theta}(X_i)}\right) = p_{\theta}^{(n)}(X^{(n)}) e^{u_{\theta}(X^{(n)})}, \quad (2.5)$$

where

$$u_{\theta}(X^{(n)}) := \sum_{i=1}^{n} \left(\log \frac{1 - \hat{D}_{n,m}^{\theta}}{1 - D_{\theta}} - \log \frac{\hat{D}_{n,m}^{\theta}}{D_{\theta}} \right)$$
(2.6)

will be further referred to as the log-posterior residual. In other words, (2.5) is a deterministic functional of auxiliary random variables $\widetilde{X}^{(m)}$ and the observed data $X^{(n)}$, and can be computed (up to a norming constant) from $\hat{D}_{n,m}^{\theta}$. The posterior density $\pi_n(\theta \mid X^{(n)})$ can be then estimated by replacing D_{θ} with $\hat{D}_{n,m}^{\theta}$ in the likelihood expression to obtain

$$\widehat{\pi}_{n,m}(\theta \mid X^{(n)}) := \exp\left(\sum_{i=1}^{n} \log \frac{1 - \widehat{D}_{n,m}^{\theta}(X_i)}{\widehat{D}_{n,m}^{\theta}(X_i)}\right) \pi(\theta) \propto \pi_n(\theta \mid X^{(n)}) \mathrm{e}^{u_{\theta}(X^{(n)})}.$$
(2.7)

Two observations ought to be made. First, the estimator (2.7) targets the posterior density only up to a norming constant. This will not be an issue in Bayesian algorithms involving posterior density ratios (such as the Metropolis-Hastings algorithm considered here). Second, the estimator (2.7) performs exponential tilting of the original posterior, where the quality of the approximation crucially depends on the statistical properties of $u_{\theta}(X^{(n)})$. Note that $u_{\theta}(X^{(n)})$ depends also on the latent data $\widetilde{X}^{(m)}_{\theta}$. We devote the entire Section 4.1 to statistical properties of $u_{\theta}(X^{(n)})$. The idea of estimating likelihood ratios via discriminative classifiers is very natural and has emerged in various contexts including hypothesis testing [14] and posterior density estimation [66].

3 Metropolis Hastings via Classification

The Metropolis-Hastings (MH) algorithm is one of the mainstays of Bayesian computation. The deployment of unbiased likelihood estimators within MH has shown great promise in models whose likelihoods are not available [3, 4, 7]. In the previous section, we have suggested how classification may be deployed to obtain estimates of likelihood ratios. This suggests a compelling question: Can we deploy these classification-based estimators within MH? This section explores this intriguing possibility and formalizes an MH variant that we further refer to as MHC, Metropolis Hastings via Classification.

Our objective is to simulate values from an (approximate) posterior distribution $\Pi_n(\cdot | X^{(n)})$ with a density $\pi_n(\theta | X^{(n)}) \propto p_{\theta}^{(n)}(X^{(n)})\pi(\theta)$ over (Θ, \mathscr{B}) using the MH routine. Recall that MH simulates a Markov chain according to the transition kernel $K(\theta, \theta') := \rho(\theta, \theta')q(\theta' | \theta) + \delta_{\theta}(\theta') \int_{\Theta} (1 - \rho(\theta, \tilde{\theta}))q(\tilde{\theta} | \theta) d\tilde{\theta}$, where

$$\rho(\theta, \theta') := \min\left\{\frac{p_{\theta'}^{(n)}(X^{(n)})\pi(\theta')}{p_{\theta}^{(n)}(X^{(n)})\pi(\theta)}\frac{q(\theta \mid \theta')}{q(\theta' \mid \theta)}, 1\right\}.$$
(3.1)

and where $q(\cdot \mid \theta)$ is a proposal density generating candidate values θ' for the next move.

It is often the case in practice that we cannot directly evaluate $p_{\theta}^{(n)}(X^{(n)})$ but have access to its (unbiased) estimator (see [19] for a recent overview). In Bayesian contexts, an unbiased likelihood estimator can be constructed using importance sampling [6] or particle filters [2, 3] via data augmentation through the introduction of auxiliary latent variables, say $\widetilde{X}_{\theta}^{(m)}$. This method has been named the pseudo-marginal approach [3]. In its simplest form (Monte Carlo Within Metropolis (MCWM) [3, 48]), this method requires independently simulating m replicates of the auxiliary data for each likelihood evaluation at each iteration. Other variants have been suggested where latent data are recycled from the previous iterations (Grouped Independence MH (GIMH) described in [7]) or using correlated latent variables for the numerator and the denominator of the acceptance ratio [16].

In this work, we propose replacing $p_{\theta}^{(n)}$ in the acceptance ratio (3.1) with the classificationbased likelihood estimator (2.5) outlined in Section 2. This estimator, similarly as with the pseudo-marginal (PM) method, also relies on the introduction of latent variables $\widetilde{X}_{\theta}^{(m)}$. However, unlike with PM methods we *do not* require an explicit hierarchical model where sampling from the conditional distribution of the latent data is feasible. Later in Section 5.2 we show an example of a generative model, where our approach fares superbly while the PM approach is not straightforward, if at all possible. As we have seen earlier, our likelihood estimator can be rewritten in terms of the estimated discriminator as

$$\hat{p}_{\theta}^{(n)}(X^{(n)}) \propto \exp\left(\sum_{i=1}^{n} \log \frac{1 - \hat{D}_{n,m}^{\theta}(X_i)}{\hat{D}_{n,m}^{\theta}(X_i)}\right).$$
(3.2)

The evaluation of $\hat{p}_{\theta}^{(n)}(X^{(n)})$ can be carried out by merely computing $\hat{D}_{n,m}^{\theta}(X_i)$ where $\hat{D}_{n,m}^{\theta}$ is a trained classifier distinguishing $X^{(n)}$ from $\widetilde{X}_{\theta}^{(m)}$. Putting the pieces together, one can replace the intractable likelihood ratio in the acceptance probability (3.1) with

$$\rho_u(\theta, \theta') := \min\left\{\frac{\widehat{p}_{\theta'}^{(n)}(X^{(n)})\pi(\theta')}{\widehat{p}_{\theta}^{(n)}(X^{(n)})\pi(\theta)}\frac{q(\theta \mid \theta')}{q(\theta' \mid \theta)}, 1\right\}.$$
(3.3)

Note that the proportionality constant in the likelihood expression (3.2) cancels out in (3.3), allowing $\rho_u(\theta, \theta')$ to be directly computable. We consider two variants. The first one, called a *fixed generator design*, assumes that the randomness of $\hat{D}_{n,m}^{\theta}$, for each given θ and $X^{(n)}$, is determined by latent variables $\widetilde{X}^{(m)}$ shared by all steps of the algorithm. This corresponds to the case when m auxiliary data points $\widetilde{X}_{\theta}^{(m)} = \{\widetilde{X}_i^{\theta}\}_{i=1}^m$ are obtained through a *deterministic* mapping $\widetilde{X}_i^{\theta} = T_{\theta}(\widetilde{X}_i)$ for some $\widetilde{X}_i \sim \widetilde{P}$ that are not changed throughout the algorithm. The second version, called a *random generator design*, assumes that the underlying latent variables variables $\widetilde{X}^{(m)} = \{\widetilde{X}_i\}_{i=1}^m$ are refreshed at each step. While the difference between these two versions is somewhat subtle, we will see important bias-variance implications (discussed in more detail below). While technically our MHC sampling procedure follows the footsteps of a standard MH algorithm, we still find it useful to summarize the computations in an algorithm box (see Table 1).

3.1 Fixed Generator MHC

It is natural to inquire whether and how the likelihood approximation affects the stationary distribution of the resulting Markov chain. Due to the exponential tilt $e^{u_{\theta}(X^{(n)})}$ in the likelihood approximation (2.5), Algorithm 1 (Table 1) does *not* yield the correct posterior $\pi_n(\theta \mid X^{(n)})$ at its steady state. Indeed, under standard assumptions (see Section 7.3.1 of

INPUT							
Draw $\widetilde{X} = \{\widetilde{X}_i\}_{i=1}^m \sim \widetilde{P}$							
Initialize $\theta^{(0)}$ and generate $\widetilde{X}_{\theta^{(0)}} = \{\widetilde{X}_i^{\theta^{(0)}}\}_{i=1}^m$ according to $\widetilde{X}_i^{\theta^{(0)}} = T_{\theta^{(0)}}(\widetilde{X}_i)$.							
LOOP							
For $t = 1, \ldots, T$ repeat steps C(1)-(3), R and U.							
Algorithm 1: Fixed Generator							
C(1): Given $\theta^{(t)}$, generate $\theta' \sim q(\cdot \mid \theta^{(t)})$.							
C(2): Generate $\widetilde{X}_{\theta'} = \{\widetilde{X}_i^{\theta'}\}_{i=1}^m$ according to $\widetilde{X}_i^{\theta'} = T_{\theta'}(\widetilde{X}_i)$.							
C(3): Compute $\hat{D}_{n,m}^{\theta'}$ from $X^{(n)}$ and $\widetilde{X}_{\theta'}$ and compute $\widehat{p}_{\theta}(X^{(n)})$ in (3.2).							
C(4) With $\rho_u(\cdot, \cdot)$ in (3.3), set $\theta^{(t+1)} = \begin{cases} \theta' & \text{with probability } \rho_u(\theta^{(t)}, \theta'), \\ \theta^{(t)} & \text{with probability } 1 - \rho_u(\theta^{(t)}, \theta'). \end{cases}$							
Algorithm 2: Random Generator							
C(1): Given $\theta^{(t)}$, generate $\theta' \sim q(\cdot \mid \theta^{(t)})$ and $\widetilde{X}' \sim \widetilde{q}(\widetilde{X}' \mid \widetilde{X}^{(t)})$							
C(2): Generate $\widetilde{X}_{\theta'} = \{\widetilde{X}_i^{\theta'}\}_{i=1}^m$ according to $\widetilde{X}_i^{\theta'} = T_{\theta'}(\widetilde{X}_i')$							
C(3): Compute $\hat{D}_{n,m}^{\theta'}$ from $X^{(n)}$ and $\widetilde{X}_{\theta'}$ and compute $\widehat{p}_{\theta}(X^{(n)})$ defined in (3.2).							
C(4): With $\rho_u(\cdot, \cdot)$ in (3.5), set $(\theta^{(t+1)}, \widetilde{X}^{(t+1)}) = \begin{cases} (\theta', \widetilde{X}') & \text{with probability } \widetilde{\rho}_u(\theta, \widetilde{X}; \theta', \widetilde{X}'), \\ (\theta^{(t)}, \widetilde{X}^{(t)}) & \text{with probability } 1 - \widetilde{\rho}_u(\theta, \widetilde{X}; \theta', \widetilde{X}'). \end{cases}$							
OUTPUT							
Samples $\theta^{(1)}, \dots, \theta^{(T)}$							

Table 1: Metropolis-Hastings via Classification.

[58]), the stationary distribution of the Markov chain, conditional on $\widetilde{X}^{(m)}$, writes as (see e.g. Theorem 7.2 in [58])

$$\pi_n^{\star}(\theta \mid X^{(n)}) = \frac{p_{\theta}^{(n)}(X^{(n)}) \times e^{u_{\theta}(X^{(n)})} \times \pi(\theta)}{\int_{\Theta} p_{\theta}^{(n)}(X^{(n)}) \times e^{u_{\theta}(X^{(n)})} \times \pi(\theta) \mathrm{d}\theta}.$$
(3.4)

We do not view this property as unsurmountable. Other approximate MH algorithms (e.g the MCWM pseudo-marginal method) also do not yield $\pi_n(\theta \mid X^{(n)})$ at stationarity [3]. However, the samples generated by Algorithm 1 will be distributed according an approximate posterior (3.4) whose statistical properties we describe in detail in Section 4. In Section 3.4, we further quantify the speed of MHC convergence in large samples under the assumption of asymptotic normality. As will be seen in Section 4, the exponential tilt induces certain bias where the pseudo-posterior (3.4) concentrates around a projection of the true parameter θ_0 . Despite the bias, the curvature of the approximate posterior can be shown to match the curvature of the actual posterior (under differentiability assumptions in Section 4.1). The random generator version, introduced in the next section, works the

other way around. It can lead to a correct location (no bias) but at the expense of enlarged variance.

3.2 Random Generator MHC

The random generator variant proceeds as Algorithm 1 but refreshes $\widetilde{X}^{(m)} \sim \widetilde{P}$ at each step before computing the acceptance ratio. For simplicity, we have dropped the subscript m in $\widetilde{X}^{(m)}$ while describing the algorithm in Table 1. The acceptance probability now also involves \widetilde{X} and writes as

$$\widetilde{\rho}_{u}(\theta, \widetilde{X}; \theta', \widetilde{X}') = \min\left\{\frac{\widehat{p}_{\theta'}^{(n)}(X^{(n)})\pi(\theta')}{\widehat{p}_{\theta}^{(n)}(X^{(n)})\pi(\theta)}\frac{q(\theta \mid \theta')}{q(\theta' \mid \theta)}\frac{\widetilde{q}(\widetilde{X} \mid \widetilde{X}')}{\widetilde{q}(\widetilde{X}' \mid \widetilde{X})}, 1\right\}.$$
(3.5)

To glean more insights into this variant, it is helpful to regard $(\theta^{(t)}, \widetilde{X}^{(t)})$ jointly as a Markov chain with an augmented proposal density $q(\theta', \widetilde{X}' \mid \theta, \widetilde{X}) = q(\theta' \mid \theta) \widetilde{q}(\widetilde{X}' \mid \widetilde{X})$ where $\widetilde{q}(\widetilde{X}' \mid \widetilde{X})$ possibly depends on \widetilde{X} . In order to make the dependence on \widetilde{X} in $u_{\theta}(X^{(n)})$ more transparent, we will denote the posterior residual defined in (2.6) with $u_{\theta}(X^{(n)}, \widetilde{X})$ going forward. It can be seen that the marginal stationary distribution of the augmented Markov chain under Algorithm 2 equals

$$\widetilde{\pi}_{n}^{\star}(\theta \mid X^{(n)}) := \int \pi_{n}^{\star}(\theta \mid X^{(n)}) \mathrm{d}\widetilde{X}, \qquad (3.6)$$

where $\pi_n^{\star}(\theta | X^{(n)})$ was defined earlier in (3.4) and depends on \widetilde{X} through $u_{\theta}(X^{(n)}, \widetilde{X})$. The following characterization will be useful for establishing statistical properties of $\widetilde{\pi}_n^{\star}(\theta | X^{(n)})$ later in Section 4. From (3.4), we can write

$$\widetilde{\pi}_{n}^{\star}(\theta \mid X^{(n)}) \propto p_{\theta}^{(n)}(X^{(n)}) \times e^{\widetilde{u}_{\theta}(X^{(n)})} \times \pi(\theta)$$
(3.7)

where

$$\widetilde{u}_{\theta}(X^{(n)}) = \log \int e^{u_{\theta}(X^{(n)},\widetilde{X})} d\widetilde{X}.$$
(3.8)

Assuming almost-sure positivity of the joint proposal density $q(\theta', \widetilde{X}' \mid \theta, \widetilde{X})$, it can be verified (e.g. from Corollary 4.1 in [67]) that the marginal distribution of $\theta^{(t)}$ after t steps of Algorithm 1 converges in total variation to $\widetilde{\pi}_n^*(\theta \mid X^{(n)})$. Interestingly, from (3.7) we can see that the stationary distribution (3.6) from Algorithm 2 has the same functional form as the

	Algorithm 3: Bias Correction
(1)	Generate a sample $\{\theta_1^{(t)}\}_{t=1}^T$ using Algorithm 1
(2)	Generate a sample $\{\theta_2^{(t)}\}_{t=1}^T$ using Algorithm 2
(3)	Debias $\{\theta_1^{(t)}\}$ using $\{\theta_2^{(t)}\}$, i.e. construct a sample $\{\theta^{(t)}\}$ by
	$\theta^{(t)} := \theta_1^{(t)} - \frac{1}{T} \sum_{s=1}^T \theta_1^{(s)} + \frac{1}{T} \sum_{s=1}^T \theta_2^{(s)}.$

Table 2: Bias Correction with Algorithm 3.

stationary distribution (3.4) from Algorithm 1. The only difference is replacing $u_{\theta}(X^{(n)})$ with an averaged-out version $\tilde{u}_{\theta}(X^{(n)})$ in (3.8). Integration may inflate the stationary distribution (3.6) by making it more spread-out compared to the fixed generator sampler. However, the exponential tilting factor $\tilde{u}_{\theta}(X^{(n)})$ is averaged out. While $u_{\theta}(X^{(n)})$ in (2.6) is *fixed* in \widetilde{X} (creating a non-vanishing bias term), $u_{\theta}(X^{(n)})$ in (3.8) can average out to 0 (depending on $\widetilde{q}(\cdot | \cdot)$), erasing the bias and yielding the actual posterior as the stationary distribution.

3.3 Debiasing

Algorithm 1 and 2 can be combined to produce a more realistic representation of the true posterior. We mentioned that Algorithm 1, under the differentiability assumptions, has the same curvature as the actual posterior but has a non-vanishing shift. Algorithm 2, on the other hand, has a reduced bias due to the averaging aspect in (3.8). We can thus diminish the bias of the fixed generator design by shifting the location towards the mean of samples obtained with the random generator. This leads to a hybrid procedure summarized in Table 2. While Algorithms 1 and 2 can be deployed as a standalone, the de-biasing variant might increase the quality of the samples. Note that if $\tilde{u}_{\theta}(X^{(n)}) = 0$, Algorithm 2 will be unbiased yielding the actual posterior as its stationary distribution. In Section 4.1 we theoretically justify Algorithm 3 by providing sufficient conditions under which it yields samples from an object which has the same limit as the actual posterior.

3.4 Mixing Properties of MHC

A critical issue for MCMC algorithms is the determination of the number of iterations needed for the result to be approximately a sample from the distribution of interest. This section sheds light on the mixing rate of Algorithm 1. Under standard assumptions on $q(\cdot | \cdot)$ (such as positivity almost surely, see Corollary 4.1 in [67]), the distribution of the MHC Markov chain after t steps will converge to $\pi_n^*(\theta | X^{(n)})$ from any initialization in Θ in total variation as $t \to \infty$. [46] derive necessary and sufficient conditions for the Metropolis algorithms (with independent or symmetric candidate distributions) to converge at a geometric rate to a prescribed continuous distribution. [8] studied the speed of convergence of MH when both $n \to \infty$ and $d \to \infty$ where $\theta \in \Theta \subset \mathbb{R}^d$.

We can reformulate their sufficient conditions for showing polynomial mixing times of MHC. Recall that the stationary distribution $\pi_n^*(\theta \mid X^{(n)})$ of the MHC sampler in (3.4) normalized to a compact set $K \subset \Theta$, writes as $\Pi_K^*(B) = \int_B \pi_n^*(\theta \mid X^{(n)}) / \int_K \pi_n^*(\theta \mid X^{(n)})$. We are interested in bounding the number of steps needed to draw a random variable from Π_K^* with a given precision. We denote with Π_K^{*t} the distribution obtained after t steps of the MHC algorithm starting from Π_K^{*0} . It is known (see e.g. [42]) that the total variation distance between Q and Q_t can be bounded by $\|\Pi_K^* - \Pi_K^{*t}\|_{TV} \leq \sqrt{M}(1 - \phi^2/2)^t$, where M is a constant which depends on the initial distribution Π_K^{*0} and ϕ is the *conductance* of the Markov chain defined, e.g., in (3.13) in [8]. To obtain bounds on the conductance, the Markov chain needs to transition somewhat smoothly (see assumption D1 and D2 in [8]). These assumptions pertain to the continuity of the transitioning measure and are satisfied by the Gaussian random walk with a suitable choice of the proposal variance (see Section 3.2.4 in [8]) The following Lemma summarizes Theorem 2 of [8] in the context of Algorithm 1 under asymptotic normality assumptions examined in more detail in Section 4.2.3.

Lemma 3.1. (Mixing Rate) Under Assumptions (7.9)-(7.10) and a Gaussian random walk $q(\cdot | \cdot)$ satisfying Lemma 4 of [8], the global conductance ϕ of the Markov chain obtained from Algorithm 1 satisfies $1/\phi = \mathcal{O}(d)$ in $P_{\theta_0}^{(n)}$ -probability. In addition, the minimal number of MCMC iterations needed to achieve $\|\Pi_K^* - \Pi_K^{*t}\|_{TV} < \epsilon$ is $\mathcal{O}(d^2 \log(M/\epsilon))$ for some suitable constant M depending on the initial distribution Π_K^{*0} .

MHC thus attains bounds on the mixing rate that are *polynomial* in d (i.e. rapid mixing) under suitable Bernstein-von Mises conditions formalized later in Section 4.2.3. This section investigates how fast the Markov chain converges to its target $\pi_n^*(\theta \mid X^{(n)})$ as the number of iterations t grows. In Section 4.2.1 (resp. Section 4.2.2), we investigate a fundamentally different question. We assess the speed at which the target $\pi_n^*(\theta \mid X^{(n)})$ shrinks around the truth θ_0 (resp. a Kullback-Leibler projection) as n grows.

4 Theory for MHC

We now shift attention from the computational aspects of MHC to its potential as a statistical inference procedure. To understand the qualitative properties of the MHC scheme, we provide an asymptotic study of its stationary distribution (convergence rates in Section 4.2 and asymptotic normality in Section 4.2.3), drawing upon its connections to empirical Bayes methods (Section 4.2.1) and Bayesian misspecification (Section 4.2.2). Before delving into the stationary distribution, however, we first derive rates of convergence for the posterior residual $u_{\theta}(X^{(n)})$ in (2.6) which plays a fundamental role.

4.1 Convergence of the Posterior Residual

We denote the sample objective function in (2.2) with $\mathbb{M}_{n,m}^{\theta}(D) := \mathbb{P}_n \log D + \mathbb{P}_m^{\theta} \log(1-D)$, where we employed the operator notation for expectation, e.g., $\mathbb{P}_n f = \frac{1}{m} \sum_{i=1}^m f(X_i)$ and $\mathbb{P}_m^{\theta} f = \frac{1}{m} \sum_{i=1}^m f(X_i^{\theta})$. Throughout this section, we will use a simplified notation u_{θ} instead of $u_{\theta}(X^{(n)})$ and similarly for p_{θ} and $p_{\theta}^{(n)}$. We denote by P the probability measure that encompasses all randomness, e.g., as $O_P(1)$.² The estimated Classifier is seen to satisfy

$$\hat{D}_{n,m}^{\theta} := \max_{D \in \mathcal{D}_n} \mathbb{M}_{n,m}^{\theta}(D)$$

where \mathcal{D}_n constitutes a sieve of classifiers that expands with the sample size and that is not too rich (as measured by the bracketing entropy $N_{[]}(\varepsilon, \mathcal{F}, d)$). In practice, the estimator $\hat{D}_{n,m}^{\theta}$ can be obtained by deploying a variety of classifiers ranging from logistic regression to deep learning (see Assumption 3 in [39] for a sieve construction using neural network classifiers). The discrepancy between two classifiers will be measured by a Hellinger-type distance

²We may think of this P as the "canonical representation" [70, Problem 1.3.4].

(see [39] and [52] for more discussion) $d_{\theta}(D_1, D_2) := \sqrt{h_{\theta}(D_1, D_2)^2 + h_{\theta}(1 - D, 1 - D_{\theta})^2}$, where $h_{\theta}(D_1, D_2) = \sqrt{(P_{\theta_0} + P_{\theta})(\sqrt{D_1} - \sqrt{D_2})^2}$. The rate of convergence of the Classifier was previously established by [39] under assumptions reviewed below. In the following, we denote with $\mathcal{D}_{n,\delta}^{\theta} := \{D \in \mathcal{D}_n : d_{\theta}(D, D_{\theta}) \leq \delta\}$ the neighborhood of the oracle classifier within the sieve.

Assumption 1. Assume that n/m converges and that an estimator $\hat{D}_{n,m}^{\theta}$ exists that satisfies $\mathbb{M}_{n,m}^{\theta}(\hat{D}_{n,m}^{\theta}) \geq \mathbb{M}_{n,m}^{\theta}(D_{\theta}) - O_P(\delta_n^2)$ for a nonnegative sequence δ_n . Moreover, assume that the bracketing entropy integral³ satisfies $J_{[]}(\delta_n, \mathcal{D}_{n,\delta_n}^{\theta}, d_{\theta}) \leq \delta_n^2 \sqrt{n}$ and that there exists $\alpha < 2$ such that $J_{[]}(\delta, \mathcal{D}_{n,\delta}^{\theta}, d_{\theta})/\delta^{\alpha}$ is decreasing in δ .

Under Assumption 1, for a given $\theta \in \Theta$, [39] conclude (see their Theorem 1) the following convergence rate result for the classifier: $d_{\theta}(\hat{D}_{n,m}^{\theta}, D_{\theta}) = O_P(\delta_n)$. While [39] focused mainly on the convergence of $\hat{D}_{n,m}^{\theta}$, here we move the investigation further by establishing the rate of convergence of $u_{\theta}(\cdot)/n$ as well as its limiting shape. To this end, we assume the following support compatibility assumption, a refinement of the bounded likelihood ratio condition in nonparametric maximum likelihood (Theorem 3.4.4 in [70] and Lemma 8.7 in [28]).

Assumption 2. There exists M > 0 such that for every $\theta \in \Theta$, $P_{\theta_0}(p_{\theta_0}/p_{\theta})$ and $P_{\theta_0}(p_{\theta_0}/p_{\theta})^2$ are bounded by M and

$$\sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} P_{\theta_0} \left(\frac{D_{\theta}}{D} \mid \frac{D_{\theta}}{D} \ge \frac{25}{16} \right) < M, \quad \sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} P_{\theta_0} \left(\frac{1 - D_{\theta}}{1 - D} \mid \frac{1 - D_{\theta}}{1 - D} \ge \frac{25}{16} \right) < M$$

for δ_n in Assumption 1. The brackets in Assumption 1 can be taken so that $P_{\theta_0}\left(\sqrt{\frac{u}{\ell}}-1\right)^2 = O(d_{\theta}(u,\ell)^2)$ and $P_{\theta_0}\left(\sqrt{\frac{1-\ell}{1-u}}-1\right)^2 = o(d_{\theta}(u,\ell)).$

The following Theorem will be crucial for understanding theoretical properties of our MHC sampling algorithm, where the rate of convergence of $u_{\theta}(\cdot)/n$ will be seen to affect the rate of convergence of the stationary distribution of our Markov chains.

Theorem 4.1. Let Assumptions 1 and 2 hold for a given $\theta \in \Theta$, then

$$\underline{u_{\theta}/n} = \mathbb{P}_n\left(\log\frac{1-\hat{D}_{n,m}^{\theta}}{1-D_{\theta}} - \log\frac{\hat{D}_{n,m}^{\theta}}{D_{\theta}}\right) = O_P(\delta_n).$$

³See the notation section

Proof. Section 7.1 in the Appendix.

One seemingly pessimistic conclusion from Theorem 4.1 is that $u_{\theta}(\cdot)$ does not vanish. [39] shows that if the true likelihood ratio has a low-dimensional representation and an appropriate neural network is used for the discriminator, the rate δ_n depends only on the underlying dimension and not on the original dimension of X_i . In spite of the non-vanishing tilting term $u_{\theta}(X^{(n)})$, it turns out that Algorithm 1 can be refined (de-biased) to produce reasonable samples as long as $\hat{D}_{n,m}^{\theta}$ estimates the score well (see Section 3.3). In the sequel, we show quadratic approximability for u_{θ} at a much faster rate than Theorem 4.1 when the model and the classifier are differentiable in some suitable sense.

Assumption 3 (Differentiability of p_{θ}). There exists $\theta_0 \in \Theta \subset \mathbb{R}^d$ such that $P_0 = P_{\theta_0}$. The model $\{p_{\theta}\}$ is differentiable in quadratic mean at θ_0 , that is, there exists a measurable function $\dot{\ell}_{\theta_0} : \mathcal{X} \to \mathbb{R}^d$ such that $\int \left[\sqrt{p_{\theta_0+h}} - \sqrt{p_{\theta_0}} - \frac{1}{2}h'\dot{\ell}_{\theta_0}\sqrt{p_{\theta_0}} \right]^2 = o(\|h\|^2)$.

This is a classical assumption (see e.g. Section 5.5 of [69]) which implies local asymptotic normality. Going back to (2.5), we write $\hat{p}_{\theta}(X^{(n)}) = \prod_{i=1}^{n} \hat{p}_{\theta}(X_i)$, where

$$\hat{p}_{\theta} = p_{\theta_0} \frac{1 - \hat{D}_{n,m}^{\theta}}{\hat{D}_{n,m}^{\theta}} \tag{4.1}$$

is an estimator of p_{θ} that is possibly unscaled so that $\int \hat{p}_{\theta}$ may not be one. The scaling constant will be denoted by $c_{\theta} := \int \hat{p}_{\theta}$. In general, \hat{p}_{θ} is not observable since p_{θ_0} is not available. From (2.6), we can see that $u_{\theta} = n\mathbb{P}_n \log \frac{1-\hat{D}_{n,m}}{\hat{D}_{n,m}} - n\mathbb{P}_n \log \frac{1-D_{\theta}}{D_{\theta}} = n\mathbb{P}_n \log \frac{\hat{p}_{\theta}}{p_{\theta_0}} - n\mathbb{P}_n \log \frac{p_{\theta}}{p_{\theta_0}}$ and, under Assumption 3, van der Vaart [69, Theorem 7.2] derives convergence of the second term above in the local neighborhood of θ_0 . In Theorem 4.2 below, we derive convergence of the first term under the a similar assumption.

Assumption 4 (Differentiability of \hat{p}_{θ}).

(i) The estimator \hat{p}_{θ} in (4.1) is differentiable in quadratic mean in probability at θ_0 with a cubic rate, that is, $\int \left[\sqrt{\hat{p}_{\theta_0+h}} - \sqrt{\hat{p}_{\theta_0}} - \frac{1}{2}h'\dot{\ell}_{\theta_0}\sqrt{\hat{p}_{\theta_0}}\right]^2 = O_P(||h||^3)$, where $\dot{\ell}_{\theta_0} : \mathcal{X} \to \mathbb{R}^d$ is the score function in Assumption 3.

⁴Integration is understood with respect to some dominating measure.

(ii) Dependence of \mathbb{P}_n and \hat{p}_{θ} is asymptotically ignorable in the sense that for every compact $K \subset \mathbb{R}^d$, in outer probability,

$$\begin{split} \sup_{h \in K} \left| n(\mathbb{P}_n - P_{\theta_0}) \left(\sqrt{\frac{\hat{p}_{\theta_0 + h/\sqrt{n}}}{\hat{p}_{\theta_0}}} - 1 - \frac{h'\dot{\ell}_{\theta_0}}{2\sqrt{n}} \right) \right| \longrightarrow 0, \\ \sup_{h \in K} \left| n(\mathbb{P}_n - P_{\theta_0}) \left(\sqrt{\frac{\hat{p}_{\theta_0 + h/\sqrt{n}}}{\hat{p}_{\theta_0}}} - 1 \right)^2 \right| \longrightarrow 0. \end{split}$$

(iii) The scaling factor is asymptotically linear in the sense that there exists a sequence of \mathbb{R}^d -valued random variables \dot{c}_{n,θ_0} such that for every compact $K \subset \mathbb{R}^d$, in outer probability, $\sup_{h \in K} \left| n \left(c_{\theta_0 + h/\sqrt{n}} - c_{\theta_0} \right) - \sqrt{n} h' \dot{c}_{n,\theta_0} \right| \to 0.$

Assumption 4 (i) requires that \hat{p}_{θ} estimates the score well and is smoother than once differentiable. If \hat{p}_{θ} is twice differentiable in θ , then it holds with $O_P(||h||^4)$. Assumption 4 (ii) requires that the dependence of \mathbb{P}_n and \hat{p}_{θ} be ignored asymptotically. If \mathbb{P}_n and \hat{p}_{θ} were independent, it would follow from Chebyshev's or Markov's inequality. Assumption 4 (iii) requires that the quadratic curvature of the scaling constant vanishes asymptotically. In general, Assumption 4 is not verifiable since the likelihood is not available. To develop intuition behind this assumption, we verify that it holds for a toy normal location-scale model example in Section 7.5 in the Appendix. With Assumption 4, the estimated log likelihood asymptotes to a quadratic function that has the oracle curvature but a different center.

Theorem 4.2. Let p_{θ} and \hat{p}_{θ} satisfy Assumptions 3 and 4 and $\int (\sqrt{\hat{p}_{\theta_0}} - \sqrt{p_{\theta_0}})^2 = O_P(\delta_n^2)$ for some $\delta_n = o(n^{-1/4})$. Then, for every compact $K \subset \mathbb{R}^d$, in outer probability,

$$\sup_{h \in K} \left| n \mathbb{P}_n \log \frac{\hat{p}_{\theta_0 + h/\sqrt{n}}}{\hat{p}_{\theta_0}} + \frac{1}{2} h' I_{\theta_0} h - \sqrt{n} \mathbb{P}_n h' \dot{\ell}_{\theta_0} + \sqrt{n} \hat{P}_{\theta_0} h' \dot{\ell}_{\theta_0} - \sqrt{n} h' \dot{c}_{n,\theta_0} \right| \longrightarrow 0$$

Proof. Section 7.2 in the Appendix.

Remark 1. Recall that the true log-likelihood ratio locally approaches a quadratic curve $-\frac{1}{2}h'I_{\theta_0}h + \sqrt{n}\mathbb{P}_n h'\dot{\ell}_{\theta_0}$. The linear term $h'\sqrt{n}(\dot{c}_{n,\theta_0} - \hat{P}_{\theta_0}\dot{\ell}_{\theta_0})$ in (4.2) shifts the center of the quadratic curve but not the curvature.

One important implication of Theorem 4.2 is linearity of u_{θ} .

Corollary 4.3. (Linear u_{θ}) Under assumptions of Theorem 4.2 we have

$$u_{\theta_0+h/\sqrt{n}} - u_{\theta_0} = h'\sqrt{n}(\dot{c}_{n,\theta_0} - \hat{P}_{\theta_0}\dot{\ell}_{\theta_0}) + o_P(1).$$
(4.2)

Proof. Follows from van der Vaart [69, Theorem 7.2] and Theorem 4.2.

We revisit linearity of u_{θ} later in Section 4.2.2 (Example 1) as one of the sufficient conditions for the Bernstein-von Mises theorem. Corollary 4.3 has a very important consequence regarding the limiting shape of the stationary distribution $\pi_n^*(\theta | X^{(n)})$ for Algorithm 1 defined in (3.4). It shows that $\pi_n^*(\theta | X^{(n)})$ approaches a *biased* normal distribution with the same variance as the true posterior. In addition, we have seen in Section 3.2 that the stationary distribution $\tilde{\pi}_n^*(\theta | X^{(n)})$ of Algorithm 1 defined in (3.7) is averaged over the bias. Therefore, if $\mathbb{E}[\dot{c}_{n,\theta_0} - \hat{P}_{\theta_0}\dot{\ell}_{\theta_0} | X] = 0$, then the stationary distribution of Algorithm 3 (in Table 2) converges to the correct normal posterior, i.e. it has the same limit as the actual posterior $\pi_n(\theta | X^{(n)})$. Theorem 4.2 thus provides a theoretical justification for de-biasing suggested in Section 3.3.

4.2 Posterior Concentration Rates

Having quantified the convergence rate of the posterior residual $u_{\theta}(X^{(n)})$ in Theorem 4.1, we are now ready to explore the convergence rate of the entire stationary distribution without necessarily imposing differentiability assumptions.

4.2.1 Empirical Bayes Lens

Recall that the MHC sampler *does not* reach $\pi_n(\theta \mid X^{(n)})$ in steady state. Recall that the stationary distribution (using the fixed generator) takes the form

$$\Pi_n^{\star}(B \mid X^{(n)}) = \frac{\int_B p_{\theta}^{(n)} / p_{\theta_0}^{(n)} \times e^{u_{\theta}} \times \pi(\theta) \mathrm{d}\theta}{\int_{\Theta} p_{\theta}^{(n)} / p_{\theta_0}^{(n)} \times e^{u_{\theta}} \times \pi(\theta) \mathrm{d}\theta}.$$
(4.3)

In the random design, we simply replace u_{θ} in (4.3) with \tilde{u}_{θ} defined in (3.8). Interestingly, (4.3) can be viewed as an actual posterior under a *tilted prior* with a density $\pi^*(\theta) \propto e^{u_{\theta}}\pi(\theta)$. This shifted prior depends on the data $X^{(n)}$ (through $u_{\theta}(X^{(n)})$) and thereby (4.3) can be loosely regarded as an empirical Bayes (EB) posterior. While EB uses plug-in estimators of prior hyper-parameters, here the data enters the prior in a less straightforward manner. We further the EB connection later in Remark 2. We first assess the quality of the posterior approximation (4.3) through its concentration rate around the *true parameter* value θ_0 using the traditional Hellinger semi-metric $d_n(\theta, \theta')$. The rate depends on the interplay between the concentration of the *actual* posterior⁵ $\Pi_n(\theta \mid X^{(n)})$ and the rate at which the residual $u_{\theta}(X^{(n)})$ in (2.6) diverges. Recall that the rate of $u_{\theta}(\cdot)/n$ was established earlier in Theorem 4.1. The following Theorem uses assumptions on prior concentration around θ_0 using the typical Kullback-Leibler neighborhood $B_n(\theta_0, \epsilon) = \left\{ \theta \in \Theta : K(p_{\theta_0}^{(n)}, p_{\theta}^{(n)}) \leq n\epsilon^2, \frac{1}{n} \sum_{i=1}^n V_2(p_{\theta_0}(X_i), p_{\theta}(X_i)) \leq \epsilon^2 \right\}.$

Theorem 4.4. Consider the pseudo-posterior distribution Π_n^* defined through (4.3). Suppose that the prior $\Pi_n(\cdot)$ satisfies conditions (3.2) and (3.4) in [29] for a sequence $\varepsilon_n \to 0$ such that $n\varepsilon_n^2 \to \infty$. In addition, let \tilde{C}_n be such that

$$P_{\theta_0}^{(n)}\left(\sup_{\theta\in\Theta}|u_{\theta}(X^{(n)})/n| > \tilde{C}_n\varepsilon_n^2\right) = o(1)$$
(4.4)

and assume that for sets $\Theta_n \subset \Theta$ the prior satisfies

$$\frac{\Pi_n(\Theta \backslash \Theta_n)}{\Pi_n(B_n(\theta_0, \varepsilon_n))} = o(\mathrm{e}^{-2(1+\widetilde{C}_n)n\varepsilon_n^2}).$$
(4.5)

Then we have, for any $M_n \to \infty$ such that $\tilde{C}_n = o(M_n)$,

$$P_{\theta_0}^{(n)} \left[\prod_n^{\star} (\theta : d_n(\theta, \theta_0) > M_n \varepsilon_n \,|\, X^{(n)}) \right] = o(1) \quad as \ n \to \infty.$$

Proof. The proof is a minor modification of Theorem 4 in [29] and is postponed until Section 7.3 in the Appendix.

Theorem 4.4 shows that the concentration rate of the pseudo-posterior nearly matches the concentration rate of the original posterior ε_n (this is implied by condition (3.2), (3.4) and a variant of (4.5) according to Theorem 4 of [29]) up to an inflation factor \tilde{C}_n which depends on the rate of $u_{\theta}(X^{(n)})/n$. If $\tilde{C}_n = \mathcal{O}(1)$ in (4.4), the rate of the actual posterior and pseudo-posterior will be the same.

⁵Using the usual notion [29], we say that the posterior $\Pi_n(\cdot | X^{(n)})$ concentrates around θ_0 at the rate ε_n (satisfying $\varepsilon_n \to 0$ and $n\varepsilon_n^2 \to \infty$) if $P_{\theta_0}\Pi_n[\theta \in \Theta : d_n(\theta, \theta_0) > M\varepsilon_n | X^{(n)}] \to 0$ as $n \to \infty$ where M possibly depends on n.

Remark 2. (Connection to Empirical Bayes) Since $\Pi_n^{\star}(\cdot | X^{(n)})$ can be regarded as an EB posterior, we could alternatively apply techniques of [18] and [61] to quantify the convergence rate in Theorem 4.4.

Remark 3. (Random Generator) Recall that the stationary distribution $\tilde{\pi}_n^*(\theta | X^{(n)})$ of the random generator MHC version can be written as (4.3) where u_{θ} is replaced with \tilde{u}_{θ} from (3.8). Theorem 4.4 holds also for the random generator where \tilde{C}_n is obtained from (4.4) with \tilde{u}_{θ} instead of u_{θ} . Due to the averaging aspect, we might expect this \tilde{C}_n to be smaller in the random generator design.

Theorem 4.4 describes the behavior of the pseudo-posterior around the truth θ_0 . We learned that the rate is artificially inflated due a bias inflicted by the likelihood approximation, where $\Pi_n^{\star}(\cdot | X^{(n)})$ may not shrink around θ_0 when ε_n is faster than the rate δ_n established in Theorem 4.1. This suggest that the truth may not be the most natural centering point for the posterior to concentrate around. A perhaps more transparent approach is to consider a different (data-dependent) centering which will allow for a more honest reflection of the contraction speed devoid of any implicit bias. We look into model misspecification for guidance about reasonable centering points.

4.2.2 Model Misspecification Lens

In Section 4.2.1, we reframed the stationary distribution (3.4) as an empirical Bayes posterior by absorbing the term $e^{u_{\theta}(X^{(n)})}$ inside the prior. This section pursues a different approach, absorbing $e^{u_{\theta}(X^{(n)})}$ inside the likelihood instead. This leads a mis-specified model $\tilde{P}_{\theta}^{(n)}$ prescribed by the following likelihood function

$$\tilde{p}_{\theta}^{(n)}(X^{(n)}) = \frac{p_{\theta}^{(n)}(X^{(n)})e^{u_{\theta}(X^{(n)})}}{C_{\theta}} \quad \text{where} \quad C_{\theta} = \int_{\mathcal{X}} p_{\theta}^{(n)}(X^{(n)})e^{u_{\theta}(X^{(n)})} \mathrm{d}X^{(n)}. \tag{4.6}$$

Defining $\tilde{\pi}(\theta) \propto \pi(\theta)C_{\theta}$, we can rewrite (3.4) as a posterior density under a mis-specified likelihood and the modified prior $\tilde{\pi}(\theta)$ as

$$\pi_n^{\star}(\theta \mid X^{(n)}) = \frac{\widetilde{p}_{\theta}^{(n)}(X^{(n)})\widetilde{\pi}(\theta)}{\int_{\Theta} \widetilde{p}_{\theta}^{(n)}(X^{(n)})\widetilde{\pi}(\theta)\mathrm{d}\theta}.$$
(4.7)

Since the model $\tilde{p}_{\theta}^{(n)}$ is mis-specified (i.e. $P_{\theta_0}^{(n)}$ is *not* of the same form as $\tilde{\mathcal{P}}^{(n)} = \{\tilde{P}_{\theta}^{(n)} : \theta \in \Theta\}$), the posterior will concentrate around the point θ^* defined as

$$\theta^* = \arg\min_{\theta\in\Theta} -P_{\theta_0}^{(n)} \log[\tilde{p}_{\theta}^{(n)}/p_{\theta_0}^{(n)}]$$
(4.8)

which corresponds to the element $\tilde{P}_{\theta^*}^{(n)} \in \tilde{\mathcal{P}}^{(n)}$ that is closest to $P_{\theta_0}^{(n)}$ in the KL sense [40]. Unlike in the iid data case studied, e.g., in [40] and [15], our likelihood (4.6) is not an independent product due to the non-separability of the function $u_{\theta}(X^{(n)})$. The following Theorem 4.5 quantifies concentration in terms of a KL neighborhoods around $\tilde{P}_{\theta^*}^{(n)}$ defined as $B(\epsilon, \tilde{P}_{\theta^*}^{(n)}, P_{\theta_0}^{(n)}) = \left\{ \tilde{P}_{\theta}^{(n)} \in \tilde{\mathcal{P}}^{(n)} : K(\theta^*, \theta_0) \le n\epsilon^2, V(\theta^*, \theta_0) \le n\epsilon^2 \right\}$, where $K(\theta^*, \theta_0) \equiv P_{\theta_0}^{(n)} \log \frac{\tilde{p}_{\theta^*}^{(n)}}{\tilde{p}_{\theta}^{(n)}}$ and $V(\theta^*, \theta_0) = P_{\theta_0}^{(n)} \left| \log \frac{\tilde{p}_{\theta^*}^{(n)}}{\tilde{p}_{\theta}^{(n)}} - K(\theta^*, \theta_0) \right|^2$.

Theorem 4.5. Denote with $Q_{\theta}^{(n)}$ a measure defined through $dQ_{\theta}^{(n)} = \frac{p_{\theta_0}^{(n)}}{\tilde{p}_{\theta^*}^{(n)}} dP_{\theta}^{(n)}$ and let $d(\cdot, \cdot)$ be a semi-metric on $\mathcal{P}^{(n)}$. Suppose that there exists a sequence $\varepsilon_n > 0$ satisfying $\varepsilon_n \to 0$ and $n\varepsilon_n^2 \to \infty$ such that for every $\epsilon > \varepsilon_n$ there exists a test ϕ_n (depending on ϵ) such that for every $J \in \mathbb{N}_0$

$$P_{\theta_0}^{(n)}\phi_n \lesssim e^{-n\epsilon^2/4} \quad and \quad \sup_{\widetilde{P}_{\theta}^{(n)}: d(\widetilde{P}_{\theta}^{(n)}, \widetilde{P}_{\theta^*}^{(n)}) > J\epsilon} Q_{\theta}^{(n)}(1-\phi_n) \le e^{-nJ^2\epsilon^2/4}.$$
(4.9)

Let $B(\epsilon, \tilde{P}_{\theta^*}^{(n)}, P_{\theta_0}^{(n)})$ be as before and let $\tilde{\Pi}_n(\theta)$ be a prior distribution with a density $\tilde{\pi}(\theta) \propto C_{\theta}\pi(\theta)$. Assume that there exists a constant L > 0 such that, for all n and $j \in \mathbb{N}$,

$$\frac{\widetilde{\Pi}_n\left(\theta\in\Theta: j\varepsilon_n < d(\widetilde{P}_{\theta}^{(n)}, \widetilde{P}_{\theta^*}^{(n)}) \le (j+1)\varepsilon_n\right)}{\widetilde{\Pi}_n\left(B(\epsilon, \widetilde{P}_{\theta^*}^{(n)}, P_{\theta_0}^{(n)})\right)} \le e^{n\varepsilon_n^2 j^2/8}.$$
(4.10)

Then for every sufficiently large constant M, as $n \to \infty$,

$$P_{\theta_0}^{(n)} \Pi_n^{\star} \left(\tilde{P}_{\theta}^{(n)} : d(\tilde{P}_{\theta}^{(n)}, \tilde{P}_{\theta^*}^{(n)}) \ge M \varepsilon_n \,|\, X^{(n)} \right) \to 0.$$

$$(4.11)$$

Proof. Section 7.4 in the Appendix.

Remark 4. For iid data, [40] introduce a condition involving entropy numbers under misspecification which implies the existence of exponential tests for a testing problem that involves non-probability measures. Since we have a non-iid situation, we assumed the existence of tests directly. **Remark 5.** (Friendlier Metrics) In parametric models indexed by θ in a metric space (Θ, d) , it is more natural to characterize the posterior concentration in terms of $d(\cdot, \cdot)$ rather than the Kullback-Leibler divergence⁶. Section 5 of [40] clarifies how Theorem 4.5 can be reformulated in terms of some metric $d(\cdot, \cdot)$ on Θ .

4.2.3 Bernstein-von Mises Theorem

The Bernstein-von Mises (BvM) theorem asserts that the posterior distribution of a parameter in a suitably regular finite-dimensional model is approximately normally distributed as the number of observations grows to infinity. More precisely, if $\theta \to p_{\theta}$ is appropriately smooth and identifiable and the prior $\Pi_n(\cdot)$ puts positive mass around the true parameter θ_0 , then the posterior distribution of $\sqrt{n}(\theta - \hat{\theta}_n)$ tends to $N(0, I_{\theta_0}^{-1})$ for most observations $X^{(n)}$, where $\hat{\theta}_n$ is an efficient estimator and I_{θ} is the Fisher information matrix of the model at θ . In this section, we want to understand the effect of the tilting factor $e^{u_{\theta}(X^{(n)})}$ on the limiting shape of the pseudo-posterior in (3.4) that is proportional to $\pi_n(\theta \mid X^{(n)})e^{u_{\theta}(X^{(n)})}$. Exponential tilting is particularly intuitive for linear $u_{\theta}(X^{(n)})$ and for Gaussian posteriors.

Example 1. (Linear u_{θ}) Suppose that the posterior $\pi_n(\theta \mid X^{(n)})$ is Gaussian with some mean μ and covariance Σ . This holds approximately in regular models according to the BvM theorem (Theorem 10.1 in [69]). Assume that there exists an invertible mapping $\tau : \Theta \to \Theta$ such that $\theta = \tau(\bar{\theta})$ where the density for $\bar{\theta}$ satisfies $\pi_n(\theta \mid X^{(n)}) e^{u_{\theta}(X^{(n)})} d\theta \propto \pi_n^*(\bar{\theta} \mid X^{(n)}) d\bar{\theta}$. Assuming the following linear form (justified in Remark 4.3)

$$u_{\theta}(X^{(n)}) = a^*(X^{(n)}) + \theta' u^*(X^{(n)})$$
(4.12)

we obtain $\bar{\theta} \sim \mathcal{N}(\mu + \Sigma u^*(X^{(n)}), \Sigma)$. In this case, the mapping τ satisfies $\theta = \tau(\bar{\theta}) = \bar{\theta} - \Sigma u^*(X^{(n)})$, implying a location shift. We had concluded a similar property below Theorem 4.2 at the end of Section 4.1.

Example 1 reveals how the behavior of $u^*(X^{(n)})$ affects the centering of the posterior limit (under linearity and Gaussianity) and how it may prevent BvM from occurring when $\hat{\theta}_n + \frac{1}{n}I_{\theta_0}^{-1}u^*(X^{(n)})$ is *not* an asymptotically efficient estimator. We now turn to more precise

⁶Hellinger neighborhoods are less appropriate for misspecified models

statements by recollecting the BvM phenomenon under misspecification in LAN models [41]. The centering and the asymptotic covariance matrix will be ultimately affected by θ^* in (4.8).

Lemma 4.6. (Bernstein von-Mises) Assume that the posterior (4.7) concentrates around θ^* at the rate ε_n^* and that for every compact $K \subset \mathbb{R}^d$

$$\sup_{h \in K} \left| \log \frac{\tilde{p}_{\theta^* + \varepsilon_n^* h}^{(n)}(X^{(n)})}{\tilde{p}_{\theta^*}^{(n)}(X^{(n)})} - h' \tilde{V}_{\theta^*} \tilde{\Delta}_{n,\theta^*} - \frac{1}{2} h' \tilde{V}_{\theta^*} h \right| \to 0 \quad in \ P_{\theta_0}^{(n)} \text{-}probability$$
(4.13)

for some random vector $\widetilde{\Delta}_{n,\theta^*}$ and a non-singular matrix \widetilde{V}_{θ^*} . Then the pseudo-posterior converges to a sequence of normal distributions in total variation at the rate ε_n^* , i.e.

$$\sup_{B} \left| \Pi_{n}^{*} \left(\varepsilon_{n}^{*-1}(\theta - \theta^{*}) \in B \mid X^{(n)} \right) - N_{\widetilde{\Delta}_{n,\theta^{*}},\widetilde{V}_{\theta^{*}}}(B) \right| \to 0 \quad in \ P_{\theta_{0}}^{(n)} \text{-probability.}$$

Proof. Follows from Theorem 2.1 of [41].

It remains to examine the assumption (4.13). For iid data, [41] derived sufficient conditions (Lemma 2.1) for (4.13) to hold. Due to the non-separability of the term $u_{\theta}(X^{(n)})$, the mis-specified model cannot be regarded as arriving from an iid experiment. In Lemma 7.4 in the Appendix (Section 7.6) we nevertheless provide intuition for when (4.13) is expected to hold if $u_{\theta}(X^{(n)})$ is linear. Recall that in Remark 4.3 we have concluded that under differentiability, the posterior residual $u_{\theta}(X^{(n)})$ does converge to a linear function in θ . In Section 7.6 in the Appendix, we formulate alternative BvM conditions which are sufficient for rapid mixing in Lemma 3.1.

5 MHC in Action

To whet reader's appetite, we present MHC performance demonstrations in three examples which we found challenging for pseudo-marginal (PM) approaches and ABC. The first one (the CIR model) exemplifies data arising as discretizations of continuous-time process for which likelihood inference can be problematic [37]. We show that, compared with PM, MHC is not only far more straightforward to implement but also more scalable. The second demonstration involves a generative model (Lotka-Volterra) for which no explicit hierarchical model exists, precluding from direct application of PM methods. We thus compare MHC with ABC, showing that ABC techniques may fall short without a very informative prior and suitable summary statistics. Lastly, we consider a Bayesian model selection example where ABC faces challenges. More examples are shown in the Appendix where we show bias-variance tradeoffs between fixed and random generator variants on a toy normal location-scale model and the Ricker model [57].

5.1 The CIR Model

The CIR model [13] is prescribed by the stochastic differential equation

$$dX_t = \beta(\alpha - X_t)dt + \sigma\sqrt{X_t}dW_t$$

where W_t is the Brownian motion, $\alpha > 0$ is a mean-reverting level, $\beta > 0$ is the speed of the process and $\sigma > 0$ is the volatility parameter. This process is an integral component of the Heston model [36] where it is deployed for modelling instantaneous variances. We want to perform Bayesian inference for the parameters $\theta = (\alpha, \beta, \sigma)'$ of this continuous-time Markov process which is observed at discrete time points $t_j = j\Delta$ for $j = 1, \ldots, T$. We will assume that there are n independent observed realizations $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{iT})'$ of this discretized series for $1 \leq i \leq n$. It has been acknowledged that if the data are recoded at discrete times, parametric inference using the likelihood can be difficult, partially due to the fact that the likelihood function is often not available [37]. A natural Bayesian inferential platform for such problems is the MH algorithm where the likelihood function can be replaced with its approximation (e.g. using the analytical closed-form likelihood approximations [1] as described in [63]). [62] perform a delicate Bayesian analysis of this model using two approximate ('pseudo-marginal') MH algorithms: the MCWM algorithm (defined in [48] and discussed in [6] and [3]) and GIMH algorithm introduced in [6]. Here, we compare MHC with the MCWM variant, referring to [62] for a detailed analysis of the CIR model using GIMH.

One common approach in the literature for Bayesian estimation of diffusion models [22, 23] is to consider estimation on the basis of discrete measurements as a classic missingdata problem (see [22] and [60] for irreducible diffusion contexts). The idea is to introduce latent observations between every two consecutive data points. The time-step interval $[0, \Delta]$ is thus partitioned into M sub-intervals, each of length $h = \Delta/M$. The granularity M should be large enough so that the grid is sufficiently fine to yield more accurate likelihood approximations. With the introduction of latent variables, the pseudo-marginal approach naturally comes to mind as a possible inferential approach. The MCWM variant (described in Section 3 of [62]) alternates between simulating θ , conditionally on the missing data blocks, say U, and then updating U, given θ . We will be using the following enumeration for the missing data $U = (u_{km}^{ij})$: we have a replicate index $1 \leq i \leq n$, a discrete time index $0 \leq j \leq T$, an index of the intermittent auxiliary series $1 \leq m \leq M$ and an index $1 \leq k \leq K$ for the number of replications inside MCWM. Given θ , one can generate the missing data using the Modified Brownian Bridge (MBB) sampler [20]. Denote with $X = [\mathbf{x}_1, \ldots, \mathbf{x}_n]'$ an $n \times (T+1)$ matrix of observations where $x_{i0} = x_0$ is the initial condition. The CIR model is an interesting test bed for both MCWM and our MHC approach, because the transition density is *actually known* (i.e. non-central χ^2 [13]). We can thereby make comparisons with an exact algorithm which constructs the likelihood from the exact transition function. The likelihood can be, however, stochastically approximated as

$$\widehat{\pi}(X|\theta) = \prod_{i=1}^{n} \prod_{j=0}^{T-1} \widehat{\pi}(x_{ij+1} \mid x_{ij}, \theta), \quad \text{where} \quad \widehat{\pi}(x_{ij+1} \mid x_{ij}, \theta) = \frac{1}{N} \sum_{k=1}^{N} R_M(u_k^{ij}), \tag{5.1}$$

where $u_k^{ij} = (u_{k0}^{ij}, \ldots, u_{kM}^{ij})' \in \mathbb{R}^{M+1}$ is the k^{th} sample of the brownian bridge (described in (3) in [62]) stretching from $u_{k0}^{ij} = x_{ij}$ and $u_{kM}^{ij} = x_{ij+1}$ and where

$$R_M(u_k^{ij}) = \frac{\prod_{m=0}^{M-1} \phi\left(u_{km+1}^{ij}; u_{km}^{ij} + h\beta(\alpha - u_{km}^{ij}), \,\sigma\sqrt{hu_{km}^{ij}}\right)}{\prod_{m=0}^{M-2} \phi\left(u_{km+1}^{ij}; u_{km}^{ij} + \frac{x_{ij+1} - u_{km}^{ij}}{M - m}, \,\sigma\sqrt{h(M - m - 1)/(M - m)u_{km}^{ij}}\right)}$$

where $\phi(x; \mu, \sigma)$ denotes the normal density with a mean μ and a standard deviation σ . Regarding the choice of M and N, [64] provide asymptotic arguments for choosing $N = M^2$ and [62] make thorough comparisons for various choices of M, N and also implement the ('exact' version having the correct stationary distribution) GIMH (see their Section 4) which recycles latent data U. There are some delicate issues regarding dependency between σ and U in GIMH and we refer the reader to [64] for further details.



Figure 1: Plots of the exact and estimated log-likelihood function (up to a constant) for MCWM (upper panel using $N = M^2 = 25$) and MHC (lower panel using fixed and random generators). Log-Likelihood slice over (Left) α keeping (β_0, σ_0), (Middle) over β fixing (α_0, σ_0) and (Right) over σ keeping (α_0, β_0).

The true data consist of n = 100 samples generated using the package sde (using the function sed.sim with 'rcCIR' initialized at $x_0 = 0.1$) using $\Delta = 1$ and T = 500 and using⁷ $\theta^0 = (0.07, 0.15, 0.07)'$. In order to implement MHC, we use the LASSO-regularized logistic regression (using an R package glmnet with a value λ chosen by 10-fold cross-validation) using the entire series \boldsymbol{x}_i as predictors. While using the entire series is useful for identifying the location parameter α , capturing more subtle aspects of the series such as speed of fluctuation and spread are needed to identify (β, σ). To this end, we add summary statistics (mean, log-variance, auto-correlations at lag 1 and 2 as well as the first 3 principal components of X) yielding the total of 507 predictors (denoted with \boldsymbol{z}_i). We consider both fixed and random generators where, for the fixed variant, we fix the random seed before generating fake data which essentially corresponds to having a deterministic generative mapping.

⁷These values are close to parameter estimates found for FedFunds data analyzed in Stramer and Bognar (2011).



Figure 2: Smoothed posterior densities obtained for the CIR model by simulation using exact MH and MHC using nrep = 1 (green) and nrep = 5 (blue). Vertical lines are the true values.

We compare the MCWM likelihood approximations obtained in MCWM (using (5.1)) with various choices $N = M^2$ with the exact one using the explicit transition distribution (top panel in Figure 1). We can see that, even for a small value of N = 2, the likelihood approximation seems to have a correct shape and is peaked close to the true values (marked by vertical dotted lines). The plots show likelihood slices along each parameter, one at a time, fixing the others at their true values. The approximation quality improves for M = 5and $N = M^2$. The lower panel in Figure 1 portrays our classification-based log-likelihood (ratio) estimates $\eta = \sum_{i=1}^{n} \log[(1 - \hat{D}(\boldsymbol{z}_i))/\hat{D}(\boldsymbol{z}_i)]$ for the fixed and random generators. The curves are nicely wrapped around the true values (perhaps even more so than for MCWM) with no visible systematic bias (even for the fixed generator). While, in the fixed case (solid lines), we would expect entirely smooth curves, recall that our classifier is based on cross-validation which introduces some randomness (thereby the wiggly estimate). The wigglyness can be alleviated by averaging over (nrep) many fake data replicates (dotted lines). The random generator (dashed lines) yields slightly more variable curves compared to the fixed design, as was expected. These plots indicate that MHC 'pseudo-likelihood' contains relevant inferential information.

To implement the exact MH, MCWM and MHC (with $nrep \in \{1, 5\}$), we adopt the same prior settings as in [62], where $\pi(\theta) = \mathbb{I}_{(0,1)}(\alpha)\mathbb{I}_{(0,\infty)}(\beta)\sigma^{-1}\mathbb{I}_{(0,\infty)}(\sigma)$. We also use their



Figure 3: Smoothed posterior densities obtained by simulation using MCWM (with $N = M^2$) for M = 2 (MCWM1 green) and M = 5 (MCWM2 blue). Vertical lines are the true values.

Method	α			β			σ			AR	Time
	$\bar{\alpha}$	1	u	$\bar{\beta}$	1	u	$\bar{\sigma}$	1	u		
MH Exact	0.0693	0.683	0.703	0.1558	0.1507	0.1608	0.07	0.696	0.704	9.1	3.3
Alg1 $(nrep = 1)$	0.0691	0.0644	0.0735	0.1505	0.1374	0.1636	0.0703	0.0669	0.0734	16.8	4.6
Alg2 $(nrep = 1)$	0.0691	0.0644	0.0741	0.1476	0.1353	0.1632	0.693	0.0667	0.0725	10.7	4.9
Alg1 $(nrep = 5)$	0.0698	0.0667	0.0725	0.1468	0.1377	0.1574	0.0699	0.676	0.725	7.8	13.9
Alg2 $(nrep = 5)$	0.0691	0.0665	0.0715	0.1468	0.1366	0.1571	0.0691	0.0674	0.0714	5.6	13.9
MCWM $(M = 2)$	0.0693	0.0658	0.0733	0.1469	0.1287	0.1632	0.067	0.0657	0.0684	13.1	15.9
MCWM $(M = 5)$	0.0694	0.0662	0.723	0.1538	0.1423	0.1634	0.0689	0.0676	0.0698	10.1	238.6

Table 3: Posterior means and 95% credible interval boundaries (lower (l) and upper (u)). *AR* is the acceptance rate and **Time** is computing time (in hours) for 10000 iterations

random walk proposals.⁸ All three algorithms are initialized at the same perturbed truth and ran for 10 000 iterations with a burnin period 1 000. Smoothed posterior densities obtained by simulation using the exact MH and MHC are in Figure 2 (random generator using $nrep \in \{1, 5\}$ where fixed generator is portrayed in Figure 18 in the Appendix). The trace-plots of 10 000 iterations are depicted in Figure 19 in the Appendix, where we can see that the random generator variant yields smaller acceptance rates (especially

⁸With probability 2/3 propose a joint move (α^*, β^*) by generating $\alpha^* \sim U(\alpha - 0.01, \alpha + 0.01)$ and $\beta^* \sim U(\beta - 0.01, \beta + 0.01)$ and with probability 1/3 propose $\sigma^* \sim U(\sigma - 0.01, \sigma + 0.01)$. To increase the acceptance rate of the exact MH algorithm, we change the window from 0.01 to 0.005.

for σ) which masks the fact that the random generator sampler generally yields more spread-out posterior approximations. Smoothing out the likelihood ratio by averaging over *nrep* repetitions reduces variance where fixed and random generators seem to yield qualitatively similar results in this example (this is why we have not used the de-biasing variant here). Histograms (together with the demarkation of 95% credible set) are in Figure 21 in the Appendix. Compared with the smoothed densities obtained from MCWM (using $N = M^2$ with $M \in \{2, 5\}$ in Figure 3) we can see that MHC yields posterior reconstructions that are wrapped more closely around the true values. Increasing M, MCWM yields posterior reconstructions that are getting closer to the actual posterior (not necessarily centered more narrowly around the truth). Recall, however, that MCWM generates Markov chains whose invariant distribution is not necessarily the exact posterior. The posterior summaries (means and 95% credible intervals) are reported in Table 3. Interestingly, both MCWM intervals for σ do not include the true value 0.07 and the MCWM computation is considerably slower relative to MHC.

5.2 Lotka-Volterra Model

The Lotka-Volterra (LV) predator-prey model [71] describes population evolutions in ecosystems where predators interact with prey. The model is deterministically prescribed via a system of first-order non-linear ordinary differential equations with four parameters $\theta = (\theta_1, \ldots, \theta_4)'$ controlling (1) the rate $r_1^t = \theta_1 X_t Y_t$ of a predator being born, (2) the rate $r_2^t = \theta_2 X_t$ of a predator dying, (3) the rate $r_3^t = \theta_3 Y_t$ of a prey being born and (4) the rate $r_4^t = \theta_4 X_t Y_t$ of a prey dying. Given the initial population sizes X_0 (predators) and Y_0 (prey) at time t = 0, the process can be simulated from exactly using the Gillespie algorithm [31]. In particular, this algorithm samples times to an event from an exponential distribution (with a rate $\sum_{j=1}^4 r_j^t$) and then picks one of the 4 reactions with probabilities proportional to their individual rates r_j^t . Despite easy to sample from, the likelihood for this model is unavailable which makes this model a natural candidate for ABC [54] and other likelihood-free methods [45, 51]. It is not entirely obvious, however, how to implement the pseudo-marginal approach since there is no explicit hierarchical model structure with a conditional likelihood, given latent data, which could be marginalized through simulation



Figure 4: Lotka-Volterra realizations for three choices of θ

to obtain a likelihood estimate.

In our experiments, each simulation is started at $X_0 = 50$ and $Y_0 = 100$ simulated over 20 time units and recorded observations every 0.1 time units, resulting in a series of T = 201observations each. We plot n = 20 time series realizations for three particular choices of θ in Figure 4 which differ in the second argument θ_2 with larger values accentuating the cyclical behavior. Slight shifts in parameters result in (often) dramatically different trajectories. Typical behaviors include (a) predators quickly eating all the prey and then slowly decaying (as in Figure 4b), (b) predators quickly dying out and then the prey population skyrocketing. For certain carefully tuned values θ , the two populations exhibit oscillatory behavior. For example, in Figure 4a and 4c we can see how the value θ_2 determines the frequency of the population renewal cycle. We rely on the ability of the discriminator to tell such different shapes apart. The real data (n = 20) is generated under the scenario (a) with $\theta^0 = (0.01, 0.5, 1, 0.01)'$.

ABC analyses of this model reported in the literature have relied on various summary statistics⁹ including the mean, log-variance, autocorrelation (at lag 1 and 2) of each series as well as their cross-correlation [51]. To see whether these summary statistics are able to

⁹In addition to the summary statistics suggested in [51], we have also considered the classification accuracy ABC summary statistic $CA = \frac{1}{n+m} \left(\sum_{i=1}^{n} \hat{D}(\boldsymbol{x}_i) + \sum_{j=1}^{m} (1 - \hat{D}(\tilde{\boldsymbol{x}}_j)) \right)$ proposed by [34]. This ABC version did not provide much better results.

capture the oscillatory behavior (at different frequencies) and distinguish it from exploding population growth, we have plotted the squared $\|\cdot\|_2$ distance of the summary statistics¹⁰ (i.e. the ABC tolerance threshold ϵ) relative to the real data for a grid of values θ_2 , fixing the rest at the true values $\theta_1^0 = 0.01, \theta_3^0 = 1, \theta_4^0 = 0.01$ (see Figure 5a). We can see a V-shaped evolution of ϵ reaching a minimum near the true value $\theta_2^0 = 0.5$, especially for nrep = 20. This creates hope that ABC based on these summary statistics has the capacity to provide a reliable posterior reconstruction. Contrastingly, in Figure 5b we have plotted the estimated log-likelihood $\eta \equiv \sum_{i=1}^{n} \log[(1 - \hat{D}(\boldsymbol{x}_i))/\hat{D}(\boldsymbol{x}_i)]$ (as a function of θ_2) where $\boldsymbol{x}_i = (X_1^i, \dots, X_T^i, Y_1^i, \dots, Y_T^i)'$ after training the LASSO-penalized logistic regression classifier on m = n fake data observations $\widetilde{\boldsymbol{x}}_i = (\widetilde{X}_1^i, \dots, \widetilde{X}_T^i, \widetilde{Y}_1^i, \dots, \widetilde{Y}_T^i)'$ for $1 \leq i \leq m$ using the cross-validated penalty λ (using the R package glmnet). Similarly as for ABC, we plot η for a single realization of fake data as well as the average η over *nrep* many replications. We can see the curve peak around the true value θ_2^0 (even for nrep = 1), indicating that the estimated log-likelihood contains relevant information which could be exploited within MH. We can also see that only a small range of values θ_2 will provide fake data that are compatible with the real data. In addition, we have seen only a small subset of parameters to give rise to the oscillatory behavior and, thus, we expect sharply peaked posteriors around the true values. This intuition is confirmed by heat-map plots of the estimated likelihood η as a function of $(\theta_2, \theta_3)'$ (Figure 6a) and as a function of $(\theta_1, \theta_4)'$ (Figure 6b), keeping the remaining parameters at the truth. In Figure 6b, we can see a sharp spike (approximating a point-mass) around the true value at $\theta_1 = \theta_4 = 0.01$ in a otherwise vastly flat landscape. This peculiar likelihood property may require a very careful consideration of initializations and proposal densities for MH and the prior domain for ABC.

In order to facilitate ABC analysis, we have used an informative uniform prior $\theta \sim U(\Xi)$ with a restricted domain $\Xi = [0, 0.1] \times [0, 1] \times [0, 2] \times [0, 0.1]$ so that the procedure does not waste time sampling from unrealistic parameter values. These values were chosen based on a visual inspection of simulated evolutions, where we have seen only a limited range of

 $^{^{10}}$ Out of curiosity, we have considered a single fake dataset as well as the average tolerance over nrep fake data replications.



Figure 5: Lotka-Volterra model. ABC discrepancy ϵ and the log-likelihood 'estimator' η .



Figure 6: Lotka-Volterra model. Estimated log-likelihood for a grid of parameters.

values to yield periodic behavior. In a pilot ABC run, we rank $M = 10\,000$ ABC samples based on ε in an ascending manner and report the histogram of the first r = 100 samples (Figure 7, the upper panel). We can see that ABC was able to narrow down the region of interest for (θ_1, θ_4) , but is still largely uninformative about parameters (θ_2, θ_3) with histograms stretching from the boundaries of the prior domain. Given how narrow the range of likely parameter values is (according to Figure 6), the likelihood of encountering such values even under the restricted uniform prior is still quite negligible. We thereby tried many more ABC samples ($M = 100\,000$ which took 47.46 hours) only to find out that



Figure 7: ABC analysis of the Lotka-Volterra model. Upper panel uses $M = 10\,000$ and r = 100 whereas the lower panel uses $M = 100\,000$ and $r = 1\,000$. Vertical red lines mark the true values.

the histograms (top $r = 1\,000$ samples) did not improve much (Figure 7, the lower panel).

The hostile likelihood landscape will create problems not only for ABC but also for Metropolis-Hastings. Indeed, initializations that are too far from the likelihood domain may result in Markov chains wandering aimlessly in the vast plateaus for a long time. Rather than competing with ABC, a perhaps more productive strategy is to combine the strengths of both. We have thereby used the pilot ABC run (the closest 100 samples out of $M = 10\,000$ which took roughly 4 hours) to obtain ABC approximated posterior means $\hat{\theta} = (0.015, 0.55, 1.31, 0.012)'$. We use these to initialize our MHC procedure to accelerate convergence (i.e. prevent painfully long burn-in). To implement MHC, we define a Gaussian random walk proposal for log-parameter values with a proposal standard deviation 0.05 and deploy the same prior as for the ABC method. We use the random generator variant here, where the fixed one can be implemented (for example) by fixing the random seed prior generating the fake data. The trace-plots after $M = 10\,000$ iterations (which took roughly 3 hours) and histograms of the parameters (after the burn-in period 1000) are portrayed

	$\theta_1^0 = 0.01$		$\theta_2^0 = 0.5$			$\theta_3 = 1$			$\theta_4 = 0.01$			
Method	$\bar{ heta}$	l	u	$\bar{ heta}$	l	u	$\bar{ heta}$	l	u	$\bar{ heta}$	l	u
ABC1	0.015	0.003	0.038	0.554	0.037	0.985	1.315	0.189	1.955	0.012	0.004	0.029
ABC2	0.016	0.003	0.042	0.604	0.087	0.980	1.259	0.205	1.971	0.013	0.003	0.024
MHC	0.01	0.008	0.014	0.531	0.41	0.685	1.029	0.791	1.301	0.010	0.007	0.014

Table 4: Posterior summary statistics using ABC1 ($M = 10\,000$ and r = 100), ABC2 ($M = 100\,000$ and $r = 1\,000$) and MHC ($M = 10\,000$ with burnin 1\,000). $\bar{\theta}$ denotes posterior mean, l and u denote the lower and upper boundaries of 95% credible intervals.

in Figure 8. We can see reasonable mixing where the acceptance probability was 0.17. The histograms report much sharper concentration around true values (compared to ABC in Figure 6) and were obtained under considerable less time (again compared to ABC with $M = 100\,000$). The posterior summaries (mean $\bar{\theta}$ and 95% credible intervals (l, u) are compared in Table 4. We can see that MHC posterior mean not only accurately estimates the true parameters, but the 95% credible intervals are much tighter and thereby perhaps more informative for inference. We believe that MHC (in collaboration with ABC pilot run) provided an inferential framework which was not attainable using neither ABC (with our choice of summary statistics), nor the pseudo-marginal method. Potentially more fruitful ABC results could be obtained by instead deploying Wasserstein distance between the empirical distributions of real and fake data [9], in particular its curve-matching variants tailored for dependent data. Other possible distances that are useful for time series include Skorokhod distance [43] or transportation distances (see [9] for related references).

5.3 Bayesian Model Selection

The performance of summary statistic-based methods is ultimately sensitive to the quality of summary statistics whose selection can be a delicate matter. One such instance is model selection, where it is known that when ABC may fail even when the summary statistic is sufficient for *each* of the models considered [59]. Our method *does not* require a summary statistic but a sieve of discriminators that can adapt to the oracle discriminator in the limit.



Figure 8: MHC analysis of the Lotka-Volterra model. Upper panel portrays MCMC trace plots (with $M = 10\,000$) where the histograms (without the burnin 1000) are in the lower panel.

This creates hope that our method can tackle model selection problems. To illustrate this point we consider a toy model choice problem considered in [59]. The actual data follows $X_i \sim N(0,1)$ for i = 1, ..., n = 500. We have two candidate models $P_{1,\mu} = N(\mu, 1)$ and $P_{2,\mu} = N(\mu, 1 + 3/\sqrt{n})$ to choose from. We let the parameters be $\theta := (m, \mu)$, where $m \in \{1, 2\}$ is the model indicator and μ is unknown mean with a prior N(0, 1). The model is assigned a uniform prior, i.e. P(m = 1) = P(m = 2) = 0.5. Following the traditional Bayesian model selection formalism, we collect evidence for model m = 1 with a Bayes factor

$$B_{12} := \frac{\pi_n(m=1 \mid X)}{\pi_n(m=2 \mid X)}.$$

The Bayes factor is the ratio of the marginal likelihoods (or posterior probabilities) of m = 1 over m = 2. The actual Bayes factor value is $B_{12} = 9$, indicating strong evidence in favor of m = 1. The Bayes factor will be estimated by the ratio of the frequencies



Figure 9: Trace plots of sampled models using: (Left) MH with the true likelihood ratio, (Middle) ABC with $s(X^{(n)}) = \bar{X}_n$ and (Right) fixed generator MHC.

of the posterior samples given by ABC or our method. Since our parameter of interest m is discrete, there is no de-biasing for this example. [59] in their Lemma 2 show that when the summary statistic is $\sum_i X_i$, the Bayes factor estimated by ABC asymptotes to 1. This is equivalent to choosing the model with a coin toss. For our method, we use the logistic regression on regressors $(1, X_i, X_i^2)$, which can mimic the oracle discriminator. The trace plots of sampled models for exact MH, MHC and ABC are provided in Figure 9. Table 5 summarizes the posterior model frequencies. The true posterior probabilities are $\pi_n(m = 1 \mid X) \approx 0.9$ and $\pi_n(m = 2 \mid X) \approx 0.1$, so the Bayes factor is 9. The "Oracle MH" is the Metropolis-Hastings with the true likelihood, in which 84.4% of the posterior draws choose model 1. Algorithms 1 and 2 choose model 1 respectively 93.2% and 70% of the times. ABC based on the sum, on the other hand, chooses the model randomly. Finally, Figure 17 in Appendix gives the estimated log-likelihood ratio for each model. In terms of μ , we again see that Algorithm 1 is slightly biased with the correct shape and Algorithm 2 is less biased but more dispersed on average.

6 Discussion

This paper develops an approximate Metropolis-Hastings (MH) posterior sampling method for when the likelihood is not tractable. By deploying a Generator and a Classifier (similarly as in Generative Adversarial Networks [32]), likelihood ratio estimators are obtained which are then plugged into the MH sampling routine. One of the main distinguishing

	Posterior	Oracle MH	Algorithm 1	Algorithm 2	ABC
Model 1	90%	422	466	350	252
Model 2	10%	78	34	150	248
Bayes factor	9.00	5.41	13.71	2.33	1.02

Table 5: "Posterior" column gives the posterior probability of each model, $\pi_n(m = j \mid X)$. Other columns give the frequencies of the corresponding sample of size 500. "Oracle MH" refers to the Metropolis-Hastings algorithm with the true likelihood. "ABC" is based on the summary statistics $s(X) = \bar{X}_n$.

features of our work (relative to other related approaches [53]) is that we consider two variants: (1) a fixed generator design yielding biased samples, and (2) a random generator yielding more dispersed samples. We provide a thorough frequentist characterization of the stationary distribution including convergence rates and asymptotic normality. Under suitable differentiability assumptions, we conclude that correct shape and location can be recovered by deploying a debiasing combination of the fixed and random generator variants. We demonstrate a very satisfactory performance on non-trivial time series examples which render existing techniques (such as PM or ABC) less practical. Along with our theoretical development, we also establish a new bound on the Kullback-Leibler divergence and variation by possibly non-divergent multiples of the Hellinger distance (Lemma 7.1). This lemma will be of independent interest to prove sharper rates of posterior contraction in models with unbounded likelihood ratios, compared to using previously known results.

Acknowledgements

Tetsuya Kaji gratefully acknowledges the support from the Richard N. Rosett Faculty Fellowship and the Liew Family Faculty Fellowship at the University of Chicago Booth School of Business. Veronika Rockova gratefully acknowledges the support from James S. Kemper Faculty Scholarship and the National Science Foundation (DMS: 1944740).

References

- Ait-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica*, 70:223–262.
- [2] Andrieu, C., Doucet, A., and Holenstein, R. (2009). Particle Markov chain Monte Carlo methods. Journal of the Royal Statistical Society (Series B), 72(2):1–33.
- [3] Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725.
- [4] Andrieu, C. and Vihola, M. (2015). Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *The Annals of Applied Probability*, 25(2):1030–1077.
- [5] Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. (2019). Discriminator rejection sampling. In *International Conference on Learning Representations*.
- [6] Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–1160.
- [7] Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035.
- [8] Belloni, A. and Chernozhukov, V. (2009). On the computational complexity of MCMCbased estimators in large samples. *The Annals of Statistics*, 37(4):2011–2055.
- [9] Bernton, E., Jacob, P., Gerber, M., and Robert, C. (2019). Approximate Bayesian computation with the Wasserstein distance. *Journal of the Royal Statistiscial Society* (Series B), pages 1–50.
- [10] Blum, M. (2010). Approximate Bayesian computation: A nonparametric perspective. Journal of the American Statistical Association, 105(491):1178–1187.
- [11] Blum, M., Nunes, M., Prangle, D., and Sisson, S. (2013). A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science*, 28(2):189–208.

- [12] Brehmer, J., Louppe, G., Pavez, J., and Cranmer, K. (2019). Mining gold from implicit models to improve likelihood-free inference. arXiv, pages 1–10.
- [13] Cox, J., Ingersoll, J., and Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica*, 53:385–407.
- [14] Cranmer, K., Pavez, J., and Louppe, G. (2016). Approximating likelihood ratios with calibrateed discriminative classifiers. arXiv, pages 1–30.
- [15] De Blasi, P. and Walker, S. (2013). Bayesian estimation of the discrepancy with misspecified parametric models. *Bayesian Analysis*, 8:781–800.
- [16] Deligiannidis, G., Doucet, A., and Pitt, M. K. (2018). The correlated pseudomarginal method. Journal of the Royal Statistical Society (Series B), 80(5):839–870.
- [17] Diggle, P. J. and Gratton, R. J. (1984). Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society (Series B)*, 46(2):193–212.
- [18] Donnet, S., Rivoirard, V., Rousseau, J., and Scricciolo, C. (2018). Posterior concentration rates for empirical Bayes procedures with applications to Dirichlet process mixtures. *Bernoulli*, 24(1):231–256.
- [19] Doucet, A., Pitt, M., Deligiannidis, G., and Kohn, R. (2014). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. arXiv, pages 1–40.
- [20] Durham, G. and Gallant, A. (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics*, 20:297–338.
- [21] Durkan, C., Murray, I., and Papamakarios, G. (2020). On contrastive learning for likelihood-free inference. Proceedings of the 37th International Conference on Machine Learning, 119:297–338.

- [22] Elerian, O., Chib, S., and Shephard, N. (2001). Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, 69:959–993.
- [23] Eraker, B. (2001). The pseudo-marginal approach for efficient Monte Carlo computations. Journal of Business & Economic Statistics, 19:177–191.
- [24] Fearnhead, P. and Prangle, D. (2011). Constructing ABC summary statistics: semiautomatic ABC. *Nature Precedings*, pages 1–10.
- [25] Frazier, D. and Drovandi, C. (2021). Robust approximate Bayesian inference with synthetic likelihood. *Journal of Computational and Graphical Statistics (to appear)*.
- [26] Frazier, D., Martin, G., Robert, C., and Rousseau, J. (2018). Asymptotic properties of approximate Bayesian computation. *Biometrika*, 105(2):593–607.
- [27] Frazier, D., Robert, C., and Rousseau, J. (2020). Model misspecification in approximate Bayesian computation: consequences and diagnostics. *Journal of the Royal Statistical Society (Series B)*, 82:421–444.
- [28] Ghosal, S., Ghosh, J. K., and van der Vaart, A. W. (2000). Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500–531.
- [29] Ghosal, S. and van der Vaart, A. W. (2007). Convergence rates of posterior distributions for noniid observations. *The Annals of Statistics*, 35(1):192–223.
- [30] Ghosal, S. and van der Vaart, A. W. (2017). Fundamentals of Nonparametric Bayesian Inference. Cambridge University Press, Cambridge.
- [31] Gillespie, D. (1977). Exact stochastic simulation of coupled chemical reactions. The Journal of Physical Chemistry, 81(25):2340–2361.
- [32] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., and Ozair, S. (2014). Generative adversarial nets. Proceedings of the 27th International Conference on Neural Information Processing Systems, 2:2672–2680.

- [33] Gutmann, M. U. and Corander, J. (2016). Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(1):4256–4302.
- [34] Gutmann, M. U., Dutta, R., Kaski, S., and Corander, J. (2018). Likelihood-free inference via classification. *Statistics and Computing*, 28:411–425.
- [35] Heermans, J., Begy, V., and Louppe, G. (2020). Likelihood-free MCMC with amortized approximate ratio estimators. Proceedings of the 37th International Conference on Machine Learning, 119(37):15112–15117.
- [36] Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bonds and currency options. *Review of Financial Studies*, 6:327–343.
- [37] Heston, S. (2004). Parametric inference for diffusion processes observed at discrete points in time: a surveys. *International Statistical Review*, 72:337–354.
- [38] Jiang, B., Wu, T., Zheng, C., and Wong, W. (2017). Learning summary statistic for approximate Bayesian computation via deep neural network. *Statistica Sinica*, 27:1595– 1618.
- [39] Kaji, T., Manresa, E., and Pouliot, G. (2020). An adversarial approach to structural estimation. *arXiv*.
- [40] Kleijn, B. and van der Vaart, A. (2006). Misspecification in infinite-dimensional Bayesian statistics. The Annals of Statistics, 34:837–877.
- [41] Kleijn, B. and van der Vaart, A. (2012). The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354 – 381.
- [42] Lovasz, L. and Simonovits, M. (1993). Random walks in a convex body and an improved volume algorithm. *Random Structures and Algorithms*, 4:359–412.
- [43] Majumdar, R. and Prabhu, V. (2015). Computing the Skorokhod distance between polygonal traces. Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, pages 1–10.

- [44] Marin, J., Pillai, N., Robert, C., and Rousseau, J. (2014). Relevant statistics for Bayesian model choice. *Journal of the Royal Statistiscial Society (Series B)*, 76(5):833– 859.
- [45] Meeds, T. and Welling, M. (2015). Optimization Monte Carlo: Efficient and embarrassingly parallel likelihood-free inference. Advances in Neural Information Processing Systems, 28(1):289–309.
- [46] Mengersen, K. and Tweedie, R. (1996). Rates of convergence of the Hastings and Metropolis algorithms. *The Annals of Statistics*, 24(1):101–121.
- [47] Mohamed, S. and Lakshminarayanan, B. (2017). Learning in implicit generative models. arXiv, pages 1–10.
- [48] O'Neil, P., Balding, D., Becker, N., Serola, M., and Mollison, D. (2000). Analyses of infectious disease data from household outbreaks by Markov chain Monte Carlo methods. *Applied Statistics*, 49:517–542.
- [49] O'Neill, P. D., Balding, D. J., Becker, N. G., Eerola, M., and Mollison, D. (2000). Analyses of infectious disease data from household outbreaks by Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society (Series C)*, 49(4):517–542.
- [50] O'Ryan, C., Harley, E. H., Bruford, M. W., Beaumont, M., Wayne, R. K., and Cherry, M. I. (1998). Microsatellite analysis of genetic diversity in fragmented South African buffalo populations. *Animal Conservation*, 1(2):85–94.
- [51] Papamakarios, G. and Murray, I. (2016). Fast ε-free inference of simulation models with Bayesian conditional density estimation. Advances in Neural Information Processing Systems, 29(1):289–309.
- [52] Patilea, V. (2001). Convex models, MLS and misspecification. The Annals of Statistics, 20:94–123.

- [53] Pham, K., Nott, D., and Chaudhuri, S. (2014). A note on approximating ABC-MCMC using flexible classifiers. The ISI's Journal for the Rapid Dissemination of Statistics Research, 3:218–227.
- [54] Prangle, D. (2017). Adapting the ABC distance function. Bayesian Analysis, 12(1):289–309.
- [55] Price, L., Drovandi, C., Lee, A., and Nott, D. (2018). Bayesian synthetic likelihood. Journal of Computational and Graphical Statistics, 27(1):1–119.
- [56] Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A., and Feldman, M. W. (1999). Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular biology and evolution*, 16(12):1791–1798.
- [57] Ricker, W. E. (1954). Stock and recruitment. Journal of the Fisheries Board of Canada, 11(5):559–623.
- [58] Robert, C. P. and Casella, G. (2004). Monte Carlo Statistical Methods. Springer, New York, second edition.
- [59] Robert, C. P., Cornuet, J.-M., Marin, J.-M., and Pillai, N. S. (2011). Lack of confidence in approximate Bayesian computation model choice. *Proceedings of the National Academy of Sciences*, 108(37):15112–15117.
- [60] Roberts, G. and Stramer, O. (2001). On inference for partially observed nonlinear diffusion models using the Metropolis-Hastings algorithms. *Biometrika*, 88:603–621.
- [61] Rousseau, J. and Szabo, B. (2017). Asymptotic behaviour of the empirical Bayes posteriors associated to maximum marginal likelihood estimator. *The Annals of Statistics*, 45(2):833–865.
- [62] Stramer, O. and Bognar, M. (2011). Bayesian inference for irreducible diffusion processes using the pseudo-marginal approach. *Bayesian Analysis*, 6:231–258.

- [63] Stramer, O., Bognar, M., and Schneider, P. (2010). Bayesian inference of discretely sampled Markov processes with closed-form likelihood expansions. *The Journal of Financial Econometrics*, 8:450–480.
- [64] Stramer, O. and Yan, J. (2007). Asymptotics of an efficient Monte Carlo estimation for the transition density of diffusion processes. *Methodology and Computing in Applied Probability*, 9:483–496.
- [65] Tavaré, S., Balding, D. J., Griffiths, R. C., and Donnelly, P. (1997). Inferring coalescence times from DNA sequence data. *Genetics*, 145(2):505–518.
- [66] Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. (2021). Likelihoodfree inference by ratio estimation. *Bayesian Analysis*, 1:1–31.
- [67] Tsvetkov, D., Hristov, L., and Angelova-Slavova, R. (2017). On the convergence of the Metropolis-Hastings Markov chains. *Serdica Math. J.*, 43(2):93–110.
- [68] Turner, R., Hung, J., Frank, E., Saatchi, Y., and Yosinski, J. (2019). Metropolis-Hastings generative adversarial networks. In *International Conference on Machine Learning*, pages 6345–6353. PMLR.
- [69] van der Vaart, A. W. (1998). Asymptotic Statistics. Cambridge University Press, Cambridge.
- [70] van der Vaart, A. W. and Wellner, J. A. (1996). Weak Convergence and Empirical Processes: With Applications to Statistics. Springer, New York.
- [71] Wilkinson, D. (2011). Stochastic Modelling for Systems Biology. Chapman and Hall/CRC, second edition.
- [72] Wood, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104.

7 Appendix

7.1 Proof of Theorem 4.1

The following lemma bounds the Kullback-Leibler divergence and variation by possibly non-diverging multiples of the Hellinger distance.¹¹ This can be used to derive sharper rates of posterior contraction in models with unbounded likelihood ratios [see 30, p. 199 and Appendix B].

Lemma 7.1. For probability measures P and P_0 such that $P_0(p_0/p) < \infty$, let $M := \inf_{c\geq 1} cP_0(\frac{p_0}{p} \mid \frac{p_0}{p} \geq [1+\frac{1}{2c}]^2)$ where $P_0(\cdot \mid A) = 0$ if $P_0(A) = 0$. For $k \geq 2$, the following hold.

(i)
$$-P_0 \log \frac{p}{p_0} \le (3+M)h(p,p_0)^2$$
.

(*ii*)
$$P_0 |\log \frac{p}{p_0}|^k \le 2^{k-1} \Gamma(k+1)(2+M)h(p,p_0)^2$$
.

(*iii*)
$$P_0 |\log \frac{p}{p_0} - P_0 \log \frac{p}{p_0}|^k \le 2^{2k-1} \Gamma(k+1)(2+M)h(p,p_0)^2.$$

(*iv*)
$$\|\frac{1}{2}\log \frac{p}{p_0}\|_{P_0,B}^2 \le (2+M)h(p,p_0)^2.$$

(v)
$$\|\frac{1}{4}(\log \frac{p}{p_0} - P_0 \log \frac{p}{p_0})\|_{P_0,B}^2 \le (2+M)h(p,p_0)^2.$$

Here, $||f||_{P,B} := \sqrt{2P(e^{|f|} - 1 - |f|)}$ is the Bernstein "norm".

Proof. (iv) Using $e^{|x|} - 1 - |x| \le (e^x - 1)^2$ for $x \ge -\frac{1}{2}$ and $e^{|x|} - 1 - |x| < e^x - \frac{3}{2}$ for $x > \frac{1}{2}$,

$$\left\|\log \sqrt{\frac{p}{p_0}}\right\|_{P_0,B}^2 \le 2P_0 \left(\sqrt{\frac{p}{p_0}} - 1\right)^2 \mathbb{1}\left\{\frac{p}{p_0} \ge \frac{1}{e}\right\} + 2P_0 \left(\sqrt{\frac{p_0}{p}} - \frac{3}{2}\right) \mathbb{1}\left\{\frac{p_0}{p} > e\right\}.$$

The first term is bounded by $2h(p, p_0)^2$. For every $c \ge 1$,

$$P_0\left(\sqrt{\frac{p_0}{p}} - \frac{3}{2}\right) \mathbb{1}\left\{\frac{p_0}{p} > e\right\} \le P_0\left(\sqrt{\frac{p_0}{p}} - 1 - \frac{1}{2c}\right) \mathbb{1}\left\{\sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right\}$$
$$= P_0\left(\sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) \left[P_0\left(\sqrt{\frac{p_0}{p}} - 1 \mid \sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) - \frac{1}{2c}\right].$$

¹¹Lemma 7.1 (iv) first appeared in Kaji et al. [39, Lemma 5]. We reproduce the proof here as it is used to prove other statements.

Since $x - \frac{1}{2c} \le \frac{c}{2}x^2$ for every x,

$$\begin{aligned} P_0\left(\sqrt{\frac{p_0}{p}} - 1 \mid \sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) - \frac{1}{2c} &\leq \frac{c}{2} \left[P_0\left(\sqrt{\frac{p_0}{p}} - 1 \mid \sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) \right]^2 \\ &\leq \frac{c}{2} P_0\left(\frac{p_0}{p} \mid \sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) P_0\left(\left[1 - \sqrt{\frac{p}{p_0}}\right]^2 \mid \sqrt{\frac{p_0}{p}} \ge 1 + \frac{1}{2c}\right) \end{aligned}$$

by the Cauchy-Schwarz inequality. Then the result follows.

(i) Write $-P_0 \log \frac{p}{p_0} = P_0(\frac{p}{p_0} - 1 - \log \frac{p}{p_0}) + P(p_0 = 0)$. With $\frac{1}{x} - 1 - \log \frac{1}{x} < 2(\sqrt{x} - \frac{3}{2})$ for every $x \ge 3$,

$$P_0\left(\frac{p}{p_0} - 1 - \log\frac{p}{p_0}\right) \le P_0\left(\frac{p}{p_0} - 1 - \log\frac{p}{p_0}\right) \mathbb{1}\left\{\frac{p}{p_0} > \frac{1}{3}\right\} + 2P_0\left(\sqrt{\frac{p_0}{p}} - \frac{3}{2}\right) \mathbb{1}\left\{\frac{p_0}{p} \ge 3\right\}$$

The second term is bounded as above. For the first term, observe that

$$\frac{P_0(\frac{p}{p_0} - 1 - \log\frac{p}{p_0})\mathbb{1}\left\{\frac{p}{p_0} > \frac{1}{3}\right\}}{P_0(1 - \sqrt{p/p_0})^2\mathbb{1}\left\{\frac{p}{p_0} > \frac{1}{3}\right\}} \le \sup_{p/p_0 > 1/3} \frac{\frac{p}{p_0} - 1 - \log\frac{p}{p_0}}{(1 - \sqrt{p/p_0})^2} < 3$$

With $P(p_0 = 0) = \int (\sqrt{p} - \sqrt{p_0})^2 \mathbb{1}\{p_0 = 0\}$ follows the result.

(ii) Since $e^x - 1 - x \ge x^k / \Gamma(k+1)$ for $k \ge 2$ and $x \ge 0, {}^{12} P_0 |\log \frac{p}{p_0}|^k \le 2^{k-1} \Gamma(k+1) ||\frac{1}{2} \log \frac{p}{p_0}||_{P_0,B}^2$. Then, apply (iv).

(iii) By the triangle and Hölder's inequalities, for $k \ge 1$, $P_0 |\log \frac{p}{p_0} - P_0 \log \frac{p}{p_0}|^k \le [(P_0 |\log \frac{p}{p_0}|^k)^{1/k} + P_0 \log \frac{p}{p_0}]^k \le 2^k P_0 |\log \frac{p}{p_0}|^k$. Then, use (ii).

(v) By the convexity of $e^{|x|} - 1 - |x|$ and Jensen's inequality, $\|\frac{1}{4}(\log \frac{p}{p_0} - P_0 \log \frac{p}{p_0})\|_{P_0,B}^2 \leq \frac{1}{2}\|\frac{1}{2}\log \frac{p}{p_0}\|_{P_0,B}^2 + \frac{1}{2}\|P_0\frac{1}{2}\log \frac{p}{p_0}\|_{P_0,B}^2 \leq \|\frac{1}{2}\log \frac{p}{p_0}\|_{P_0,B}^2$. With (iv) follows the result.

Proof of Theorem 4.1. For $D \in \mathcal{D}$, write $\mathbb{P}_n(\log \frac{1-D}{1-D_{\theta}} - \log \frac{D}{D_{\theta}})$ as

$$P_0 \log \frac{1-D}{1-D_{\theta}} - P_0 \log \frac{D}{D_{\theta}} + (\mathbb{P}_n - P_0) \log \frac{1-D}{1-D_{\theta}} - (\mathbb{P}_n - P_0) \log \frac{D}{D_{\theta}}.$$

Let $W_1 := \sqrt{\frac{D}{D_{\theta}}} - 1$, $W_2 := \sqrt{\frac{1-D}{1-D_{\theta}}} - 1$, and $\delta := d_{\theta}(D, D_{\theta})$. By Taylor's theorem, $\log(1+x) = x - \frac{1}{2}x^2 + \frac{1}{2}x^2R(x)$ where R(x) = O(x) as $x \to 0$. Therefore, $P_0 \log \frac{D}{D_{\theta}} = 2P_0W_1 - P_0W_1^2 + P_0W_1^2R(W_1)$. Note that $P_0W_1^2 = P_0(\sqrt{D/D_{\theta}} - 1)^2 = h_{\theta}(D, D_{\theta})^2$. Since $W_1^2 \ge 0$, this implies that $W_1(X_i)^2 = O_P(\delta^2)$ and $W_1(X_i) = o_P(1)$. Also,

$$2P_0W_1 = \left[2P_0\frac{\sqrt{D(p_0+p_\theta)}}{\sqrt{p_0}} - \int D(p_0+p_\theta) - \int p_0\right] + (P_0+P_\theta)(D-D_\theta)$$
$$= -h_\theta(D,D_\theta)^2 + (P_0+P_\theta)(D-D_\theta).$$

 ${}^{12}\Gamma(k-1) \ge \int_x^\infty y^{k-2} e^{-y} dy \ge x^{k-2} e^{-x} \text{ implies } \frac{d^2}{dx^2} (e^x - 1 - x) \ge \frac{d^2}{dx^2} x^k / \Gamma(k+1).$

Note that $(P_0 + P_\theta)|D - D_\theta| \le (P_0 + P_\theta)(\sqrt{D} + \sqrt{D_\theta})|\sqrt{D} - \sqrt{D_\theta}| \le 2\sqrt{2}h_\theta(D, D_\theta)$ by the Cauchy-Schwarz inequality. Next, for $1/5 \le c < 1$,

$$\begin{aligned} |P_0 W_1^2 R(W_1)| &\leq P_0 W_1^2 |R(W_1)| \mathbb{1}\{W_1 \leq -c\} + P_0 W_1^2 |R(W_1)| \mathbb{1}\{W_1 > -c\} \\ &\leq P_0 (-R(W_1) \mathbb{1}\{W_1 \leq -c\}) + P_0 W_1^2 |R(-c) \lor R(W_1)|. \end{aligned}$$

Since R(x) < 1 and $R(W_1) = o_P(1)$, the second term is $o(\delta^2)$ for every c by the dominated convergence theorem. By the diagonal argument, there exists a sequence $c \to 1$ for given $D \to D_{\theta}$ such that the second term remains $o(\delta^2)$. Since $0 < -R(x) < -2\log(1+x)$ for $x \leq -\frac{1}{5}$,

$$P_{0}(-R(W_{1})\mathbb{1}\{W_{1} \leq -c\}) \leq P_{0}\left(\log \frac{D_{\theta}}{D}\mathbb{1}\{W_{1} \leq -c\}\right) \\ = P_{0}\left(\frac{D}{D_{\theta}}\log \frac{D_{\theta}}{D} \cdot \frac{D_{\theta}}{D}\mathbb{1}\{W_{1} \leq -c\}\right) \leq \sup_{x \geq (1-c)^{-2}}\left|\frac{1}{x}\log x\right| \cdot P_{0}\left(\frac{D_{\theta}}{D}\mathbb{1}\{W_{1} \leq -c\}\right).$$

The first term is o(1) as $c \to 1$. The second term is bounded by $P_0(\frac{D_{\theta}}{D}\mathbb{1}\{W_1 \leq -\frac{1}{5}\}) = P_0(W_1 \leq -\frac{1}{5})P_0(\frac{D_{\theta}}{D} \mid \frac{D_{\theta}}{D} \geq \frac{25}{16}) \leq P_0(W_1 \leq -\frac{1}{5})M$ by Assumption 2. By Markov's inequality, $P_0(W_1 \leq -\frac{1}{5}) \leq 25P_0W_1^2 = O(\delta^2)$. Thus, $|P_0W_1^2R(W_1)| = o(\delta^2)$. Altogether, we have $P_0\log\frac{D}{D_{\theta}} = O(\delta)$.

Next, write $P_0 \log \frac{1-D}{1-D_{\theta}} = 2P_0W_2 - P_0W_2^2 + P_0W_2^2R(W_2)$. By the Cauchy-Schwarz inequality,

$$P_0 W_2 \leq \sqrt{P_0 \frac{p_0}{p_\theta}} \cdot h_\theta (1 - D, 1 - D_\theta) \leq \sqrt{M} \delta,$$
$$P_0 W_2^2 \leq \sqrt{(P_0 + P_\theta) \left(\frac{p_0}{p_\theta}\right)^2 (\sqrt{1 - D} - \sqrt{1 - D_\theta})^2} \cdot h_\theta (1 - D, 1 - D_\theta)$$

Since D and D_{θ} are bounded by 0 and 1,

$$(P_0 + P_{\theta}) \left(\frac{p_0}{p_{\theta}}\right)^2 (\sqrt{1 - D} - \sqrt{1 - D_{\theta}})^2 \le P_0 \left(\frac{p_0}{p_{\theta}}\right)^2 + P_0 \frac{p_0}{p_{\theta}} \le 2M.$$

Therefore, by the dominated convergence theorem, the LHS is o(1) and hence $P_0W_2^2 = o(\delta)$. This also implies $W_2(X_i) = o_P(1)$, $W_2^2(X_i) = o_P(\delta)$, and $R(W_2(X_i)) = o_P(1)$. Next, similarly as before, for $1/5 \le c < 1$,

$$|P_0 W_2^2 R(W_2)| \le P_0(-R(W_2) \mathbb{1}\{W_2 \le -c\}) + P_0 W_2^2 |R(-c) \lor R(W_2)|.$$

There exists a sequence $c \to 1$ such that the second term is $o(\delta)$. Also,

$$P_0(-R(W_2)\mathbb{1}\{W_2 \le -c\}) \le \sup_{x \ge (1-c)^{-2}} |\frac{1}{x} \log x| \cdot P_0(\frac{1-D_\theta}{1-D}\mathbb{1}\{W \le -c\}).$$

The first term is o(1) as $c \to 1$. The second term is bounded by $P_0(W_2 \leq -\frac{1}{5})M$ by Assumption 2. By Markov's inequality, $P_0(W_2 \leq -\frac{1}{5}) \leq 25P_0W_2^2 = O(\delta)$. Thus, $|P_0W_2^2R(W_2)| = o(\delta)$. Altogether, we have $P_0 \log \frac{1-D}{1-D_{\theta}} = O(\delta)$.

Next, we bound $\mathbb{E}^* \sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} |\sqrt{n}(\mathbb{P}_n - P_0) \log \frac{D}{D_{\theta}}|$. Under Assumption 2, an analogous argument as Lemma 7.1 (iv) yields that $\left\|\frac{1}{2}\log \frac{D}{D_{\theta}}\right\|_{P_0,B}^2 \leq 2(1+M)h_{\theta}(D, D_{\theta})^2 = O(\delta^2)$. By van der Vaart and Wellner [70, Lemma 3.4.3], we have

$$\mathbb{E}^* \sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} \left| \sqrt{n} (\mathbb{P}_n - P_0) \log \frac{D}{D_{\theta}} \right| \lesssim J \left(1 + \frac{J}{\delta^2 \sqrt{n}} \right).$$

for $J = J_{[]}(\delta, \{\log \frac{D}{D_{\theta}} : D \in \mathcal{D}_{n,\delta_n}^{\theta}\}, \|\cdot\|_{P_0,B})$. Note that a δ -bracket in $\mathcal{D}_{n,\delta_n}^{\theta}$ induces a $C\delta$ bracket in $\{\log \frac{D}{D_{\theta}}\}$ for some constant C since $\left\|\log \frac{u}{D_{\theta}} - \log \frac{\ell}{D_{\theta}}\right\|_{P_0,B}^2 \leq 4P_0\left(\sqrt{u/\ell} - 1\right)^2 = O(d_{\theta}(u,\ell)^2)$ by Assumption 2. Therefore, $J \leq J_{[]}(\delta, \mathcal{D}_{n,\delta_n}^{\theta}, d_{\theta})$ and hence $J(1+\frac{J}{\delta^2\sqrt{n}}) \lesssim \delta^2\sqrt{n}$ by Assumption 1.

Finally, we bound $\mathbb{E}^* \sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} |\sqrt{n}(\mathbb{P}_n - P_0) \log \frac{1-D}{1-D_{\theta}}|$. As in Lemma 7.1 (iv), we obtain $\rho^2 := \left\| \frac{1}{2} \log \frac{1-D}{1-D_{\theta}} \right\|_{P_0,B}^2 \leq 2(1+M)P_0W_2^2 = o(\delta)$. Therefore, by van der Vaart and Wellner [70, Lemma 3.4.3], $\mathbb{E}^* \sup_{D \in \mathcal{D}_{n,\delta_n}^{\theta}} |\sqrt{n}(\mathbb{P}_n - P_0) \log \frac{1-D}{1-D_{\theta}}| \leq J\left(1 + \frac{J}{\delta^2\sqrt{n}}\right)$ for $J = J_{[]}(\rho, \{\log \frac{1-D}{1-D_{\theta}} : D \in \mathcal{D}_{n,\delta_n}^{\theta}\}, \|\cdot\|_{P_0,B})$. With a δ -bracket in $\mathcal{D}_{n,\delta_n}^{\theta}$, Assumption 2 implies $\left\|\log \frac{1-\ell}{1-D_{\theta}} - \log \frac{1-u}{1-D_{\theta}}\right\|_{P_0,B}^2 \leq 4P_0\left(\sqrt{(1-\ell)/(1-u)} - 1\right)^2 = o(\delta)$. Therefore, the expectation of the supremum is of order $o(\delta\sqrt{n})$.

7.2 Proof of Theorem 4.2

Let h_n be a bounded sequence and $\theta_n := \theta_0 + \frac{h_n}{\sqrt{n}}$ and $W_n := \sqrt{\hat{p}_{\theta_n}/\hat{p}_{\theta_0}} - 1$. Since $\log(1+x) = x - \frac{1}{2}x^2 + \frac{1}{2}x^2R(x)$ for R(x) = O(x), $n\mathbb{P}_n \log \frac{\hat{p}_{\theta_n}}{\hat{p}_{\theta_0}} = 2n\mathbb{P}_n W_n - n\mathbb{P}_n W_n^2 + n\mathbb{P}_n W_n^2 R(W_n)$. By Assumption 4 (ii) and $P_{\theta_0}\dot{\ell}_{\theta_0} = 0$, $2n\mathbb{P}_n W_n - n\mathbb{P}_n W_n^2 = 2nP_{\theta_0} W_n + \sqrt{n}\mathbb{P}_n h'_n \dot{\ell}_{\theta_0} - nP_{\theta_0} W_n^2 + \frac{1}{2}nP_{\theta_0} W_n + \sqrt{n}P_n h'_n \dot{\ell}_{\theta_0} - nP_{\theta_0} W_n^2 + \frac{1}{2}nP_{\theta_0} W_n + \sqrt{n}P_n h'_n \dot{\ell}_{\theta_0} - nP_{\theta_0} W_n^2 + \frac{1}{2}nP_n W_n + \frac{1}{2}nP_n W_n + \frac{1}{2}nP_n W_n^2 + \frac{1}{2}nP_n W_n + \frac{1}{2}$ $o_P(1)$. Observe that $nP_{\theta_0}W_n^2 = \frac{1}{4}h'_n I_{\theta_0}h_n + o_P(1)$ and

$$2nP_{\theta_0}W_n = 2n\hat{P}_{\theta_0}W_n + 2n(P_{\theta_0} - \hat{P}_{\theta_0})W_n$$

= $-n\int \left(\sqrt{\hat{p}_{\theta_n}} - \sqrt{\hat{p}_{\theta_0}}\right)^2 + n(c_{\theta_n} - c_{\theta_0}) - \sqrt{n}\hat{P}_{\theta_0}h'_n\dot{\ell}_{\theta_0}$
 $+ 2n\int \left(\sqrt{p_{\theta_0}} - \sqrt{\hat{p}_{\theta_0}}\right) \left(\sqrt{p_{\theta_0}} + \sqrt{\hat{p}_{\theta_0}}\right) \left(W_n - \frac{h'_n\dot{\ell}_{\theta_0}}{2\sqrt{n}}\right).$

By Assumption 4 (i) and $\int (\sqrt{\hat{p}_{\theta_0}} - \sqrt{p_{\theta_0}})^2 = O_P(\delta_n^2), n \int (\sqrt{\hat{p}_{\theta_n}} - \sqrt{\hat{p}_{\theta_0}})^2 = \frac{1}{4}h'_n I_{\theta_0} h_n + o_P(1).$ By the Cauchy-Schwarz inequality,

$$\begin{split} \left| \int \left(\sqrt{p_{\theta_0}} - \sqrt{\hat{p}_{\theta_0}} \right) \left(\sqrt{p_{\theta_0}} + \sqrt{\hat{p}_{\theta_0}} \right) \left(W_n - \frac{h'_n \dot{\ell}_{\theta_0}}{2\sqrt{n}} \right) \right| \\ & \leq \left[\int \left(\sqrt{p_{\theta_0}} - \sqrt{\hat{p}_{\theta_0}} \right)^2 \int \left(\sqrt{p_{\theta_0}} + \sqrt{\hat{p}_{\theta_0}} \right)^2 \left(W_n - \frac{h'_n \dot{\ell}_{\theta_0}}{2\sqrt{n}} \right)^2 \right]^{1/2}, \end{split}$$

which is $O_P(\delta_n n^{-3/4}) = o_P(n^{-1})$ under Assumption 4 (i) and $\delta_n = o(n^{-1/4})$.

Since $|n\mathbb{P}_n W_n^2 R(W_n)| \leq |n\mathbb{P}_n W_n^2| \max_{1 \leq i \leq n} |R(W_n(X_i))|$, it remains to show that the maximum is $o_P(1)$. Write $V_n := W_n - \frac{h'_n \ell_{\theta_0}}{2\sqrt{n}}$. Then,

$$\max_{i} |W_n(X_i)| \le \max_{i} \left| \frac{1}{2\sqrt{n}} h'_n \dot{\ell}_{\theta_0}(X_i) \right| + \max_{i} |V_n(X_i)|.$$

By Markov's inequality,

$$P\left(\max_{1\leq i\leq n} \left| \frac{1}{\sqrt{n}} h'_n \dot{\ell}_{\theta_0}(X_i) \right| > \varepsilon\right) \leq nP\left(\left| \frac{1}{\sqrt{n}} h'_n \dot{\ell}_{\theta_0}(X_i) \right| > \varepsilon \right)$$
$$\leq \varepsilon^{-2} P_{\theta_0}((h'_n \dot{\ell}_{\theta_0})^2 \mathbb{1}\{(h'_n \dot{\ell}_{\theta_0})^2 > n\varepsilon^2\}),$$

which converges to zero as $n \to \infty$ for every $\varepsilon > 0$. Thus, $\max_i \left| \frac{1}{\sqrt{n}} h'_n \dot{\ell}_{\theta_0}(X_i) \right|$ converges to zero in probability. Since Assumption 4 (ii) and (i) imply that $n \mathbb{P}_n V_n^2 = n P_{\theta_0} V_n^2 + o_P(1) = o_P(1)$, we have $\max_i V_n^2(X_i) = o_P(1)$ and hence $\max_i |V_n(X_i)| = o_P(1)$. Conclude that $\max_i |W_n(X_i)|$ converges to zero in probability and so does $\max_i |R(W_n(X_i))|$.

7.3 Proof of Theorem 4.4

We will prove Theorem 4.4 under weaker assumptions. In particular, we slightly relax Assumption 4.4 by considering the aggregate behavior of $u_{\theta}(X^{(n)})$ around θ_0 with respect to the prior $\Pi_n(\cdot)$. Instead, we assume

$$P_{\theta_0}^{(n)}\left(I_n(\Pi_n, X^{(n)}, \varepsilon_n) \le e^{-\widetilde{C}_n n \varepsilon_n^2}\right) = o(1)$$

where

$$I_n(\Pi_n, X^{(n)}, \epsilon) = \int_{B_n(\theta_0, \epsilon)} e^{u_\theta(X^{(n)})} d\Pi_n(\theta)$$
(7.1)

and, at the same time,

$$P_{\theta_0}^{(n)} \left[\sup_{\substack{\Theta_n^c \cup d_n(\theta, \theta_0) > \epsilon}} |u_{\theta}(X^{(n)})| > \tilde{C}_n n \varepsilon_n^2 \right] = o(1)$$

for any $\epsilon > \varepsilon_n$. Assumption (4.5) is not needed if one is only interested in the concentration inside Θ_n . Alternatively, we could also replace Assumption (4.4) with the following condition to lower-bound the denominator in (4.3)

$$\sup_{\theta \in B_n(\theta_0,\varepsilon_n)} P_{\theta_0}^{(n)} \left[\ln(p_{\theta}^{(n)}/p_{\theta_0}^{(n)}) + u_{\theta} < -n\varepsilon_n^2 \right] = o(n\varepsilon_n^2).$$

Instead of relying on the existence of exponential tests (through Lemma 9 in [29]), we could then directly assume that for any $\epsilon > \varepsilon_n$ and for all $\theta \in \Theta_n$ such that $d(\theta, \theta_0) > j\epsilon$ for any $j \in \mathbb{N}$ there exists a test $\phi_n(\theta)$ satisfying

$$P_{\theta_0}^{(n)}\phi_n \lesssim e^{-n\epsilon^2/2}$$
 and $\int_{\mathcal{X}} (1-\phi_n) p_{\theta}^{(n)} e^{u_{\theta}} \leq e^{-j^2 n\epsilon^2/2}.$

We will use the following Lemma (an analogue of Lemma 10 [29]).

Lemma 7.2. Recall the definition $I_n(\Pi_n, X^{(n)}, \epsilon)$ in (7.1) and define $q_{\theta}^{(n)} = p_{\theta}^{(n)} / p_{\theta_0}^{(n)} e^{u_{\theta}}$. Then we have for any $C, \varepsilon > 0$

$$P_{\theta_0}^{(n)}\left(\int_{B(\theta_0,\varepsilon)} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) \le \mathrm{e}^{-(1+C)n\varepsilon^2} \times I_n(\Pi_n, X^{(n)}, \epsilon)\right) \le \frac{1}{C^2 n\varepsilon^2}$$

Proof. Define a changed prior measure $\Pi_n^{\star}(\cdot)$ through $d\Pi_n^{\star}(\theta) = \frac{e^{u_{\theta}(X^{(n)})}}{\int e^{u_{\theta}(X^{(n)})} d\theta} d\Pi_n(\theta)$. Lemma 10 of [29] then yields

$$P_{\theta_0}^{(n)} \left(\int_{B(\theta_0,\varepsilon)} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) \le \mathrm{e}^{-(1+C)n\varepsilon^2} I_n(\Pi_n, X^{(n)}, \epsilon) \right)$$
$$= P_{\theta_0}^{(n)} \left(\int_{B(\theta_0,\varepsilon)} p_{\theta}^{(n)} / p_{\theta_0}^{(n)} \mathrm{d}\Pi_n^{\star}(\theta) \le \Pi_n^{\star}(B(\theta_0,\varepsilon)) \mathrm{e}^{-(1+C)n\varepsilon^2} \right) \le \frac{1}{C^2 n\epsilon^2}.$$

Recall the definition $I_n(\Pi_n, X^{(n)}, \varepsilon_n) = \int_{B(\theta_0, \varepsilon_n)} e^{u_\theta(X^{(n)})} d\Pi_n(\theta)$ and define an event

$$\mathcal{A}_n = \left\{ X^{(n)} : \int_{B(\theta_0,\varepsilon_n)} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) > \mathrm{e}^{-2n\varepsilon_n^2} I_n(\Pi_n, X^{(n)}, \varepsilon_n) \right\}$$

where $q_{\theta}^{(n)} = p_{\theta}^{(n)} / p_{\theta_0}^{(n)} e^{u_{\theta}}$. From our assumptions, there exists a sequence $\tilde{C}_n > 0$ such that the complement of the set

$$\mathcal{B}_n = \left\{ X^{(n)} : I_n(\Pi_n, X^{(n)}, \varepsilon_n) > e^{-\widetilde{C}_n n \varepsilon_n^2} \text{ and } \sup_{\Theta_n^c \cup d_n(\theta, \theta_0) > \varepsilon_n} |u_\theta(X^{(n)})| \le \widetilde{C}_n n \varepsilon_n^2 \right\}$$

has a vanishing probability. Lemma 7.2 then yields $P_{\theta_0}^{(n)}[\mathcal{A}_n^c \cup \mathcal{B}_n^c] = o(1)$ as $n \to \infty$. The following calculations are thus conditional on the set $\mathcal{A}_n \cap \mathcal{B}_n$. On this set, we can lower-bound the denominator of (4.3) as follows

$$\int_{\Theta} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) > \int_{B(\theta_0,\varepsilon_n)} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) > \mathrm{e}^{-2n\varepsilon_n^2} I_n(\Pi_n, X^{(n)}, \varepsilon_n) \ge \mathrm{e}^{-(2+\widetilde{C}_n)n\varepsilon_n^2}$$

We first show that $P_{\theta_0}^{(n)}[\Pi_n^{\star}(\Theta \setminus \Theta_n \mid X^{(n)})] = o(1)$ as $n \to \infty$. On the set $\mathcal{A}_n \cap \mathcal{B}_n$ we have from (4.5) and from the Fubini's theorem

$$P_{\theta_0}^{(n)} \left[\Pi_n^{\star}(\Theta \setminus \Theta_n \mid X^{(n)}) \right] = P_{\theta_0}^{(n)} \left[\frac{\int_{\Theta \setminus \Theta_n} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta)}{\int_{\Theta} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta)} \right] \le \mathrm{e}^{2n\varepsilon_n^2} \frac{\Pi_n^{\star}(\Theta \setminus \Theta_n)}{\Pi_n^{\star}(B_n(\theta_0, \varepsilon_n))}$$
$$= \mathrm{e}^{2(1+\widetilde{C}_n)n\varepsilon_n^2} \frac{\Pi_n(\Theta \setminus \Theta_n)}{\Pi_n(B_n(\theta_0, \varepsilon_n))} = o(1).$$

For some J > 0 (to be determined later) we define the complement of the ball around the truth as a union of shells

$$U_n = \{\theta \in \Theta_n : d_n(\theta, \theta_0) > MJ\varepsilon_n\} = \bigcup_{j \ge J} \Theta_{n,j}$$

where each shell equals

$$\Theta_{n,j} = \{ \theta \in \Theta_n : M j \varepsilon_n < d_n(\theta, \theta_0) \le M (j+1) \varepsilon_n \}.$$

We now invoke the local entropy Assumption (3.2) in [29] which guarantees (according to Lemma 9 in [29]) that there exist tests ϕ_n (for each n) such that

$$P_{\theta_0}^{(n)}\phi_n \lesssim e^{n\varepsilon_n^2 - nM^2\varepsilon_n/2} \quad \text{and} \quad P_{\theta}^{(n)}(1 - \phi_n) \le e^{-nM^2\varepsilon_n^2 j^2/2}$$
(7.2)

for all $\theta \in \Theta_n$ such that $d_n(\theta, \theta_0) > M \varepsilon_n j$ and for every $j \in \mathbb{N} \setminus \{0\}$ and M > 0. One can then write

$$P_{\theta_{0}}^{(n)}\Pi\left(\theta \in \Theta : d(\theta,\theta_{0}) > MJ\varepsilon_{n} \mid X^{(n)}\right) \leq P_{\theta_{0}}^{(n)}\Pi(\Theta_{n}^{c} \mid X^{(n)}) + P_{\theta_{0}}^{(n)}\phi_{n} + P_{\theta_{0}}^{(n)}(\mathcal{A}_{n}^{c}) + P_{\theta_{0}}^{(n)}(\mathcal{B}_{n}^{c}) + \sum_{j \geq J} P_{\theta_{0}}^{(n)}[\Pi(\Theta_{n,j} \mid X^{(n)})(1-\phi_{n})\mathbb{I}(\mathcal{A}_{n} \cap \mathcal{B}_{n})]$$

For the last term above, we recall that $\Pi(\Theta_{n,j} | X^{(n)}) = \frac{\int_{\Theta_{n,j}} q_{\theta}^{(n)} d\Pi_n(\theta)}{\int_{\Theta} q_{\theta}^{(n)} d\Pi_n(\theta)}$. We bound the denominator as before. Regarding the numerator, on the event \mathcal{B}_n we have from (7.2) and from the Fubini's theorem

$$P_{\theta_0}^{(n)} \int_{\Theta_{n,j}} q_{\theta}^{(n)} \mathrm{d}\Pi_n(\theta) (1 - \phi_n) \le \mathrm{e}^{-nM^2 \varepsilon_n^2 j^2 / 2 + \widetilde{C}_n n \varepsilon_n^2} \Pi_n(\Theta_{n,j})$$
(7.3)

Putting the pieces together, we obtain

$$P_{\theta_0}^{(n)}[\Pi(\Theta_{n,j} \mid X^{(n)})(1-\phi_n)\mathrm{I}(\mathcal{A}_n \cap \mathcal{B}_n)] \le \mathrm{e}^{-nM^2\varepsilon_n^2 j^2/2 + 2(1+\widetilde{C}_n)n\varepsilon_n^2} \frac{\Pi_n(\Theta_{n,j})}{\Pi_n[B_n(\theta_0,\varepsilon_n)]}$$

Assumption (3.4) of [29] writes as

$$\frac{\Pi_n(\Theta_{n,j})}{\Pi_n[B_n(\theta_0,\varepsilon_n)]} \le e^{nM^2\varepsilon_n^2j^2/4}$$
(7.4)

which yields

$$P_{\theta_0}^{(n)}\Pi\left(\theta\in\Theta: d(\theta,\theta_0) > MJ\varepsilon_n \,|\, X^{(n)}\right) \le o(1) + \sum_{j\ge J} e^{-n\varepsilon_n^2(M^2j^2/4 - 2 - 2\widetilde{C}_n)}.$$

The right hand side converges to zero as long as $J = J_n \to \infty$ fast enough so that $\tilde{C}_n = o(J_n)$ and $n\varepsilon_n^2$ is bounded away from zero.

7.4 Proof of Theorem 4.5

We define the event

$$\mathcal{A} = \left\{ X^{(n)} \in \mathcal{X} : \int \frac{\widetilde{p}_{\theta}^{(n)}}{\widetilde{p}_{\theta^*}^{(n)}} \mathrm{d}\widetilde{\Pi}_n(\theta) > \mathrm{e}^{-(1+C)n\epsilon^2} \widetilde{\Pi}_n[B(\epsilon, \widetilde{P}_{\theta^*}^{(n)}, P_{\theta_0}^{(n)})] \right\}.$$

The following lemma shows that $P_{\theta_0}^{(n)}[\mathcal{A}^c] = o(1)$ as $n \to \infty$.

Lemma 7.3. For $k \ge 2$, every $\epsilon > 0$ and a prior measure $\widetilde{\Pi}_n(\theta)$ on Θ , we have for every C > 0

$$P_{\theta_0}^{(n)}\left(\int \frac{\widetilde{p}_{\theta}^{(n)}}{\widetilde{p}_{\theta^*}^{(n)}} \mathrm{d}\widetilde{\Pi}_n(\theta) \le \mathrm{e}^{-(1+C)n\epsilon^2} \widetilde{\Pi}_n[B(\epsilon, \widetilde{P}_{\theta^*}^{(n)}, P_{\theta_0}^{(n)})]\right) \le \frac{1}{C^2 n\epsilon^2}$$

Proof. This follows directly from Lemma 10 in [29].

We now define $U_n(\epsilon) = \prod_n (\theta \in \Theta : d(\tilde{P}_{\theta}^{(n)}, \tilde{P}_{\theta^*}^{(n)}) > \epsilon \mid X^{(n)})$. For every $n \ge 1$ and $J \in \mathbb{N} \setminus \{0\}$, we can decompose

$$P_{\theta_0}^{(n)}U_n(JM\varepsilon_n) = P_{\theta_0}^{(n)}[U_n(JM\varepsilon_n)\phi_n] + P_{\theta_0}^{(n)}[U_n(JM\varepsilon_n)(1-\phi_n)\mathbb{I}(\mathcal{A}^c)] + P_{\theta_0}^{(n)}[U_n(JM\varepsilon_n)(1-\phi_n)\mathbb{I}(\mathcal{A})].$$

The first term is bounded (from the assumption (4.9)) as

$$P_{\theta_0}^{(n)}[U_n(JM\varepsilon_n)\phi_n] \le P_{\theta_0}^{(n)}\phi_n \lesssim e^{-n\varepsilon_n^2 J^2 M^2}$$

The second term can be bounded by $P_{\theta_0}^{(n)}[\mathbb{I}(\mathcal{A}^c)] \leq \frac{1}{C^2 J^2 M^2 n \varepsilon_n^2}$ which converges to zero as $n \varepsilon_n^2 \to \infty$. The last term satisfies

$$P_{\theta_{0}}^{(n)}[U_{n}(JM\varepsilon_{n})(1-\phi_{n})\mathbb{I}(\mathcal{A})] = P_{\theta_{0}}^{(n)}\left[(1-\phi_{n})\mathbb{I}(\mathcal{A})\frac{\int_{\theta:d(\widetilde{P}_{\theta}^{(n)},\widetilde{P}_{\theta^{*}}^{(n)})>JM\varepsilon_{n}}{\widetilde{p}_{\theta^{*}}^{(n)}}\frac{\widetilde{p}_{\theta^{*}}^{(n)}}{\widetilde{p}_{\theta^{*}}^{(n)}}\widetilde{\Pi}_{n}(\theta)\mathrm{d}\theta}\right]$$

$$\leq \frac{\mathrm{e}^{(1+C)n\epsilon^{2}}}{\widetilde{\Pi}_{n}[B(\epsilon,\widetilde{P}_{\theta^{*}}^{(n)},P_{\theta_{0}}^{(n)})]}\int_{\theta:d(\widetilde{P}_{\theta}^{(n)},\widetilde{P}_{\theta^{*}}^{(n)})>JM\varepsilon_{n}}\left[\int_{\mathcal{X}}(1-\phi_{n})p_{\theta_{0}}^{(n)}\frac{\widetilde{p}_{\theta}^{(n)}}{\widetilde{p}_{\theta^{*}}^{(n)}}\right]\widetilde{\Pi}_{n}(\theta)\mathrm{d}\theta$$

$$\leq \frac{\mathrm{e}^{(1+C)n\epsilon^{2}}}{\widetilde{\Pi}_{n}[B(\epsilon,\widetilde{P}_{\theta^{*}}^{(n)},P_{\theta_{0}}^{(n)})]}\sum_{j\geq J}\int_{U_{n,j}}Q_{\theta}^{(n)}(1-\phi_{n})\mathrm{d}\widetilde{\Pi}_{n}(\theta),$$

where $U_{n,j} = \{\theta : jM\varepsilon_n < d(\tilde{P}_{\theta}^{(n)}, \tilde{P}_{\theta^*}^{(n)}) \le (j+1)M\varepsilon_n)\}$. The tests (from the assumption (4.9)) satisfy $Q_{\theta}^{(n)}(1-\phi_n) \le e^{-nj^2M^2\varepsilon_n^2/4}$ uniformly on $U_{n,j}$. Then we find (using the assumption (4.10))

$$P_{\theta_0}^{(n)}[U_n(JM\varepsilon_n)(1-\phi_n)\mathbb{I}(\mathcal{A})] \le e^{(1+C)n\varepsilon_n^2} \sum_{j\ge J} e^{-nj^2M^2\varepsilon_n^2/4+nj^2M^2\varepsilon_n^2/8}.$$

The sum converges to zero when $n\varepsilon_n^2$ is bounded away from zero and $J \to \infty$.

7.5 Normal Location-Scale Example

Let $X_i \sim P_0 = N(0,1)$ and $P_{\theta} = N(\mu, \sigma^2)$ where $\theta = (\mu, \sigma^2)$ are the unknown parameters and $\theta_0 = (0,1)$ are the true values. This model satisfies Assumption 3 with the score $\dot{\ell}_{\theta_0}(x) = \begin{bmatrix} x \\ (x^2-1)/2 \end{bmatrix}$ and the Fisher information matrix $I_{\theta_0} = \begin{bmatrix} 1 & 0 \\ 0 & 1/2 \end{bmatrix}$. The oracle discriminator of P_0 from P_{θ} is $D_{\theta}(x) = \begin{bmatrix} 1 + \exp(-\frac{1}{2}\log\sigma^2 + \frac{x^2}{2} - \frac{(x-\mu)^2}{2\sigma^2}) \end{bmatrix}^{-1}$. Let us use the logistic regression using regressors $(1, x, x^2)$ to estimate D_{θ} , i.e.,

$$D_{\theta}(x) = [1 + \exp(-\beta_0 - \beta_1 x - \beta_2 x^2)]^{-1}.$$

Thus, the true parameter for the logistic regression is $\beta = (\beta_0, \beta_1, \beta_2) = (\frac{1}{2} \log \sigma^2 + \frac{\mu^2}{2\sigma^2}, -\frac{\mu}{\sigma^2}, \frac{1}{2\sigma^2} - \frac{1}{2})$. Let $\hat{\beta} = (\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2)$ be the estimator of β . Then,

$$\hat{p}_{\theta}(x) = \frac{\exp\left(-\frac{x^2}{2} - \hat{\beta}_0 - \hat{\beta}_1 x - \hat{\beta}_2 x^2\right)}{\sqrt{2\pi}} \quad \text{and} \quad c_{\theta} = \frac{\exp\left(-\hat{\beta}_0 + \frac{1}{2}\frac{\hat{\beta}_1^2}{1 + 2\hat{\beta}_2}\right)}{\sqrt{1 + 2\hat{\beta}_2}}.$$

Being a MLE, $\hat{\beta}$ is regular and efficient, so $\sqrt{n}(\hat{\beta} - \beta) = \Delta + o_P(1)$ for a normal vector Δ . Moreover, if we generate X_i^{θ} through $X_i^{\theta} = \mu + \sigma \widetilde{X}_i$, $\widetilde{X}_i \sim N(0, 1)$, there is one-to-one correspondence between $X_i^{\theta_1}$ and $X_i^{\theta_2}$ for every θ_1 and θ_2 , so the dependence of Δ on θ disappears as $n \to \infty$ for otherwise a more efficient estimator exists to contradict efficiency. Therefore, the formula for \hat{p}_{θ} implies that Assumption 4 (i) is satisfied with the oracle score function $\dot{\ell}_{\theta_0}$; since \hat{p}_{θ} is twice differentiable, it holds with a faster rate of $O_P(||h||^4)$. Finally, we check Assumption 4 (ii) and (iii) by simulation. Figure 10 shows the supremands of Assumption 4 (ii) and (iii) as functions of $\theta = (\mu, \sigma^2)$. The black lines plot $n(c_{\theta} - c_{\theta_0})$ as we change θ ; they are linear and its quadratic curvatures are ignorable. The blue lines represent $n(\mathbb{P}_n - P_{\theta_0}) \left(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1 - (\theta - \theta_0)'\dot{\ell}_{\theta_0}/2\right)$ and the red lines $n(\mathbb{P}_n - P_{\theta_0}) \left(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1\right)^2$; compared to the values of $n(c_{\theta} - c_{\theta_0})$, both are uniformly ignorable.

Since this model with the logistic classifier satisfies Assumptions 3 and 4, it is susceptible to Theorem 4.2. This is supported by a diagnostics plot in Figure 11 which portrays true and estimated likelihood ratios. In Figure 11a, μ is varied with σ^2 fixed at σ_0^2 while, in Figure 11b, σ^2 is varied with μ held at μ_0 . The difference between the estimated log likelihood (blue) and the quadratic approximation (dashed red) is negligible, demonstrating that the validity of Theorem 4.2 is justifiable. Compared to the oracle log likelihood (black),





(a) The black line $n(c_{\theta} - c_{\theta_0})$; the blue line $n(\mathbb{P}_n - P_{\theta_0})(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1 - (\theta - \theta_0)'\dot{\ell}_{\theta_0}/2)$; the red line $n(\mathbb{P}_n - P_{\theta_0})(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1)^2$. σ^2 is fixed at σ_0^2 .

(b) The black line $n(c_{\theta} - c_{\theta_0})$; the blue line $n(\mathbb{P}_n - P_{\theta_0})(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1 - (\theta - \theta_0)'\dot{\ell}_{\theta_0}/2)$; the red line $n(\mathbb{P}_n - P_{\theta_0})(\sqrt{\hat{p}_{\theta}/\hat{p}_{\theta_0}} - 1)^2$. μ is fixed at μ_0 .

Figure 10: Illustration of Assumption 4 (ii–iii) in the normal location-scale example with n = m = 5000.

the estimated log likelihood is shifted by the random term $\sqrt{n}(\dot{c}_{n,\theta_0} - \hat{P}_{\theta_0}\dot{\ell}_{\theta_0})$. The curvature, however, is the same as oracle since the red line curves by the Fisher information I_{θ_0} . Thus, we expect Algorithm 1 to produce a biased sample and Algorithm 2 a dispersed sample. Note that we can compute $\sqrt{n}\hat{P}_{\theta_0}\dot{\ell}_{\theta_0} = c_{\theta_0}\sqrt{n}\left[-\frac{\hat{\beta}_1}{1+2\hat{\beta}_2}, -\frac{1}{2} + \frac{1}{2(1+2\hat{\beta}_2)} + \frac{\hat{\beta}_1^2}{2(1+2\hat{\beta}_2)^2}\right]'$, which is asymptotically linear in Δ by the delta method. It is then reasonable to expect that this term has mean zero when averaged over \widetilde{X} since $\hat{\beta}$ is asymptotically unbiased. If \dot{c}_{n,θ_0} also has mean zero, then Algorithm 2 is unbiased and Algorithm 3 recovers the exact normal posterior.

To see that this is indeed the case, we impose a conjugate normal-inverse-gamma prior, $\theta \sim N\Gamma^{-1}(\mu_0, \nu, \alpha, \beta)$, that is, the marginal prior of σ^2 is the inverse-gamma $\Gamma^{-1}(\alpha, \beta)$ and the conditional prior of μ given σ^2 is $N(\mu_0, \frac{\sigma^2}{\nu})$. The posterior is then analytically calculated as (for $\bar{X}_n = \frac{1}{n} \sum_i X_i$)

$$\theta \mid X \sim N\Gamma^{-1} \Big(\frac{\nu\mu_0 + n\bar{X}_n}{\nu + n}, \nu + n, \alpha + \frac{n}{2}, \beta + \frac{1}{2} \sum_i (X_i - \bar{X}_n)^2 + \frac{n\nu}{\nu + n} \frac{(\bar{X}_n - \mu_0)^2}{2} \Big).$$

Figure 12 shows the histograms of Algorithm 1, 2 and 3 after K = 500 MCMC steps. Since the estimated log likelihood has a rightward bias (as seen from Figure 11), Algorithm 1 produces a sample that is shifted to the right (Figures 12a and 12c). Algorithm 2, on the



(a) True log likelihood, estimated log likelihood, and quadratic approximation by Theorem 4.2. $\sigma^2 = \sigma_0^2$.



(b) True log likelihood, estimated log likelihood, and quadratic approximation by Theorem 4.2. $\mu = \mu_0$.

Figure 11: Illustration of Theorem 4.2 in the normal mean-scale example with n = m = 5000.

other hand, gives a sample that is more dispersed than the posterior but is correctly placed, indicating that the random bias has mean zero. Consequently, Algorithm 3 generates a sample that is placed and shaped correctly (Figures 12b and 12d).

7.6 BvM Conditions

Below, we provide sufficient conditions for the LAN assumption 4.13, relaxing slightly Lemma 2.1 in [41]. The assumptions in Lemma 7.4 are closely related to the ones in Theorem 4.2. The main difference is that Lemma 7.4 is concerned with the behavior of the (misspecified) likelihood around θ^* as opposed to θ^0 .

Lemma 7.4. Assume that $P_{\theta_0}^{(n)} = P_{\theta_0}^n$ with a density $\prod_{i=1}^n p_{\theta_0}(x_i)$ where the function $\theta \to \log p_{\theta}(x)$ is differentiable at θ^* with a derivative $\dot{\ell}_{\theta}$. Assume there exists an open neighborhood U of θ^* such that $\left|\log \frac{p_{\theta_1}(x)}{p_{\theta_2}(x)}\right| \le m_{\theta^*} \|\theta_1 - \theta_2\| P_{\theta_0} - a.s. \forall \theta_1, \theta_2 \in U$ where m_{θ} is a square integrable function. Assume that the log-likelihood has a 2nd order Taylor expansion around θ^* (i.e. (7.7) holds). Assume that u_{θ} is asymptotically linear around θ^* (i.e. (7.8) holds), then (4.13) holds with $\varepsilon_n^* = 1/\sqrt{n}$ and

$$\widetilde{V}_{\theta} = V_{\theta} \quad and \quad \widetilde{\Delta}_{n,\theta} = V_{\theta}^{-1} \left[\frac{\dot{C}_{\theta}}{\sqrt{n}} + \sqrt{n} \mathbb{P}_n \dot{\ell}_{\theta} + \frac{u^*(X^{(n)})}{\sqrt{n}} \right]$$
(7.5)



Figure 12: Histograms of the MHC samples of μ and σ^2 in the normal location-scale model. Algorithm 1 (resp. 2) yield more biased (resp. dispersed) samples compared to the true posterior (black curve). Algorithm 3 (on the right) tracks the black curve more closely.

Proof. We can write

$$\log \frac{\widetilde{p}_{\theta^*+\varepsilon_n h}^{(n)}}{\widetilde{p}_{\theta^*}^{(n)}} = \log \frac{C_{\theta^*+\varepsilon_n h}}{C_{\theta^*}} + \log \frac{p_{\theta^*+\varepsilon_n h}^{(n)}}{p_{\theta^*}^{(n)}} + u_{\theta^*+\varepsilon_n h} - u_{\theta^*}.$$
(7.6)

This yields, from Lemma 19.31 in [69], that

$$\mathbb{G}_n\left(\sqrt{n}\log\frac{p_{\theta^*+h/\sqrt{n}}}{p_{\theta^*}^*}-h'\dot{\ell}_{\theta^*}\right)\to 0 \quad \text{in } P_0,$$

where $\mathbb{G}_n = \sqrt{n}(\mathbb{P}_n - P_{\theta_0})$ is the empirical process. Assuming that

$$P_{\theta_0} \log\left(\frac{p_\theta}{p_{\theta^*}}\right) = P_{\theta_0} \dot{\ell}'_{\theta^*}(\theta - \theta^*) + \frac{1}{2}(\theta - \theta^*)' V_{\theta^*}(\theta - \theta^*) + o(\|\theta - \theta^*\|^2) \quad \text{as } \theta \to \theta^* \quad (7.7)$$

one obtains

$$\log \frac{p_{\theta^* + h/\sqrt{n}}^{(n)}}{p_{\theta^*}^{(n)}} = n \mathbb{P}_n \log \frac{p_{\theta^* + h/\sqrt{n}}}{p_{\theta^*}} = o_P(1) + \mathbb{G}_n h' \dot{\ell}_{\theta^*} + n P_{\theta_0} \log \frac{p_{\theta^* + h/\sqrt{n}}}{p_{\theta^*}}$$
$$= o_P(1) + \mathbb{G}_n h' \dot{\ell}_{\theta^*} + \frac{h'_n V_{\theta^*} h}{2} + \sqrt{n} P_{\theta_0} h' \dot{\ell}_{\theta}$$

If we assume asymptotic linearity of u_{θ} around θ^* , i.e.

$$u_{\theta^*+h/\sqrt{n}}(X^{(n)}) - u_{\theta^*}(X^{(n)}) = \frac{1}{\sqrt{n}}h'u^*(X^{(n)}) + o_P(1)$$
(7.8)

for some $u^{\star}(X^{(n)})$ and

$$\log \frac{C_{\theta^* + h_n/\sqrt{n}}}{C_{\theta^*}} = \frac{\dot{C}_{\theta^*}'h_n}{\sqrt{n}} + o(1)$$

then (4.13) holds with (7.5).

Related BvM conditions have been characterized in [8]. We restate these conditions utilizing the localized re-parametrization $h = \sqrt{n}(\theta - \theta_0) - s$, where $s = \sqrt{n}(\hat{\theta} - \theta_0)$ is a zero-mean vector where $\hat{\theta}$ is some suitable estimator. We first define a localized criterion function $\ell(h) \equiv \frac{\tilde{p}_{\hat{\theta}+h/\sqrt{n}}(X^{(n)})\tilde{\pi}(\hat{\theta}+h/\sqrt{n})}{\tilde{p}_{\hat{\theta}}(X^{(n)})\tilde{\pi}(\hat{\theta})}$, which corresponds to the normalized pseudoposterior $\pi^*(\theta | X^{(n)})/\pi^*(\hat{\theta} | X^{(n)})$. [8] impose a centered variant of (4.13) requiring that $\ell(h)$ approaches a quadratic form on a closed ball K (such that $\Lambda \equiv \sqrt{n}(\Theta - \theta_0) - s = K \cup K^c$) in the sense that

$$\left|\log \ell(h) - (-h'Jh)/2\right| \le \epsilon_1 + \epsilon_2 \times h'Jh/2 \quad \forall h \in K,$$
(7.9)

for some matrix J > 0 with eigenvalues bounded away from zero. If

$$\epsilon_1 = o(1)$$
 and $\epsilon_2 \times \lambda_{max}^2(J)(\sup_{h \in K} ||h||)^2 = o(1)$ in $P_{\theta_0}^{(n)}$ -probability. (7.10)

Theorem 1 of [8] shows that $\ell(h)/\int_{\Lambda}\ell(h)dh$ approaches the standard normal density in $P_{\theta_0}^{(n)}$ -probability as $n, d \to \infty$. The condition (7.9) (a) allows for mild deviations from smoothness and log-concavity, (b) involves also the prior (unlike (4.13)) but, (c) requires the existence of a \sqrt{n} -consistent estimator $\hat{\theta}$. Lemma 4.6 is more general, where the rate ε_n^* does not need to be $1/\sqrt{n}$ and where the posterior is allowed to have a non-vanishing bias. The requirement (7.10) imposes certain restrictions on $u_{\theta}(X^{(n)})$. For example, in the linear case (4.12) one would need $u^*(X^{(n)}) = o(\sqrt{n})$ in $P_{\theta_0}^{(n)}$ -probability from (7.10).

 $^{13}\int_{K}\ell(h)\mathrm{d}h/\int_{\Lambda}\ell(h)\mathrm{d}h \geq 1 - o_{P_{\theta_0}}(1)$ and $\int_{K}\phi(h)\mathrm{d}h$ for $\phi(\cdot)$ standard Gaussian density



Figure 13: Estimated log likelihood ratio for the Ricker model: (Left) function of log r fixing $\sigma = \sigma_0$ and $\varphi = \varphi_0$, (Middle) function of σ^2 fixing $r = r_0$ and $\varphi = \varphi_0$, (Right) function of φ fixing $\sigma = \sigma_0$ and $r = r_0$.

7.7 Example: Ricker Model

The Ricker model is a classic discrete model that describes partially observed population dynamics of fish and animals in ecology. The latent population $N_{i,t}$ follows

$$\log N_{i,t+1} = \log r + \log N_{i,t} - N_{i,t} + \sigma \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim N(0,1),$$

where r denotes the intrinsic growth rate and σ is the dispersion of innovations. The index t represents time and runs through 1 to T = 20. The index i represents independent observations and runs through 1 to n = 300. The initial population $N_{i,0}$ may be set as 1 or set randomly after some burn-in period. We observe $X_{i,t}$ such that

$$X_{i,t} \mid N_{i,t} \sim \text{Poisson}(\varphi N_{i,t}),$$

where φ is a scale parameter. The objective is to make inference on $\theta := (\log r, \sigma^2, \varphi)$. Each time sequence $X_i := (X_{i,1}, \ldots, X_{i,T})$ constitutes an observation, where *i* runs through *n*. In our notation, we can define the underlying data-generating process as $\widetilde{X}_{i,t} := (U_{i,t}, \varepsilon_{i,t})$ for $U_{i,t} \sim U[0, 1]$ and set the function T_{θ} to map ε_i to N_i and then (U_i, N_i) to X_i through the Poisson inverse transform sampling of $U_{i,t}$ into $X_{i,t}$. We set the true parameter as $(\log r_0, \sigma_0^2, \varphi_0) = (3.8, 1, 10)$ and employ an improper, flat prior. Note that our method can accommodate an improper prior, unlike ABC.

There is no obvious sufficient statistic for this model, and the likelihood is intractable due to the nontrivial time dependence of $N_{i,t}$. We use an average of neural network discrim-



Figure 14: MHC samples for the Ricker model.

inators to adapt to the unknown likelihood ratio. First, we estimate D_{θ} by a neural network with one hidden layer with 50 nodes, each of which is equipped with the hyperbolic tangent sigmoid activation function. Then, we compute the log likelihood of the data $\sum_i \log \frac{1-\hat{D}_{\theta}}{\hat{D}_{\theta}}$. We repeat this for 20 times with independently drawn \widetilde{X} and take the average of the log likelihood. This specification produces approximately quadratic likelihood-ratio curves (Figure 13). Unlike the location-scale normal model, the fixed design does not produce entirely smooth curves due to the averaging aspect over many discriminators. The quadratic shape is nevertheless recovered here, implying that the differentiability assumptions from Section 4.1 are not entirely objectionable.

Figure 14 shows the marginal histograms of the MHC samples (500 MCMC iterations). The proposal distribution is independent across parameters; $\log r$ uses the normal distribution, σ^2 the inverse-gamma distribution, and φ the gamma distribution; each of them has the mean equal to the previous draw and variance 1/n. The vertical dashed lines indicate the true parameter θ_0 . Note that the posterior is asymptotically centered at the MLE, not θ_0 . However, the blue histograms on the left (Algorithm 1) seem too far away from θ_0 relative to the widths of the histograms. On the other hand, the red histograms (Algorithm 1) are more dispersed but located closer to θ_0 . These observations confirm our theoretical findings. Histograms of Algorithm 3 (Figure 15) look reasonable as a posterior sample, center around the true values.

Figure 15 and 16 compare our method with the MCWM pseudo-marginal Metropolis-



Figure 15: MHC samples for the Ricker model (Algorithm 3)



Figure 16: Posterior samples for the Ricker model using the pseudo-marginal MCWM method

Hastings algorithm [3]. We have implemented the default pseudo-marginal method which deploys an average of conditional likelihoods for X_i , given N_i ,

$$\hat{p}(X_i) = \frac{1}{K} \sum_{k=1}^K \prod_{t=1}^T p(X_{i,t} \mid N_{i,t,k}) = \frac{1}{K} \sum_{k=1}^K \prod_{t=1}^T \frac{(\varphi N_{i,t,k})^{X_{i,t}} e^{-\varphi N_{i,t,k}}}{X_{i,t}!}$$

as the likelihood approximation, where K is some positive integer and where $N_{i,t,k}$ are independently drawn across k = 1, ..., K. In our comparisons, we let K = 20n. Figure 15 shows that the two methods produce posterior draws that are located at similar places, and the widths of the histograms are also comparable. We would like to point out, again, that our method does not require that a tractable conditional likelihood is available nor that a user-specified summary statistic is supplied.



(a) Estimated log likelihood for m = 1. The vertical dashed line indicates $\mu_0 = 0$.

(b) Estimated log likelihood for m = 2. The vertical dashed line indicates $\mu_0 = 0$.

Figure 17: Estimated log likelihood for models 1 and 2. The figures indicate that it is smooth in μ and have the same curvature as the true log likelihood.

7.8 Bayesian Model Selection: Further Details

Figure 17 shows true likelihood ratio and and classification-based estimates for fixed and random designs for the Bayesian model selection example from Section 5.3. Under the fixed design, the curve is smooth and slightly biased with a similar shape to the true log-likelihood. For the random design, there is no smoothness (due to the fake data refreshing aspect).

7.9 The CIR Model: Further Details

This section presents additional plots for the CIR analysis from Section ??. Figure 18 shows smoothed posterior samples for MHC (fixed generator) and $nrep \in \{1, 5\}$. These plots look qualitatively similar to the random generator results presented in Figure 8. Next, Figure 19 and 20 show trace-plots of the MHC samples. We can see that (1) using larger nrepreduces variance, (2) random generators have smaller acceptance rates for the same proposal distribution. Lastly, histograms of the posterior samples together with demarkations of the 95% credible intervals are in Figure 21 and 22.



Figure 18: Smoothed posterior densities obtained by simulation using the exact MH and MHC fixed generator using $nrep \in \{1, 5\}$



Figure 19: Trace-plots of 10 000 MHC iterations with nrep=1



Figure 20: Trace-plots of 10 000 MHC iterations with nrep = 5



Figure 21: Histogram of 9 000 MHC iterations (after 1 000 burnin) with nrep = 1



Figure 22: Histogram of 9 000 MHC iterations (after 1 000 burnin) with nrep = 5