

Approximate Bayesian Computation via Classification

Yuexi Wang

YUEXI.WANG@CHICAGOBOOTH.EDU

Tetsuya Kaji

TETSUYA.KAJI@CHICAGOBOOTH.EDU

Veronika Rockova

VERONIKA.ROCKOVA@CHICAGOBOOTH.EDU

*Booth School of Business
University of Chicago
Chicago, IL 60637, USA*

Editor: Sayan Mukherjee

Abstract

Approximate Bayesian Computation (ABC) enables statistical inference in simulator-based models whose likelihoods are difficult to calculate but easy to simulate from. ABC constructs a kernel-type approximation to the posterior distribution through an accept/reject mechanism which compares summary statistics of real and simulated data. To obviate the need for summary statistics, we directly compare empirical distributions with a Kullback-Leibler (KL) divergence estimator obtained via contrastive learning. In particular, we blend flexible machine learning classifiers within ABC to automate fake/real data comparisons. We consider the traditional accept/reject kernel as well as an exponential weighting scheme which does not require the ABC acceptance threshold. Our theoretical results show that the rate at which our ABC posterior distributions concentrate around the true parameter depends on the estimation error of the classifier. We derive limiting posterior shape results and find that, with a properly scaled exponential kernel, asymptotic normality holds. We demonstrate the usefulness of our approach on simulated examples as well as real data in the context of stock volatility estimation.

Keywords: Approximate Bayesian Computation, Classification, Likelihood-free Inference, Kullback-Leibler Divergence, Posterior Concentration

1. Introduction

We consider a collection of i.i.d. observations $\mathbf{X} = (X_1, \dots, X_n)'$ where each $X_i \in \mathcal{X}$ is realized from a parametric model $\{P_{\theta_0} : \theta_0 \in \Theta \subset \mathbb{R}^d\}$. We assume that P_{θ} , for each $\theta \in \Theta$, admits a density p_{θ} . We are interested in Bayesian inference about θ_0 based on the posterior distribution

$$\pi_n(\theta | \mathbf{X}) \propto p_{\theta}^{(n)}(\mathbf{X})\pi(\theta) \quad (1)$$

prescribed by the likelihood $p_{\theta}^{(n)}(\mathbf{X})$ and the prior density $\pi(\theta)$. We are particularly interested in simulator-based models whose the likelihood function cannot be directly expressed/evaluated (such as discretely observed diffusions (Sørensen, 2004) or generative models) but can be sampled from.

Simulator-based models are often called implicit models because the the likelihood function p_{θ} cannot be numerically evaluated (Diggle and Gratton, 1984). Fortunately, it may still be possible to simulate synthetic datasets from the model. The ability to simulate from

the likelihood has opened up new opportunities for simulating from the posterior. For example, Approximate Bayesian Computation (ABC) (Pritchard et al., 1999; Beaumont et al., 2002) emerged as a default likelihood-free Bayesian inferential tool. It is an accept/reject posterior sampling mechanism which obviates likelihood evaluations. Each iteration proceeds by (1) simulating prior parameter guesses and fake data from the likelihood, and then (2) accepting those parameter values whose fake data were close to the observed data. A big challenge with ABC has been gauging the similitude between observed and fake data.

Measures of similarity between data sets have traditionally been based on summary statistics (see Blum et al. (2013) for an overview within the ABC context). In other words, two datasets are considered similar if their summary statistics are close. In the absence of expert knowledge, however, constructing effective summary statistics can be challenging (Joyce and Marjoram, 2008; Nunes and Balding, 2010; Blum et al., 2013) and one may need to resort to automated strategies. One possibility is regressing parameter values onto (functionals of) fake data in a pilot ABC run to train a flexible mapping which can be substituted for summary statistics (Fearnhead and Prangle, 2012; Jiang et al., 2017; Akesson et al., 2021). Another possibility, related to indirect inference, is to construct summary statistics from an auxiliary model (Drovandi et al., 2011; Wood, 2010). One can also choose a subset of candidate summary statistics that satisfy some optimality criterion (Joyce and Marjoram, 2008; Nunes and Balding, 2010) or find an optimal projection of a set of summary statistics onto a lower-dimensional map (Fearnhead and Prangle, 2012). Alternatively, one can directly use a discrepancy between the empirical distributions of the observed and synthetic data sets inside ABC (such as Kullback-Leibler (Jiang et al., 2018) or Wasserstein (Bernton et al., 2019) or Maximum Mean (MM) discrepancy (Park et al., 2016)) or Energy Statistics (ES) (Nguyen et al., 2020)). See Drovandi and Frazier (2022) for a nice comprehensive review of the distribution-style ABC discrepancies. Our work fortifies this ABC point of view by focusing on the Kullback-Leibler discrepancy estimated via classification.

The KL divergence is one of the most widely used discrepancy metrics. It expresses the average information per observation to discriminate between two probabilistic models (Kullback, 1958). In large deviations, for example, it characterizes the exponential decay rate at which empirical measures converge to their probabilities (see Sanov’s theorem in Den Hollander (2008)) and the rate of decay of the probability of error in a binary hypothesis testing problem (see Stein’s Lemma in Cover and Thomas (1991)). KL also naturally connects to maximum likelihood estimation through its interpretation as the expectation of the log-likelihood ratio. There exist many methods for estimating the KL divergence. For example, Wang et al. (2009) proposed nearest-neighbor techniques to obtain a mean-square consistent estimator. Wang et al. (2005) proposed a histogram-based KL estimator based on partitioning of the space into statistically equivalent intervals. Silva and Narayanan (2007) and Silva and Narayanan (2010) went a step further and proposed using data-driven partitions (including multivariate recursive partitioning) and formulated sufficient consistency conditions. Alternatively, Nguyen et al. (2007) proposed a variational approach by turning KL estimation into a penalized convex risk minimization problem. Our work is different from the approaches above as we adopt a KL estimator based on classification.

We suggest embedding a machine learning classifier inside ABC to determine whether or not fake and observed data are similar and, thereby, whether or not the underlying parameter value should be kept in the ABC reference table. The fundamental premise of

this proposal is as follows: parameter values that yield indistinguishable simulated datasets can be deemed close. Bayesian inference via classification has been suggested before. Kaji and Ročková (2022) developed a version of the Metropolis-Hastings algorithm, called MHC, based on classification-based estimators of likelihood ratios. Thomas et al. (2022) derived a marginal approach by contrasting two fake datasets generated from the marginal and conditional likelihoods. Gutmann et al. (2018) proposed a classification strategy related to ours using a different discrepancy metric. Our paper reframes the method of Gutmann et al. (2018) as a genuine ABC algorithm with a KL divergence discriminator and provides supporting theory which justifies its inferential potential.

In particular, we study statistical properties of the approximate posterior which, in part, depend on the properties of the KL divergence estimator. We consider the traditional accept/reject ABC version (with a uniform kernel) as well as an exponential kernel variant which does not require the ABC tolerance threshold. Similar to Frazier et al. (2018), we show that the choice of the ABC acceptance threshold ϵ plays a critical role in the convergence rate and in the limiting posterior shape. In practice, it is often not obvious what the optimal threshold ϵ should be. Motivated by the connections with the MHC algorithm of Kaji and Rockova (2021), we propose an exponential kernel which yields ABC posteriors that correspond to the stationary distribution of MHC. Our ABC kernel method can be thus regarded as a parallelizable counterpart to the sequential MHC sampling, targeting the same posterior approximation. The concentration and asymptotic shape behavior of the ABC posterior, which can be derived from Kaji and Ročková (2022), theoretically justify our exponential weighting scheme. Finally, our classification-based ABC approach provides a viable computational strategy for obtaining coarsened posteriors for Bayesian robust inference (Miller and Dunson, 2018). Our ABC approach leverages machine learning but does so in a perhaps more traditional way than the recent sequential neural likelihood and mixture density network approaches for learning posteriors (Papamakarios and Murray, 2016; Papamakarios et al., 2019).

The remaining of the paper is structured as follows. In Section 2, we flesh out the basic idea of ABC and introduce our framework with classification. In Section 3, we investigate the posterior concentration and limiting shape behaviors of the ABC posteriors. Section 4 shows performance on simulated datasets and Section 5 further highlights the practical value of our approach on real data. In Section 6, we conclude with a discussion.

Notation. We use the shorthand notation $p_0 = p_{\theta_0}$ and $P_0 = P_{\theta_0}$. We employ the operator notation for expectation, e.g., $P_0 f = \int f dP_0$. The ϵ -bracketing number $N_{[]}(\epsilon, \mathcal{F}, d)$ 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 $J_{[]}(\delta, \mathcal{F}, d) = \int_0^\delta \sqrt{1 + \log N_{[]}(\epsilon, \mathcal{F}, d)} d\epsilon$. Next, $K(f, g) = \int f \log(f/g) d\mu$ denotes the Kullback-Leibler divergence between two density functions and $V_2(f, g) = \int f |\log(f/g)|^2 d\mu$. For real-valued sequences $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, $a_n \lesssim b_n$ means that $a_n \leq C b_n$ for some generic constant $C > 0$, $a_n \asymp b_n$ means that $a_n \lesssim b_n \lesssim a_n$, and $a_n \gg b_n$ indicates a greater order of magnitude. For a sequence of random variables x_n , $x_n = o_P(a_n)$ if $\lim_{n \rightarrow \infty} P(|x_n/a_n| \geq C) = 0$ for every $C > 0$, and $x_n = O_P(a_n)$ if for every $C > 0$ there exists a finite $M > 0$ and

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

a finite N such that $P(|x_n/a_n| \geq M) \leq C$ for all $n > N$. All limits are taken as $n \rightarrow \infty$. Take $\|\cdot\|$ to be the Euclidean norm.

2. ABC without Summary Statistics

The now default ABC method for Bayesian likelihood-free inference constructs a nested kernel-type approximation to the posterior distribution. The first approximation occurs when the data is distilled into summary statistics to obtain $\pi(\theta | S_X) \propto \pi(S_X | \theta)\pi(\theta)$, where $S_X = S(\mathbf{X})$ is a vector of summary statistics. The quality of this approximation depends crucially on the informativeness of S_X . The actual ABC approximation to the posterior (1) is then constructed via a kernel function as $\pi_{ABC}(\theta | S_X) = \int \pi(\theta, S | S_X) dS$ with $\pi(\theta, S) \propto T_\epsilon(\|S - S_Y\|)\pi(S | \theta)\pi(\theta)$, where $\|\cdot\|$ is a general norm to be specified later by the user and where $T_\epsilon(\|u\|) = T(\|u\|/\epsilon)$ is a standard (smoothing) kernel with a scale parameter $\epsilon > 0$. The key two challenges with ABC are (1) deriving low-dimensional summary statistics with a minimal loss of information and (2) selecting the kernel and its the tolerance level ϵ . To remediate the reliance of ABC on summary statistics, we focus on viewing the observed and fake data as empirical distributions and gauge the discrepancy between them. See Drovandi and Frazier (2022) for an overview of other discrepancy-based ABC methods. Regarding the choice of aggregation kernels, we consider the traditional uniform kernel yielding an accept/reject algorithm and a smoothing kernel free from ϵ -tuning.

We are interested in regimes where ABC is most effective (Jiang et al., 2017), i.e. settings where the sample size n of \mathbf{X} is moderately high and the dimension of Θ is low to ensure that we can hit the high ABC-posterior region with a reasonable prior probability. The number of observations n has an impact on the effectiveness of the embedded classifier. The nonparametric neural network classifiers usually demand n to be somewhat large so that the estimation errors are manageable.

2.1 ABC with KL Divergence

Instead of summary statistics, we use the estimated KL divergence inside the ABC algorithm. Our interest in the KL divergence as a discrepancy metric stems partially from the following connection to the generalized Bayesian inference (Bissiri et al., 2016). The posterior distribution (1) can be rewritten as a generalized posterior $\pi_n(\theta | \mathbf{X}) \propto \pi(\theta) \exp(-n \times KL(p_0^{(n)}, p_\theta^{(n)}))$ where the parameter θ is linked to data through the empirical Kullback-Liebler (KL) divergence $KL(p_0^{(n)}, p_\theta^{(n)}) \equiv \frac{1}{n} \sum_{i=1}^n \log(p_0/p_\theta)(X_i)$. For when the KL divergence cannot be easily evaluated, we consider various estimators in the next section. We denote a generic KL divergence estimator obtained from observed data $\mathbf{X} \sim P_0^{(n)}$ and pseudo-data $\tilde{\mathbf{X}}^\theta \sim P_\theta^{(n)}$ as $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)$. Adopting $\hat{K}(\cdot, \cdot)$ as the ABC discrepancy, we consider a simple accept/reject ABC mechanism detailed in Algorithm 1 below. While Jiang et al. (2017) used a nearest-neighbor estimator of the KL divergence, we devise a different estimator based on classification in Section 2.3.

Algorithm 1: KL-ABC with Accept-Reject

- For a pre-determined tolerance level $\epsilon > 0$ repeat for $j = 1, \dots, N$:
1. Simulate θ_j from $\pi(\theta)$.
 2. Simulate $\tilde{\mathbf{X}}^{\theta_j} = (\tilde{X}_1^{\theta_j}, \dots, \tilde{X}_m^{\theta_j})'$ through i.i.d. sampling from the model p_{θ_j} .
 3. Construct $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^{\theta_j})$ by training a classifier distinguishing \mathbf{X} and $\tilde{\mathbf{X}}^{\theta_j}$ as in (8).
 4. Accept θ_j when $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^{\theta_j}) \leq \epsilon$.

Algorithm 1 simulates pairs of parameter values and pseudo-data $\{\theta, \tilde{\mathbf{X}}^\theta\}$ from the joint posterior density

$$\hat{\pi}^{AR}(\theta, \tilde{\mathbf{X}}^\theta \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon) = \frac{\pi(\theta)p_\theta^{(n)}(\tilde{\mathbf{X}}^\theta)\mathbb{I}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon)}{\int \pi(\theta)p_\theta^{(n)}(\tilde{\mathbf{X}}^\theta)\mathbb{I}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon)d\tilde{\mathbf{X}}^\theta d\theta}, \quad (2)$$

which margins towards the following approximate (Accept/Reject) posterior density

$$\begin{aligned} \hat{\pi}_\epsilon^{AR}(\theta \mid \mathbf{X}) &= \int \hat{\pi}^{AR}(\theta, \tilde{\mathbf{X}}^\theta \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon) d\tilde{\mathbf{X}}^\theta \\ &\equiv \frac{\pi(\theta)P_\theta^{(n)}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon)}{\int \pi(\theta)P_\theta^{(n)}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon)d\theta}. \end{aligned} \quad (3)$$

The inferential potential of the approximation (3) will be scrutinized theoretically later in Section 3.2. In particular, we will later see that the convergence rate of (3) around θ_0 depends on the choice of ϵ (Frazier et al., 2018) as well as the quality of the discriminator. It is interesting to note that the ABC posterior (3) is mathematically equivalent to the c -posterior proposed by Miller and Dunson (2018) for robust inference in mis-specified (tractable) models. The computation of the c -posteriors has relied on powered-likelihood approximations and MCMC sampling. While we instead view (3) as an approximate posterior in models with intractable likelihoods, our ABC algorithm can be nevertheless used to compute coarsened posteriors in broader scenarios when MCMC sampling may not be available (see Remark 2.2 below). Algorithm 1 uses the uniform kernel which corresponds to the indicator function $T_\epsilon(\|u\|) = \mathbb{I}(\|u\| \leq \epsilon)$. In practice, it is difficult to balance out conflicting demands of smaller ϵ (yielding good approximability) and larger acceptance rates (yielding more posterior samples). As a remedy, we propose a way to aggregate the ABC samples through a scaled exponential kernel motivated by the connection between KL and the log-likelihood ratio. This ABC variant requires no ad-hoc thresholding and is summarized in Algorithm 2.

Algorithm 2: KL-ABC with Exponential Weighting

- Repeat for $j = 1, \dots, N$:
1. Simulate θ_j from $\pi(\theta)$.
 2. Simulate $\tilde{\mathbf{X}}^{\theta_j} = (\tilde{X}_1^{\theta_j}, \dots, \tilde{X}_m^{\theta_j})'$ through i.i.d. sampling from the model p_{θ_j} .
 3. Construct $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^{\theta_j})$ by training a classifier distinguishing \mathbf{X} and $\tilde{\mathbf{X}}^{\theta_j}$ as in (8).
 4. Assign θ_j a weight proportional to $\exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^{\theta_j}))$.

Algorithm 2 generates draws for the pair $\{\theta, \tilde{\mathbf{X}}^\theta\}$ from a joint posterior density

$$\hat{\pi}^{EK}(\theta, \tilde{\mathbf{X}}^\theta | \mathbf{X}) = \frac{\pi(\theta)p_\theta^{(n)}(\tilde{\mathbf{X}}^\theta) \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))}{\int \pi(\theta)p_\theta^{(n)}(\tilde{\mathbf{X}}^\theta) \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))d\tilde{\mathbf{X}}^\theta d\theta}, \quad (4)$$

which leads to the approximated Bayesian posterior as

$$\hat{\pi}^{EK}(\theta | \mathbf{X}) = \int \hat{\pi}^{EK}(\theta, \tilde{\mathbf{X}}^\theta | \mathbf{X})d\tilde{\mathbf{X}}^\theta = \frac{\pi(\theta)P_\theta^{(n)} \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))}{\int \pi(\theta)P_\theta^{(n)} \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))d\theta}. \quad (5)$$

Remark 1 (Generating Fake Data) *We assume $\tilde{\mathbf{X}}^\theta = g_\theta(\tilde{\mathbf{X}})$, where $\tilde{\mathbf{X}} \in \mathbb{R}^m$ are random variables arriving from $\tilde{P}^{(m)}$ and where $g_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a deterministic mapping. Generating random variable draws by passing $\tilde{\mathbf{X}}$ through some mapping is commonly done in practice, also known as the reparameterization trick (Kingma and Welling, 2013). For example, Gaussian random variables $\tilde{\mathbf{X}} = \{\tilde{X}_i^\theta\}_{i=1}^m$ that follow i.i.d. $N(\mu, \sigma^2)$ distribution can be obtained by transforming $\{\tilde{X}_i\}_{i=1}^m \stackrel{i.i.d.}{\sim} N(0, 1)$ via $\tilde{X}_i^\theta = \mu + \sigma\tilde{X}_i$. In other cases, one can use uniform draws $\tilde{\mathbf{X}}$ and the inverse transform sampling.*

Smooth kernels have been used inside ABC before to rescale the acceptance probability (e.g. Beaumont et al. (2002) employed the Epanechnikov kernel). Wilkinson (2013) and Sisson et al. (2018) provide a thorough overview and comparisons of the commonly used kernels. Our smoothed weights are directly interpretable due to their linkage between the KL divergence and the log-likelihood ratio. Algorithm 2 can be regarded as a version of Importance Sampling ABC (see Nguyen et al. (2020) for a variant using energy statistics and Park et al. (2016) for minimal description length ABC). We later show in Section 3.2 that, with the scaled exponential kernel, the ABC posterior corresponds to the stationary distribution of the MHC algorithm of Kaji and Ročková (2022) and can be regarded as a posterior under a misspecified model. Computational comparisons of the sequential MHC sampler with our parallelizable ABC sampler are performed in the Appendix Section C where we show benefits of the ABC strategy when convergence issues may arise for MHC due to initialization. In Section 3.4, we perform comparisons of the accept/reject (AR) and exponential kernels under model misspecification where we show that the AR kernel is far more robust.

The classifier needs to be trained for each ABC draw, which may incur additional computational cost compared to traditional ABC where the summary statistics and their distance can be computed without optimization. We provide comparisons of computation times in Appendix (Section F). Although the computation costs of our methods are higher when the data dimensionality d is relatively small, we are less disadvantageous when d is large compared to other discrepancies like Wasserstein distance or Maximum Mean Discrepancy. In addition, nonparametric discriminator classes such as neural network classifiers can efficiently benefit when there is an inherent low-dimensional structure in the data (Kaji et al., 2020). Additionally, training can be accelerated if one initiates the training at some pre-trained neural networks.

2.2 Connection to Robust Bayesian Inference

Our algorithms can be regarded as “robust ABC” algorithms that estimate (relative-entropy coarsened posteriors) c -posteriors introduced in Miller and Dunson (2018). Note that our focus of robustness is different from the robustness to different distance measures, but more relevant to data perturbation or model misspecification (discussed later in Section 3.4). The c -posteriors yield robust inference by conditioning on the event that the observed data \mathbf{X} is sufficiently close (in terms of the KL divergence) to the data generated by the model. A similar case as the Huber-type data contamination is considered in γ -divergence ABC Fujisawa et al. (2021). The proposed computation of c -posteriors in Miller and Dunson (2018) is made feasible only through asymptotic approximations (Section 3.1 in Miller and Dunson (2018)), e.g. with powered posteriors that are computable using conjugate priors. Our ABC methods can compute them without any approximation and for a broader class of priors. In particular, if in Algorithm 1 we draw $\epsilon \sim \text{Exp}(\alpha)$ for some $\alpha > 0$ and accept θ if $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) < \epsilon$, then our ABC posterior coincidentally approximates the relative-entropy c -posterior proportional to $\pi(\theta)P_\theta^{(n)}e^{-\eta K_n}$ where $K_n = \mathbb{P}_n \log \frac{p_0}{p_\theta}$. Algorithm 2 corresponds to the case $\alpha = n$ in Miller and Dunson (2018) without any approximation. Interestingly, the degree of robustness corresponds to the acceptance rate of ABC. For example, if we let $\alpha \ll n$, the c -posterior puts larger weight on the prior and robustifies the model, which corresponds to accepting many draws (more than proportional to n) and the draws reflecting the shape of the prior. On the contrary, if we let $\alpha \gg n$, the c -posterior puts most weight on a narrow neighborhood of the observed data, which corresponds to accepting very few draws for which the Kullback–Leibler divergence is the smallest. From an ABC’s perspective, probably the most interesting case is when the acceptance rate is roughly fixed throughout $n \rightarrow \infty$, in which case α is comparable with n and our algorithms produce a correct c -posterior without utilizing the approximation in Miller and Dunson (2018) which stands on $n \gg \alpha$ or $n \ll \alpha$. In addition, another advantage of our method is that it allows us to calculate the c -posterior for different α easily. If we use an MCMC with tempering (i.e. the powered likelihood approximation), we might need to run separate MCMC chains for different α . On the other hand, our ABC-based algorithm lets us calculate the c -posteriors by filtering out independent samples of candidate draws according to various α . This may be advantageous in applications when no ex-ante preference is available on the degree of robustness and when one wants to see how the concentration of the c -posterior varies with it.

2.3 Estimating KL Divergence via Classification

We adopt the ‘ $-\log D$ ’ trick to estimate the KL divergence (Goodfellow et al., 2014). More precisely, a flexible discriminator D (such as a neural network or logistic regression) is trained to maximize

$$\mathbb{M}_{n,m}^\theta(D) = \mathbb{P}_n \log D + \mathbb{P}_m^\theta \log(1 - D), \quad (6)$$

where we employ the operator notation for expectation, e.g., $\mathbb{P}_n f = \frac{1}{n} \sum_{i=1}^n f(X_i)$ and $\mathbb{P}_m^\theta f = \frac{1}{m} \sum_{i=1}^m f(\tilde{X}_i^\theta)$. This can also be regarded as a classification problem where we label $\{X_i\}_{i=1}^n$ (‘real’ data) with 1 and $\{\tilde{X}_i^\theta\}_{i=1}^m$ (‘fake’ data) with 0. The oracle maximizer to (6)

can be shown to be (Goodfellow et al., 2014, Proposition 1)

$$D_\theta(X) = \frac{p_0(X)}{p_0(X) + p_\theta(X)}. \quad (7)$$

The functional form of the oracle solution in (7) naturally suggests the following KL estimator obtained from a trained discriminator $\hat{D}_{n,m}^\theta$ (Thomas et al., 2020)

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \mathbb{P}_n \log \frac{\hat{D}_{n,m}^\theta}{1 - \hat{D}_{n,m}^\theta} = \frac{1}{n} \sum_{i=1}^n \log \frac{\hat{D}_{n,m}^\theta(X_i)}{1 - \hat{D}_{n,m}^\theta(X_i)}. \quad (8)$$

Later we show that our classification-based KL estimator (8) converges to a well-defined limit counterpart $K(p_0, p_\theta)$ under mild conditions in Section 3.1.

2.4 Other KL Estimators

Our ABC results, presented later in Section 3, can be extended to other types of KL estimators if similar estimation error results as in Theorem 5 can be shown. Since the rate $1/(nu^2)$ stems from the estimation error of the empirical KL divergence, the fundamental difference between our classification-based KL estimator and other KL estimators lies in the rate δ_n . One example is the k-Nearest Neighbor (kNN) estimator proposed in Pérez-Cruz (2008). Wang et al. (2009) showed that this estimator is asymptotically unbiased and mean-square consistent and they propose a data-dependent choice of k which can improve the convergence speed. Jiang et al. (2018) assess data discrepancy inside ABC with the special case of 1-nearest neighbor, which is defined as

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \frac{d}{n} \sum_{i=1}^n \log \frac{\min_j \|X_i - \tilde{X}_j^\theta\|}{\min_{j \neq i} \|X_i - X_j\|} + \log \frac{m}{n-1}. \quad (9)$$

where d is the number of covariates in \mathbf{X} . Zhao and Lai (2020) provide convergence rate of the bias for this kNN estimator is bounded by $n^{-2\gamma/(d+2)} \log n$, where γ is a parameter characterizing the tail behavior of the target distribution. Note that the kNN estimator is not applicable to cases where \mathbf{X} arises from a discrete distribution.

Another route to estimate the KL divergence is via (data-dependent) partitioning methods. Wang et al. (2005) proposed to estimate the Radon-Nikodym derivative dP_0/dP_θ using frequency counts on a statistically equivalent partition of \mathbb{R}^d . However, the computational complexity of their method is exponential in d and the estimation accuracy deteriorates quickly as the dimension increases. Silva and Narayanan (2010) further contributed to multivariate data-driven partition schemes by using a Barron-type histogram-based density estimate. They provide sufficient conditions on the partitions scheme to make the estimator strongly consistent.

Lastly, Nguyen et al. (2007) adopted a variational approach to estimate KL by reframing the estimation problem as a penalized convex risk minimization problem, where they restrict the estimate to a bounded subset of a Reproducing Kernel Hilbert Space (RKHS). Convergence rates are then obtained from empirical process theory on nonparametric M-estimators (van de Geer, 2000). In an independent contribution, Ghimire et al. (2020) used

a discriminator in RKHS to estimate KL using a similar approach to ours. They showed that the estimator error bound is related to the complexity of the discriminator in RKHS. We compare computational complexities of these methods and our approach in the Appendix Section G.

Beyond the forward KL divergence, our classification framework allows us to consider other discrepancy metrics. Alternatively to (8), we could instead estimate the reversed KL divergence

$$\hat{K}_{\text{reverse}}(\tilde{\mathbf{X}}^\theta, \mathbf{X}) = \frac{1}{m} \sum_{i=1}^m \log \frac{1 - \hat{D}_{n,m}^\theta(\tilde{X}_i^\theta)}{\hat{D}_{n,m}^\theta(\tilde{X}_i^\theta)}$$

which converges to $K(p_\theta, p_0)$ and which is still uniquely minimized at $p_\theta = p_0$. One can show that the estimation error of this reversed KL estimator is still $O_{P^*}(\delta_n)$ by following the same techniques used in Theorem 3. The reversed KL divergence is widely used in variational inference (Jordan et al., 1999; Wainwright and Jordan, 2008). Forward and reversed KL’s could perform differently when the function class inside the variational approach is not rich enough. The reversed KL is zero-forcing/mode-seeking, while the forward KL is mass-covering/mean-seeking (Bishop, 2006). Another related metric, deployed by Gutmann et al. (2018), is the classification accuracy (CA) defined as

$$\text{CA}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \frac{1}{n+m} \left(\sum_{i=1}^n \hat{D}_{n,m}^\theta(X_i) + \sum_{i=1}^m (1 - \hat{D}_{n,m}^\theta(\tilde{X}_i^\theta)) \right). \quad (10)$$

Since $\hat{D}(\cdot)$ and $\log \frac{\hat{D}}{1-\hat{D}}(\cdot)$ are linked by a logistic transformation which is Lipschitz-continuous, CA can be roughly regarded as a weighted average of the forward KL divergence and the reversed KL divergence. Our framework is connected to other GAN-style discrepancy metrics which are also used in ABC literature. We provide a discussion in Appendix D.

3. Frequentist Analysis of the ABC Posterior

One way to assess the quality of the posterior distribution is through the speed at which it contracts around the truth θ_0 as $n \rightarrow \infty$. While the ABC posterior is ultimately an approximation, it might still concentrate about θ_0 at a reasonable rate. In this section, we look into theoretical properties of both Algorithm 1 and Algorithm 2. First, we develop a tail bound result quantifying how fast the classification-based estimator $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)$ converges to the true KL divergence $K(p_0, p_\theta)$ conditional on approximability of the discriminator class. The tail bound analysis is crucial in our convergence analysis. Next, we show that the convergence rate of the accept-reject ABC in Algorithm 1 is determined jointly by the accept-reject threshold ϵ_n , the estimation error δ_n (with respect to $\mathbb{P}_n \log p_0/p_\theta$), and the rate $n^{-1/2}$ of estimation between $\mathbb{P}_n \log p_0/p_\theta$ and $K(p_0, p_\theta)$. Further, the typical posterior distribution converges to a uniform distribution over an ellipse when the acceptance threshold ϵ_n dominates the other two. On the other hand, the exponentially weighted posterior in Algorithm 2 can be viewed as the posterior under a “misspecified” model. The convergence rate is then determined by the contraction rate of the true posterior and the estimation error, where the ABC posterior is asymptotically normal around the KL projection under LAN conditions (Kaji and Ročková, 2022).

3.1 Convergence Rate of Estimation Errors

We assume that the set of considered classifiers \mathcal{D} resides in a sieve \mathcal{D}_n that expands with the sample size and its size is measured by bracketing entropy $N_{[]}(\epsilon, \mathcal{F}, d)$. Kaji et al. (2020) prove the rate of convergence of such a classifier (under assumptions reviewed later) with a Hellinger-type distance defined as

$$d_\theta(D_1, D_2) = \sqrt{h_\theta(D_1, D_2)^2 + h_\theta(1 - D_1, 1 - D_2)^2},$$

where $h_\theta(D_1, D_2) = \sqrt{(P_0 + P_\theta)(\sqrt{D_1} - \sqrt{D_2})^2}$.

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

Assumption 1 requires three conditions. First, the synthetic data sample size m should be at least as large as the actual data size n , which can be assured. Second, the discriminator class needs to be expressive enough so we can find a sufficiently good maximizer approximating the oracle discriminator D_θ . Lastly, the entropy of the sieve should be moderate to prevent overfitting.

The following theorem states that the sequence δ_n in Assumption 1 determines the convergence rate of \hat{D} . The speed at which δ_n converges to 0 depends on the choice of the sieve and smoothness of the model. When a nonparametric estimator is employed, δ_n is often slower than $n^{-1/2}$. In Appendix A, we give a specific expression of δ_n for a neural network classifier and give a few examples in which δ_n vanishes faster than $n^{-1/4}$.

Lemma 2 (Kaji et al., 2020, Theorem S.1) *Under Assumption 1, we have $d_\theta(\hat{D}_{n,m}^\theta, D_\theta) = O_P^*(\delta_n)$.*²

To establish the rate of convergence of our approximated posterior, we have the following assumption.

Assumption 2 *There exists $\Lambda > 0$ such that for every $\theta \in \Theta$, $P_0(p_0/p_\theta)$ and $P_0(p_0/p_\theta)^2$ are bounded by Λ and*

$$\sup_{D \in \mathcal{D}_{n,\delta_n}^\theta} P_0 \left(\frac{D_\theta}{D} \middle| \frac{D_\theta}{D} \geq \frac{25}{16} \right) < \Lambda, \quad \sup_{D \in \mathcal{D}_{n,\delta_n}^\theta} P_0 \left(\frac{1 - D_\theta}{1 - D} \middle| \frac{1 - D_\theta}{1 - D} \geq \frac{25}{16} \right) < \Lambda,$$

for δ_n in Assumption 1. The brackets in Assumption 1 can be taken so that $P_0(\sqrt{u/l} - 1)^2 = O(d_\theta(u, l)^2)$ and $P_0(\sqrt{(1-l)/(1-u)} - 1)^2 = o(d_\theta(u, l))$.

2. We use P^* to denote outer expectation (see Section 1.2 of van der Vaart and Wellner (1996)), here is the expectation of a “a smallest measurable function g that dominates $d_\theta(\hat{D}_{n,m}^\theta, D_\theta)$ ”.

Assumption 2 constrains the tail behavior of the discriminator so that the residual of the cross-entropy loss in (6) can be circumscribed by the bracketing entropy. See Kaji and Ročková (2022) for a discussion of how this can be reasonably satisfied for logistic discriminator and neural network discriminators that use sigmoid activation functions.

The following theorem quantifies the rate of convergence of our estimator (8) towards the *empirical* KL divergence $\mathbb{P}_n \log \frac{p_0}{p_\theta}$.

Lemma 3 (Convergence Rate of Estimation Errors) *Under Assumptions 1 and 2,*

$$\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - \mathbb{P}_n \log \frac{p_0}{p_\theta} \right| = O_{P^*}(\delta_n). \quad (11)$$

Proof. Since the estimation error can be rewritten as

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - \mathbb{P}_n \log \frac{p_0}{p_\theta} = -\mathbb{P}_n \left(\log \frac{1 - \hat{D}_{n,m}^\theta}{1 - D_\theta} - \log \frac{\hat{D}_{n,m}^\theta}{D_\theta} \right),$$

it follows from Kaji and Ročková (2022, Theorem 4.1).

Remark 4 (Uniform Convergence Rate) *The rate of estimation in Theorem 2 and Theorem 3 is characterized as point-wise. While for each $\theta \in \Theta$ the estimation error is shrinking at the rate δ_n , the multiplication constant in front the rate could potentially depend on θ . Assuming a compact parameter support (which is not a-typical in deep learning models; see e.g. Schmidt-Hieber (2020) or Polson and Rockova (2018) and Wang and Rockova (2020)) and continuity of the multiplication constant, we can essentially regard the rate as uniform. Hereafter, we thereby abuse the notation and tacitly assume that δ_n is the worst rate over all $\theta \in \Theta \subset \mathbb{R}^d$, i.e. the rate with the largest multiplication constant.*

Next, we investigate convergence around the *actual* KL divergence $K(p_0, p_\theta)$. The next lemma will be utilized later in the proof of Theorem 7. However, it is of independent interest as it shows how the speed at which the joint error probability (accounting for randomness of both the observed and fake data $(\mathbf{X}, \tilde{\mathbf{X}})$) decays in terms of the estimation error. Below, the probability P corresponds to $P_0^{(n)} \otimes \tilde{P}^{(m)}$, where $\tilde{P}^{(m)}$ is the measure for $\tilde{\mathbf{X}}$ (see Remark 1).

Theorem 5 *For a given $\theta \in \Theta$, we define for $u > 0$ and $\delta_n > 0$ as in Lemma 2 and for an arbitrarily slowly increasing sequence $C_n > 0$*

$$\rho_{n,\theta}(u; C_n; \delta_n) \equiv P \left(\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > 2u, d_\theta(\hat{D}_{n,m}^\theta, D^\theta) \leq C_n \delta_n \right). \quad (12)$$

Under Assumptions 1 and 2, we then have

$$\rho_n(u; C_n; \delta_n) \equiv \sup_{\theta \in \Theta} \rho_{n,\theta}(u; C_n; \delta_n) = O \left(\frac{C_n \delta_n}{u} + \frac{1}{nu^2} \right).$$

The proof can be found in Appendix H.1. Note that the intersecting event above has probability converging to one, according to Lemma 2.

Remark 6 (Neural Network Sieve) *The results discussed in this section apply to any nonparametric sieve discriminator that satisfies the entropy conditions. To facilitate understanding of how the rate of convergence δ can be affected by the dimension of the data space d , the smoothness of the density, and the choice of the classifier, we provide a discussion on neural network sieve specifically here.*

Borrowing from the idea of Bauer and Kohler (2019), the convergence rate of the appropriately configured neural network discriminator depends only on the low underlying dimension of the oracle discriminator, however large the ostensible dimension of the data is. Intuitively, the low-dimensional structure can be described as follows. The log-likelihood ratio $\log(p_0/p_\theta)$ takes a d -dimensional input X as an argument, which may be large. If this function admits a representation as a nested composition of smooth functions, each of which takes a possibly smaller number d^ of arguments, the neural network sieve can adapt to this underlying structure and converges faster than the traditionally proven rate.*

In particular, if $d^ < 2p$, we have $\delta_n = o_P(n^{-1/4})$, which is often the desired rate for the nonparametric estimator of a nuisance parameter. Additionally we show that one can obtain $\delta_n \lesssim n^{-2/5}$ for the binary choice model with logistic errors and δ_n arbitrarily close to $n^{-1/2}$ for the discretely sampled Brownian motion model. The detailed characterizations of the low underlying dimension d^* and two examples can be found in Appendix A.*

Finally, with Assumption 2, the convergence rate δ_n translates into the convergence rate of the Kullback-Leibler estimator in Theorem 3. Thus, in smooth low-dimensional hierarchical models, our Kullback-Leibler estimator converges reasonably fast even when the nominal dimension of the data is large.

3.2 Posterior Concentration Rate

Concentration rates are typically quantified in terms of a prior concentration (measured in terms of a combination of the KL divergence and the KL variation) and the entropy of the model. We have a similar prior mass condition (see (8.4) in Section 8.2 in Ghosal and Van der Vaart (2017)). We denote the KL-neighborhood of p_0 by

$$B_2(p_0, \epsilon) = \{\theta : K(p_0, p_\theta) < \epsilon^2\}. \quad (13)$$

Assumption 3 (Prior Mass) *There exist some constants $\kappa > 0$ and $\xi > 0$ such that for every $0 < \epsilon < \xi$ and some constant $C > 0$, the prior probability satisfies $\Pi[B_2(p_0, \epsilon)] \geq C\epsilon^\kappa$.*

Next, we assume that the parameter θ is identifiable in the sense that the KL divergence is locally compatible with the Euclidean norm. This assumption is adopted from Assumption 3(ii) of Frazier et al. (2018).

Assumption 4 (Identification) *The density function p_θ is continuous in θ and for every θ in some open neighborhood of θ_0 satisfies*

$$\|\theta - \theta_0\| \leq L \times K(p_0, p_\theta)^\alpha$$

for some $L > 0$ and $\alpha > 0$.

Similarly as in (5.1) in Kleijn and van der Vaart (2006), Assumption 4 ensures posterior concentration around θ_0 when $K(p_0, p_\theta) \rightarrow 0$. This holds for many distributions. For example, for the exponential distribution with a rate parameter θ , we have $K(p_0, p_\theta) = \frac{\theta}{\theta_0} - \log(\frac{\theta}{\theta_0}) - 1$. Since $\log(1+x) = x - \frac{x^2}{2} + o(x^2)$ when $x \rightarrow 0$, we have $K(p_0, p_\theta) \geq \frac{1}{2}\theta_0^{-2}(\theta - \theta_0)^2$. For multivariate normal distribution with a known variance Σ and an unknown location μ , we have $K(p_0, p_\theta) = \frac{1}{2}(\mu - \mu_0)\Sigma^{-1}(\mu - \mu_0) \geq \frac{1}{2}\rho(\Sigma)^{-1}\|\mu - \mu_0\|^2$, where $\rho(\Sigma)$ is the spectral radius, i.e., the largest eigenvalue, of a matrix Σ .

First, we focus on the uniform kernel $T_\epsilon(x) = \mathbb{I}(|x| \leq \epsilon)$ used in Algorithm 1. Recall (from Remark 1) that $\tilde{\mathbf{X}}^\theta = g_\theta(\tilde{\mathbf{X}})$ for some suitable mapping $g_\theta(\tilde{\mathbf{X}})$ where $\tilde{\mathbf{X}} \sim \tilde{P}^{(m)}$. The ABC joint posterior (2) is a weighted aggregation of uniform kernels, i.e

$$\hat{\pi}^{AR}(\theta, \tilde{\mathbf{X}} \mid \hat{K}[\mathbf{X}, g_\theta(\tilde{\mathbf{X}})] \leq \epsilon) = \frac{\tilde{\pi}(\tilde{\mathbf{X}})\pi(\theta)\mathbb{I}(\hat{K}[\mathbf{X}, g_\theta(\tilde{\mathbf{X}})] \leq \epsilon)}{\int \tilde{\pi}(\tilde{\mathbf{X}})\pi(\theta)\mathbb{I}(\hat{K}[\mathbf{X}, g_\theta(\tilde{\mathbf{X}})] \leq \epsilon)d\tilde{\mathbf{X}}d\theta}, \quad (14)$$

which yields the following ABC posterior distribution

$$\hat{\Pi}_{\epsilon_n}^{AR}(A \mid \mathbf{X}) = \frac{\int_{\theta \in A} \pi(\theta)\tilde{P}^{(m)}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n)d\theta}{\int_{\Theta} \pi(\theta)\tilde{P}^{(m)}(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n)d\theta} \quad \text{for a Borel-measurable } A \subset \Theta. \quad (15)$$

The following theorem (a modification of Theorem 1 in Frazier et al. (2018)) quantifies the concentration rate in terms of the tolerance threshold ϵ_n as well as the rate at which the classification-based KL estimator can estimate $\mathbb{P}_n \log(p_0/p_\theta)$ (as formulated in (11)).

Theorem 7 *Let Assumptions 1, 2 and 3 hold and take δ_n as in (11) in Theorem 3. Then, as $n \rightarrow \infty$ and with $\epsilon_n = o(1)$ such that $n\epsilon_n^2 \rightarrow \infty$ and $C_n\delta_n = o(\epsilon_n)$ for some arbitrarily slowly increasing sequence $C_n > 0$ we have*

$$P_0^{(n)}\Pi[K(p_0, p_\theta) > \lambda_n \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n] = o(1), \quad (16)$$

where $\lambda_n = \epsilon_n + M_n C_n \delta_n \epsilon_n^{-\kappa} + \sqrt{M_n} n^{-1/2} \epsilon_n^{-\kappa/2}$ for some arbitrarily slowly increasing sequence $M_n > 0$. Moreover, if Assumption 4 also holds, as $n \rightarrow \infty$, we have

$$P_0^{(n)}\Pi[\|\theta - \theta_0\| > L\lambda_n^\alpha \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n] = o(1). \quad (17)$$

The proof of the theorem is provided in Appendix H.2. Thus, the convergence rate of our ABC posterior depends on three components: the accept-reject threshold ϵ_n , the estimation error of the KL estimator δ_n and the rate of discrepancy $n^{-1/2}$ between the empirical and true KL divergence. Since δ_n will typically be greater than the parametric rate $n^{-1/2}$, the overall convergence rate is then driven by $\lambda_n = \epsilon_n + \tilde{M}_n \delta_n \epsilon_n^{-\kappa}$, where \tilde{M}_n is an arbitrarily slowly increasing sequence.

In practice, it is unclear how to properly choose ϵ_n . In Algorithm 2, we proposed to weight the draws using a scaled exponential kernel $\exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))$. We denote the ABC posterior under the exponential kernel as $\hat{\Pi}^{EK}(\cdot \mid \mathbf{X})$ where

$$\hat{\Pi}^{EK}(A \mid \mathbf{X}) = \frac{\int_A \tilde{P}^{(m)} \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))\pi(\theta)d\theta}{\int_{\Theta} \tilde{P}^{(m)} \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta))\pi(\theta)d\theta}. \quad (18)$$

To gain more insights into the ABC posterior behavior under the exponential kernel, we take a closer look at the “likelihood function” above

$$\tilde{P}^{(m)} \exp(-n\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)) = \frac{p_\theta^{(n)}}{p_0^{(n)}} \tilde{P}^{(m)} e^{u_\theta},$$

where $u_\theta(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = -n \times \left(\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - \mathbb{P}_n \log \frac{p_\theta}{p_0} \right)$. From the equations above, we can write

$$\hat{\pi}^{EK}(\theta | \mathbf{X}) \propto p_\theta^{(n)}(\mathbf{X}) \times e^{\hat{u}_\theta(\mathbf{X})} \times \pi(\theta) \quad \text{with} \quad \hat{u}_\theta(\mathbf{X}) = \log \int e^{u_\theta(\mathbf{X}, \tilde{\mathbf{X}}^\theta)} d\tilde{P}^{(m)}(\tilde{\mathbf{X}}). \quad (19)$$

When $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)$ is the classification-based estimator, $\hat{u}_\theta(\mathbf{X})$ can be related to the random generator setting of the Metropolis-Hastings MHC algorithm in Kaji and Ročková (2022) which has (19) as its stationary distribution. Similarly as in Appendix Section 5 of Kaji and Ročková (2022), we can regard the posterior approximation in (19) as a posterior $\hat{\pi}^{EK}(\theta | \mathbf{X}) \propto q_\theta^{(n)} \tilde{\pi}(\theta)$ under a misspecified likelihood

$$q_\theta^{(n)} = \frac{p_\theta^{(n)}(\mathbf{X}) e^{\hat{u}_\theta(\mathbf{X})}}{C_\theta} \quad \text{where} \quad C_\theta = \int_{\mathcal{X}} p_\theta^{(n)}(\mathbf{X}) e^{\hat{u}_\theta(\mathbf{X})} d\mathbf{X} \quad (20)$$

and a modified prior $\tilde{\pi}(\theta) \propto \pi(\theta) C_\theta$. Since the likelihood is misspecified, the ABC posterior concentrates around a projection point θ^* defined as

$$\theta^* = \arg \min_{\theta \in \Theta} -P_0^{(n)} \log[q_\theta^{(n)}/p_0^{(n)}], \quad (21)$$

which corresponds to the mis-specified model that is closest to $P_0^{(n)}$ in the KL sense (Kleijn and van der Vaart, 2006). Kaji and Ročková (2022) study posterior concentration of (19). We recall their Theorem 4.5 in Appendix B. Unlike in Theorem 7, the posterior concentration rate here depends both on the estimation error δ_n and the actual concentration rate of the true posterior, not the acceptance threshold.

Remark 8 (Vanishing Bias) *To better understand the severity of the centering bias of the misspecified model, we note*

$$\begin{aligned} -P_0^{(n)} \log[q_{\theta^*}^{(n)}/p_0^{(n)}] &\leq -P_0^{(n)} \log[q_{\theta_0}^{(n)}/p_0^{(n)}] = P_0^{(n)} \log \frac{p_0^{(n)}}{p_0^{(n)} e^{\hat{u}_{\theta_0}(\mathbf{X})} / C_{\theta_0}} \\ &= \log C_{\theta_0} - P_0^{(n)} \hat{u}_{\theta_0}(\mathbf{X}) = \log P_0^{(n)} e^{\hat{u}_{\theta_0}(\mathbf{X})} - P_0^{(n)} \hat{u}_{\theta_0}(\mathbf{X}). \end{aligned}$$

This is essentially the Jensen gap. If we have this Jensen gap vanishing when $n \rightarrow \infty$, then we can conclude that the centering bias is also vanishing, and the ABC posterior in (18) will eventually concentrate at the right location.

3.3 Shape of the Limiting ABC Posterior Distribution

We now analyze the limiting shape of $\hat{\Pi}_{\epsilon_n}^{AR}(\cdot \mid \mathbf{X})$ defined in (15). We focus on the case when $\epsilon_n \gg \delta_n^{1/(\kappa+1)}$, where κ was defined in Assumption 3, since the posterior is not guaranteed to converge when the decision threshold ϵ_n is smaller than the estimation error δ_n of the KL estimator.

Assumption 5 *Assume that for every $\varepsilon > 0$, we have $\inf_{\|\theta - \theta_0\| > \varepsilon} K(p_0, p_\theta) > 0$. In addition, assume that $\log p_\theta$ is twice differentiable with respect to θ and that, for every θ in some neighborhood of θ_0 , the remainder of the second order Taylor expansion of $K(p_0, p_\theta) = P_0 \log \frac{p_0}{p_\theta}(x)$ around θ_0 is comparatively small relative to the second-order term, i.e.*

$$\begin{aligned} K(p_0, p_\theta) &= \nabla_{\theta=\theta_0} K(p_0, p_\theta)(\theta - \theta_0) + \frac{1}{2}(\theta - \theta_0)' \nabla_{\theta=\theta_0}^2 K(p_0, p_\theta)(\theta - \theta_0) \{1 + o(1)\} \\ &= \frac{1}{2}(\theta - \theta_0)' I(\theta_0)(\theta - \theta_0) \{1 + o(1)\}, \end{aligned}$$

where $I(\theta) = \nabla_\theta^2 K(p_0, p_\theta) = P_0[(\nabla_\theta \log p_\theta)^2]$ is the Fisher information matrix.

When Assumption 5 is satisfied, the condition in Assumption 4 is immediately satisfied as well and the identification of θ_0 is guaranteed. In the open-neighborhood of θ_0 , since we have

$$(\theta - \theta_0)' I(\theta_0)(\theta - \theta_0) = 2K(p_0, p_\theta) \{1 + o(1)\},$$

and $I(\theta_0)$ is positive definite, the convergence in the bilinear form $(\theta - \theta_0)' I(\theta_0)(\theta - \theta_0)$ ensures the convergence of the parameters.

Theorem 9 *Assume that the prior function $\pi(\cdot)$ is continuous around θ_0 . Then, under Assumptions 1, 2, 3 and 5, if $\lim_n \delta_n / \epsilon_n^{\kappa+1} \rightarrow 0$, the average posterior distribution of $\epsilon_n^{-1/2}(\theta - \theta_0)$ converges to the uniform distribution over the ellipse $\{w : w' I(\theta_0) w \leq 2\}$ where $I(\theta)$ is the Fisher information matrix defined in Assumption 5. In particular, as $n \rightarrow 0$, we have*

$$P_0^{(n)} \int f\left(\epsilon_n^{-1/2}(\theta - \theta_0)\right) \hat{\Pi}_{\epsilon_n}^{AR}(\theta \mid \mathbf{X}) \rightarrow \int_{w' I(\theta_0) w \leq 2} f(u) du / \int_{w' I(\theta_0) w \leq 2} du$$

for every continuous and bounded function $f(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$.

The proof is provided in Appendix H.3.

Remark 10 *Theorem 9 is adapted from the case (i) in Theorem 2 of Frazier et al. (2018). We only consider situations when $\epsilon_n \gg \delta_n^{1/(\kappa+1)}$ with the prior shrinkage parameter κ defined in Assumption 3. In other words, we assume that the ABC decision threshold ϵ_n is dominating both the estimation error δ_n and the asymptotic error $n^{-1/2}$ and, thereby, determines the posterior concentration rate. It is not entirely obvious how the posterior would behave when the threshold ϵ_n shrinks faster than the estimation error δ_n .*

For the asymptotic behavior of the ABC posterior (19) induced by the exponential kernel, we resort to BvM characterizations under misspecification in LAN models (Kleijn and van der Vaart, 2012). When the posterior (19) concentrates around θ^* in (21) at the rate ϵ_n^* , one can show under a suitable LAN condition that the approximate posterior converges to a sequence of normal distributions in total variation at the rate ϵ_n^* . The centering and the asymptotic covariance matrix both depend on θ^* . The formal statement and the proof is in Kaji and Ročková (2022).

Although we only consider the case where the model is correctly specified in our paper, our results can be extended to the mis-specified model along the lines of Frazier et al. (2019).

3.4 ABC Kernels and Model Misspecification

Although the exponential kernel ABC (Algorithm 2) obviates the need for the threshold ϵ_n and performs very well in our examples, the accept/reject kernel ABC (Algorithm 1) may perform better under mis-specification (see Remark 2.2). When the model is mis-specified, the posterior under Algorithm 1 will converge to the “pseudo-true” value, which is the point that minimizes the distance between summary statistics within the mis-specified class. Since our summary statistic is replaced with KL divergence, this point will coincide with the KL projection in our case. The exponential kernel will also concentrate around a certain KL projection but its bias will now be compounded by the influence of both $P_0 \notin \mathcal{P} = \{\theta \in \Theta : P_\theta\}$ and the exponential tilt $e^{\hat{u}_\theta(\mathbf{X})}$ arising from the approximation error of the KL estimator. We would thereby expect a bigger bias from the exponential kernel when the model is misspecified.

We illustrate the intuition above with a toy example. We use the simple example proposed in Frazier et al. (2019) where the assumed data-generating process (DGP) is i.i.d. $\mathcal{N}(\theta, 1)$, but the actual DGP for \mathbf{X} is i.i.d. $\mathcal{N}(\theta, \sigma^2)$ for $\sigma^2 \neq 1$. In other words, the assumed DGP has an incorrect specification of the variance of the observed data. We consider the oracle logistic classifier built on \mathbf{X} and the quadratic term \mathbf{X}^2 for our KL estimator. We fix $\theta = 1$ and simulate \mathbf{X} (using $n = 100$) with respect to different values of σ^2 , ranging from 0.5 to 5 with evenly spaced increments of 0.05. Using one common set of latent variables $\nu_i \sim \mathcal{N}(0, 1)$, the observed data is generated as $X_i = 1 + \nu_i\sigma$ for each σ^2 . The prior belief is $\theta \sim \mathcal{N}(0, 25)$, and we implement ABC methods with $N = 100\,000$ pseudo datasets. The parameter draws and the latent variable datasets are the same across the different values of σ^2 . To explore how the tail behavior influences the ABC bias, we also include misspecified lognormal distributions using the same setup.

Figure 1(a) compares the posterior mean of the accept-reject ABC (AR) and the ABC with exponential weighting (EW) across different values of σ^2 . We can see that both AR and EW have a relatively small bias in estimating θ when the misspecification level in σ^2 varies. Since the logistic regression classifier on \mathbf{X} and \mathbf{X}^2 is almost the “oracle” discriminator for gaussian distributions, the error of the KL estimator should be minimal. Nevertheless, the posterior mean of EW exhibits a downward moving trend when the level of misspecification increases. For the heavy-tailed lognormal distribution shown in Figure 1(b), although the posterior mean of AR does shift away from true value $\theta = 1$ as the degree of




Figure 1: ABC posterior under misspecified models.

model misspecification increases, the posterior mean of EW shifts away from $\theta = 1$ at a faster speed.

4. Simulations

In this section, we illustrate our approach and make comparisons with other likelihood-free inference techniques. Within our KL-ABC framework, we include two types of KL estimators. One is obtained with the logit discriminator score, which we refer to as KL estimation via classification (KLC), and the other one is estimated via the kNN method (kNN) with $k = 1$ (Jiang et al., 2018). For both estimators, we aggregate ABC samples with the accept-reject kernel as in Algorithm 1 and the exponential kernel as in Algorithm 2. The latter will be denoted with a suffix ‘exp’, e.g. KLC-exp or kNN-exp. The discriminator used for each dataset will be specified later. We provide more discussions on discriminator calibrations in Appendix F. Besides the two examples presented here, we have another analysis on the g-and-k distribution in Appendix E.

The ABC discrepancy metrics we choose for comparisons are (1) the classification accuracy (CA) (Gutmann et al., 2018) defined as (10); (2) the 2-Wasserstein (W2) distance under the Euclidean metric (Bernton et al., 2019) defined as $W2(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \min_\gamma [\sum_{i=1}^n \sum_{j=1}^m \gamma_{ij} \|X_i - \tilde{X}_j^\theta\|^2]^{1/2}$ s.t. $\gamma \mathbf{1}_m = \mathbf{1}_n, \gamma' \mathbf{1}_n = \mathbf{1}_m$ with $0 \leq \gamma_{ij} \leq 1$; (3) ℓ_2 -distance between summary statistics (SS) and we use the semi-automatic (SA) method (Fearnhead and Prangle, 2012) if no candidate summary statistics are given; (4) approximated posterior mean of the parameters predicted by trained deep neural network (DNN) (Jiang et al., 2017); (5) Maximum Mean (MM) discrepancy (Park et al., 2016) defined as $MM(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \frac{1}{n(n-1)} \sum_{i \neq j} k(X_i, X_j) + \frac{1}{m(m-1)} \sum_{i \neq j} k(\tilde{X}_i^\theta, \tilde{X}_j^\theta) - \frac{2}{nm} \sum_{i,j} k(X_i, \tilde{X}_j^\theta)$ where $k(\cdot, \cdot)$ is a Gaussian kernel with the bandwidth being the median of $\{\|X_i - X_j\| : i \neq j\}$; (6) a V-statistic estimator of Energy Statistics (ES) proposed by Nguyen et al. (2020); (7) auxiliary likelihood $AL = \frac{1}{m} \ln p_A(\tilde{\mathbf{X}}^\theta | \hat{\phi}(\tilde{\mathbf{X}}^\theta)) - \frac{1}{m} \ln p_A(\tilde{\mathbf{X}}^\theta | \hat{\phi}(\mathbf{X}))$ proposed by Drovandi et al. (2011), where $p_A(x | \phi)$ is a d -dimensional Gaussian distribution with ϕ being the sample mean

and covariance. For the classification accuracy (CA), we use the same discriminator as the one in KL estimation. For the DNN approach, we deploy a 3-layer DNN with 100 neurons and hyperbolic tangent (tanh) activation on each hidden layer. The model is trained on 10^6 samples and validated on 10^5 samples, with early stopping once the validation error starts to increase. In each experiment, unless otherwise noted, we set the tolerance threshold ϵ adaptively such that 1 000 of 100 000 (i.e. the top 1%) proposed ABC samples are accepted.

4.1 M/G/1-Queuing Model

Because queuing models are usually easy to simulate from, but have no tractable likelihoods, they have been frequently used as test cases in the ABC literature, see e.g. Fearnhead and Prangle (2012) and Bernton et al. (2019). Here, we choose the same setup as in Jiang et al. (2017). Each datum is a 5-dimensional vector consisting of the first five inter-departure times $x_i = (x_{i1}, x_{i2}, x_{i3}, x_{i4}, x_{i5})'$. In the model, the service times u_{ik} follow a uniform distribution $U[\theta_1, \theta_2]$, and the arrival times w_{ik} are exponentially distributed with the rate θ_3 . We only observe the interdeparture times x_i , given by the process $x_{ik} = u_{ik} + \max(0, \sum_{j=1}^k w_{ij} - \sum_{j=1}^{k-1} x_{ij})$. We perform ABC on $n = 500$ observed samples which are generated from the true parameter $\theta_0 = (1, 5, 0.2)$. The prior on $(\theta_1, \theta_2 - \theta_1, \theta_3)$ is uniform on $[0, 10]^2 \times [0, 0.5]$.³

Regarding the choice of the discriminator, we consider both the Random Forest (RF) classifier and a ℓ_1 -penalized logistic classifier (LRD). For the former, we use the default setting in the R package `randomForest`. We also denote CA calculated from RF classifier as RF-CA. For the latter, we implement the discriminator with R package `glmnet`, and the model is built on degree-2 polynomials of the data, including quadratic and interaction terms, with penalty parameter λ is selected via 5-fold cross-validation. We find that RF outperforms LRD.

The shape of the ABC posteriors is given in Figure 2. The plot reveals that, among the three parameters, θ_1 is the hardest one to estimate where all methods, except for RF and its variants, gave relatively flat posterior estimates. Regarding θ_2 , all methods seem to center well around the truth. RF-exp provides the tightest estimation. Regarding θ_3 , all except kNN return spiky posterior. Next, we repeat the experiments on 10 different datasets, and summarize the average squared estimation errors and the width of the 95% credible intervals in Table 1. Overall, we see that RF and RF-exp are able to correctly identify the right locations of the parameters and outperform the kNN estimator, with RF-exp providing the tightest credible interval. The method RF-CA tends to give very similar results to RF, which is not entirely unexpected since they are derived from the same discriminators.

4.2 Lotka-Volterra Model

The Lotka-Volterra (LV) predator-prey model (Wilkinson, 2018) describes population evolutions in ecosystems where predators interact with prey. It is one of the classical stochastic kinetic network model examples. The state of the population is prescribed deterministically via a system of ordinary differential equations (ODEs). Inference for such models is chal-

3. We place the uniform prior on $\theta_2 - \theta_1$ instead of θ_2 , since θ_2 must be larger than θ_1 . This is used in Jiang et al. (2017).




Figure 2: ABC reconstructed posteriors under M/G/1-queuing model with $\theta_0 = (1, 5, 0.2)'$. The vertical black lines mark the true values. The blue curves in kNN and RF boxes mark a smoothed density calculated from the exponential kernel.

Method (scale)	$\theta_1 = 1$		$\theta_2 = 5$		$\theta_3 = 0.2$	
	$(\hat{\theta}_1 - \theta_1)^2$	95% CI width	$(\hat{\theta}_2 - \theta_2)^2$	95% CI width	$(\hat{\theta}_3 - \theta_3)^2$ (10^{-4})	95% CI width
RF	0.008	2.044	0.038	4.304	0.571	0.084
RF -exp	0.035	0.959	0.055	1.473	0.307	0.028
LRD	0.197	3.116	0.217	4.599	0.308	0.064
LRD-exp	0.169	2.851	0.312	3.708	0.234	0.030
kNN	0.525	3.135	0.106	3.986	3.659	0.094
kNN-exp	1.057	2.664 (0.8)	0.431	3.331 (0.9)	2.634	0.072
DNN	0.150	3.350	1.100	7.361	81.931	0.328
ES	0.255	3.321	0.087	5.280	1.176	0.070
CA	0.191	3.107	0.235	4.602	0.280	0.064
RF-CA	0.014	2.194	0.036	4.104	0.404	0.084
AL	0.259	3.423	0.057	5.651	0.575	0.044
SA	0.180	2.457	0.355	3.514	45.297	0.446
W2	0.595	3.631	0.039	4.871	3.846	0.052 (0.8)

Table 1: ABC performance on the M/G/1 queuing model over 10 repetitions with top 1% ABC samples selected. Most of the 95% CIs have full coverage with the rest having their coverage marked in brackets. The bold fonts mark the best model in each metric.




Figure 3: Estimated $\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)$ (via Classification) over a grid of (θ_1, θ_4) values. The other two parameters are fixed as $\theta_2 = 0.5, \theta_3 = 1$.

lenging because the transition density is intractable. However, simulation from the model is possible, which makes it a natural candidate for ABC methods.

The model monitors population sizes of predators X_t and prey Y_t over time t . The changes in states are determined by four parameters $\theta = (\theta_1, \dots, \theta_4)'$ 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; (4) the rate $r_4^t = \theta_4 X_t Y_t$ of a prey dying. Given the initial population sizes (X_0, Y_0) at time $t = 0$, the dynamics can be simulated using the Gillespie algorithm (Gillespie, 1977). The algorithm samples times to an event from an exponential distribution (with a rate $\sum_{j=1}^4 r_j^t$) and picks one of the four reactions with probabilities proportional to their individual rates r_j^t .

We use the same simulation setup as Kaji and Ročková (2022). Each simulation is started at $X_0 = 50$ and $Y_0 = 100$ and state observations are recorded every 0.1 time units for a period of 20 time units, resulting in a series of $T = 201$ observations each. The real data ($n = 20$ time series) are generated with true values $\theta_0 = (0.01, 0.5, 1, 0.01)'$. The predator-prey interaction dynamic is very sensitive to parameter changes. For example, Figure 4 of Kaji and Ročková (2022) shows that a slight perturbation in θ_2 leads to significant changes in the population renewal cycle. We rely on the ability of the discriminator to tell such different patterns apart. The sensitivity of the model to minor parameter changes is confirmed with heat-map plots of the estimated KL divergence as a function of $(\theta_1, \theta_4)'$ in Figure 3. Figure 3(a) provides a plot of the estimated KL over the region $[0, 0.1]^2$ where, apparently, the majority of the region is flat and uninformative with a sharp spike around the true values at $\theta_1 = \theta_4 = 0.01$. We thus narrow the investigation down to a smaller region $[0, 0.02]^2$ in Figure 3(b). Again, the curvature in the estimated KL around the truth is quite steep. This may pose some issues for Metropolis-Hasting algorithms, since the majority of the prior region is uninformative and improper initialization could lead to extremely slow convergence.

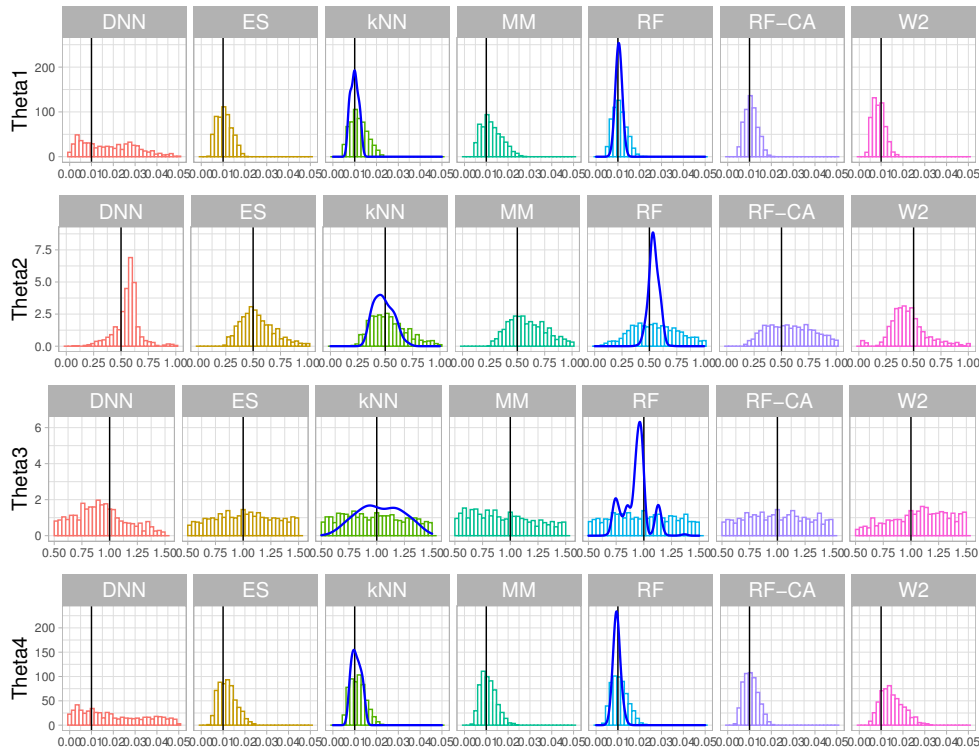


Figure 4: ABC posterior densities for the Lotka-Volterra model with $\theta_0 = (0.01, 0.5, 1, 0.01)'$. The black vertical lines mark the true parameter values. The blue curves in kNN and RF boxes represent the smoothed density calculated from the exponential kernel. The ABC posteriors are plotted from the top 1% out of 10^5 samples.

Previous ABC analyses of this model suggested various summary statistics including the mean, log-variance, autocorrelation (at lag 1 and 2) of each series as well as their cross-correlation (Papamakarios and Murray, 2016). For the discriminator of our method, we choose the ℓ_1 -penalized (LASSO) logistic regression classifier (LRD) with $m = n$ and with a penalty λ selected via 5-fold cross-validation (as implemented in the R package `glmnet`), as well as a random forest classifier (RF).

Similar to Kaji and Ročková (2022), we use an informative prior $\theta \in U(\Xi)$ with a restricted domain $\Xi = [0, 0.1] \times [0, 1] \times [0, 2] \times [0, 0.1]$ so that the computation is more economic and efficient. A typical snapshot of the ABC posteriors is plotted in Figure 4. From the plot, we can see that the difficulty in estimating different parameters varies a lot and θ_3 is the most challenging among all. While other methods give relatively flat posteriors for θ_3 , our method (RF) combined with the scaled exponential kernel identifies the correct location of the parameter with a much tighter posterior. For the DNN approach, the posteriors seem to be very flat and the estimation for θ_2 is biased. Considering its heavy computation costs shown in Table 7, we exclude this method in the repetition experiment. The average performance of ABC methods under this model is summarized in Table 2. We can see that RF with the exponential kernel gives the tightest CI most of the time, while maintaining relatively small estimation errors.

Method (scale)	$\theta_1 = 0.01$		$\theta_2 = 0.5$		$\theta_3 = 1$		$\theta_4 = 0.01$	
	$(\hat{\theta}_1 - \theta_1)^2$ (10^{-5})	95% CI width	$(\hat{\theta}_2 - \theta_2)^2$	95% CI width	$(\hat{\theta}_3 - \theta_3)^2$ (10^{-2})	95% CI width	$(\hat{\theta}_4 - \theta_4)^2$ (10^{-5})	95% CI width
RF	0.118	1.116	0.389	0.719	0.037	0.938	0.051	1.290
RF-exp	0.015	0.354	0.092	0.191	0.272	0.483	0.037	0.493
LRD	42.949	4.474	1.688	0.626	0.119	0.945	17.159	4.720
LRD-exp	0.066	0.354	0.247	0.285	0.405	0.643	0.021	0.626
KL	0.563	1.462	0.107	0.602	0.054	0.934	0.203	1.327
KL-exp	0.049	0.544 (0.9)	0.044	0.243 (0.9)	0.346	0.599 (0.9)	0.066	0.616 (0.9)
CA	36.477	4.375	1.602	0.629	0.118	0.945	15.684	4.622
RF-CA	0.141	1.061	0.699	0.700	0.042	0.933	0.056	1.181
ES	0.101	1.233	0.203	0.573	0.045	0.923	0.341	1.443
W2	1.265	1.666	0.478	0.599	0.011	0.917	4.080	2.051
SS	1.639	1.777	2.260	0.721	1.019	0.901	3.767	1.660
MM	0.824	1.727	0.959	0.601	0.414	0.932	0.074	1.294

Table 2: ABC performance evaluations on the Lotka-Volterra Model, averaged over 10 repetitions, with the top 1% ABC samples selected. Most of the 95% CIs have full coverage with the rest having their coverage marked in brackets. The bold fonts mark the best model in each metric.

5. Empirical Analysis

We further demonstrate our approach on the nontrivial problem of estimating stock volatility using merely daily observations on high, low and closing prices. All of these price observations are typically available to investors. We use a similar data generating process as in Magdon-Ismail and Atiya (2003), assuming that the assets follow a Brownian motion with a constant drift and volatility. In particular, suppose that the log-price processes $X_j(t), i = 1, \dots, d$, are correlated Brownian motions, that is $E[X_i(s)X_j(t)] = \sigma_{ij} \min\{s, t\}$, and that the joint movement of the log-price processes $\mathbf{X}(t) = (X_1(t), \dots, X_d(t))'$ follows a multivariate Brownian motion as

$$d\mathbf{X}(t) = \boldsymbol{\mu}dt + \Sigma d\mathbf{W}(t), \tag{22}$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_d)'$ and $\Sigma = [\sigma_{ij}]_{1 \leq i, j \leq d}$ denote the drift and the volatility of the log processes, respectively. We write

$$H_j = \max_{0 \leq t \leq 1} X_j(t), \quad L_j = \min_{0 \leq t \leq 1} X_j(t), \quad S_j = X_j(1),$$

for the high, low and final log price, respectively, over a fixed time interval $[0, 1]$. We want to estimate the drift $\boldsymbol{\mu}$ and the volatility matrix Σ merely from observing these three prices over a period of time.

We impose a normal-inverse-Wishart prior $(\boldsymbol{\mu}, \Sigma) \sim NIW(\boldsymbol{\mu}_0, \lambda, \Phi, \nu)$. This distribution can be sampled from in two steps: (1) sample Σ from an inverse Wishart distribution $\Sigma \mid \Phi, \nu \sim W^{-1}(\Phi, \nu)$; (2) sample $\boldsymbol{\mu}$ from a multivariate normal distribution $\boldsymbol{\mu} \mid \boldsymbol{\mu}_0, \lambda, \Sigma \sim N(\boldsymbol{\mu}_0, \frac{1}{\lambda}\Sigma)$. Since Σ is a semi-positive definite matrix, we model the parameters through its Cholesky root $\Sigma^{1/2}$. Without loss of generality, we only consider the case $W(0) = (0, \dots, 0)'$, since the closing prices on the previous day or the opening prices of today are usually known.

5.1 Synthetic Data

To compare various likelihood-free estimators, we first generate synthetic data for 1 000 trading days. For each day (of length $t = 1$), we simulate the Brownian motion using 500

time steps to obtain the high, low and closing price data for each particular window. We first restrict our attention to the case of just two assets, which leaves us with 5 parameters to estimate: μ_1, μ_2 and the upper triangular root of Σ , denoted with $L = \begin{bmatrix} l_{11} & 0 \\ l_{12} & l_{22} \end{bmatrix}$. We first illustrate how the covariance parameter σ_{12} impacts the co-movement of asset prices. Holding $\mu_1 = \mu_2 = 0$ and $\sigma_{11} = \sigma_{22} = 1$ fixed, we plot time series realizations of the closing prices for three particular choices of σ_{12} in Figure 5. The patterns are as expected where the prices tend to co-fluctuate when σ_{12} is closer to one. The success of our method depends on how well the discriminator can tell apart these trajectories.




Figure 5: Closing prices time series for three choices of σ_{12} ($\mu_1 = \mu_2 = 0, \sigma_{11} = \sigma_{22} = 1$).

For our simulation, we set $\mu = (0, 0)'$ and $\sigma_{11} = \sigma_{22} = 1, \sigma_{12} = 0.5$. We choose relatively non-informative prior hyper-parameters $\mu_0 = (0, 0)'$, $\lambda = 1, \Phi = I_d$ and $\nu = d$. Posterior distributions reconstructed with different ABC methods are given in Figure 6 and the averaged performance over 10 repetitions is summarized in Table 3. We explore two discriminators built on the prices and the quadratic/interaction terms of the prices: (1) a lasso classifier with penalty term λ selected from a 5-fold cross-validation (LRD); (2) a random forest classifier (RF). We find out that for the Lotka-Volterra model, the linear classifier performs better than the nonparametric random forest, as reflected in Figure 6. We observe a smaller estimation bias and the computation of LRD is much shorter than RF. Thus, we only include LRD in the repetitions. It is clear that our exponential kernel methods place more mass around the true location of the parameters, and that the shape of kNN-exp is slightly less regular than LRD-exp. Although LRD-exp induces a larger bias in estimating the drift (μ_1, μ_2) , it does a better job at capturing the correct location of the volatilities.

5.2 Real Data Analysis

Following the example in Rogers and Zhou (2008), we examine a small data set of stock prices focusing on two stocks: Boeing (BA) and Proctor & Gamble (PG). The prices were obtained from NYSE (Yahoo Finance), starting from 3rd January 2011 and consisting of 1 000 trading days. Since the off-market trades follow a different mechanism than the market trade, we only model price changes from the opening price to the closing price each day

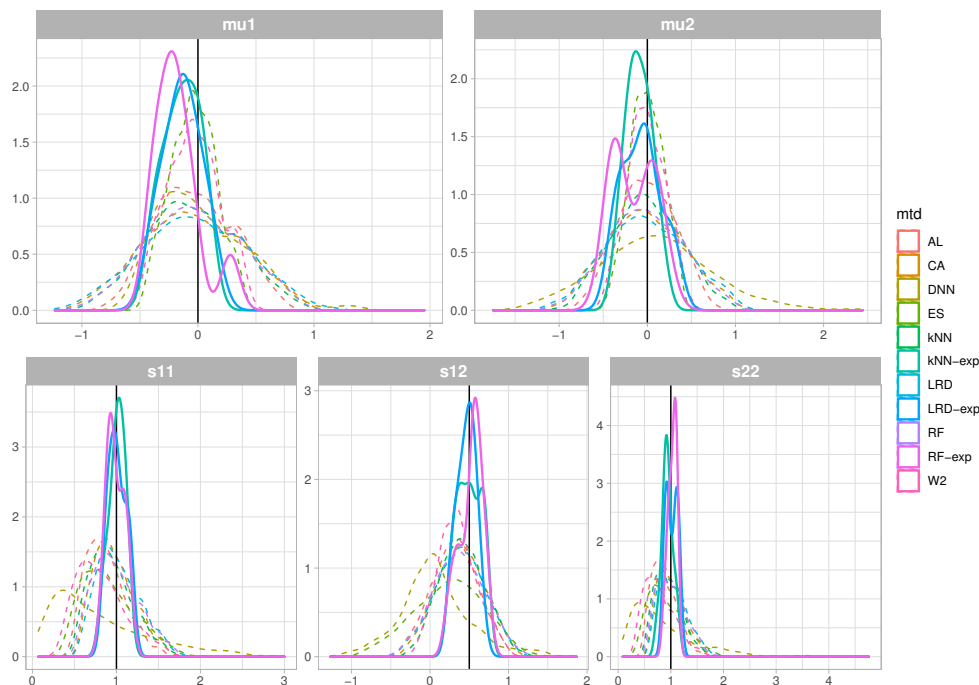


Figure 6: Posterior densities on simulated volatility data ($d = 2$)

		LRD	LRD-exp	kNN	kNN-exp	ES	CA	W2	AL	SS	DNN
$\mu_1 = 0$	MSE($\times 10^{-4}$)	2.004	2.113	0.638	0.159	1.484	1.497	0.170	0.628	1.216	1.838
	95% CI width	1.760	0.331	1.480	0.442	0.740	1.660	0.831	1.240	1.734	3.079
$\mu_2 = 0$	MSE($\times 10^{-4}$)	0.734	7.439	0.363	0.259	0.050	0.142	0.277	0.089	0.021	7.029
	95% CI width	1.628	0.431	1.435	0.538	0.737	1.564	0.813	1.247	1.569	2.527
$\sigma_{11} = 1$	MSE($\times 10^{-3}$)	0.311	0.111	0.167	0.460	0.834	0.002	2.634	1.669	1.404	3.502
	95% CI width	0.963	0.371	0.954	0.329	1.231	0.896	1.140	0.910	0.763	6.233
$\sigma_{12} = 0.5$	MSE($\times 10^{-3}$)	1.035	0.103	0.427	0.107	6.657	0.861	1.528	3.337	0.109	18.427
	95% CI width	1.193	0.355	1.092	0.324	1.700	1.122	0.988	1.114	0.628	5.530
$\sigma_{22} = 1$	MSE($\times 10^{-3}$)	0.416	0.304	0.124	0.018	0.776	0.007	4.244	1.679	3.687	1.882
	95% CI width	0.959	0.252	0.974	0.314	1.274	0.920	1.157	0.922	0.783	6.693

Table 3: Performance on the stock volatility estimation example, averaged over 10 repetitions, with top 1% selected. All 95% CIs have full coverage of the true parameters. The bold fonts mark the best model in each row.

where the log price differences $X(t)$ are all computed based on the opening prices of that day.

We observe that the fluctuations in the prices are much smaller than in our simulated time series and we thereby choose the hyperparameters based on the mean and covariance of the closing prices. Figure 7 gives the ABC posterior distributions estimated from different methods and the corresponding summaries of the distributions are provided in Table 4. We again observe that our methods with the exponential kernels return much narrower posterior distributions. The estimates on the volatilities are quite close except for the ABC

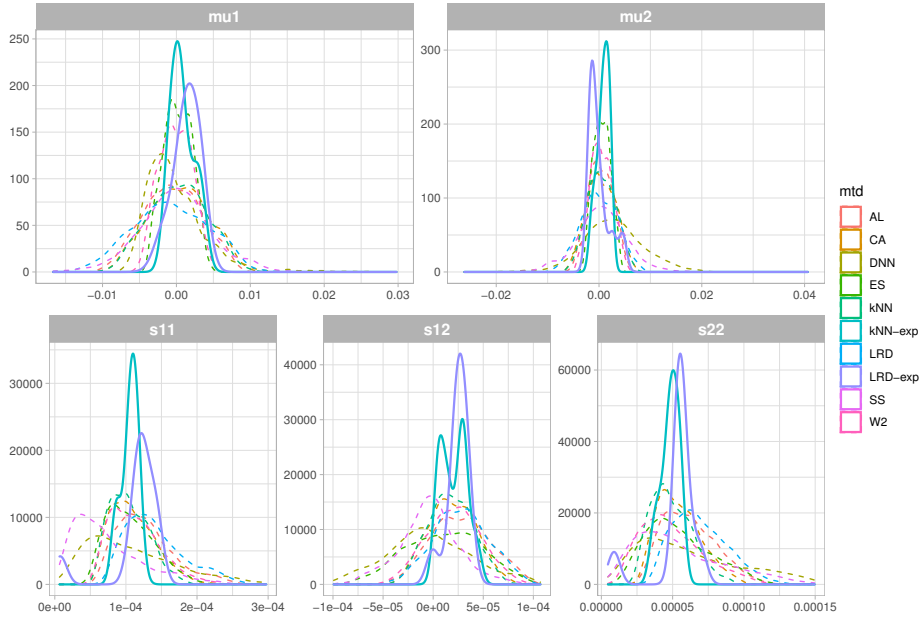


Figure 7: Posterior densities estimated from log-prices of BA and PG

with summary statistics (SS). The drifts of both BA and PG are not significantly different from zero.

	LRD	LRD-exp	kNN	kNN-exp	CA	ES	W2	AL	DNN	SS	
$\mu_1 (\times 10^{-3})$	$\bar{\mu}_1$	-0.774	1.354	0.250	0.779	0.342	0.173	-0.126	-0.500	-0.091	-0.071
	l	-10.380	-1.981	-6.816	-1.145	-7.099	-3.201	-3.987	-7.603	-5.415	-9.532
	u	8.202	3.514	7.229	3.385	7.395	3.443	3.736	6.810	9.335	10.121
$\mu_2 (\times 10^{-3})$	$\bar{\mu}_2$	0.009	-0.039	0.491	0.822	0.445	0.545	0.453	0.316	2.793	0.450
	l	-6.591	-1.718	-4.329	-1.406	-4.604	-2.364	-3.491	-4.103	-8.913	-9.831
	u	6.531	4.563	5.503	2.166	5.538	3.616	4.450	4.690	14.701	10.171
$\sigma_{11} (\times 10^{-4})$	$\bar{\sigma}_{11}$	1.406	1.142	1.049	1.054	1.101	1.070	1.130	1.273	1.224	0.738
	l	0.824	1.142	0.633	0.844	0.634	0.577	0.626	0.689	0.243	0.155
	u	2.362	1.142	1.640	1.202	1.819	1.818	1.938	2.259	3.616	1.976
$\sigma_{12} (\times 10^{-5})$	$\bar{\sigma}_{12}$	2.811	2.446	1.767	2.084	1.906	1.074	2.052	2.759	-0.530	-0.376
	l	-2.379	2.446	-2.359	2.084	-2.390	-6.152	-3.137	-2.656	-25.243	-6.660
	u	8.226	2.446	6.037	2.084	6.442	8.008	7.253	8.568	23.343	6.065
$\sigma_{22} (\times 10^{-4})$	$\bar{\sigma}_{22}$	6.922	5.263	4.924	4.820	5.470	5.419	4.947	6.179	18.485	5.521
	l	4.069	5.263	2.766	4.820	3.149	2.005	1.734	3.322	1.571	1.326
	u	11.049	5.263	8.020	4.820	8.826	10.258	9.844	10.223	93.218	12.852

Table 4: Posterior estimates on analysis of BA and PG. For each parameter, we report three summary statistics, the posterior means, the lower limit of the 95% CI intervals (l) and the upper limit of the 95% CI intervals (u).

6. Discussion

This paper develops an ABC variant using a classification-based KL estimator as a discrepancy measure. By deploying a flexible classifier, the empirical KL divergence can be

estimated with a vanishing error. In addition, inspired by the connection between the KL divergence and the log-likelihood ratio, we propose a scaled exponential kernel to aggregate ABC samples. This smoothing variant avoids the need for choosing the ad hoc threshold ϵ_n and fully utilizes information returned from all ABC samples. Under mild conditions, we show that the posterior concentration rate of the accept-reject ABC depends on the estimation error δ_n and the accept-reject threshold ϵ_n , while the rate of the smooth version depends on the estimation error δ_n and the contraction rate of the actual posterior distribution. Our methodology can also be related to many other likelihood-free inference methods, including ABC with Classification Accuracy (Gutmann et al., 2018), Wasserstein distance ABC (Bernton et al., 2019), and Generalized Posteriors (Schmon et al., 2020). Our methods coincide with the c -posterior (Miller and Dunson, 2018), which is robust to perturbations. In particular, the accept-reject ABC can be shown to be robust under model misspecification (see Section 3.4). In addition, the exponential kernel can be motivated as an instantiation of General Bayesian Inference (GBI) (Bissiri et al., 2016) with the KL estimator as the loss function. See Thomas and Corander (2019) and Thomas et al. (2020) for examples of conducting robust inference using probabilistic classifiers under the generalized Bayes update setup. Along with our theoretical investigations, we demonstrate competitive performance of our methods on benchmark examples. Our theoretical analysis provides theoretical justifications for the method of Gutmann et al. (2018).

Acknowledgements

The authors gratefully acknowledge the support from the James S. Kemper Faculty Fund, the Liew Family Junior Faculty Fellowship, and the Richard Rosett Faculty Fellowship at the University of Chicago Booth School of Business and the National Science Foundation (Grant No. NSF DMS-1944740).

References

- Mattias Akesson, Prashant Singh, Fredrik Wrede, and Andreas Hellander. Convolutional neural networks as summary statistics for approximate Bayesian computation. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2021.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, 2017.
- Benedikt Bauer and Michael Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47(4):2261–2285, 2019.
- Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- Espen Bernton, Pierre E Jacob, Mathieu Gerber, and Christian P Robert. Approximate Bayesian computation with the Wasserstein distance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(2):235–269, 2019.

- Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, Berlin, Heidelberg, 2006.
- Pier Giovanni Bissiri, Chris C Holmes, and Stephen G Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 78(5):1103, 2016.
- Michael GB Blum. Regression approaches for ABC. In *Handbook of Approximate Bayesian Computation*, pages 71–85. Chapman and Hall/CRC, 2018.
- Michael GB Blum, Maria Antonieta Nunes, Dennis Prangle, and Scott A Sisson. A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science*, 28(2):189–208, 2013.
- Rainer Burkard, Mauro Dell’Amico, and Silvano Martello. *Assignment problems: revised reprint*. SIAM, 2012.
- Thomas M Cover and Joy A Thomas. Entropy, relative entropy and mutual information. *Elements of Information Theory*, 2(1):12–13, 1991.
- Frank Den Hollander. *Large deviations*, volume 14. American Mathematical Soc., 2008.
- Peter J Diggle and Richard J Gratton. Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 46(2):193–212, 1984.
- Christopher Drovandi and David T Frazier. A comparison of likelihood-free methods with and without summary statistics. *Statistics and Computing*, 32(3):1–23, 2022.
- Christopher C Drovandi, Anthony N Pettitt, and Malcolm J Faddy. Approximate Bayesian computation using indirect inference. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 60(3):317–337, 2011.
- Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate Bayesian computation: Semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):419–474, 2012.
- David T Frazier, Gael M Martin, Christian P Robert, and Judith Rousseau. Asymptotic properties of approximate Bayesian computation. *Biometrika*, 105(3):593–607, 2018.
- David T Frazier, Christian Robert, and Judith Rousseau. Model misspecification in ABC: Consequences and diagnostics. *Journal of the Royal Statistical Society: Series B*, 2019.
- Masahiro Fujisawa, Takeshi Teshima, Issei Sato, and Masashi Sugiyama. γ -abc: Outlier-robust approximate Bayesian computation based on a robust divergence estimator. In *International Conference on Artificial Intelligence and Statistics*, pages 1783–1791. PMLR, 2021.
- Sandesh Ghimire, Prashna K Gyawali, and Linwei Wang. Reliable estimation of Kullback-Leibler divergence by controlling discriminator complexity in the reproducing kernel Hilbert space. *arXiv preprint arXiv:2002.11187*, 2020.

- Subhashis Ghosal and Aad Van der Vaart. *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press, 2017.
- Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- Michael U Gutmann, Ritabrata Dutta, Samuel Kaski, and Jukka Corander. Likelihood-free inference via classification. *Statistics and Computing*, 28(2):411–425, 2018.
- Bai Jiang, Tung-yu Wu, Charles Zheng, and Wing H Wong. Learning summary statistic for approximate Bayesian computation via deep neural network. *Statistica Sinica*, pages 1595–1618, 2017.
- Bai Jiang, Tung-Yu Wu, and Wing Hung Wong. Approximate Bayesian computation with Kullback-Leibler divergence as data discrepancy. In *International Conference on Artificial Intelligence and Statistics*, pages 1711–1721. PMLR, 2018.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Paul Joyce and Paul Marjoram. Approximately sufficient statistics and Bayesian computation. *Statistical applications in Genetics and Molecular Biology*, 7(1), 2008.
- Tetsuya Kaji and Veronika Ročková. Metropolis–Hastings via classification. *Journal of the American Statistical Association*, pages 1–15, 2022.
- Tetsuya Kaji, Elena Manresa, and Guillaume Pouliot. An adversarial approach to structural estimation. *arXiv preprint arXiv:2007.06169*, 2020.
- Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Bas JK Kleijn and Aad W van der Vaart. Misspecification in infinite-dimensional Bayesian statistics. *The Annals of Statistics*, 34(2):837–877, 2006.
- Bas JK Kleijn and Aad W van der Vaart. The Bernstein-von-Mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.
- Solomon Kullback. *Information Theory and Statistics*. Wiley, New York, 1958.
- Malik Magdon-Ismail and Amir F Atiya. A maximum likelihood approach to volatility estimation for a Brownian motion using high, low and close price data. *Quantitative Finance*, 3(5):376, 2003.
- Jeffrey W Miller and David B Dunson. Robust Bayesian inference via coarsening. *Journal of the American Statistical Association*, 2018.

- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, pages 429–443, 1997.
- Hien Duy Nguyen, Julyan Arbel, Hongliang Lü, and Florence Forbes. Approximate Bayesian computation via the energy statistic. *IEEE Access*, 8:131683–131698, 2020.
- XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, pages 1089–1096, 2007.
- Matthew A Nunes and David J Balding. On optimal selection of summary statistics for approximate Bayesian computation. *Statistical applications in Genetics and Molecular Biology*, 9(1), 2010.
- George Papamakarios and Iain Murray. Fast ε -free inference of simulation models with Bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*, pages 1028–1036, 2016.
- George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848. PMLR, 2019.
- Mijung Park, Wittawat Jitkrittum, and Dino Sejdinovic. K2-ABC: Approximate Bayesian computation with kernel embeddings. In *Artificial Intelligence and Statistics*, pages 398–407. PMLR, 2016.
- Fernando Pérez-Cruz. Kullback-Leibler divergence estimation of continuous distributions. In *2008 IEEE international Symposium on Information Theory*, pages 1666–1670. IEEE, 2008.
- Nicholas G Polson and Veronika Rockova. Posterior concentration for sparse deep learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 938–949, 2018.
- Jonathan K Pritchard, Mark T Seielstad, Anna Perez-Lezaun, and Marcus W Feldman. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 1999.
- Leonard CG Rogers and Fanyin Zhou. Estimating correlation from high, low, opening and closing prices. *The Annals of Applied Probability*, 18(2):813–823, 2008.
- Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with ReLU activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.
- Sebastian M Schmon, Patrick W Cannon, and Jeremias Knoblauch. Generalized posteriors in approximate Bayesian computation. In *Third Symposium on Advances in Approximate Bayesian Inference*, 2020.

- Lorraine Schwartz. On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 4(1):10–26, 1965.
- Jorge Silva and Shrikanth Narayanan. Universal consistency of data-driven partitions for divergence estimation. In *2007 IEEE International Symposium on Information Theory*, pages 2021–2025. IEEE, 2007.
- Jorge Silva and Shrikanth S Narayanan. Information divergence estimation based on data-dependent partitions. *Journal of Statistical Planning and Inference*, 140(11):3180–3198, 2010.
- Scott A Sisson, Yanan Fan, and Mark A Beaumont. Overview of ABC. In *Handbook of Approximate Bayesian Computation*, pages 3–54. Chapman and Hall/CRC, 2018.
- Helle Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a survey. *International Statistical Review*, 72(3):337–354, 2004.
- Owen Thomas and Jukka Corander. Diagnosing model misspecification and performing generalized Bayes’ updates via probabilistic classifiers. *arXiv preprint arXiv:1912.05810*, 2019.
- Owen Thomas, Henri Pesonen, and Jukka Corander. Generalised Bayes updates with f -divergences through probabilistic classifiers. *arXiv preprint arXiv:2007.04358*, 2020.
- Owen Thomas, Ritabrata Dutta, Jukka Corander, Samuel Kaski, and Michael U Gutmann. Likelihood-free inference by ratio estimation. *Bayesian Analysis*, 17(1):1–31, 2022.
- Sara A van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- Aad W van der Vaart and Jon A Wellner. *Weak convergence and empirical processes: with applications to statistics*. Springer, 1996.
- Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Martin J Wainwright and Michael Irwin Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- Qing Wang, Sanjeev R Kulkarni, and Sergio Verdú. Divergence estimation of continuous distributions based on data-dependent partitions. *IEEE Transactions on Information Theory*, 51(9):3064–3074, 2005.
- Qing Wang, Sanjeev R Kulkarni, and Sergio Verdú. Divergence estimation for multidimensional densities via k -nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55(5):2392–2405, 2009.
- Yuexi Wang and Veronika Rockova. Uncertainty quantification for sparse deep learning. In *International Conference on Artificial Intelligence and Statistics*, pages 298–308. PMLR, 2020.
- Darren J Wilkinson. *Stochastic modeling for systems biology*. CRC press, 2018.

Richard David Wilkinson. Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. *Statistical applications in Genetics and Molecular Biology*, 12(2):129–141, 2013.

Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.

Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.

Puning Zhao and Lifeng Lai. Analysis of k nearest neighbor KL divergence estimation for continuous distributions. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2562–2567. IEEE, 2020.

Appendix A. Convergence Rate of Estimation Errors with Neural Network Sieves

To develop a precise definition of the low underlying dimension d^* described in Remark 6, we borrow the smoothness notion of Bauer and Kohler (2019).

Definition 1 ((p, C)-Smoothness) Let $p = q + s$ for some $q \in \mathbb{N}_0$ and $0 < s \leq 1$. A function $m : \mathbb{R}^d \rightarrow \mathbb{R}$ is called (p, C)-smooth if for every $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}_0^d$ with $\sum_{j=1}^d \alpha_j = q$, the partial derivative $\frac{\partial^q m}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}}$ exists and satisfies

$$\left| \frac{\partial^q m}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}}(x) - \frac{\partial^q m}{\partial x_1^{\alpha_1} \dots \partial x_d^{\alpha_d}}(z) \right| \leq C \|x - z\|^s$$

for every $x, z \in \mathbb{R}^d$ where $\|\cdot\|$ denotes the Euclidean norm.

With this, the nested composition structure is defined as follows.

Definition 2 (Generalized Hierarchical Interaction Model) Let $d \in \mathbb{N}$, $d^* \in \{1, \dots, d\}$, and $m : \mathbb{R}^d \rightarrow \mathbb{R}$. We say that the function m satisfies a generalized hierarchical interaction model of order d^* and level 0, if there exist $a_1 \in \mathbb{R}^d, \dots, a_{d^*} \in \mathbb{R}^d$, and $f : \mathbb{R}^{d^*} \rightarrow \mathbb{R}$ such that

$$m(x) = f(a'_1 x, \dots, a'_{d^*} x)$$

for every $x \in \mathbb{R}^d$. We say that m satisfies a generalized hierarchical interaction model of order d^* and level $l + 1$ with K components if there exist $g_k : \mathbb{R}^{d^*} \rightarrow \mathbb{R}$ and $f_{1,k}, \dots, f_{d^*,k} : \mathbb{R}^d \rightarrow \mathbb{R}$ ($k = 1, \dots, K$) such that $f_{1,k}, \dots, f_{d^*,k}$ ($k = 1, \dots, K$) satisfy a generalized hierarchical model of order d^* and level l and

$$m(x) = \sum_{k=1}^K g_k(f_{1,k}(x), \dots, f_{d^*,k}(x))$$

for every $x \in \mathbb{R}^d$. We say that the generalized hierarchical interaction model is (p, C)-smooth if all functions occurring in its definition are (p, C)-smooth.

For example, a conditional binary choice model satisfies a generalized hierarchical interaction model of order $d^* \leq 3$ and level 0, irrespectively of the dimension of the covariates.

Example 1 (Binary Choice Model) Let $y_i = \mathbb{I}\{x'_i \alpha + \varepsilon_i > 0\}$, $\varepsilon_i \sim P_\varepsilon$, be the true DGP and $y_i^\beta = \mathbb{I}\{x'_i \beta + \tilde{\varepsilon}_i > 0\}$, $\tilde{\varepsilon}_i \sim \tilde{P}_\varepsilon$, be the generative model. Then,

$$\log \frac{p_0(y, x)}{p_\theta(y, x)} = y \log \frac{1 - P_\varepsilon(-x' \alpha)}{1 - \tilde{P}_\varepsilon(-x' \beta)} + (1 - y) \log \frac{P_\varepsilon(-x' \alpha)}{\tilde{P}_\varepsilon(-x' \beta)}.$$

Therefore, we can write this as $f(a'_1 z, a'_2 z, a'_3 z)$ where $z = (y, x')'$, $a_1 = (1, 0, \dots, 0)'$, $a_2 = (0, -\alpha)'$, $a_3 = (0, -\beta)'$, and $f(y, x_1, x_2) = y[\log(1 - P_\varepsilon(x_1)) - \log(1 - \tilde{P}_\varepsilon(x_2))] + (1 - y)[\log P_\varepsilon(x_1) - \log \tilde{P}_\varepsilon(x_2)]$. If P_ε and \tilde{P}_ε are logistic distributions, then f is (p, C)-smooth for $(p, C) = (8, 1)$ or for $(p, C) = (11, 10)$, for example. Therefore, it satisfies the generalized hierarchical interaction model of order $d^* = 3$ and level $l = 0$.

Example 2 (Diffusion Process) Let $X_{i,t/d}$, $t = 1, \dots, d$ be discretely sampled observations of a Brownian motion $dX_{it} = \mu dt + \sigma dW_{it}$ with $X_{i0} = 0$, where W_{it} is a standard Brownian motion independent across i . Let the generative model be $X_{i,t/d}^\theta$, $t = 1, \dots, d$ from $X_{it}^\theta = mdt + s d\tilde{W}_{it}$ and $\theta = (m, s)$. Then, the log likelihood ratio is

$$\begin{aligned} \log \frac{p_0(x_{1/d}, \dots, x_1)}{p_\theta(x_{1/d}, \dots, x_1)} &= d \log \frac{s}{\sigma} - \frac{d}{2} \left(\frac{1}{\sigma^2} - \frac{1}{s^2} \right) \sum_{j=1}^d (x_{j/d} - x_{(j-1)/d})^2 \\ &\quad + \left(\frac{\mu}{\sigma^2} - \frac{m}{s^2} \right) x_1 - \frac{1}{2} \left(\frac{\mu^2}{\sigma^2} - \frac{m^2}{s^2} \right). \end{aligned}$$

Letting $z = ((x_{1/d} - x_0)^2, \dots, (x_1 - x_{(d-1)/d})^2, x_1)'$, we can write this as $f(a'z)$ where $a = (-\frac{d}{2}(\frac{1}{\sigma^2} - \frac{1}{s^2}), \dots, -\frac{d}{2}(\frac{1}{\sigma^2} - \frac{1}{s^2}), \frac{\mu}{\sigma^2} - \frac{m}{s^2})'$ and $f(y) = d \log \frac{s}{\sigma} + y - \frac{1}{2}(\frac{\mu^2}{\sigma^2} - \frac{m^2}{s^2})$. Then, f is (p, C) -smooth with $(p, C) = (\infty, 1)$, and the log likelihood ratio satisfies the hierarchical model with $d^* = 1$ and level $l = 0$.

Next, we define the configuration of the neural network appropriate for estimating a generalized hierarchical interaction model.

Definition 3 (Hierarchical Neural Network) Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a q -admissible activation function. For $M^* \in \mathbb{N}$, $d \in \mathbb{N}$, $d^* \in \{1, \dots, d\}$, and $\alpha > 0$, let $\mathcal{F}_{M^*, d^*, d, \alpha}$ be the class of functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$f(x) = \sum_{i=1}^{M^*} \mu_i \sigma \left(\sum_{j=1}^{4d^*} \lambda_{i,j} \sigma \left(\sum_{v=1}^d \theta_{i,j,v} x_v + \theta_{i,j,0} \right) + \lambda_{i,0} \right) + \mu_0$$

for some $\mu_i, \lambda_{i,j}, \theta_{i,j,v} \in \mathbb{R}$, where $|\mu_i| \leq \alpha$, $|\lambda_{i,j}| \leq \alpha$, and $|\theta_{i,j,v}| \leq \alpha$. For $l = 0$, define the set of neural networks with two hidden layers by $\mathcal{H}_{M^*, d^*, d, \alpha}^{(0)} = \mathcal{F}_{M^*, d^*, d, \alpha}$; for $l > 0$, define the set of neural networks with $2l + 2$ hidden layers by

$$\mathcal{H}_{M^*, d^*, d, \alpha}^{(l)} = \left\{ h : \mathbb{R}^d \rightarrow \mathbb{R} : h(x) = \sum_{k=1}^K g_k(f_{1,k}(x), \dots, f_{d^*,k}(x)), g_k \in \mathcal{F}_{M^*, d^*, d^*, \alpha}, f_{j,k} \in \mathcal{H}^{(l-1)} \right\}.$$

With these definitions, the convergence rate of the neural network sieve is characterized as follows. If the log likelihood ratio satisfies the generalized hierarchical model of order d^* and the corresponding neural network sieve is used, the convergence rate of the discriminator depends only on d^* , not on d . The precise assumption is formulated as follows.

Assumption 6 (Neural Network Discriminator) Let P_0 and P_θ have subexponential tails and finite first moments.⁴ Let $\log(p_0/p_\theta)$ satisfy a (p, C) -smooth generalized hierarchical interaction model of order d^* and finite level l with K components for $p = q + s$, $q \in \mathbb{N}_0$,

4. We say that P on \mathbb{R}^d has subexponential tails if $\log P(\|X\|_\infty > a) \lesssim -a$ for large a .

and $s \in (0, 1]$. Let $\mathcal{H}_{M^*, d^*, d, \alpha}^{(l)}$ be the class of neural networks with the Lipschitz activation function with Lipschitz constant 1 for

$$M_* = \left[\binom{d^* + q}{d^*} (q + 1) \left(\left[\frac{(\log \delta_n)^{2(2q+3)}}{\delta_n} \right]^{\frac{1}{p}} + 1 \right)^{d^*} \right],$$

$$\alpha = \left[\frac{(\log \delta_n)^{2(2q+3)}}{\delta_n} \right]^{\frac{d^* + p(2q+3) + 1}{p}} \frac{\log n}{\delta_n^2},$$

and $\delta_n = [(\log n)^{\frac{p+2d^*(2q+3)}{p}} / n]^{\frac{p}{2p+d^*}}$. Denote by $\mathcal{D}_n = \{\Lambda(f) : f \in \mathcal{H}_{M^*, d^*, d, \alpha}^{(l)}\}$ the sieve of neural network discriminators for the standard logistic cdf Λ .

Assumption 6 gives a sufficient condition for the entropy condition in Assumption 1. With this, we obtain the “classification counterpart” of Bauer and Kohler (2019, Theorem 1) as below.

Theorem 11 (Kaji et al., 2020, Proposition S.3) *Suppose Assumption 6 holds. If n/m converges and an estimator $\hat{D}_{n,m}^\theta$ exists that satisfies $\mathbb{M}_{n,m}(\hat{D}_{n,m}^\theta) \geq \mathbb{M}_{n,m}^\theta(D_\theta) - O_P(\delta_n^2)$ for the δ_n in Assumption 6, then $d_\theta(\hat{D}_{n,m}^\theta, D_\theta) = O_P^*(\delta_n)$.*

This theorem, combined with the explicit expression of δ_n in Assumption 6, tells that if $d^* < 2p$, we have $\delta_n = o_P(n^{-1/4})$, which is often the desired rate for the nonparametric estimator of a nuisance parameter. For the binary choice model in Example 1, the discriminator converges much faster than $n^{-1/4}$.

Example 1 (Continued) *The binary choice model with logistic errors satisfies a (p, C) -smooth hierarchical model with $(p, C) = (8, 1)$ and $d^* = 3$. We can substitute these numbers (and $q = p - 1$) into δ_n in Assumption 6 and obtain $\delta_n = \left(\frac{(\log n)^{13.75}}{n}\right)^{8/19} \lesssim n^{-2/5}$.*

Example 2 (Continued) *The discretely sampled Brownian motion model satisfies a (p, C) -smooth hierarchical model with $d^* = 1$, $C = 1$, and arbitrarily large p . Therefore, δ_n can be arbitrarily close to $n^{-1/2}$, however large the sampling frequency is.*

When there is no low-dimensional structure, Definition 2 simply reduces to the smoothness of m and d^* is equal to d . Therefore, the convergence rate in Theorem 11 reduces to the traditionally proven rate that deteriorates quickly with d .

We note that the hierarchical structure is not the only way to proving the superior adaptivity of a neural network discriminator. Similar results can possibly be deduced with other smoothness assumptions and network configurations, such as Schmidt-Hieber (2020) and Yarotsky (2017).

Appendix B. Frequentist’s Analysis on the Exponential Kernel

We thus study the concentration in terms of a KL neighborhood around $Q_{\theta^*}^{(n)}$ defined as

$$B(\epsilon, Q_{\theta^*}^{(n)}; P_0^{(n)}) = \{Q_{\theta^*}^{(n)} \in \mathcal{Q}^{(n)} : \tilde{K}(\theta^*, \theta) \leq n\epsilon^2, \tilde{V}(\theta^*, \theta) \leq n\epsilon^2\}, \quad (23)$$

where $\tilde{K}(\theta^*, \theta) \equiv P_0^{(n)} \log \frac{q_{\theta^*}^{(n)}}{q_{\theta}^{(n)}}$ and $\tilde{V}(\theta^*, \theta) \equiv P_0^{(n)} \left| \log \frac{q_{\theta^*}^{(n)}}{q_{\theta}^{(n)}} - \tilde{K}(\theta^*, \theta) \right|^2$.

The following corollary is directly adopted from Theorem 5.1 of Kaji and Ročková (2022).

Corollary 12 *Denote with $\tilde{Q}_{\theta}^{(n)}$ a measure defined through $d\tilde{Q}_{\theta}^{(n)} = (p_0^{(n)}/q_{\theta^*}^{(n)})dP_{\theta}^{(n)}$ and let $d(\cdot, \cdot)$ be a semi-metric on $\mathcal{P}^{(n)}$. Suppose that there exists a sequence $\epsilon_n > 0$ satisfying $\epsilon_n \rightarrow 0$ and $n\epsilon_n^2 \rightarrow \infty$ such that for every $\epsilon > \epsilon_n$ there exists a test ϕ_n (depending on ϵ) such that for every $J \in \mathbb{N}_0$*

$$P_0^{(n)} \phi_n \lesssim e^{-n\epsilon^2/4} \quad \text{and} \quad \sup_{Q_{\theta}^{(n)}: d(Q_{\theta}^{(n)}, P_{\theta^*}^{(n)}) > J\epsilon} \tilde{Q}_{\theta}^{(n)}(1 - \phi_n) \leq e^{-nJ^2\epsilon^2/4}. \quad (24)$$

Let $B(\epsilon, Q_{\theta^*}^{(n)}; P_0^{(n)})$ be as in (23) and let $\tilde{\Pi}_n(\theta)$ be a prior distribution with a density $\tilde{\pi}(\theta) \propto C_{\theta}\pi(\theta)$ with C_{θ} as in (20). Assume that there exists a constant $L > 0$ such that, for all n and $j \in \mathbb{N}$,

$$\frac{\tilde{\Pi}_n \left(\theta \in \Theta : j\epsilon_n < d(Q_{\theta}^{(n)}, P_{\theta^*}^{(n)}) \leq (j+1)\epsilon_n \right)}{\tilde{\Pi}_n \left(B(\epsilon_n, Q_{\theta^*}^{(n)}; P_0^{(n)}) \right)} \leq e^{n\epsilon_n^2 j^2 / 8}. \quad (25)$$

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

$$P_0^{(n)} \Pi^{EK} \left(Q_{\theta}^{(n)} : d(Q_{\theta}^{(n)}, P_{\theta^*}^{(n)}) \geq M\epsilon_n \mid X^{(n)} \right) \rightarrow 0. \quad (26)$$

Next, we want to show the shape of the posterior is actually asymptotically gaussian around θ^* . The following corollary follows from Theorem 2.1 of Kleijn and van der Vaart (2012) and Lemma 8.1 of Kaji and Ročková (2022).

Corollary 13 *(Bernstein von-Mises) Assume that the posterior (20) concentrates around θ^* at the rate ϵ_n^* and that for every compact $K \in \mathbb{R}^d$*

$$\sup_{h \in K} \left| \log \frac{q_{\theta^* + \epsilon_n^* h}^{(n)}(\mathbf{X})}{q_{\theta^*}^{(n)}(\mathbf{X})} - h' \tilde{V}_{\theta^*} \tilde{\delta}_{n, \theta^*} - \frac{1}{2} h' \tilde{V}_{\theta^*} h \right| \rightarrow 0 \quad (27)$$

for some random vector $\tilde{\delta}_{n, \theta^*}$ and a non-singular matrix \tilde{V}_{θ^*} . Then the approximated posterior $\Pi^{EK}(\cdot)$ converges to a sequence of normal distributions in total variation at the rate ϵ_n^* , i.e.

$$\sup_B \left| \Pi^{EK} \left(\epsilon_n^{*-1}(\theta - \theta^*) \in B \mid \mathbf{X} \right) - N_{\tilde{\delta}_{n, \theta^*}, \tilde{V}_{\theta^*}^{-1}(B)} \right| \rightarrow 0 \quad \text{in } P_0^{(n)} \text{ - probability.} \quad (28)$$

Appendix C. ABC versus MHC

One may wonder how the performance of the exponentially weighted ABC (Algorithm 2) differs from the sequential MHC version Kaji and Ročková (2022). The target distribution is the same but one may expect convergence issues of the sequential version. The major difference are: (1) ABC does not need a proposal distribution for state changes; (2) ABC

is parallelizable while MHC can only be computed in a sequential fashion; (3) ABC is less sensitive to how the θ 's are initialized. For the scenarios where the likelihood is spiky (as in Figure 3), priors can be adjusted for ABC focus on the most promising region (Blum, 2018). Meanwhile, the performance of MHC can be largely influenced by the initializations. When the acceptance ratio is low (or the effective sample size is small), tuning can be difficult.

We use the Lotka-Volterra model to illustrate how MHC could suffer when the initialization point lands in a non-informative region. We randomly sampled a few different initialization points from the prior, and run the MHC for 100 000 iterations with the same proposal distribution. The posterior means, effective sample sizes and the acceptance rates are reported in Table 5 and trace-plots are provided in Figure 8.

	%accept		θ_1	θ_2	θ_3	θ_4
#1	11.934%	initial	0.041	0.197	1.411	0.009
		posterior mean	0.013	0.463	0.900	0.011
		ESS	3.483	21.114	109.692	36.124
#2	0.013%	initial	0.047	0.508	0.709	0.050
		posterior mean	0.042	0.663	0.846	0.047
		ESS	2.980	2.474	4.074	2.508
#3	0.019%	initial	0.016	0.355	0.761	0.042
		posterior mean	0.016	0.417	0.582	0.041
		ESS	3.537	49.821	6.926	49.189
#4	0.035%	initial	0.031	0.710	0.777	0.033
		posterior mean	0.026	0.477	1.183	0.030
		ESS	2.012	1.861	1.069	4.323
#5	14.828%	initial	0.022	0.208	1.398	0.012
		posterior mean	0.010	0.527	0.935	0.010
		ESS	45.167	123.381	236.331	164.153

Table 5: MHC from different initializations for the Lotka-Volterra model with $\theta_0 = (0.01, 0.5, 1, 0.01)'$. For each case, we report the initializations (sampled randomly from the prior), the acceptance ratio and the effective sample sizes (ESS) over a chain of 100 000 iterations, the posterior mean with the first 10 000 iterations treated as burn-ins.

Obviously, the performance and mixing speed of MHC is largely contingent on the starting point for this model. Even for the case #1, whose acceptance ratio is relatively high, it took more than 25 000 iterations for the chain to slowly navigate itself towards the informative region, and the effective sample sizes are still small since the MHC samples are highly correlated. For the case #3, the convergence of the chain seems hopeless after 100 000 iterations. The ABC method, on the contrary, is robust to the initializations since the parameters are sampled completely independently identically from the prior, as shown in the repetition experiment in Table 2.



Figure 8: Trace plots of MHC from different initializations. The tests 1, 2 and 5 are corresponding to the cases #1, #2 and #5 in Table 5. The black solid horizontal lines mark the locations of the true parameters.

Appendix D. Other GAN-style Discrepancy Metrics

There are many other divergences that can be estimated using different configurations of (6). Given a set \mathcal{F} of functions from \mathcal{X} to \mathbb{R} , we can define

$$d_{\mathcal{F}}(p_0, p_{\theta}) = \sup_{f \in \mathcal{F}} \mathbb{P}_n f(X_i) - \mathbb{P}_m^{\theta} f(\tilde{X}_i^{\theta}), \quad (29)$$

as the Integral Probability Metric (IPM) (Müller, 1997) associated with the function class \mathcal{F} . Arjovsky et al. (2017) show that different choices of the function class \mathcal{F} can dramatically change the topology, as we discuss below, and the regularity of $d_{\mathcal{F}}(p_0, p_{\theta})$ as a loss function. Several widely used discrepancy metrics can be obtained from IPMs. For example, the 1-Wasserstein distance or the Earth-Mover distance (Villani, 2009) are obtained when \mathcal{F} is the set of 1-Lipschitz functions. Arjovsky et al. (2017) proposed Wasserstein GANs, which approximate the 1-Wasserstein distance with gradient-clipped neural networks. Another example is the Total Variation (TV) distance, obtained when \mathcal{F} is the set of all measurable functions bounded between -1 and 1. Next, the Maximum Mean (MM) discrepancy is obtained when $\mathcal{F} = \{f \in \mathcal{H} : \|f\|_{\infty} \leq 1\}$ for \mathcal{H} some Reproducing Kernel Hilbert Space (RKHS) associated with a given kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Our classification-

based GAN framework thus provides an alternative route towards implementing ABCs with the Wasserstein distance (Bernton et al., 2019) as well as MM (Park et al., 2016).

One important aspect of these various metrics is that they induce topologies of different strengths. A divergence makes it easier for a sequence of distributions to converge when it induces a weaker topology. Arjovsky et al. (2017) show that KL induces the strongest one, followed by JS and TV, with Wasserstein distance being the weakest. However, the computation costs of the Wasserstein distance are high, i.e. $O((n+m)^3 \log(n+m))$ to calculate it exactly (Burkard et al., 2012). The GAN version approximation also requires weight clipping to enforce the Lipschitz constraint which can be problematic to implement. Our KL estimator, on the other hand, is more scalable and places no constraints on the discriminator. The neural network discriminator can converge faster than traditional nonparametric methods if some low-dimensionality assumption is satisfied (Kaji et al., 2020; Bauer and Kohler, 2019). In addition, the KL estimator has a natural interpretation through its link to the log-likelihood ratio. The KL divergence appears in many theoretical results in Bayesian non-parametrics. For example, Schwartz’s Theorem (Schwartz, 1965) requires that, for posterior consistency, the prior should assign positive probability to any KL neighborhood of p_0 . By adopting the KL estimator within our framework, we can relate ABC to classical posterior convergence rate results.

Appendix E. Simulations: Univariate g -and- k Distributions

Another classical example in the ABC literature is the univariate g -and- k distribution. It is defined implicitly by its inverse distribution function

$$F^{-1}(x) = A + B \left[1 + 0.8 \frac{1 - e^{-gz_x}}{1 + e^{-gz_x}} \right] (1 + z_x^2)^k z_x,$$

where z_x is the x -th quantile of the standard normal distribution, and parameters A, B, g, k are related to location, scale, skewness and kurtosis. The probability density function has no analytical form but can be numerically calculated with high precision since it only involves one-dimensional inversions and differentiations of the quantile function. Therefore, Bayesian inference can be carried out.

We generate $n = 500$ observations from the model by letting $A = 3, B = 1, g = 2, k = 0.5$. Among the four parameters, g is the hardest to identify, as observed in previous analyses. From a pilot run, narrow the range of our uniform priors to $A \sim U([2, 4]), B \sim U([0, 0.25]), g \sim U([0, 10])$ and $k \sim U([0, 2.5])$.

For our methods, two classifiers are explored for the model, both built on observed data \mathbf{X} and the higher-order terms \mathbf{X}^2 and \mathbf{X}^3 .⁵ The first one is a random forests (RF) classifier, and the second one is a neural network (NND) with two hidden layers with 10 nodes each. The first layer is activated with the rectified linear unit (ReLU) function, while the second layer is equipped with a hyperbolic tangent function (tanh).

We compare the posteriors obtained from different ABC methods in Figure 9. We can see that, indeed, the parameter g is the most difficult one to estimate. While the majority of methods return a flat posterior, KLC and kNN with the exponential kernel still place

5. We are trying to manually create more features for this univariate data. We also explore even higher order terms and find the including up to \mathbf{X}^3 works the best in practice.

most of the posterior mass around the true value. AL and SA misidentify the location of one or two parameters. While CA and KLC give relatively similar ranking on the ABC sampled parameters, KLC with the exponential kernel can utilize all the samples and re-adjust the shape according to weights. In addition, RF-exp returns more spiky posteriors than kNN-exp. This is also reflected in the performance summaries in Table 6 which show that RF-exp has a narrower CI than kNN-exp on average. The NND method may not ideal for this univariate g -and- k distribution.

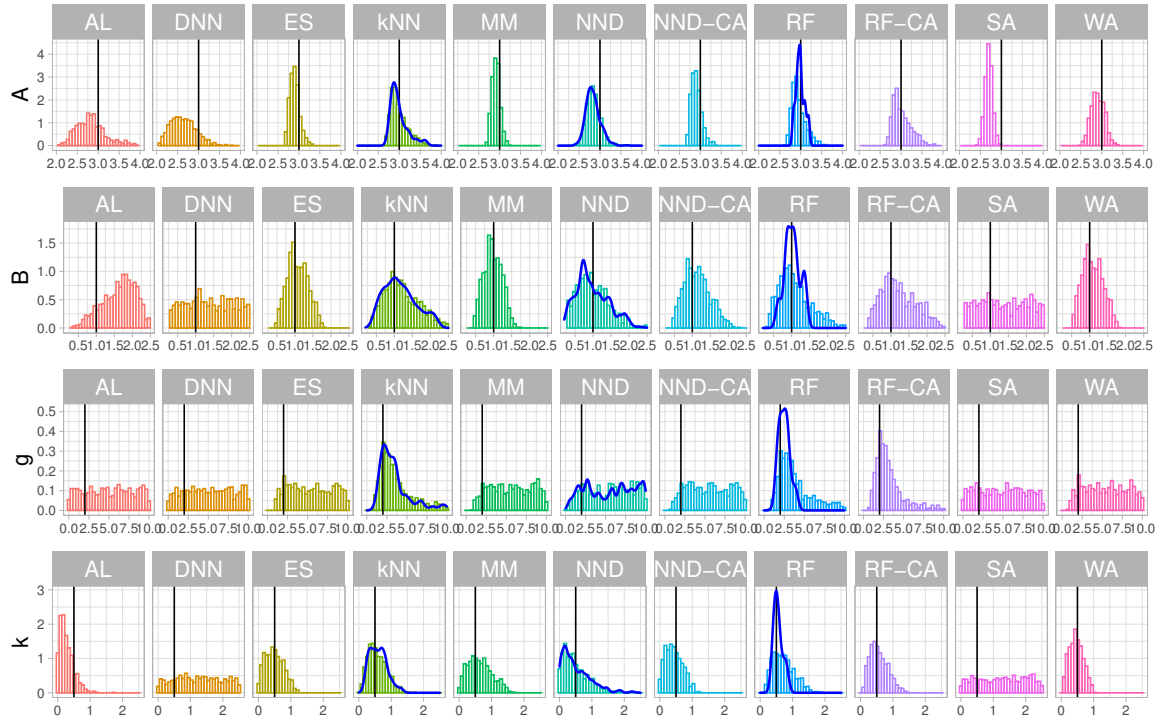


Figure 9: Posterior densities for the g -and- k distributions ($A = 3, B = 1, g = 2, k = 0.5$). The black vertical lines mark the true parameter values. The blue curves in kNN, NND, RF boxes represent the weighted density calculated from the exponential kernel. The ABC posteriors were plotted with the top 1% of 10^5 samples.

Appendix F. Discriminator Calibration

We provide more intuition for how the choice of the discriminator, including the m/n ratio and the number of sets of latent variables needed to approximate the KL divergence, affects our results. We also discuss possible modifications we make to the KL estimation to eliminate bias. We use the M/G/1-queueing model as an example.

First, we want to investigate how the choice of the discriminator and the m/n ratio could impact the estimation of the KL divergence. We consider three discriminators: (1) logistic regression on degree-2 polynomial terms of \mathbf{X} (LRD); (2) a neural network with one layer of 10 ReLU nodes and two layers of 10 tanh nodes with input \mathbf{X} (NND1); (3)

Method	$A = 3$		$B = 1$		$g = 2$		$k = 0.5$	
	$(\hat{A} - A)^2$	95% CI width	$(\hat{B} - B)^2$	95% CI width	$(\hat{g} - g)^2$	95% CI width	$(\hat{k} - k)^2$	95 %CI width
RF	0.0036	0.758	0.0316	1.699	2.545	7.748	0.02550	1.170
RF-exp	0.0016	0.394	0.0281	0.913	0.131	1.373	0.00541	0.486 (0.9)
NND	0.0143	0.535	0.0369	1.528	12.840	8.356	0.00341	1.044
NND-exp	0.0156	0.455	0.0339	1.229	12.516	8.191	0.00800	0.784
kNN	0.0024	0.843	0.0486	1.770	2.695	7.558	0.00421	1.086
kNN-exp	0.0032	0.610	0.0410	1.333	1.673	4.109	0.00480	0.805
AL	0.0176	1.325	0.2211	1.593	8.704	9.480	0.04451	0.835
NND-CA	0.0125	0.503	0.0466	1.568	13.077	8.162	0.00574	1.080
RF-CA	0.0016	0.843	0.0414	1.719	1.293	7.184	0.00286	1.022
SA	0.1191	0.424 (0)	0.0598	2.379	8.800	9.478	0.56399	2.379
WA	0.0145	0.627	0.0240	1.034	11.621	8.566	0.01060	0.776
DNN	0.0560	1.166	0.0622	2.373	8.976	9.456	0.56306	2.374
ES	0.0140	0.487	0.0154	1.023	11.015	8.497	0.00298	0.957
MM	0.0120	0.386	0.0102	0.916	13.379	8.340	0.01459	1.248

Table 6: Performance summaries on g -and- k distributions over 10 repetitions with top 1% ABC samples selected. Most of the 95% CIs have full coverage with the rest having their coverage marked after the CI width. The bold fonts mark the best model in each metric.

a neural network with one layer of 10 ReLU nodes and one layer of 10 tanh nodes with input as the degree-2 polynomial terms of \mathbf{X} (NND2). For each discriminator, we do a grid search over $m/n \in \{1, 2, 3, 5\}$. We plot the estimated KL divergence (rescaled) in Figure 10. From the plot we can see it is quite obvious that LRD best identifies the location of the true values, and the fluctuations/variance in its estimations are the smallest. It also worth noting that the difficulties in estimating different parameters are significantly different here. All discriminators are able to learn the location of θ_3 with little variance, but only LRD is able to learn θ_1 and θ_2 well. The m/n ratio does not seem to have significant impact for LRD, but we observe for NND1, NND2 and some other models that the variance in the estimation tends to get smaller when $m/n > 1$ but not overly large to avoid bias. We continue our experiments with LRD and $m/n = 3$.

Next, we want to know how many different sets of latent variables are needed to control the variance in estimation. We focus on the best discriminator we have previously found (i.e. LRD with $m/n = 3$). We include the cases where $nlatent \in \{1, 5, 10\}$ here. The results are shown in Figure 11. It is expected that with more sets of latent variables to average over, the variance in estimation gets smaller. The main challenge for the M/G/1 model lies in estimating θ_1 correctly and stably. With only one or five sets of latent variable, it seems hopeless to identify θ_1 . To trade off variance and computation costs, we set $nlatent = 10$ in our experiments.

Appendix G. Computational Complexities

We list the computation complexity and the computation time (in hours) of the three simulation examples in Table 7. The computation complexities are considered in the multivariate settings, and the costs of ES, WA, MM, kNN are adopted from Jiang et al. (2018); Nguyen et al. (2020). For logistic regression, RF and DNN, the cost depends on how fast the iterations actually converge and we provide only a rough approximation here. From the three simulation examples, g -and- k examples and $m/g/1$ -queuing examples has small d but large

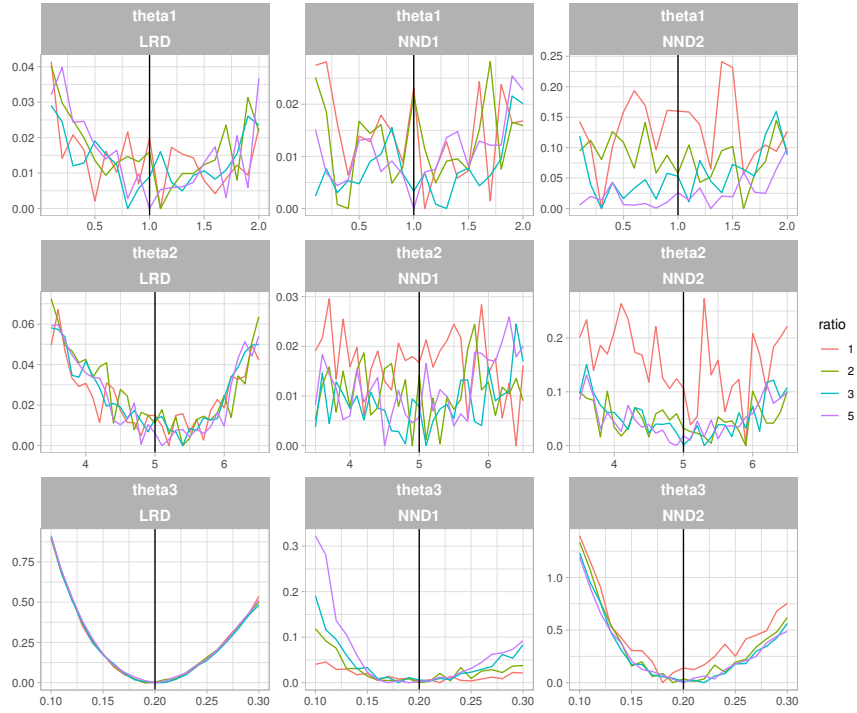


Figure 10: Estimated KL divergence (rescaled) under different discriminators ($n_{\text{latent}}=10$). We plot the marginal changes in estimated KL divergence with respect to different parameters and different m/n ratios here. The estimated values are rescaled by subtracting the minimum in each scenario so the changes are more comparable.




Figure 11: Estimated KL divergence (rescaled) using LRD and $m/n = 3$. We plot the marginal changes in estimated KL divergence with respect to different parameters and different number of sets of latent variables. The estimated values are rescaled by subtracting the minimum in each scenario so the changes are more comparable.

n , while the Lotka-Volterra example is large in d but small in n . We see that conventional summary statistics calculation usually costs least amount of time. AL is also fast but it becomes almost infeasible when the dimension of X , i.e., d , is large and calculating the

density function of the multivariate normals has numerical issues. For the KL estimators, kNN is easiest to compute and it has the computation advantage when d is small, but both its accuracy and complexity suffer when d is large. Though our classifier costs longer time on examples with small d , the computation costs are relatively similar for the LV example whose dimension d is large.

	Complexity	Computation Time (in hours)		
		g-and-k	m/g/1	LV
kNN	$\mathcal{O}((n+m)\log(n+m))$	0.017	0.150	11.618
Logistic	$\mathcal{O}(n+m)$	\	12.052	12.214
NN	$\mathcal{O}((n+m) \times \#weight \times \#epoch)$	28.624	\	\
RF	$\mathcal{O}((n+m)\log(n+m))$	4.926	8.268	12.382
WA	$\mathcal{O}((n+m)^2)$	0.009	0.706	11.706
MM	$\mathcal{O}((n+m)^2)$	10.009	24.617	11.882
ES	$\mathcal{O}((n+m)^2)$	0.048	0.180	12.192
SA/SS		0.961	1.986	2.517
DNN	$\mathcal{O}(10^6 \times 10^6 \times \Theta \times \#epoch)$	1.968	2.378	34.556
AL		0.009	0.120	\

Table 7: Computational Complexities/Runtime Comparisons

Appendix H. Proofs

H.1 Proof of Theorem 5

Denote $K_n = \mathbb{P}_n \log \frac{p_0}{p_\theta}$. Using Chebyshev's inequality and Assumption 2 as

$$\begin{aligned}
 P_0^{(n)}[|K_n - K(p_0, p_\theta)| > u] &= P_0^{(n)}\left(\left|(\mathbb{P}_n - P_0) \log \frac{p_0}{p_\theta}\right| > u\right) \leq \frac{1}{u^2} P_0^{(n)}\left[\left|(\mathbb{P}_n - P_0) \log \frac{p_0}{p_\theta}\right|^2\right] \\
 &= \frac{1}{u^2} P_0^{(n)}\left[\left|\mathbb{P}_n\left(\log \frac{p_0}{p_\theta}\right) - P_0 \log \frac{p_0}{p_\theta}\right|^2\right] = \frac{1}{nu^2} P_0\left[\left|\log \frac{p_0}{p_\theta} - P_0 \log \frac{p_0}{p_\theta}\right|^2\right] \\
 &\leq \frac{16(2 + \Lambda)h(p_0, p_\theta)}{nu^2},
 \end{aligned}$$

where the last inequality follows from Lemma 2.1 (iii) of Kaji and Ročková (2022). Next, note

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K_n = -\mathbb{P}_n\left(\log \frac{1 - \hat{D}_{n,m}^\theta}{1 - D_\theta} - \log \frac{\hat{D}_{n,m}^\theta}{D_\theta}\right).$$

Recall the outer expectation P^* in Theorem 2 and the set of classifiers $D_{n,\delta}^\theta$ defined in Assumption 1, we can bound

$$\begin{aligned} P\left(\left|\mathbb{P}_n\left(\log\frac{1-\hat{D}_{n,m}^\theta}{1-D_\theta}-\log\frac{\hat{D}_{n,m}^\theta}{D_\theta}\right)\right|>u, d_\theta(\hat{D}_{n,m}^\theta, D_\theta)\leq C_n\delta_n\right) \\ \leq P^*\left(\sup_{D\in\mathcal{D}_{C_n\delta_n}^\theta}\left|\mathbb{P}_n\left(\log\frac{1-D}{1-D_\theta}-\log\frac{D}{D_\theta}\right)\right|>u\right) \\ \leq\frac{1}{u}\mathbb{E}^*\sup_{D\in\mathcal{D}_{C_n\delta_n}^\theta}\left|\mathbb{P}_n\left(\log\frac{1-D}{1-D_\theta}-\log\frac{D}{D_\theta}\right)\right| \end{aligned}$$

by Markov's inequality. The proof of Theorem 4.1 of Kaji and Ročková (2022) shows that the expectation is $O(C_n\delta_n)$. Using the triangle inequality and the Bonferroni inequality, since $h(p_0, p_\theta)\leq 2$, we can then write

$$\begin{aligned} P\left(\left|\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)-K(p_0, p_\theta)\right|>2u, d_\theta(\hat{D}_{n,m}^\theta, D^\theta)\leq C_n\delta_n\right) \\ \leq P\left(\left|\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)-K_n\right|+|K_n-K(p_0, p_\theta)|>2u, d_\theta(\hat{D}_{n,m}^\theta, D_\theta)\leq C_n\delta_n\right) \\ \leq P\left(\left|\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)-K_n\right|>u, d_\theta(\hat{D}_{n,m}^\theta, D_\theta)\leq C_n\delta_n\right)+P_0^{(n)}[|K_n-K(p_0, p_\theta)|>u] \\ \leq O\left(\frac{C_n\delta_n}{u}\right)+\frac{32(2+\Lambda)}{nu^2}. \end{aligned}$$

H.2 Proof of Theorem 7

Throughout, we continue to assume that $\tilde{\mathbf{X}}^\theta = g_\theta(\tilde{\mathbf{X}})$ and we denote with $P = P_0^{(n)} \otimes \tilde{P}^{(m)}$ the joint measure for $(\mathbf{X}, \tilde{\mathbf{X}})$. Below, we will be using the notation $\Pi(\cdot)$ to denote the generic probability, i.e. for $(\theta, \tilde{\mathbf{X}})$ or for the conditional probability θ given $\tilde{\mathbf{X}}$. Later, we will define a high-probability event $\Omega_n(C, \varepsilon_n)$ such that $P_0^{(n)}[\Omega_n(C, \varepsilon_n)^c] = o(1)$ for some $C \in (0, 1)$. Given $\delta_n > 0$ from our assumptions, we can write for every $\lambda_n > 0$ and $\varepsilon_n > 0$ and for every arbitrarily slowly increasing sequence $C_n > 0$

$$\begin{aligned} P_0^{(n)}\Pi\left(K(p_0, p_\theta)>\lambda_n\left|\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)\leq\varepsilon_n\right.\right)\leq\Pi_1+o(1)+ \\ P_0^{(n)}\Pi\left(K(p_0, p_\theta)>\lambda_n\left|\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta)\leq\varepsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta)\leq C_n\delta_n\right.\right)\mathbb{I}[\Omega_n(C, \varepsilon_n)], \quad (30) \end{aligned}$$

where, using Lemma 2 and the fact that the rate δ_n is uniform (see Remark 4),

$$\Pi_1 \equiv P_0^{(n)}\Pi(d(\hat{D}_{n,m}^\theta, D_\theta)>C_n\delta_n)\leq\sup_{\theta\in\Theta}P(d(\hat{D}_{n,m}^\theta, D_\theta)>C_n\delta_n)=o(1).$$

Consider the joint event, for some $\delta' > 0$,

$$A_{\varepsilon_n}(\delta') = \{(\tilde{\mathbf{X}}, \theta) : \hat{K}(\mathbf{X}, g_\theta(\tilde{\mathbf{X}})) \leq \varepsilon_n\} \cap \{K(p_0, p_\theta) > \delta'\}.$$

For every $(\tilde{\mathbf{X}}, \theta) \in A_{\varepsilon_n}(\delta')$ we have

$$K(p_0, p_\theta) \leq \hat{K}(\mathbf{X}, g_\theta(\tilde{\mathbf{X}})) + \left|\hat{K}(\mathbf{X}, g_\theta(\tilde{\mathbf{X}})) - K(p_0, p_\theta)\right| \leq \varepsilon_n + \left|\hat{K}(\mathbf{X}, g_\theta(\tilde{\mathbf{X}})) - K(p_0, p_\theta)\right|.$$

Hence $(\tilde{\mathbf{X}}, \theta) \in A_{\epsilon_n}(\delta')$ implies that

$$\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \delta' - \epsilon_n,$$

and choosing $\delta' \geq \epsilon_n + t_\epsilon$ leads to

$$\Pi[A_{\epsilon_n}(\delta')] \leq \int_{\Theta} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > t_\epsilon \right] d\Pi(\theta).$$

Using (30), we now focus on the conditional probability, given $d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n$,

$$\begin{aligned} & \Pi \left(K(p_0, p_\theta) > \epsilon_n + t_\epsilon \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right) \\ & \leq \frac{\int_{\Theta} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > t_\epsilon \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta)}{\int_{\Theta} \tilde{P}^{(m)} \left[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta)}. \end{aligned} \quad (31)$$

We now find a lower bound for the denominator. Recall the KL neighborhood $B_2(p_0, \epsilon_n)$ defined in (13). Since

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq K(p_0, p_\theta) + \left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| \leq \epsilon_n/2 + K(p_0, p_\theta),$$

provided that $\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| \leq \epsilon_n/2$. The denominator can be then bounded by

$$\begin{aligned} & \int_{\Theta} \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\Pi(\theta) \\ & \geq \int_{B_2(p_0, \epsilon_n/2)} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| \leq \epsilon_n/2 \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta) \\ & \geq \Pi[B_2(p_0, \epsilon_n/2)] - \int_{B_2(p_0, \epsilon_n/2)} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n/2 \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta). \end{aligned}$$

Denoting

$$Z(\mathbf{X}) \equiv \int_{B_2(p_0, \epsilon_n/2)} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n/2 \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta)$$

we can write, for every $C > 0$, using Fubini's theorem and Markov's inequality

$$\begin{aligned} P_0^{(n)}(Z(\mathbf{X}) > C) & \leq \frac{1}{C} \int_{B_2(p_0, \epsilon_n/2)} P \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n/2 \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta) \\ & = \frac{\Pi(B_2(p_0, \epsilon_n/2)) \sup_{\theta \in \Theta} P \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n/2, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right]}{C \sup_{\theta \in \Theta} P \left[d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right]} \\ & = \frac{\Pi(B_2(p_0, \epsilon_n/2))}{C(1 + o(1))} \sup_{\theta \in \Theta} \rho_{n,\theta}(\epsilon_n/2; C_n; \delta_n), \end{aligned}$$

where we have used Theorem 5 and the fact that $P(d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n) = 1 + o(1)$ for every $\theta \in \Theta$ from Lemma 2. We now define an event, for some $0 < C < 1$ and $\epsilon_n > 0$,

$$\Omega_n(C, \epsilon_n) = \{\mathbf{X} : Z(\mathbf{X}) \leq C \times \Pi(B_2(p_0, \epsilon_n/2))\}.$$

Using Theorem 5, we have

$$\rho_{n,\theta}(\epsilon_n/2; C_n; \delta_n) = O\left(\frac{C_n \delta_n}{\epsilon_n} + \frac{1}{n \epsilon_n^2}\right) \quad \text{for every } \theta \in \Theta.$$

Choosing $\epsilon_n > 0$ such that $\epsilon_n = o(1)$ and $n \epsilon_n^2 \rightarrow \infty$ and $C_n \delta_n = o(\epsilon_n)$ we have $P_0^{(n)}[\Omega_n(C, \epsilon_n)^c] = o(1)$ for every $C \in (0, 1)$. On the event $\Omega_n(C, \epsilon_n)$, for some $0 < C < 1$, we can lower-bound the denominator with

$$\int_{\Theta} \tilde{P}^{(m)}[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\Pi(\theta) > (1 - C) \times \Pi(B_2(p_0, \epsilon_n/2)).$$

Using this bound and applying Fubini's theorem, we can further write

$$\begin{aligned} P_0^{(n)} \Pi \left(K(p_0, p_\theta) > \epsilon_n + t_\epsilon \mid \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right) \mathbb{I}[\Omega_n(C, \epsilon_n)] \\ \leq \frac{\int_{\Theta} P \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > t_\epsilon \mid d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\Pi(\theta)}{(1 - C) \times \Pi(B_2(p_0, \epsilon_n/2))} \end{aligned} \quad (32)$$

Using Theorem 5 again, we obtain an upper bound for the display above with

$$\frac{\rho_n(t_\epsilon; C_n; \delta_n)}{(1 - C) \times \Pi(B_2(p_0, \epsilon_n/2))(1 + o(1))}.$$

Using the prior Assumption 3, we can choose t_ϵ such that

$$\left(\frac{C_n \delta_n}{t_\epsilon} + \frac{1}{n t_\epsilon^2} \right) / \epsilon_n^\kappa = 1/M_n$$

for some arbitrarily slowly increasing sequence $M_n > 0$. We choose $t_\epsilon = M_n C_n \delta_n / \epsilon_n^\kappa + \sqrt{M_n} n^{-1/2} / \epsilon_n^{\kappa/2}$. Since $\delta_n \gtrsim n^{-1/2}$ and $\epsilon_n^{-\kappa} \geq \epsilon_n^{-\kappa/2}$, the overall rate is then driven by $\epsilon_n + t_\epsilon = \epsilon_n + \widetilde{M}_n \delta_n \epsilon_n^{-\kappa}$, where $\widetilde{M}_n = M_n C_n$.

H.3 Proof of Theorem 9

Recall from Theorem 7 that with the accept-reject strategy, the true KL divergence $K(p_0, p_\theta)$ is contracting at the rate $\lambda_n = \epsilon_n + \widetilde{M}_n \delta_n \epsilon_n^{-\kappa}$, where \widetilde{M}_n is a slowly increasing sequence that diverges faster than C_n . Consider the case when $\epsilon_n \gg \widetilde{M}_n \delta_n \epsilon_n^{-\kappa}$ or, equivalently, $\epsilon_n \gg \delta_n^{1/(\kappa+1)}$. Denote

$$x(\theta) = \epsilon_n^{-1} K(p_0, p_\theta) \quad \text{and} \quad f_n(\theta - \theta_0) = f(\epsilon_n^{-1/2}(\theta - \theta_0)).$$

We express the ABC posterior expectation of $f_n(\theta - \theta_0)$ for a non-negative and bounded function $f_n(\cdot)$ by

$$\begin{aligned}
 P_0^{(n)} E_{\hat{\Pi}_{\epsilon_n}^{AR}} [f_n(\theta - \theta_0)] &= P_0^{(n)} \int f_n(\theta - \theta_0) d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X}) \\
 &= \underbrace{P_0^{(n)} \int f_n(\theta - \theta_0) \mathbb{I}[K(p_0, p_\theta) \leq \lambda_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X})}_{\text{(I)}} \\
 &\quad + \underbrace{P_0^{(n)} \int f_n(\theta - \theta_0) \mathbb{I}[K(p_0, p_\theta) \leq \lambda_n, d(\hat{D}_{n,m}^\theta, D_\theta) > C_n \delta_n] d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X})}_{\text{(II)}} \\
 &\quad + \underbrace{P_0^{(n)} \int f_n(\theta - \theta_0) \mathbb{I}[K(p_0, p_\theta) > \lambda_n] d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X})}_{\text{(III)}}
 \end{aligned}$$

where the term (III) can be controlled using Fubini's theorem and the concentration result in Theorem 7 as follows

$$\begin{aligned}
 \text{(III)} &= \int f_n(\theta - \theta_0) P_0^{(n)} \mathbb{I}[K(p_0, p_\theta) > \lambda_n] d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X}) \\
 &\leq \|f\|_\infty P_0^{(n)} \Pi(K(p_0, p_\theta) > \lambda_n | \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n) = o(1).
 \end{aligned}$$

The second term (II) can be bounded similarly as

$$\begin{aligned}
 \text{(II)} &= P_0^{(n)} \int_{K(p_0, p_\theta) \leq \lambda_n} f_n(\theta - \theta_0) \mathbb{I}[d(\hat{D}_{n,m}^\theta, D_\theta) > C_n \delta_n] d\hat{\Pi}_{\epsilon_n}^{AR}(\theta | \mathbf{X}) \\
 &\leq \|f\|_\infty \int_{K(p_0, p_\theta) \leq \lambda_n} P(d(\hat{D}_{n,m}^\theta, D_\theta) > C_n \delta_n) d\Pi(\theta) \\
 &\leq \|f\|_\infty \sup_{\theta \in \Theta} P(d(\hat{D}_{n,m}^\theta, D_\theta) > C_n \delta_n) = o(1)
 \end{aligned}$$

where we use the fact that $\sup_{\theta \in \Theta} P(d(\hat{D}_{n,m}^\theta, D_\theta) > C_n \delta_n) = o(1)$ from Theorem 2.

Thus, the asymptotic behavior is mainly determined by the term (I). Combined with the continuity of $\pi(\theta)$ at θ_0 we can re-write (I) as

$$\begin{aligned}
 \text{(I)} &= P_0^{(n)} \frac{\int_{K(p_0, p_\theta) \leq \lambda_n} \pi(\theta) f_n(\theta - \theta_0) \tilde{P}^{(m)}[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta}{\int_{K(p_0, p_\theta) \leq \lambda_n} \pi(\theta) \tilde{P}^{(m)}[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta} \\
 &= P_0^{(n)} \frac{\int_{K(p_0, p_\theta) \leq \lambda_n} f_n(\theta - \theta_0) \tilde{P}^{(m)}[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta}{\int_{K(p_0, p_\theta) \leq \lambda_n} \tilde{P}^{(m)}[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta} (1 + o(1))
 \end{aligned}$$

We have $x(\theta) \geq 0$ for all $\theta \in \Theta$ and since

$$\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) = \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) + \epsilon_n x(\theta),$$

we can write

$$\begin{aligned}
 & \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] \\
 & \geq \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| \leq \epsilon_n(1 - x(\theta)), d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] \\
 & = 1 - \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n(1 - x(\theta)), d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right]. \quad (33)
 \end{aligned}$$

Denoting with

$$\tilde{Z}(\mathbf{X}) = \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} \tilde{P}^{(m)} \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n(1 - x(\theta)), d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\theta.$$

Using Markov's inequality and Fubini's theorem we have, for every $\tilde{C} > 0$,

$$\begin{aligned}
 & P_0^{(n)}(\tilde{Z}(\mathbf{X}) > \tilde{C}) \\
 & \leq \frac{1}{\tilde{C}} \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} P \left[\left| \hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) - K(p_0, p_\theta) \right| > \epsilon_n(1 - x(\theta)), d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n \right] d\theta \\
 & \leq \frac{1}{\tilde{C}} \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} \rho_n(\epsilon_n(1 - x(\theta)); C_n; \delta_n) d\theta \\
 & \leq \frac{1}{\tilde{C}} \times \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta \times \sup_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} \rho_n(\epsilon_n(1 - x(\theta)); C_n; \delta_n) \\
 & \leq \frac{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta}{\tilde{C}} \times \rho_n(M_n C_n \delta_n; C_n; \delta_n) \\
 & = \frac{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta}{\tilde{C}} \times O\left(\frac{1}{M_n} + \frac{1}{n M_n^2 C_n^2 \delta_n^2}\right).
 \end{aligned}$$

where $\rho_n(\cdot; C_n; \delta_n)$ is defined in Theorem 5. Thus, we can define a set $\tilde{\Omega}_n(\tilde{C}_n)$, for some for some arbitrarily slowly increasing sequence $\tilde{C}_n > 0$, and $\tilde{C}_n = O(1/M_n)$, as

$$\tilde{\Omega}_n(\tilde{C}_n) = \left\{ \mathbf{X} : \tilde{Z}(\mathbf{X}) \leq \tilde{C}_n \times \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta \right\}.$$

Then we have that $P_0^{(n)}(\tilde{\Omega}_n(\tilde{C}_n)^c) = o(1)$. Recall the inequality in (33), on $\tilde{\Omega}_n(\tilde{C}_n)$, we have that

$$\begin{aligned}
 & \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta \\
 & \geq \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta - \tilde{Z}(\mathbf{X}) = \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta \times (1 - o(1))
 \end{aligned}$$

And this quantity is upper-bounded by $\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta$. Therefore, we can conclude that

$$\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta = \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta (1 + o(1)).$$

Note that $x(\theta) \leq 1 - M_n C_n \delta_n / \epsilon_n$ implies that $K(p_0, p_\theta) \leq \epsilon_n - M_n C_n \delta_n$. On this set $\tilde{\Omega}_n(\tilde{C}_n)$, we can further lower bound the denominator as

$$\begin{aligned} & \int_{K(p_0, p_\theta) \leq \lambda_n} \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta \\ &= \underbrace{\int_{K(p_0, p_\theta) \leq \epsilon_n - M_n C_n \delta_n} d\theta (1 + o(1))}_{D_1} \\ &+ \underbrace{\int_{\epsilon_n - M_n C_n \delta_n < K(p_0, p_\theta) \leq \lambda_n} \tilde{P}^{(m)} [\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta}_{D_2}. \end{aligned}$$

Next, we show that the second term D_2 is $o(D_1)$. Let $u = \epsilon_n^{-1/2}(\theta - \theta_0)$. Under Assumption 5, we have $K(p_0, p_\theta) = \frac{1}{2}(\theta - \theta_0)' I(\theta_0) (\theta - \theta_0) \{1 + o(1)\}$, which is $x(\theta_0 + \sqrt{\epsilon_n} u) = \frac{1}{2} u' I(\theta_0) u$. Since $I(\theta_0)$ is positive definite we can write

$$\frac{D_2}{D_1} \leq \frac{\int_{1 - M_n C_n \delta_n / \epsilon_n < x(\theta) \leq 1 + \tilde{M}_n \delta_n / \epsilon_n^{\kappa+1}} d\theta}{\int_{x(\theta) \leq 1 - M_n C_n \delta_n / \epsilon_n} d\theta} \quad (34)$$

$$\leq \frac{\int_{2(1 - M_n C_n \delta_n / \epsilon_n) \leq u' I(\theta_0) u \leq 2(1 + \tilde{M}_n \delta_n / \epsilon_n^{\kappa+1})} d\theta}{\int_{u' I(\theta_0) u \leq 2(1 - M_n C_n \delta_n / \epsilon_n)} d\theta} \lesssim \frac{M_n C_n \delta_n}{\epsilon_n} + \frac{\tilde{M}_n \delta_n}{\epsilon_n^{\kappa+1}} = o(1). \quad (35)$$

where the approximation follows from the fact that $\epsilon_n \gg \delta_n^{1/(\kappa+1)}$ and because the denominator is approximating the integral $\int_{u' I(\theta_0) u \leq 2} d\theta$ and the numerator is the length of shrinking intervals. Combining the above results, we find that, on the set $\tilde{\Omega}_n(\tilde{C}_n)$, the denominator can be lower-bounded by $\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta \{1 + o(1)\}$.

This implies

$$\begin{aligned} \text{(I)} &= \frac{\int_{K(p_0, p_\theta) \leq \lambda_n} f_n(\theta - \theta_0) P[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta}{(1 + o(1)) \int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta} (1 + o(1)) + o(1) \\ &= (N_1 + N_2) \{1 + o(1)\} + o(1). \end{aligned} \quad (36)$$

with

$$N_1 \equiv \frac{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} f\left(\epsilon_n^{-1/2}(\theta - \theta_0)\right) d\theta}{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta} \quad (37)$$

and

$$N_2 \equiv \frac{\int_{K(p_0, p_\theta) \leq \lambda_n} \mathbb{I}\left[x(\theta) > 1 - \frac{M_n C_n \delta_n}{\epsilon_n}\right] f\left(\epsilon_n^{-1/2}(\theta - \theta_0)\right) P[\hat{K}(\mathbf{X}, \tilde{\mathbf{X}}^\theta) \leq \epsilon_n, d(\hat{D}_{n,m}^\theta, D_\theta) \leq C_n \delta_n] d\theta}{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta}, \quad (38)$$

where the second equality follows from the fact that $x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}$ leads to $K(p_0, p_\theta) \leq \epsilon_n - M_n C_n \delta_n$ and then $K(p_0, p_\theta) \leq \lambda_n$ is trivially satisfied and where the last $o(1)$ comes from the set $\tilde{\Omega}_n(\tilde{C}_n)^c$.

Since we have $\epsilon_n \gg \delta_n$, with $u = \epsilon_n^{-1/2}(\theta - \theta_0)$, the first term is approximately equal to

$$N_1 = \frac{\int_{K(p_0, p_{\theta_0 + \sqrt{\epsilon_n} u}) \leq \epsilon_n} f(u) du}{\int_{K(p_0, p_{\theta_0 + \sqrt{\epsilon_n} u}) \leq \epsilon_n} du} (1 + o(1)).$$

Using Assumption 5 again, we have

$$\int_{K(p_0, p_{\theta_0 + \sqrt{\epsilon_n} u}) \leq \epsilon_n} du = \int_{\frac{1}{2}\sqrt{\epsilon_n} u' I(\theta) \sqrt{\epsilon_n} u \leq \epsilon_n} du + o(1) = \int_{u' I(\theta_0) u \leq 2} du + o(1).$$

This leads to

$$N_1 = \frac{\int_{u' I(\theta_0) u \leq 2} f(u) du}{\int_{u' I(\theta_0) u \leq 2} du} (1 + o(1)).$$

Next, we show that the fraction N_2 converges to 0. The numerator can be simplified as an integral over $1 - \frac{M_n C_n \delta_n}{\epsilon_n} < x(\theta) \leq 1 + \frac{\tilde{M}_n \delta_n}{\epsilon_n^{\kappa+1}}$, which can be bounded by (34) as

$$\begin{aligned} N_2 &\leq \frac{\|f\|_\infty \int_{1 - \frac{M_n C_n \delta_n}{\epsilon_n} < x(\theta) \leq 1 + \frac{\tilde{M}_n \delta_n}{\epsilon_n^{\kappa+1}}} d\theta}{\int_{x(\theta) \leq 1 - \frac{M_n C_n \delta_n}{\epsilon_n}} d\theta} \\ &\leq \|f\|_\infty \frac{D_2}{D_1} = o(1). \end{aligned}$$

Since we have $\epsilon_n^{\kappa+1} \gg \delta_n$, putting all the terms together, we obtain that the $P_0^{(n)}$ -averaged ABC posterior distribution of $\epsilon_n^{-1/2}(\theta - \theta_0)$ is asymptotically uniform over the ellipsoid $\{u : u' I(\theta_0) u \leq 2\}$.