arXiv:1704.04315v3 [stat.ME] 7 Oct 2020

INFORMATION CRITERION FOR BOLTZMANN APPROXIMATION PROBLEMS

Youngjun Choe[†], Yen-Chi Chen[‡], and Nick Terry[†]

[†]*Department of Industrial & Systems Engineering and* [‡]*Department of Statistics*

University of Washington, Seattle

Abstract: This paper considers the problem of approximating a density when it can be evaluated up to a normalizing constant at a limited number of points. We call this problem the Boltzmann approximation (BA) problem. The BA problem is ubiquitous in statistics, such as approximating a posterior density for Bayesian inference and estimating an optimal density for importance sampling. Approximating the density with a parametric model can be cast as a model selection problem. This problem cannot be addressed with traditional approaches that maximize the (marginal) likelihood of a model, for example, using the Akaike information criterion (AIC) or Bayesian information criterion (BIC). We instead aim to minimize the cross-entropy that gauges the deviation of a parametric model from the target density. We propose a novel information criterion called the cross-entropy information criterion (CIC) and prove that the CIC is an asymptotically unbiased estimator of the cross-entropy (up to a multiplicative constant) under some regularity conditions. We propose an iterative method to approximate the target density by minimizing the CIC. We demonstrate that the

proposed method selects a parametric model that well approximates the target density.

Key words and phrases: cross-entropy information criterion, density estimation, importance sampling, Kullback-Leibler divergence, parametric mixture model.

1. Introduction

This paper considers the problem of approximating a target density $q^*(\mathbf{x}) = r(\mathbf{x})/\rho$ when the normalizing constant ρ is unknown but we can evaluate the nonnegative function r at a limited number of points. We call this problem the *Boltzmann approximation problem* (BA problem) due to its relation to the Boltzmann distribution, detailed in Section 2. This problem is ubiquitous in statistics. Here we list three notable examples where this problem occurs.

Example 1: Simulation-based inference. Stochastic simulation models [Choe et al., 2015, Pratola and Chkrebtii, 2018] are widely used to estimate the mean $\mathbb{E}(V)$, where the simulation output of interest, V, depends on the input $\mathbf{X}$ from a known density $p(\mathbf{x})$. When the simulator is computationally expensive, the importance sampling estimator $\frac{1}{n}\sum_{i=1}^{n} V_i \frac{p(\mathbf{X}_i)}{q(\mathbf{X}_i)}$ is widely used because it is unbiased and has the minimum variance when using the output V_i corresponding to the input $\mathbf{X}_i$,

$i = 1, \ldots, n$, sampled from the optimal density [Chen and Choe, 2019]

$$q^*(\mathbf{x}) \propto \sqrt{\mathbb{E}(V^2|\mathbf{X} = \mathbf{x})}p(\mathbf{x}).$$

By identifying $r(\mathbf{x}) = \sqrt{\mathbb{E}(V^2|\mathbf{X} = \mathbf{x})}p(\mathbf{x})$, we encounter the BA problem.

Example 2: Causal inference. Consider a simple causal inference problem where we have a response variable Y and a binary treatment $A \in \{0, 1\}$. Under the potential outcome model [Rubin, 2005], the causal effect from A on Y can be analyzed via studying how the two potential outcomes, $Y(1)$ and $Y(0)$, differ. To investigate this, we need to know four distributions: Two of them, $p(Y(0) = y_0|A = 0)$ and $p(Y(1) = y_1|A = 1)$, can be directly identified from the data while the other two, $p(Y(0) = y_0|A = 1)$ and $p(Y(1) = y_1|A = 0)$, cannot. Luckily, Tukey's factorization shows that

$$p(Y(0) = y_0|A = 1) \propto \underbrace{\frac{P(A = 1|Y(0) = y_0)}{P(A = 0|Y(0) = y_0)}}_{O(y_0)} p(Y(0) = y_0|A = 0)$$

and similarly for $p(Y(1) = y_1|A = 0)$. With a model on $O(y_0)$, one approach to approximate $p(Y(0) = y_0|A = 1)$ is to sample from its density implied by the above factorization. In this case, we only have access to a function that is proportional to the density we want to sample from. By identifying

$q^*(\cdot) = p(Y(0) = \cdot|A = 1)$ and $r(\cdot) = O(\cdot)p(Y(0) = \cdot|A = 0)$, we run into the same BA problem.

Example 3: Bayesian inference. The most prominent BA problem is approximating a posterior density for Bayesian inference. Let $\pi(\theta)$ be the prior density of the parameter θ and $L(\theta|X_1, \cdots, X_n)$ be the likelihood function. The posterior (density) is $\pi(\theta|X_1, \cdots, X_n) \propto L(\theta|X_1, \cdots, X_n) \cdot \pi(\theta)$. In this case, $q^*(\cdot) = \pi(\cdot|X_1, \cdots, X_n)$ is the posterior density and $r(\cdot) = \pi(\cdot)L(\cdot|X_1, \cdots, X_n)$ is the prior density times the likelihood of a model. Then, ρ is the model evidence, which can be compared between models for Bayesian model selection [Knuth et al., 2015].

When r is computationally light to evaluate, Markov chain Monte Carlo methods have been used extensively in practice, such as the Metropolis-Hastings algorithm that probabilistically accepts or rejects samples [Metropolis et al., 1953, Hastings, 1970], mainly because of theoretical guarantee on asymptotically exact sampling from q^*. Also, for exact sampling, the acceptance-rejection sampling algorithm is widely used if an upper-bounding function of r is known.

But there is growing interest in the BA problems where evaluating r is computationally heavy. In Bayesian inference, r may involve computing the likelihood for a large dataset [Blei et al., 2017] or a complex Bayesian

model [Beaumont, 2010, Sunnåker et al., 2013]. For the causal inference, if the model involves covariates (which is often the case), we need to sample counterfactual variables several times per each possible covariate value. For simulation-based inference, running a simulator is often computationally expensive [Choe et al., 2015, Pratola and Chkrebtii, 2018]. In such scenarios, one can only afford a relatively small number (e.g., hundreds or thousands) of evaluations of r and does not want to discard any of them. This paper focuses on this scenario where we use every evaluation of r to approximate q^*.

While the approximation of q^* is of primary interest, this study also emphasizes the estimation of the normalizing constant ρ because this quantity is, in some applications, even more important to estimate than the density q^* itself. ρ is the model evidence for Bayesian inference and the estimand for importance sampling. In particular, importance sampling requires sampling from the approximate density of q^* and evaluating r at the sampled points to unbiasedly estimate $\rho = \mathbb{E}_\mu[r]$ (with respect to the measure μ) with a minimal variance. Therefore, a sampling-based approach to approximate q^* is a key to the BA problem and is the focus of this study to ensure desirable probabilistic guarantees on estimating ρ.

The existing literature on approximating q^* in the BA problem can be

generally grouped into parametric approaches [Rubinstein, 1999, Wang and Zhou, 2015, Choe et al., 2015] and nonparametric approaches [Zhang, 1996, Neddermeyer, 2009, Chen and Choe, 2019]. Chen and Choe [2019] compare both approaches theoretically and empirically.

This paper considers a class of parametric approaches where we posit a parametric family of densities and find its member closest to q^* in terms of the closeness measured by the Kullback-Leibler (KL) divergence [Kullback and Leibler, 1951]. This parametric framework itself is very general and includes maximum likelihood estimation.

For importance sampling, this framework is first used for the so-called cross-entropy method [Rubinstein, 1999]. In this framework, the parametric approximation of q^* takes two steps, namely, 1) choosing a parametric family and 2) minimizing the KL divergence from a member density in the chosen family to the target density q^*. The latter is a well-studied optimization problem and the former is an under-explored model selection problem and the focus of this paper.

To tackle this problem, we combine two key ideas: importance sampling and information criterion. Importance sampling [Kahn and Marshall, 1953] allows us to generate points from distributions different from q^* and still be able to estimate an expectation with respect to q^* up to a constant as long

as we have access to r. With this, we are able to compute the information criterion to perform parameter estimation and model selection.

To the best of our knowledge, no rigorous framework of combining importance sampling and information criterion is proposed in the literature yet, despite its critical role in choosing an importance sampling density and in addressing BA problems. We propose a novel information criterion called the cross-entropy information criterion (CIC) and prove that the criterion is an asymptotically unbiased estimator of the KL divergence (up to a multiplicative constant and an additive constant). We justify that the minimization of the criterion leads to a good model selection.

We also show that the CIC is reduced to the Akaike information criterion (AIC)[Akaike, 1974] when data are directly generated from the distribution to estimate. Rigorous theoretical analysis of the AIC is a long-standing problem in the literature. Theoretical analyses of information criteria akin to the AIC often impose uniform integrability conditions directly on estimators to make the model complexity penalty to be expressed in the free parameter dimension d (see Conditions A7–A8 in Donohue et al. [2011], Theorem 1 in Claeskens and Consentino [2008], and more references cited in Bhansali and Papangelou [1991, p.1157] and Findley and Wei [2002, p.416]). For autoregressive models, sufficient conditions for such uniform

integrability conditions are established much later than the seminal paper of the AIC [Akaike, 1974] where the uniform integrability is not established rigorously [Bhansali, 1986, Findley and Wei, 2002]. For general parametric models, establishing general versions of the sufficient conditions has been an open problem. This paper establishes such sufficient conditions to validate the CIC, which is more general than AIC.

The remainder of this paper is organized as follows. Section 2 briefly reviews the relevant background. Section 3 proposes the cross-entropy information criterion. Section 4 explains how the proposed criterion can be used in practice with a numerical example for approximating an optimal density for importance sampling. Section 5 concludes the paper.

2. Background

This section briefly explains the origin of the BA problem and reviews the KL divergence, the maximum likelihood estimator (MLE), and the Akaike information criterion (AIC) to a) introduce the minimum cross-entropy estimator (MCE), which is a generalized version of MLE, and b) pave the way for generalizing the AIC to the cross-entropy information criterion (CIC).

The name of the BA problem originates from statistical physics where the nonnegative function $r(\mathbf{x})$ is what we can evaluate and express as

$r(\mathbf{x}) = e^{-\phi(\mathbf{x})}$, where $\phi(\mathbf{x})$ is often called the energy of a state $\mathbf{x}$ that can be calculated using theories in physics. Using this notation, the target density $q^*(\mathbf{x}) \propto e^{-\phi(\mathbf{x})}$ is called the Boltzmann distribution and $\rho = \int r(\mathbf{x})\,\mathrm{d}\mathbf{x} = \int e^{-\phi(\mathbf{x})}\,\mathrm{d}\mathbf{x}$ is called the partition function. The BA problem often has one practical constraint: evaluating r is (computationally) expensive so we want to make good use of all evaluations.

The KL divergence [Kullback and Leibler, 1951] is commonly used to gauge the difference between two distributions in statistical inference. Consider two probability measures, Q^* and Q, on a common measurable space such that Q^* is absolutely continuous with respect to Q (written $Q^* \ll Q$). Then, the KL divergence from Q to Q^* is defined as

$$\mathbb{D}(Q^*||Q) := \mathbb{E}_{Q^*}\left[\log \frac{\mathrm{d}Q^*}{\mathrm{d}Q}\right], \tag{2.1}$$

where $\mathbb{E}_{Q^*}$ denotes the expectation with respect to Q^* and $\mathrm{d}Q^*/\mathrm{d}Q$ is the Radon—Nikodym derivative of Q^* with respect to Q. If Q^* and Q are absolutely continuous with respect to a dominating measure μ (e.g., counting or Lebesgue measure), their respective densities, q^* and q, exist by the Radon—Nikodym theorem and the KL divergence in (2.1) can be expressed

as

$$\begin{aligned}\mathbb{D}(Q^*||Q) &= \mathbb{E}_{Q^*}\left[\log \frac{q^*}{q}\right] \\ &= \int q^* \log\left(\frac{q^*}{q}\right) \mathrm{d}\mu.\end{aligned}$$

The KL divergence is non-negative and takes zero if and only if $q^* = q$ almost everywhere (a.e.). If Q is an approximation of Q^*, the choice of Q minimizing $\mathbb{D}(Q^*||Q)$ leads to a good approximation.

The MLE is a prominent example of using the KL divergence. When the data $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are drawn from an unknown distribution Q^*, we can approximate Q^* by a distribution in a parametric family $\{Q_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \boldsymbol{\Theta}_d \subset \mathbb{R}^d\}$ by minimizing the KL divergence from $Q_{\boldsymbol{\theta}}$ to Q^* over $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$. Suppose $Q^* \ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$ so that the KL divergence is well defined over $\boldsymbol{\Theta}_d$. Also, suppose $Q_{\boldsymbol{\theta}} \ll \mu$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$ and $Q^* \ll \mu$ so that densities $q_{\boldsymbol{\theta}} = \mathrm{d}Q_{\boldsymbol{\theta}}/\mathrm{d}\mu$ and $q^* = \mathrm{d}Q^*/\mathrm{d}\mu$ exist. Then, the KL divergence is

$$\begin{aligned}\mathbb{D}(Q^*||Q_{\boldsymbol{\theta}}) &= \int q^* \log\left(\frac{q^*}{q_{\boldsymbol{\theta}}}\right) \mathrm{d}\mu \\ &= \int q^* \log q^* \, \mathrm{d}\mu - \int q^* \log q_{\boldsymbol{\theta}} \, \mathrm{d}\mu. \qquad (2.2)\end{aligned}$$

Note that only the second term in (2.2), called cross-entropy, depends on

$\boldsymbol{\theta}$. Therefore, minimizing the KL divergence over $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$ is equivalent to minimizing the cross-entropy over $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$. Because q^* is unknown, the cross-entropy should be estimated based on $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q^*$. An unbiased, consistent estimator of the cross-entropy is

$$-\frac{1}{n}\sum_{i=1}^{n} \log q_{\boldsymbol{\theta}}(\mathbf{X}_i), \tag{2.3}$$

which is the average of negative log-likelihoods for the observed data. Therefore, the MLE of $\boldsymbol{\theta}$, denoted by $\widehat{\boldsymbol{\theta}}_n$, is the minimizer of the cross-entropy estimator in (2.3).

Another example of using the KL divergence or cross-entropy is the Akaike information criterion (AIC) [Akaike, 1974]. As the free parameter dimension d of the parameter space $\boldsymbol{\Theta}_d$ (or equivalently, the model degrees of freedom) increases, $Q_{\widehat{\boldsymbol{\theta}}_n}$ may become a better approximation of Q^*. To compare the different approximating distributions (or models), we could use a plug-in estimator

$$-\frac{1}{n}\sum_{i=1}^{n} \log q_{\widehat{\boldsymbol{\theta}}_n}(\mathbf{X}_i) \tag{2.4}$$

of the cross-entropy, but this is problematic because of the downward bias created from using the data twice (once for $\widehat{\boldsymbol{\theta}}_n$ and another for estimating

the cross-entropy). While complex models (or distributions with larger d's) tend to have a smaller cross-entropy estimate, they are subject to the over-fitting problem. The AIC remedies this issue by correcting the asymptotic bias of the estimator in (2.4). The AIC is defined (up to a multiplicative constant) as

$$-\frac{1}{n}\sum_{i=1}^{n}\log q_{\widehat{\boldsymbol{\theta}}_n}(\mathbf{X}_i)+\frac{d}{n}, \tag{2.5}$$

where the bias correction term d/n penalizes the model complexity, balancing it with the goodness-of-fit represented by the first term. Minimizing the AIC can be interpreted as minimizing an asymptotically unbiased estimator of the cross-entropy. Therefore, both the MLE and AIC aim at minimizing the cross-entropy from an approximate distribution $Q_{\boldsymbol{\theta}}$ to the unknown distribution Q^* generating the data.

3. Cross-Entropy Information Criterion for the BA problem

This section develops an information criterion for the BA problem. We introduce an approximation of the cross-entropy in the BA problem that is suitable for parameter estimation and an information criterion that can be used for model selection.

3.1 Approximate Cross-Entropy for the BA problem

Analogous to the MLE and AIC, our approximation task in this paper considers minimizing the KL divergence (or cross-entropy) from a parametric distribution $Q_{\boldsymbol{\theta}}$ to the target distribution Q^* over $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$ (to well-define the KL divergence, hereafter assume $Q^* \ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$) when the target density q^* is proportional to a nonnegative function r, i.e., $q^* = r/\rho$ for a positive unknown constant $\rho = \int r \, \mathrm{d}\mu$. Minimizing the KL divergence in (2.2) over $\boldsymbol{\theta}$ is equivalent to minimizing the cross-entropy $-\int q^* \log q_{\boldsymbol{\theta}} \, \mathrm{d}\mu = -\frac{1}{\rho} \int r \log q_{\boldsymbol{\theta}} \, \mathrm{d}\mu$ and equivalent to minimizing

$$\mathcal{C}(\boldsymbol{\theta}) := -\int r \log q_{\boldsymbol{\theta}} \, \mathrm{d}\mu, \tag{3.6}$$

which is unknown in practice because r can be evaluated only at observed data points. Using the idea of importance sampling, we approximate $\mathcal{C}(\boldsymbol{\theta})$ in (3.6) by an unbiased and consistent estimator

$$\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta}) := -\frac{1}{n} \sum_{i=1}^{n} \frac{r(\mathbf{X}_i)}{q_{\boldsymbol{\eta}}(\mathbf{X}_i)} \log q_{\boldsymbol{\theta}}(\mathbf{X}_i) \tag{3.7}$$

where $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q_{\boldsymbol{\eta}}$ with any parameter $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$ as long as $q_{\boldsymbol{\eta}}$ covers the support of q^*. Because the estimator in (3.7) is to approximate the

cross-entropy in the BA problem, we call it the *approximate cross-entropy* (ACE).

Therefore, by minimizing $\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ in (3.7) over $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$, we can approximately minimize the KL divergence (or cross-entropy) from $Q_{\boldsymbol{\theta}}$ to Q^*. Thus, we call

$$\widehat{\boldsymbol{\theta}}_n := \operatorname*{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta}) \tag{3.8}$$

the *minimum cross-entropy estimator* (MCE) because it minimizes the ACE. Note that if the random sample is directly drawn from the target distribution (i.e., $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q_{\boldsymbol{\eta}} = Q^*$), then the MCE reduces to the MLE because minimizing (3.7) is equivalent to minimizing (2.3) due to $r/q^* = \rho$.

In the importance sampling literature, the density $q_{\boldsymbol{\eta}}$ in (3.7) is called the importance sampling (or proposal) density and q^* the optimal importance sampling density. The density q^* is optimal because the variance of the following importance sampling estimator of ρ is reduced to zero if $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are drawn from $q_{\boldsymbol{\eta}} = q^*$:

$$\widehat{\rho}_{\mathrm{IS}} = \frac{1}{n} \sum_{i=1}^{n} \frac{r(\mathbf{X}_i)}{q_{\boldsymbol{\eta}}(\mathbf{X}_i)}. \tag{3.9}$$

In practice, q^* is unknown and thus approximated by $q_{\boldsymbol{\eta}}$. Therefore, finding $q_{\boldsymbol{\eta}}$ closest to q^* is of primary interest.

3.2 Properties of the MCE

The MCE is a minimizer of (3.7) so it is an M-estimator [Huber, 1964], which allows us to use the M-estimation theory [Van der Vaart, 1998]. Hereafter, we assume the parameter minimizing the true cross-entropy,

$$\boldsymbol{\theta}^* := \operatorname*{argmin}_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d} \mathcal{C}(\boldsymbol{\theta}),$$

is unique. Then, the MCE $\widehat{\boldsymbol{\theta}}_n$ is a strongly consistent estimator of $\boldsymbol{\theta}^*$ (Lemma 1) and has asymptotic normality (Lemma 2) under standard regularity conditions (see assumptions in Appendix A).

Lemma 1 (Strong consistency of the MCE)**.** *Suppose that assumptions (A1) and (B1-2) hold and $\boldsymbol{\Theta}_d$ is compact. Then, for any $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$, the MCE $\widehat{\boldsymbol{\theta}}_n := \operatorname*{argmin}_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ converges almost surely to $\boldsymbol{\theta}^*$ as $n \to \infty$.*

Lemma 2 (Asymptotic normality of the MCE)**.** *Suppose that assumptions (A1-3) and (B3-4) hold and that for any $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$, $\widehat{\boldsymbol{\theta}}_n$ converges in probability to $\boldsymbol{\theta}^*$ as $n \to \infty$. Then, for any $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$, $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right)$ converges in distribution to $N\left(0, \boldsymbol{\Gamma}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}\boldsymbol{\Gamma}^{-1}\right)$ as $n \to \infty$, where $\boldsymbol{\Gamma} := -\mathbb{E}_\mu[r\nabla^2_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}]$*

and $\mathbf{\Lambda}_{\boldsymbol{\eta}} := \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*) \nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)^T\right]$ *with*

$$h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}) := \frac{r(\mathbf{X})}{q_{\boldsymbol{\eta}}(\mathbf{X})} \log q_{\boldsymbol{\theta}}(\mathbf{X}). \tag{3.10}$$

The above two lemmas are common results from the standard assumptions, see, e.g., Theorem A1 and A2 in Rubinstein and Shapiro [1993].

3.3 Iterative procedure for minimizing the ACE

Although the ACE $\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ in (3.7) is an unbiased estimator of the cross-entropy (up to a multiplicative constant) from $Q_{\boldsymbol{\theta}}$ to Q^*, a bad choice of sampling density $q_{\boldsymbol{\eta}}$ may lead to a large variance. Inspired by the cross-entropy method [Rubinstein, 1999], we propose to update $\boldsymbol{\eta}$ iteratively by the closest estimate of $\boldsymbol{\theta}^*$ as described in the procedure in Figure 1. We use the same sample size n for each iteration for notational simplicity without loss of generality.

An important property of the procedure in Figure 1 is that the iterative update of $\boldsymbol{\eta}$ by the best estimate of $\boldsymbol{\theta}^*$ leads to an information criterion that mimics the AIC under good regularity conditions. If $\boldsymbol{\eta}$ is fixed, we will not obtain such a nice property.

The iterative approach can be used to estimate ρ, the normalizing con-

ITERATIVE PROCEDURE FOR APPROXIMATING Q^*

Inputs: iteration counter $t = 1$, the number of iterations τ, the sample size n, and the initial parameter $\widehat{\boldsymbol{\theta}}_n^{(0)} = \boldsymbol{\eta} \in \boldsymbol{\Theta}_d$.

1. Sample $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)}$ from $Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$.
2. Find the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d}{\operatorname{argmin}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$, where

$$\begin{aligned} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}) &:= -\frac{1}{n} \sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\boldsymbol{\theta}}\left(\mathbf{X}_i^{(t-1)}\right), \\ &= -\frac{1}{n} \sum_{i=1}^{n} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}\right) \end{aligned} \tag{3.11}$$

3. If $t = \tau$, output the approximate distribution $Q_{\widehat{\boldsymbol{\theta}}_n^{(\tau)}}$. Otherwise, increment t by 1 and go to Step 1.

Figure 1: Iterative procedure for approximating Q^* by minimizing the estimator of cross-entropy from a parametric distribution $Q_{\boldsymbol{\theta}}$ to Q^*.

stant, by modifying the estimator in (3.9) as follows:

$$\widehat{\rho}^{(t-1)} = \frac{1}{n} \sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)}. \tag{3.12}$$

Furthermore, if $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = q^*$, the estimator in (3.12) is the optimal importance sampling estimator having zero variance [Kahn and Marshall, 1953]. Because the iterative procedure refines $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$ to be closer to q^*, $\widehat{\rho}^{(t-1)}$ will generally have a smaller variance as t gets larger.

The fact that we can estimate ρ with a small variance as a byproduct of approximating q^* is particularly desirable, because ρ, the normalizing constant of $q^* = r/\rho$, is often a quantity of interest as discussed in Section 1.

3.4 Approximating AIC

To simplify the model complexity penalty term in the AIC, Akaike [1974] assume that the true data generating distribution belongs to the parametric distribution family being considered. We make a similar assumption $Q^* = Q_{\boldsymbol{\theta}^*}$ (i.e., assumption (A1) in Appendix A) to simplify the asymptotic bias of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in estimating $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$.

In what follows we describe our main result as Theorem 1 that quantifies the asymptotic bias. Before deriving the asymptotic bias, we first introduce a useful lemma that characterizes the limiting behavior of any MCE. The assumptions and proofs are deferred to Appendix A. We note that as n tends to infinity, the number of iterations $\tau \geq 2$ remains fixed and the parameter dimension d is also fixed.

Lemma 3. *Suppose that assumptions (A1-5) and (B1-5) hold and let* $\widehat{\boldsymbol{\theta}}_n$ *be the MCE. Then* $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right) = Z_{\boldsymbol{\eta}} + \epsilon_{n,\boldsymbol{\eta}}$ *where*

$$Z_{\boldsymbol{\eta}} := -\boldsymbol{\Gamma}^{-1}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\nabla_{\boldsymbol{\theta}}h(\mathbf{X}_i;\boldsymbol{\eta},\boldsymbol{\theta}^*), \quad \sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\mathbb{E}(\|\epsilon_{n,\boldsymbol{\eta}}\|) = O(1/n^{1/4})$$

and $\mathbb{E}(Z_{\boldsymbol{\eta}}) = 0, \quad \mathsf{Cov}(Z_{\boldsymbol{\eta}}) = \boldsymbol{\Gamma}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}\boldsymbol{\Gamma}^{-1}, \quad \boldsymbol{\Lambda}_{\boldsymbol{\eta}} = \rho\boldsymbol{\Gamma} + O(\|\boldsymbol{\eta} - \boldsymbol{\theta}^*\|).$

Lemma 3 is different from the conventional asymptotic normality (e.g., Lemma 2) in two senses. First, we require a bound on the expectation of the remainder term $\mathbb{E}(\|\epsilon_{n,\boldsymbol{\eta}}\|)$, which implies the rate of $O_P(1/n^{1/4})$ and the asymptotic normality. We need the expectation because the derivation of the asymptotic bias in the CIC analysis requires an expectation bound. Second, the expectation of the remainder term is bounded uniformly for all $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$. This is in contrast to the conventional analysis that only requires the result at $\boldsymbol{\eta} = \boldsymbol{\theta}^*$. We need the uniform bound over $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$ because the data are generated from a measure $Q_{\boldsymbol{\eta}}$ where $\boldsymbol{\eta}$ is a random vector. In our analysis, we generate a sequence of estimators $\widehat{\boldsymbol{\theta}}_n^{(1)}, \widehat{\boldsymbol{\theta}}_n^{(2)}, \cdots, \widehat{\boldsymbol{\theta}}_n^{(\tau)}$ and the data used to estimate $\widehat{\boldsymbol{\theta}}_n^{(t)}$ are generated from $Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$, the measure based on the previous set of data. The uniform bound ensures that the remainder term $\epsilon_{n,\boldsymbol{\eta}}$ is small even if $\boldsymbol{\eta}$ is a random quantity. For any sequential/iterative sampling procedure, we would need a similar bound to control the remainder terms.

Theorem 1 (Asymptotic bias of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in estimating $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$)**.** *Sup-*

pose that assumptions (A1-6) and (B1-5) hold. Then

$$\mathbb{E}\Big[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big) - \mathcal{C}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)\Big] = -\rho\frac{d}{n} + o\bigg(\frac{1}{n}\bigg)$$

for each $t = 2, \ldots, \tau$.

The asymptotic bias, $-\rho d/n$, is proportional to the free parameter dimension d of the parameter space $\boldsymbol{\Theta}_d$, similar to the penalty term of the AIC in (2.5). In practice, $\rho = \mathbb{E}_\mu[r]$ is unknown, but we can use a consistent estimator of ρ to estimate the asymptotic bias, such as the estimator in (3.12).

As a bias-corrected estimator of the cross-entropy $\mathcal{C}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)$ (up to a multiplicative constant), we define the *cross-entropy information criterion* (CIC) as the follows:

Definition 1 (Cross-entropy information criterion (CIC))**.**

$$\mathrm{CIC}^{(t)}(d) = \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big) + \widehat{\rho}\frac{d}{n} \tag{3.13}$$

for $t = 1, \ldots, \tau$, where $\widehat{\boldsymbol{\theta}}_n^{(t)} := \operatorname*{argmin}_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ with $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\cdot)$ in (3.11). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimator in (3.12).

We note that the CIC reduces to the AIC up to an additive $o_p(1/n)$ if the samples are all drawn from the target distribution, that is, $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)} \sim$

$Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = Q^*$ for $t = 1, \ldots, \tau$ in Figure 1. If so, the first term of the CIC in (3.13) becomes

$$\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) := -\frac{1}{n}\sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right) \tag{3.14}$$

$$= -\frac{1}{n}\sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q^*\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right) \tag{3.15}$$

$$= -\frac{\rho}{n}\sum_{i=1}^{n} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right), \tag{3.16}$$

because $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = q^*$ in (3.14) and $r/q^* = \rho$ in (3.15). Plugging the expression in (3.16) into the CIC in (3.13) shows that the CIC is equal to ρ times the AIC in (2.5) up to an additive $o_p(1/n)$.

The asymptotic bias expression in Theorem 1 holds only for $t \geq 2$, because when $t = 1$, the initial sample is drawn from $Q_{\widehat{\boldsymbol{\theta}}_n^{(0)}} = Q_{\boldsymbol{\eta}}$, which is not a distribution converging to $Q_{\boldsymbol{\theta}^*}$. If we want to select a reasonable parameter dimension d to use at the first iteration, it is still necessary to penalize the model complexity. Therefore, we define the CIC even for $t = 1$.

3.5 The CIC based on cumulative data

If we use the equal sample size n for each iteration, the model dimension d for later iterations may vary only a little from the earlier iterations. Al-

ternatively, we can aggregate the samples gathered through iterations to obtain a cumulative version of the CIC as discussed in this subsection.

In the t^{th} iteration, the cumulative version uses all the observed data up to the current iteration to estimate $\mathcal{C}(\boldsymbol{\theta})$, instead of using only the current iteration's data $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)} \sim Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$ (recall Figure 1). The benefit of the cumulative version is the tendency of the aggregated estimator of $\mathcal{C}(\boldsymbol{\theta})$ to have a smaller variance than the non-aggregated estimator $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ in (3.11). This approach, in turn, can reduce the variance of the MCE as well, which minimizes the aggregated estimator of $\mathcal{C}(\boldsymbol{\theta})$.

For more flexibility, we can allocate a different sample size for each iteration, that is, n_t for the t^{th} iteration, $t = 0, 1, \ldots, \tau$ (for example, a large n_0 for the initial sample to broadly cover the support of $Q_{\boldsymbol{\eta}}$ and equal sample sizes $n_1 = \ldots = n_\tau$ for the later iterations). Then, we can find the MCE

$$\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d}{\operatorname{argmin}} \, \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta}), \tag{3.17}$$

where the aggregated estimator of $\mathcal{C}(\boldsymbol{\theta})$ is denoted as

$$\begin{aligned}\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta}) &:= \frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} n_s \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}^{(s)}}(\boldsymbol{\theta}) \\ &= -\frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} \sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)} \log q_{\boldsymbol{\theta}}\left(\mathbf{X}_i^{(s)}\right) \\ &= -\frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} \sum_{i=1}^{n_s} h\left(\mathbf{X}_i^{(s)}, \widehat{\boldsymbol{\theta}}^{(s)}, \boldsymbol{\theta}\right) \end{aligned} \tag{3.18}$$

for $t = 1, \ldots, \tau$ with $\widehat{\boldsymbol{\theta}}^{(0)} := \boldsymbol{\eta}$. We note that $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.18) is an unbiased estimator of $\mathcal{C}(\boldsymbol{\theta})$.

By using all data gathered up to the t^{th} iteration, we can determine the model parameter dimension d at the t^{th} iteration with the following CIC:

Definition 2 (CIC: Cumulative version)**.**

$$\overline{\text{CIC}}^{(t)}(d) = \bar{\mathcal{C}}^{(t-1)}\left(\widehat{\boldsymbol{\theta}}^{(t)}\right) + \widehat{\rho}\frac{d}{\sum_{s=0}^{t-1} n_s} \tag{3.19}$$

for $t = 1, \ldots, \tau$, where $\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d}{\operatorname{argmin}} \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.17). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimators in (3.20) and (3.21).

As t increases, the accumulated sample size $\sum_{s=0}^{t-1} n_s$ increases so that the free parameter dimension d can increase. Thus, the cumulative version of the CIC allows the use of a highly complex model if it can better approxi-

mate Q^*.

As a consistent and unbiased estimator of ρ, we can use

$$\widehat{\rho}^{(0)} = \frac{1}{n_0} \sum_{i=1}^{n_0} \frac{r\left(\mathbf{X}_i^{(0)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(0)}}\left(\mathbf{X}_i^{(0)}\right)} \tag{3.20}$$

at the 1^{st} iteration. At the t^{th} iteration for $t = 2, \ldots, \tau$, we can use

$$\widehat{\rho}^{(t-1)} = \frac{1}{\sum_{s=1}^{t-1} n_s} \sum_{s=1}^{t-1} \sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}, \tag{3.21}$$

where we do not use the data $\mathbf{X}_1^{(0)}, \ldots, \mathbf{X}_{n_0}^{(0)}$ from the initial distribution $Q_{\widehat{\boldsymbol{\theta}}^{(0)}} = Q_{\boldsymbol{\eta}}$ because they could potentially increase the variance of the resulting estimator if $Q_{\boldsymbol{\eta}}$ is too different from Q^*. The estimator in (3.21) is an importance sampling estimator of $\rho = \mathbb{E}_\mu[r]$. The potential for the increased variance has been well studied in the importance sampling literature [e.g., Hesterberg, 1995, Owen and Zhou, 2000].

4. Application of the Cross-Entropy Information Criterion

This section details how the CIC can be useful in practice. We first present how the CIC can help choose the number of components, k, for a mixture model in conjunction with an expectation-maximization (EM) algorithm

that finds the MCE for a given k in the t^{th} iteration, $t = 1, \ldots, \tau$. Then, we present the summary of how to use the cumulative version of CIC to iteratively approximate a target distribution. Lastly, we provide a numerical example to illustrate the use of the CIC for approximating an optimal importance sampling distribution.

4.1 Mixture model and an EM algorithm

To approximate a target distribution, we can consider a parametric mixture model whose parameter dimension d determines the model complexity. Parametric mixture models are often used to approximate a posterior density for Bayesian inference [Gutmann and Corander, 2016, Blei et al., 2017] and an optimal importance sampling density [Botev et al., 2013, Kurtz and Song, 2013, Wang and Zhou, 2015]. The density approximation quality hinges on the number of mixture components, k (or equivalently, the model dimension d). Prior studies either assume that k is given [Botev et al., 2013, Kurtz and Song, 2013] or use a rule of thumb to choose k based on "some understanding of the structure of the problem at hand" [Wang and Zhou, 2015].

We can use the CIC to select k for any parametric mixture model, considering various parametric component families. For example, exponential

families are especially convenient because the MCE can be found by using an expectation-maximization (EM) algorithm [Dempster et al., 1977]. In this paper, we use the Gaussian mixture model (GMM) for illustration. Appendix B details our version of the EM algorithm to find the MCE.

Figure 2 illustrates our EM algorithm in action for the first outer iteration ($t = 1$), where a GMM (gray-scale filled countour plot) with three component densities (white countour lines) is updated over EM iterations to approximate an unknown target density. We can see that in contrast to the conventional EM algorithm that maximizes the likelihood of a model (i.e., goodness-of-fit) to approximate the distribution of observed data, our algorithm uses the data (yellow dots), $\mathbf{X}_1^{(0)}, \ldots, \mathbf{X}_{1000}^{(0)} \sim Q_{\boldsymbol{\eta}}$, to estimate and minimize the cross-entropy from the approximate density to the target density. Out of the 1000 observations (yellow dots) in Figure 2(a) (note that the same data are plotted in (b)–(g) as well), only the small portion of them that fall *above* the red dashed line contribute to the cross-entropy estimate (in (3.18) because $r(\mathbf{x})$ is zero *below* the red dashed line for the numerical example in Section 4.3). See Appendix C for more details of the EM algorithm implemented for the numerical example in Section 4.3.

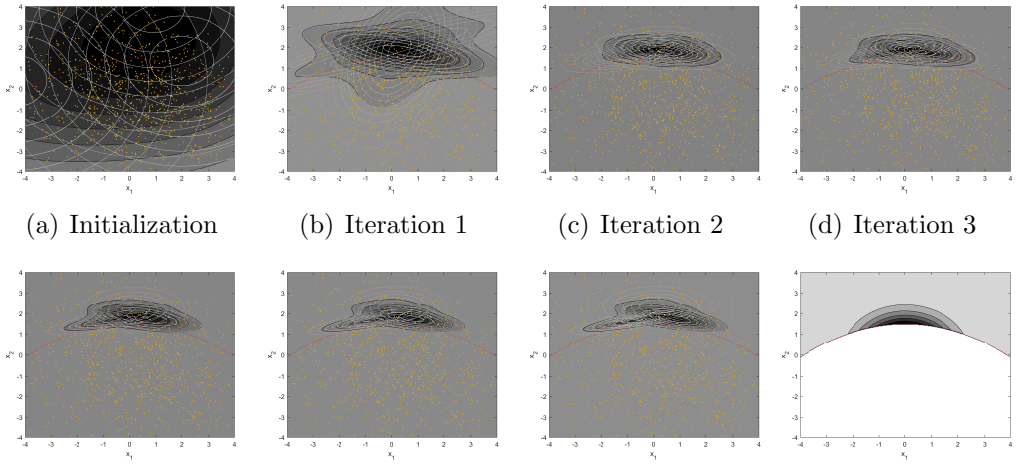

(a) Initialization (b) Iteration 1 (c) Iteration 2 (d) Iteration 3

(e) Iteration 4 (f) Iteration 5 (g) Iteration 6 (h) Target density

Figure 2: Illustration of our EM algorithm (for $k = 3$ at $t = 1$) that updates a randomly initialized density in (a) through Iterations 1-6 in (b)–(g) to approximate the unknown target density in (h). *Yellow dots* in (a)–(g) are the 2-dimensional pilot data $\mathbf{X}_i^{(0)}, i = 1, \ldots, 1000$ sampled from the initial distribution $Q_{\widehat{\boldsymbol{\theta}}^{(0)}} = Q_{\boldsymbol{\eta}}$. *Gray-scale filled contour plots* represent the Gaussian mixture density with $k = 3$ components in (a)–(g) and the target density in (h). *White contour line plots* in (a)–(g) represent the three component densities of the Gaussian mixture density. *Red dashed line* is the reference line that marks the shape of the target density in (h). The target density is the optimal importance sampling density in the numerical example (with $b = 1.5$) in Section 4.3.

4.2 Summary of the CIC-based distribution approximation procedure

This subsection summarizes how we can use the CIC to approximate a target distribution in practice. Using the EM algorithm in Section 4.1 for different k's (or d's) in the t^{th} iteration for $t = 1, \ldots, \tau$, we can find the MCE in (3.17) and calculate the CIC in (3.19). At the minimum of the

CIC, we can then find the best number of components, $k^{*(t)}$ (or the best model dimension $d^{*(t)}$) to use in the t^{th} iteration.

The CIC tends to decrease and then slowly increase as d increases, subject to the randomness of the data. Figure 3 shows such a pattern, where k is the number of mixture components in the GMM with unconstrained means and covariances. Note that k is proportional to the free parameter dimension $d = (k-1) + k(p + p(p+1)/2)$, with p denoting the dimension of the GMM density support.

Within the t^{th} iteration, we calculate the CIC over k as shown in Figure 3 for $t = 1, 4, 7$ for the numerical example (with $b = 1.5$) in Section 4.3. As the iteration counter t increases, $\overline{\text{CIC}}^{(t)}(d)$ in (3.19) uses a larger sample that accumulated data over iterations to more accurately estimate the cross-entropy.

At $t = 1$, the GMM with $k = 7$ in Figure 4(b) achieves the minimum CIC (as shown in Figure 3), while $k = 1$ and $k = 10$ result in seemingly over-simplified and over-complicated densities in Figure 4(a) and Figure 4(c), respectively, for the given sample size, 1000 (note that the effective sample size is much smaller because only a small portion of the data fall above the red dashed line as explained earlier with Figure 2). The right choice of k (neither too small nor too large) for the given sample size at the current

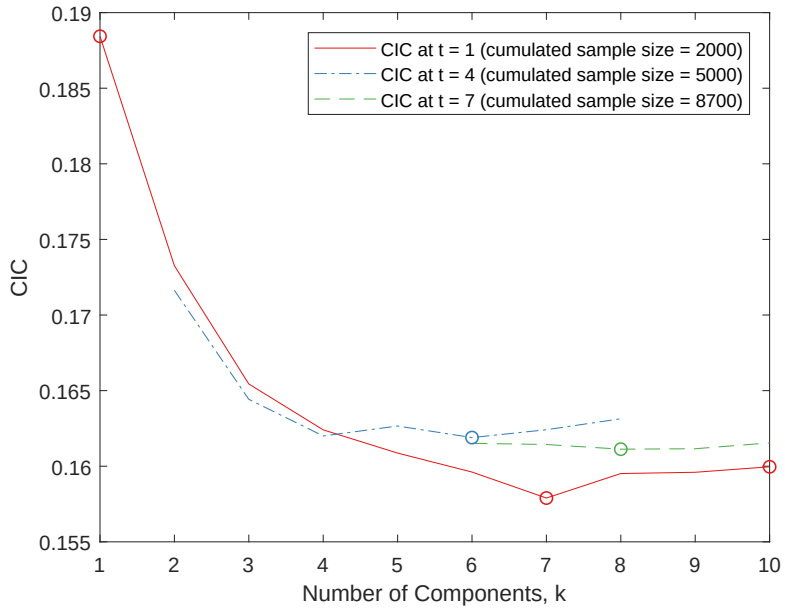

Figure 3: Plot of the CIC (cumulative version), $\overline{\mathrm{CIC}}^{(t)}(d)$, in (3.19) versus the number of components k, which determines the model dimension d, of the Gaussian mixture model. As the iteration counter t increases from 1 (red solid line) to 4 (black dash-dot line) to 7 (green dashed line), the CIC is calculated using a larger sample. The CIC is minimized at $k = 7$ for $t = 1$, $k = 6$ for $t = 4$, and $k = 8$ for $t = 7$ in this example. The circles in the plot correspond to the approximate densities shown in Figures 4(a)–(e).

iteration helps subsequent iterations by preventing sampling from an overly simplified/complicated distribution that could misguide the later iterations. Thus, it is beneficial to refine the approximate distribution *proportionally* (neither too much nor too little) for the given data size. Over iterations, sampled data (yellow dots in Figure 2) should increasingly cover the entire support of the target density. The CIC-minimizing densities at $t = 4$ in

Figure 4(d) and $t = 7$ in Figure 4(e) capture the overall shape of the target density in Figure 4(f).

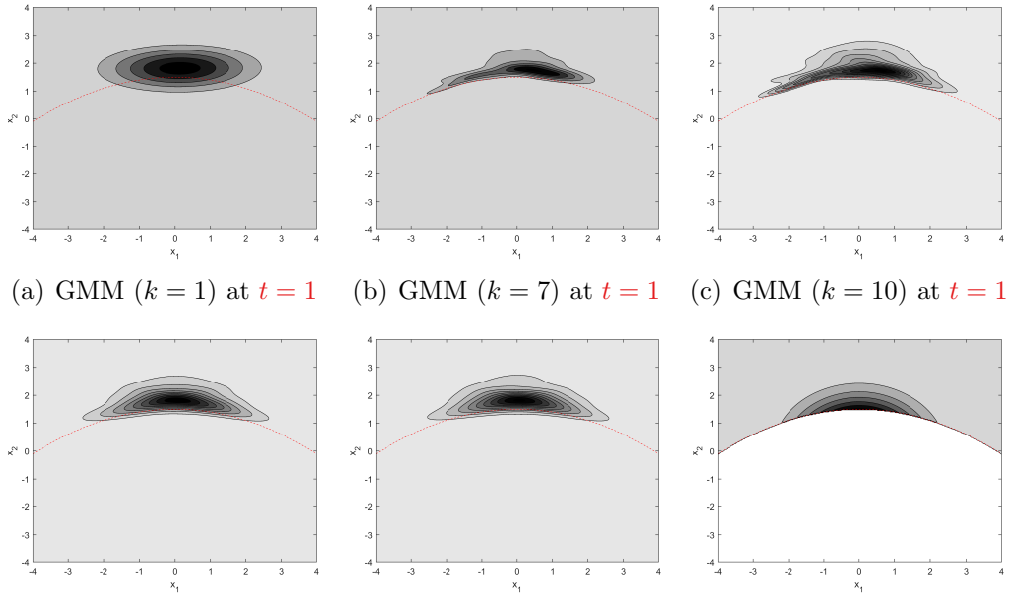

(a) GMM ($k = 1$) at $t = 1$ (b) GMM ($k = 7$) at $t = 1$ (c) GMM ($k = 10$) at $t = 1$

(d) GMM ($k = 6$) at $t = 4$ (e) GMM ($k = 8$) at $t = 7$ (f) Target density

Figure 4: Gaussian mixture models (GMMs) with k components in (a)–(e) correspond to the circles in Figure 3 and are compared with the target density in (f), which is the optimal importance sampling density in the numerical example (with $b = 1.5$) in Section 4.3.

Figure 5 summarizes the CIC-based distribution approximation procedure. Note that in addition to approximating the target distribution Q^*, if we want to estimate a quantity of interest such as ρ, we can sample $\mathbf{X}_1^{(\tau)}, \ldots, \mathbf{X}_{n_\tau}^{(\tau)} \sim Q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$ and use the estimator such as $\widehat{\rho}^{(\tau)}$ in (3.21). The following numerical example in Section 4.3 uses this additional step for importance sampling.

CIC-Based Approximation of the Target Distribution Q^*

Inputs: iteration counter $t = 1$, the number of iterations τ, the sample size per iteration $n_t, t = 1, \ldots, \tau$, the initial parameter dimension $d^{(0)}$, and the initial parameter $\widehat{\boldsymbol{\theta}}^{(0)} = \boldsymbol{\eta} \in \boldsymbol{\Theta}_{d^{(0)}}$.

1. Sample $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_{n_{t-1}}^{(t-1)} \sim Q_{\widehat{\boldsymbol{\theta}}^{(t-1)}}$.
2. Find the best model dimension $d^{*(t)} := \underset{d \geq 1}{\operatorname{argmin}} \overline{\mathrm{CIC}}^{(t)}(d)$ to use, where $\overline{\mathrm{CIC}}^{(t)}(d) := \bar{\mathcal{C}}^{(t-1)}\left(\widehat{\boldsymbol{\theta}}^{(t)}\right) + \widehat{\rho}\frac{d}{\sum_{s=0}^{t-1} n_s}$ is in (3.19) with the MCE $\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d}{\operatorname{argmin}} \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.17). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimators in (3.20) and (3.21).
3. If $t = \tau$, output the approximate distribution $Q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$. Otherwise, increment t by 1 and go to Step 1.

Figure 5: CIC-based approximation of the target distribution Q^*

4.3 Numerical example

This subsection empirically demonstrates how the CIC can be used to approximate the optimal distribution of importance sampling for estimating ρ. We expect that the better the approximation is, the smaller the standard error of the estimator $\widehat{\rho}^{(\tau)}$ in (3.21) will be. We call this importance sampling method the CIC-IS.

As benchmarks, we consider two methods to estimate ρ. First, the crude Monte Carlo (CMC) method is the most straightforward and most widely used Monte Carlo method. Second, we implement the importance

sampling method in Kurtz and Song [2013], which represents a state-of-the-art, cross-entropy method using a mixture model. Their method, which is called cross-entropy-based adaptive importance sampling using Gaussian mixture (CE-AIS-GM) uses the GMM with a *pre-specified* value for the number of mixture components, k. The GMM parameters are updated *once* within each iteration (as in Figure 2(b)) instead of using an EM algorithm.

We note that this empirical comparison is to *illustrate* how the CIC can improve importance sampling (by improving the approximation of the optimal density), *not* to comprehensively test the empirical performance of the CIC-IS itself. Interested readers are referred to the work of Cao and Choe [2019], which applies the CIC-IS to stochastic simulation models and studies its empirical performance using extensive numerical experiments and a case study on reliability evaluation of a wind turbine.

We use a classical example in the structural safety literature, which is also used in Kurtz and Song [2013]. With $\mathbf{X}$ following the bivariate Gaussian density $\phi(\mathbf{x})$ with zero mean and identity covariance matrix, a system of interest fails when $\mathbf{X}$ falls on the region $\{\mathbf{x} \in \mathbb{R}^2 : g(\mathbf{x}) \leq 0\}$, where $g(\mathbf{x}) = b - x_2 - \kappa\,(x_1 - e)^2$. We vary the parameter $b = 1.5, 2.0, 2.5$ to test three different failure thresholds. We fix the other two parameters, $\kappa = 0.1$ and $e = 0$, to maintain the shape of the failure boundary $\{\mathbf{x} \in$

$\mathbb{R}^2 : g(\mathbf{x}) = 0\}$ (the red dashed line in Figure 4). Note that $g(\mathbf{x})$ *represents* a computationally expensive function to evaluate, such as a finite element model in structural engineering.

The quantity of interest is the probability of the failure event, $\rho = \mathbb{E}_\mu[r]$, where $r(\mathbf{x}) = \phi(\mathbf{x})\mathbb{I}(g(\mathbf{x}) \leq 0)$. Here, $\mathbb{I}(\cdot)$ denotes the indicator function. The CMC estimator of ρ is

$$\widehat{\rho}_{\text{CMC}} = \frac{1}{n_{\text{CMC}}} \sum_{i=1}^{n_{\text{CMC}}} \mathbb{I}(g(\mathbf{X}_i) \leq 0), \tag{4.22}$$

where $\mathbf{X}_i, i = 1, \ldots, n_{\text{CMC}}$, are sampled from the density $\phi(\mathbf{x})$. For both importance sampling methods, we use the same sample size in Kurtz and Song [2013], namely, the total of 8700 replications: $n_t = 1000$ for $t = 0, \ldots, 6$ and $n_\tau = 1700$ for $\tau = 7$. Because the standard error of the CMC estimator can be analytically calculated as $\rho(1-\rho)/n_{\text{CMC}}$, we estimate the standard error instead of running CMC simulations.

We set the CE-AIS-GM to use $k = 30$ and to estimate ρ based only on the last ($\tau^{(th)}$) iteration data as in Kurtz and Song [2013]. The CIC-IS adaptively chooses k within the algorithm described in Figure 5 and uses the data from multiple iterations ($t = 2, \ldots, \tau$) to estimate ρ by $\widehat{\rho}^{(\tau)}$ in (3.21), as described in Section 4.2. Because the CIC helps find a distribution fairly

close to the optimal distribution throughout all the iterations, the CIC-IS can use the accumulated data to estimate ρ.

Table 1 shows the estimation results based on 500 experiment repetitions. As the parameter b increases, the estimand ρ decreases. Regardless of b, the CIC-IS obtains at least 50% smaller standard errors than the CE-AIS-GM. The smaller standard errors translate into larger computational savings for estimating ρ at a desired accuracy. Specifically, to compare both importance sampling methods with CMC, we analytically calculate 'CMC Ratio', which is the number of replications used in each row's method (that is, 8700) divided by the number of replications necessary for the CMC estimator in (4.22) (that is, n_{CMC}) to achieve the same standard error in the row. Although CE-AIS-GM saves significantly compared to CMC, CIC-IS saves even more by 4 to 6 times. The good performance of the CIC-IS can be attributed to a) closeness of the approximate densities $q_{\widehat{\boldsymbol{\theta}}^{(t)}}$, $t = 1, \ldots, \tau$ to the optimal density q^* as illustrated in Figure 4, and b) use of the data from multiple iterations ($t = 2, \ldots, \tau$) to estimate ρ, thanks to good approximate densities throughout the iterations.

Table 1: Comparison between CE-AIS-GM and CIC-IS

b	Method	Mean	Standard Error	CMC Ratio
1.5	CE-AIS-GM	0.082902	0.001145	15.00%
	CIC-IS	0.082911	0.000506	2.93%
2.0	CE-AIS-GM	0.030174	0.000526	8.23%
	CIC-IS	0.030173	0.000213	1.35%
2.5	CE-AIS-GM	0.008908	0.000211	4.39%
	CIC-IS	0.008910	0.000099	0.97%

Note: 'Mean' and 'Standard Error' are the sample mean and standard error of the estimates, respectively. The 'CMC Ratio' is $n_{\text{Total}}/n_{\text{CMC}}$, where $n_{\text{Total}} = 8700$ and $n_{\text{CMC}} = \bar{\rho}(1-\bar{\rho})/(S.E.)^2$. $\bar{\rho}$ is the CIC-IS's sample mean and $S.E.$ is the standard error of the method in the row. The smaller the CMC Ratio, the larger the computational saving of the method over CMC.

5. Conclusion

This paper proposed a novel information criterion, called the cross-entropy information criterion (CIC), to find a parametric density that has the asymptotically minimum cross-entropy to a target density to approximate. The CIC is the sum of two terms: an estimator of the cross-entropy (up to a multiplicative constant) from the parametric density to the target density, and a model complexity penalty term that is proportional to the free parameter dimension of the parametric density. Under certain regularity conditions, we proved that the penalty term corrects the asymptotic bias of the first term in estimating the true cross-entropy. Empirically, we demonstrated that minimizing the CIC leads to a density that well approximates an optimal importance sampling density. The CIC allowed us to develop

a principled algorithm to automatically approximate an unknown density that can be evaluated up to a normalizing constant at a limited number of points. This opens up the possibility to create an off-the-shelf software package for importance sampling and other similar applications.

Supplementary Materials

The online supplement includes Appendix A (Assumptions and Proofs), Appendix B (EM Algorithm for Minimizing the Cross-Entropy), and Appendix C (Implementation Details of the Numerical Example).

Acknowledgements

This work was supported by the National Science Foundation (NSF grant DMS-1952781).

References

Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

Mark A. Beaumont. Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41:379–406, 2010.

Rajendra J. Bhansali. A derivation of the information criteria for selecting autoregressive models. *Advances in Applied Probability*, 18(2):360–387, 1986.

Rajendra J. Bhansali and Fredos Papangelou. Convergence of moments of least squares estimators for the coefficients of an autoregressive process of unknown order. *The Annals of Statistics*, 19(3):1155–1162, 1991.

David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

Zdravko I. Botev, Dirk P. Kroese, Reuven Y. Rubinstein, and Pierre L'Ecuyer. The cross-entropy method for optimization. *Machine Learning: Theory and Applications, V. Govindaraju and C. R. Rao, Eds, Chennai: Elsevier*, 31:35–59, 2013.

Quoc Dung Cao and Youngjun Choe. Cross-entropy based importance sampling for stochastic simulation models. *Reliability Engineering & System Safety*, 191:106526, 2019.

Yen-Chi Chen and Youngjun Choe. Importance sampling and its optimality for stochastic simulation models. *Electronic Journal of Statistics*, 13(2):3386–3423, 2019.

Youngjun Choe, Eunshin Byon, and Nan Chen. Importance sampling for reliability evaluation with stochastic simulation models. *Technometrics*, 57(3):351–361, 2015.

Gerda Claeskens and Fabrizio Consentino. Variable selection with incomplete covariate data. *Biometrics*, 64(4):1062–1069, 2008. doi: 10.1111/j.1541-0420.2008.01003.x.

Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

Michael C. Donohue, Rosanna Overholser, Ronghui Xu, and Florin Vaida. Conditional Akaike information under generalized linear and proportional hazards mixed models. *Biometrika*, 98(3):685–700, 2011.

Mario A. T. Figueiredo and Anil K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):381–396, 2002.

David F. Findley and Ching-Zong Wei. AIC, overfitting principles, and the boundedness of moments of inverse matrices for vector autotregressions and related models. *Journal of Multivariate Analysis*, 83(2):415–450, 2002.

Michael U. Gutmann and Jukka Corander. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(125):1–47, 2016.

Wilfred K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

Tim Hesterberg. Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–194, 1995.

Peter J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.

Herman Kahn and Andy W. Marshall. Methods of reducing sample size in Monte Carlo computations. *Journal of the Operations Research Society of America*, 1(5):263–278, 1953.

Kevin H. Knuth, Michael Habeck, Nabin K. Malakar, Asim M. Mubeen, and Ben Placek. Bayesian evidence and model selection. *Digital Signal Processing*, 47:50–67, 2015.

Solomon Kullback and Richard A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

Nolan Kurtz and Junho Song. Cross-entropy-based adaptive importance sampling using gaussian mixture. *Structural Safety*, 42:35–44, 2013.

Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

Jan C. Neddermeyer. Computationally efficient nonparametric importance sampling. *Journal of the American Statistical Association*, 104(486):788–802, 2009.

Art Owen and Yi Zhou. Safe and effective importance sampling. *Journal of the American Statistical Association*, 95(449):135–143, 2000.

Mathew T Pratola and Oksana Chkrebtii. Bayesian calibration of multistate stochastic simulators. *Statistica Sinica*, 28:693–719, 2018.

Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.

Reuven Y. Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability*, 1(2):127–190, 1999.

Reuven Y. Rubinstein and Alexander Shapiro. *Discrete event systems: Sensitivity analysis and stochastic optimization by the score function method*. Chichester: John Wiley & Sons Ltd, 1993.

Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate Bayesian computation. *PLOS Computational Biology*, 9(1):1–10, 01 2013.

Aad W. Van der Vaart. *Asymptotic statistics*. New York: Cambridge University Press, 1998.

Hui Wang and Xiang Zhou. A cross-entropy scheme for mixtures. *ACM Transactions on Modeling and Computer Simulation*, 25(1):6:1–6:20, 2015.

Ping Zhang. Nonparametric importance sampling. *Journal of the American Statistical Association*, 91(435):1245–1253, 1996.

Youngjun Choe

Department of Industrial and Systems Engineering

University of Washington

Seattle, WA 98195-2650, USA

E-mail: ychoe@uw.edu

Yen-Chi Chen

Department of Statistics

University of Washington

Seattle, WA 98195-4322, USA

E-mail: yenchic@uw.edu

Nick Terry

Department of Industrial and Systems Engineering

University of Washington

Seattle, WA 98195-2650, USA

E-mail: pnterry@uw.edu

Appendix A: Assumptions and Proofs

We recall that the functions, $\mathcal{C}(\boldsymbol{\theta}) := -\mathbb{E}_\mu[r \log q_{\boldsymbol{\theta}}]$ and $h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}) := \frac{r(\mathbf{X})}{q_{\boldsymbol{\eta}}(\mathbf{X})} \log q_{\boldsymbol{\theta}}(\mathbf{X})$, are defined in (3.6) and (3.10), respectively. Below, $\nabla_{\boldsymbol{\theta}}$ and $\nabla^2_{\boldsymbol{\theta}}$ denote the gradient and the Hessian with respect to $\boldsymbol{\theta}$, respectively. For example, $\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)$ denotes the gradient of $h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ at $\boldsymbol{\theta}^*$.

Assumptions

(A) For $\boldsymbol{\theta}^* := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_d}{\operatorname{argmin}} \mathcal{C}(\boldsymbol{\theta})$,

(A1) the optimal distribution $Q^* = Q_{\boldsymbol{\theta}^*}$.

(A2) $\boldsymbol{\theta}^*$ is an interior point of $\boldsymbol{\Theta}_d$.

(A3) $\boldsymbol{\Gamma} := -\mathbb{E}_\mu[r \nabla^2_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}]$ is nonsingular.

(A4) $\nabla_{\boldsymbol{\theta}}(1/q_{\boldsymbol{\theta}}(\mathbf{x}))$ is continuous in a neighborhood (under the Euclidean norm) of $\boldsymbol{\theta}^*$ for a.e. $\mathbf{x}$ under μ.

(A5) $\left\| \mathbb{E}_\mu \left[\nabla_{\boldsymbol{\theta}}(1/q_{\boldsymbol{\theta}}) \, \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \right)^T \right] \right\|$ is bounded for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$ and $\boldsymbol{\Theta}_d$ is a compact set.

(A6) There exists a measurable function $g_0(\mathbf{x})$ such that $\int g_0 \, \mathrm{d}\mu < \infty$ and

$$\max\{ \|\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}}(\mathbf{x})\| , \left\|\nabla^2_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}}(\mathbf{x})\right\| \} \leq g_0(\mathbf{x})$$

for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$.

(B) For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$,

(B1) $h(\mathbf{x}, \boldsymbol{\eta}, \cdot)$ is continuous on $\boldsymbol{\Theta}_d$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$.

(B2) there exists a measurable function $g_{\boldsymbol{\eta}}$ such that $\int |g_{\boldsymbol{\eta}}| \, \mathrm{d}Q_{\boldsymbol{\eta}} < \infty$ and $|h(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\theta})| \leq g_{\boldsymbol{\eta}}(\mathbf{x})$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$ and all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$.

(B3) $h(\mathbf{x}, \boldsymbol{\eta}, \cdot)$ is thrice continuously differentiable on $\boldsymbol{\Theta}_d$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$.

(B4) there exist measurable functions $g_{\boldsymbol{\eta}}^{(i)}$ such that $\int |g_{\boldsymbol{\eta}}^{(i)}|^4 \, \mathrm{d}Q_{\boldsymbol{\eta}} < \infty$ and $\|\nabla_{\boldsymbol{\theta}}^i h(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\theta})\|_{\max} \leq g_{\boldsymbol{\eta}}^{(i)}(\mathbf{x})$ for $i = 1, 2, 3$ and all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$. The norm is the elementwise maximal norm.

(B5) $\boldsymbol{\Lambda}_{\boldsymbol{\eta}} := \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*) \nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)^T\right]$ exists and $\int \|\nabla_{\boldsymbol{\theta}}^2 h(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)\|_{\max}^2 \, \mathrm{d}Q_{\boldsymbol{\eta}}(\mathbf{x}) < \infty$.

Assumptions (A1-3) and (B1-4) are the standard regularity conditions to establish consistency (see Lemma 1) and asymptotic normality (see Lemma 2) of an M-estimator. Assumptions (B3) and (B4) imply Assumptions (B1) and (B2), respectively, as Lemma 2 requires stronger assumptions than Lemma 1.

To establish Lemma 3 and Theorem 1, additional assumptions (A4-6) and (B5) are needed because we need to control the expectation of the

smaller order terms in the asymptotic analysis. Assumptions (A4-6) ensure that we can exchange the expectation and derivative of the function $h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta})$. Assumption (B5) bounds the second moment of the Hessian matrix of h. This will be needed when deriving the convergence of sample Hessian matrix to the population Hessian matrix.

A crucial difference from the conventional analysis on the asymptotic normality is that to derive the AIC and CIC, we need expectation bound on the smaller order terms. The asymptotic normality only requires convergence in probability and distribution, which are not enough for the AIC and CIC. Also, the bounded derivatives up to the third (not second) order in (B4) is an additional requirement that allows us to bound the Taylor's remainder term. This is needed because there is no multivariate mean value theorem.

Note that assumptions (B4-5) imply the uniform integrability conditions that are commonly used in the literature for information criteria akin to the AIC to make the model complexity penalty to be expressed in the free parameter dimension d (see Conditions A7–A8 in Donohue et al. [2011], Theorem 1 in Claeskens and Consentino [2008], and more references cited in Bhansali and Papangelou [1991, p.1157] and Findley and Wei [2002, p.416]). Essentially, the uniform integrability condition is to make sure the expec-

tation of the smaller order terms is still a smaller order term. Here we use bounds on derivatives and Taylor remainder theorem to handle this.

Below, we provide the proofs of Lemma 3 and Theorem 1.

Proof of Lemma 3. We define the directional derivative with respect to variable $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d \subset \mathbb{R}^d$ in the direction $v \in \mathbb{R}^d$ as

$$\nabla_{\boldsymbol{\theta},v} := v^T \nabla_{\boldsymbol{\theta}}.$$

Recall that $\boldsymbol{\theta}^* := \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d}{\operatorname{argmin}} \mathcal{C}(\boldsymbol{\theta})$ is the population minimizer and

$$\begin{aligned}
\widehat{\boldsymbol{\theta}}_n &= \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d}{\operatorname{argmin}} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta}) \\
&= \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d}{\operatorname{argmin}} \frac{1}{n}\sum_{i=1}^{n} h(\mathbf{X}_i;\boldsymbol{\eta},\boldsymbol{\theta}) \\
&= \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d}{\operatorname{argmin}} h_n(\boldsymbol{\eta},\boldsymbol{\theta})
\end{aligned}$$

is the estimator, where $\mathbf{X}_i \sim Q_{\boldsymbol{\eta}}$ and

$$h_n(\boldsymbol{\eta},\boldsymbol{\theta}) := \frac{1}{n}\sum_{i=1}^{n} h(\mathbf{X}_i;\boldsymbol{\eta},\boldsymbol{\theta}).$$

The expectation of h_n is

$$\mathbb{E}_{Q_{\boldsymbol{\eta}}}(h(\mathbf{X}_i;\boldsymbol{\eta},\boldsymbol{\theta})) = \mathcal{C}(\boldsymbol{\theta})$$

for all $\boldsymbol{\eta} \in \boldsymbol{\Theta}_d$.

Because $\boldsymbol{\theta}^*$ and $\widehat{\boldsymbol{\theta}}_n$ are the minimizers, they solve the score equations

$$0 = \nabla_{\boldsymbol{\theta}} \mathcal{C}(\boldsymbol{\theta}^*) = \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta}, \widehat{\boldsymbol{\theta}}_n).$$

Thus, by the Taylor's remainder theorem,

$$\begin{aligned}
-\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) &= \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\widehat{\boldsymbol{\theta}}_n) - \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) \\
&= \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) + \int_{\epsilon=0}^{\epsilon=1} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^* + \epsilon(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*))\,\mathrm{d}\epsilon \\
&= \boldsymbol{\Gamma}_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) + R_{1,n}(\boldsymbol{\eta}),
\end{aligned}$$

where

$$\begin{aligned}
\boldsymbol{\Gamma}_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) &:= \nabla^2_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*), \\
R_{1,n}(\boldsymbol{\eta}) &:= \int_{\epsilon=0}^{\epsilon=1} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^* + \epsilon(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*))\,\mathrm{d}\epsilon.
\end{aligned}$$

$R_{1,n}(\boldsymbol{\eta})$ represents the remainder terms.

Thus, when $\mathbf{\Gamma}_n(\boldsymbol{\eta}, \boldsymbol{\theta}^*)$ is invertible, the difference in the estimator can be expressed as

$$\sqrt{n}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) = -\sqrt{n}\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta}, \boldsymbol{\theta}^*) - \sqrt{n}\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)R_{1,n}(\boldsymbol{\eta}). \quad \text{(S.23)}$$

The first quantity is how we establish the asymptotic normality and the second part (remainder term) is something we need to bound its expectation.

Bounding the remainder term $R_{1,n}(\boldsymbol{\eta})$. To see how we obtain this, first note that

$$\mathbb{E}(\|\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)R_{1,n}(\boldsymbol{\eta})\|) \leq \sqrt{\mathbb{E}(\|\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)\|^2)\mathbb{E}(\|R_{1,n}(\boldsymbol{\eta})\|^2)}. \quad \text{(S.24)}$$

Here both norms are the 2-norm for matrices and vectors. So it suffices to bound the two terms individually.

Bounding $\mathbb{E}(\|\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)\|^2)$. To bound the inverse matrix part, recall that $\mathbf{\Gamma} = \nabla^2_{\boldsymbol{\theta}}\mathcal{C}(\boldsymbol{\theta}^*)$ and note that since $\mathbb{E}(\mathbf{\Gamma}_n(\boldsymbol{\eta}, \boldsymbol{\theta}^*)) = \mathbf{\Gamma}$ due to assumption (A6) and $\mathbf{\Gamma}_n(\boldsymbol{\eta}, \boldsymbol{\theta}^*)$ is essentially a sample mean matrix so the variance shrinks at rate $O(1/n)$ due to assumption (B5). Thus, we can decompose

$$\mathbf{\Gamma}_n(\boldsymbol{\eta}, \boldsymbol{\theta}^*) = \mathbf{\Gamma} + \Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta}) \quad \text{(S.25)}$$

with $\mathbb{E}(\|\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\|) = O(1/\sqrt{n})$. Moreover, because assumption (B5) holds for all $\boldsymbol{\eta} \in \mathbf{\Theta}_d$,

$$\sup_{\boldsymbol{\eta}\in\mathbf{\Theta}_d} \mathbb{E}(\|\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\|) = O(1/\sqrt{n}).$$

For the inverse matrix, because

$$\begin{aligned}
\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*) &= [\mathbf{\Gamma} + \Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})]^{-1} \\
&= [\mathbf{\Gamma}(I + \mathbf{\Gamma}^{-1}\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta}))]^{-1} \\
&= (I - \mathbf{\Gamma}^{-1}\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta}) + O(\|\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\|^2))\mathbf{\Gamma}^{-1} \\
&= \mathbf{\Gamma}^{-1} - \mathbf{\Gamma}^{-1}\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\mathbf{\Gamma}^{-1} + O(\|\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\|^2).
\end{aligned}$$

By the fact that $\mathbf{\Gamma}$ is invertible and $\sup_{\boldsymbol{\eta}\in\mathbf{\Theta}_d} \mathbb{E}(\|\Delta_{n,\mathbf{\Gamma}}(\boldsymbol{\eta})\|) = O(1/\sqrt{n})$, we conclude that

$$\sup_{\boldsymbol{\eta}\in\mathbf{\Theta}_d} \mathbb{E}(\|\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)\|) = \|\mathbf{\Gamma}^{-1}\| + O(1/\sqrt{n}) \tag{S.26}$$

and

$$\sup_{\boldsymbol{\eta}\in\mathbf{\Theta}_d} \mathbb{E}(\|\mathbf{\Gamma}_n^{-1}(\boldsymbol{\eta}, \boldsymbol{\theta}^*)\|^2) = \|\mathbf{\Gamma}^{-1}\|^2 + O(1/\sqrt{n}). \tag{S.27}$$

Bounding $\mathbb{E}(\|R_{1,n}(\boldsymbol{\eta})\|^2)$. By the definition of $R_{1,n}(\boldsymbol{\eta})$,

$$R_{1,n}(\boldsymbol{\eta}) = \int_{\epsilon=0}^{\epsilon=1} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta},\widehat{\boldsymbol{\theta}}_n-\boldsymbol{\theta}^*} \nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta}, \boldsymbol{\theta}^* + \epsilon(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*)) \, \mathrm{d}\epsilon$$

so its norm can be bounded by

$$\begin{aligned}
\|R_{1,n}(\boldsymbol{\eta})\| &\leq \|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^2 \sup_{\boldsymbol{\theta}\in\boldsymbol{\Theta}_d} \sup_{\|v\|=\|u\|=\|w\|=1} \|\nabla_{\boldsymbol{\theta},v} \nabla_{\boldsymbol{\theta},u} \nabla_{\boldsymbol{\theta},w} h_n(\boldsymbol{\eta}, \boldsymbol{\theta})\| \\
&\leq \|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^2 d^{3/2} \int g_{\boldsymbol{\eta}}^{(3)} \, \mathrm{d}\widehat{Q}_{\boldsymbol{\eta}}
\end{aligned}$$

using assumption (B4) and d is the dimension of $\boldsymbol{\theta}$ and $\widehat{Q}_{\boldsymbol{\eta}}$ is the empirical measure formed by $\mathbf{X}_1, \cdots, \mathbf{X}_n$.

Because assumption (B4) holds for all $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$,

$$\begin{aligned}
\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} \mathbb{E}\|R_{1,n}(\boldsymbol{\eta})\|^2 &\leq \sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} d^3 \mathbb{E}\left(\|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^4 \left(\int g_{\boldsymbol{\eta}}^{(3)} \, \mathrm{d}\widehat{Q}_{\boldsymbol{\eta}} \right)^2 \right) \\
&\leq d^3 \sqrt{\mathbb{E}\left(\|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^8 \right) \sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} \mathbb{E}\left(\left(\int g_{\boldsymbol{\eta}}^{(3)} \, \mathrm{d}\widehat{Q}_{\boldsymbol{\eta}} \right)^4 \right)} \\
&= O\left(\sqrt{\mathbb{E}\|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^8} \right) = O(1/n^2).
\end{aligned} \tag{S.28}$$

Putting it altogether, we obtain

$$\begin{aligned}\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\sqrt{n}\mathbb{E}(\|\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)R_{1,n}(\boldsymbol{\eta})\|) &\leq \sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\sqrt{n}\sqrt{\mathbb{E}(\|\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)\|^2)\mathbb{E}(\|R_{1,n}(\boldsymbol{\eta})\|^2)} \\ &= \sqrt{n}(\|\boldsymbol{\Gamma}^{-1}\| + O(1/n^{1/4}))\cdot O\left(\sqrt{\mathbb{E}\|\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\|^4}\right) \\ &= O(1/\sqrt{n}) = o(1).\end{aligned}$$

Analysis of the first order term in equation (S.23)**.** Note that the difference between the first order term and $Z_{\boldsymbol{\eta}}$ is

$$R_{2,n}(\boldsymbol{\eta}) := \sqrt{n}\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)\nabla_{\boldsymbol{\theta}}h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) - Z_{\boldsymbol{\eta}} = \sqrt{n}(\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*) - \boldsymbol{\Gamma}^{-1})\nabla_{\boldsymbol{\theta}}h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)$$

so we need to show that

$$\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\mathbb{E}\|R_{2,n}(\boldsymbol{\eta})\| = o(1).$$

We can decompose the above bound via

$$\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\mathbb{E}\|R_{2,n}(\boldsymbol{\eta})\| \leq \sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d}\sqrt{\mathbb{E}\left\|\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*) - \boldsymbol{\Gamma}^{-1}\right\|^2 \cdot n\mathbb{E}\left\|\nabla_{\boldsymbol{\theta}}h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)\right\|^2}.$$

Using the derivation of equation (S.26), we can easily show that

$$\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} \mathbb{E}\|\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*) - \boldsymbol{\Gamma}^{-1}\|^2 = O(1/\sqrt{n})$$

and the second part

$$n\mathbb{E}\|\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)\|^2 = O(1)$$

because the covariance matrix $\mathsf{Cov}(\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)) = \frac{1}{n}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}$, where

$$\boldsymbol{\Lambda}_{\boldsymbol{\eta}} = \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X},\boldsymbol{\eta},\boldsymbol{\theta}^*\right)\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X},\boldsymbol{\eta},\boldsymbol{\theta}^*\right)^T\right]$$

is a bounded matrix for all $\boldsymbol{\eta}$ due to assumption (B5) and $\mathbb{E}\|\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)\|^2 = \mathsf{Tr}(\mathsf{Cov}(\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)))$. $\mathsf{Tr}(\boldsymbol{A})$ is the trace of the matrix $\boldsymbol{A}$.

Equation (S.23) also implies a useful result about the limiting behavior of $\sqrt{n}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*)$. First, because $\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)$ converges in probability to $\boldsymbol{\Gamma}^{-1}$ and $\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*)$ is a sample average quantity, we have a central limit theory about $\sqrt{n}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*)$ in the sense that

$$\sqrt{n}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) \xrightarrow{D} N(0, \boldsymbol{\Sigma}_{\boldsymbol{\eta}}),$$

where $\boldsymbol{\Sigma}_{\boldsymbol{\eta}} := \boldsymbol{\Gamma}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}\boldsymbol{\Gamma}^{-1}$ and $\boldsymbol{\Sigma}_{\boldsymbol{\eta}}$ is the asymptotic variance of $\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*$, i.e.,

$$n\mathsf{Var}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) \to \boldsymbol{\Sigma}_{\boldsymbol{\eta}}.$$

Putting it altogether, we conclude that

$$\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} \mathbb{E}\|R_{2,n}(\boldsymbol{\eta})\| = O(1/n^{1/4})$$

and we can rewrite equation (S.23) as

$$\begin{aligned}\sqrt{n}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) &= -\sqrt{n}\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)\nabla_{\boldsymbol{\theta}} h_n(\boldsymbol{\eta},\boldsymbol{\theta}^*) - \sqrt{n}\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)R_{1,n}(\boldsymbol{\eta}) \\ &= -Z_{\boldsymbol{\eta}} \underbrace{-\sqrt{n}\boldsymbol{\Gamma}_n^{-1}(\boldsymbol{\eta},\boldsymbol{\theta}^*)R_{1,n}(\boldsymbol{\eta}) - R_{2,n}(\boldsymbol{\eta})}_{:=\epsilon_{n,\boldsymbol{\eta}}}\end{aligned}$$

with $\sup_{\boldsymbol{\eta}\in\boldsymbol{\Theta}_d} \mathbb{E}(\|\epsilon_{n,\boldsymbol{\eta}}\|) = O(1/n^{1/4})$.

The covariance of $Z_{\boldsymbol{\eta}}$ is

$$\mathsf{Cov}(Z_{\boldsymbol{\eta}}) = \boldsymbol{\Gamma}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}\boldsymbol{\Gamma}^{-1}$$

and we have

$$
\begin{aligned}
\boldsymbol{\Lambda}_{\boldsymbol{\eta}} &= \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*\right) \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*\right)^T\right] \\
&= \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\frac{r^2(\mathbf{X})}{q_{\boldsymbol{\eta}}^2(\mathbf{X})} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X}) \left(\nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X})\right)^T\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r^2}{q_{\boldsymbol{\eta}}} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}\right)^T\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r^2}{q_{\boldsymbol{\eta}} q_{\boldsymbol{\theta}^*}^2} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}\right)^T\right] \\
&= \rho^2 \mathbb{E}_{\mu}\left[\frac{1}{q_{\boldsymbol{\eta}}} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}\right)^T\right] \qquad \text{(S.29)} \\
&= \rho^2 \int \left(\frac{1}{q_{\boldsymbol{\theta}^*}(\mathbf{x})} + R_{3,n}(\mathbf{x}; \boldsymbol{\eta})\right) \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x}) \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x})\right)^T \, \mathrm{d}\mu(\mathbf{x}) \qquad \text{(S.30)} \\
&= \rho \boldsymbol{\Gamma} + \rho^2 \int R_{3,n}(\mathbf{x}; \boldsymbol{\eta}) \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x}) \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x})\right)^T \, \mathrm{d}\mu(\mathbf{x}) \\
&\overset{(A5)}{=} \rho \boldsymbol{\Gamma} + O(\|\boldsymbol{\eta} - \boldsymbol{\theta}^*\|), \qquad \text{(S.31)}
\end{aligned}
$$

where

$$
R_{3,n}(\mathbf{x}; \boldsymbol{\eta}) := \int_{\epsilon=0}^{\epsilon=1} \left(\boldsymbol{\eta} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}} \left(\frac{1}{q_{\boldsymbol{\theta}^* + \epsilon(\boldsymbol{\eta} - \boldsymbol{\theta}^*)}(\mathbf{x})}\right) \mathrm{d}\epsilon
$$

is the Taylor's remainder term.

Therefore,

$$
n\mathsf{Var}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*) = \boldsymbol{\Sigma}_{\boldsymbol{\eta}} + o(1) = \rho \boldsymbol{\Gamma}^{-1} + o(1) + O(\|\boldsymbol{\eta} - \boldsymbol{\theta}^*\|).
$$

□

Proof of Theorem 1. We first simplify $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ and $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$. We then use the simplified expressions to derive the bias of interest. The first quantity is a population objective function while the second quantity is a sample objective function.

Analysis on the population objective function $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$.

$$\begin{aligned}\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= -\int r \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}} \,\mathrm{d}\mu \\ &= -\int \frac{r}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}} q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}} \,\mathrm{d}\mu \\ &= -\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right],\end{aligned} \tag{S.32}$$

where $\mathbf{X} \sim Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$. The equation in (S.32) holds under the condition $Q^* \ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}_d$, because $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\mathbf{x}) = 0$ implies $r(\mathbf{x}) = 0$ for any $\mathbf{x}$.

We take a second-order expansion of $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ about $\boldsymbol{\theta}^*$ and use the Taylor's remainder theorem:

$$\begin{aligned}\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) = &-\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] - \left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] \\ &- \frac{1}{2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right]\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right) \\ &+ R_{4,n},\end{aligned} \tag{S.33}$$

where $R_{4,n}$ is the Taylor's remainder term at rate $\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^3$ and has a bounded expectation of rate $\mathbb{E}(|R_{4,n}|) = O(\mathbb{E}\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^3)$ due to assumption (B4) on the third derivatives. In (S.33), the zeroth-order term is $\mathcal{C}(\boldsymbol{\theta}^*)$ by definition and the first-order term is zero because

$$\begin{aligned}
\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] &= \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\frac{r(\mathbf{X})}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\mathbf{X})} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r(\mathbf{X})}{q_{\boldsymbol{\theta}^*}(\mathbf{X})} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \\
&= \rho \mathbb{E}_{\mu}\left[\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \qquad \text{(S.34)} \\
&= \rho \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\mu}\left[q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \qquad \text{(S.35)} \\
&= 0,
\end{aligned}$$

where the equation in (S.34) holds under assumption (A1) and the interchange of expectation and differentiation in (S.35) holds under assumption (A6).

Let $\boldsymbol{\delta}_n := \sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)$. Using the fact that $\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] = \mathcal{C}(\boldsymbol{\theta}^*)$ implies

$$\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] = -\boldsymbol{\Gamma}$$

under assumption (A6), we simplify the expression of $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in (S.33) as

$$\begin{aligned}\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= \mathcal{C}(\boldsymbol{\theta}^*) - \frac{1}{2n}\boldsymbol{\delta}_n^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right]\boldsymbol{\delta}_n + R_{4,n} \\ &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{1}{2n}\boldsymbol{\delta}_n^T \boldsymbol{\Gamma}\boldsymbol{\delta}_n + R_{4,n}.\end{aligned}$$

Using Lemma 3, we can rewrite $\boldsymbol{\delta}_n$ as

$$\boldsymbol{\delta}_n = \sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right) = Z_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} + \epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}},$$

where $E\|\epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\| = o(1)$ since we have the uniform bound. And the covariance

$$\begin{aligned}\mathsf{Cov}(\boldsymbol{\delta}_n) &= \mathbb{E}\left(\mathsf{Cov}\left(\boldsymbol{\delta}_n | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\right)\right) + \mathsf{Cov}\left(\mathbb{E}\left(\boldsymbol{\delta}_n | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\right)\right) \\ &= \boldsymbol{\Gamma}^{-1}\mathbb{E}\left(\boldsymbol{\Lambda}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\right)\boldsymbol{\Gamma}^{-1} + o(1).\end{aligned} \tag{S.36}$$

Equation (S.36) is derived as follows. We first bound the quantity $\mathsf{Cov}\left(\mathbb{E}\left(\boldsymbol{\delta}_n | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\right)\right)$. Note that for a random vector $\mathbf{X} = (X_1, X_2)$, its

covariance can be bounded by

$$\begin{aligned}
|\mathsf{Cov}(X_1, X_2)| &\leq \sqrt{\mathsf{Var}(X_1)\mathsf{Var}(X_2)} \\
&\leq \max\{\mathsf{Var}(X_1), \mathsf{Var}(X_2)\} \\
&\leq \max\{\mathbb{E}(X_1^2), \mathbb{E}(X_2^2)\} \\
&\leq \mathbb{E}(\max\{X_1^2, X_2^2\}) \\
&\leq \mathbb{E}(\|\mathbf{X}\|^2).
\end{aligned}$$

The same bound can be established for a random vector of length d. By setting $\mathbf{X} = \mathbb{E}\Big(\epsilon_{n,\boldsymbol{\eta}=\widehat{\boldsymbol{\theta}}_n^{(t-1)}} | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)$ and using the fact that $\mathbb{E}\Big(Z_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big) = 0$, each element in the covariance matrix $\mathsf{Cov}\Big(\mathbb{E}\Big(\boldsymbol{\delta}_n | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\Big) = \mathsf{Cov}\Big(\mathbb{E}\Big(\epsilon_{n,\boldsymbol{\eta}=\widehat{\boldsymbol{\theta}}_n^{(t-1)}} | \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\Big)$ is less than or equal to

$$\begin{aligned}
\mathbb{E}\left\{\left\|\mathbb{E}\Big(\epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}} \mid \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\right\|^2\right\} &\leq \mathbb{E}\left\{\mathbb{E}^2\Big(\left\|\epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\right\| \mid \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\right\} \\
&\leq \sup_{\widehat{\boldsymbol{\theta}}_n^{(t-1)} \in \boldsymbol{\Theta}_d} \mathbb{E}^2\Big(\left\|\epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\right\| \mid \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big) \\
&= \left(\sup_{\widehat{\boldsymbol{\theta}}_n^{(t-1)} \in \boldsymbol{\Theta}_d} \mathbb{E}\Big(\left\|\epsilon_{n,\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\right\| \mid \widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\right)^2 \\
&\leq O(1/\sqrt{n}).
\end{aligned}$$

Note that the last inequality is due to Lemma 3.

For another quantity, equation (S.31) implies

$$\mathbb{E}\left\|\mathbf{\Lambda}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} - \rho\mathbf{\Gamma}\right\| = O\Big(\mathbb{E}\left\|\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\right\|\Big) = o(1).$$

So we conclude

$$\mathsf{Cov}(\boldsymbol{\delta}_n) = \rho\mathbf{\Gamma}^{-1} + o(1). \tag{S.37}$$

This establishes the bound in equation (S.36).

Let

$$\tilde{\boldsymbol{\delta}}_n := \mathbf{\Gamma}^{1/2}\boldsymbol{\delta}_n = \sqrt{n}\mathbf{\Gamma}^{1/2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right).$$

Then we have

$$\begin{aligned}
\frac{1}{2n}\mathbb{E}\big[\boldsymbol{\delta}_n^T\mathbf{\Gamma}\boldsymbol{\delta}_n\big] &= \frac{1}{2n}\mathbb{E}\Big[\tilde{\boldsymbol{\delta}}_n^T\tilde{\boldsymbol{\delta}}_n\Big] \\
&= \frac{1}{2n}\mathsf{Tr}(\mathsf{Cov}(\tilde{\boldsymbol{\delta}}_n)) + o\left(\frac{1}{n}\right) \\
&= \frac{1}{2n}\mathsf{Tr}(\mathbf{\Gamma}^{1/2}\mathsf{Cov}(\boldsymbol{\delta}_n)\mathbf{\Gamma}^{1/2}) + o\left(\frac{1}{n}\right) \\
&= \frac{1}{2n}\rho d + o\left(\frac{1}{n}\right).
\end{aligned} \tag{S.38}$$

As a result, we conclude that

$$\begin{aligned}\mathbb{E}\Big[\mathcal{C}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)\Big] &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{1}{2n}\mathbb{E}\big[\boldsymbol{\delta}_n^T \boldsymbol{\Gamma} \boldsymbol{\delta}_n\big] + \mathbb{E}(R_{4,n}) \\ &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{\rho d}{2n} + o\Big(\frac{1}{n}\Big). \end{aligned} \tag{S.39}$$

Analysis of the sample objective function $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)$.

We take a second-order expansion of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)$ about $\boldsymbol{\theta}^*$ with a Taylor remainder theorem on the third-order:

$$\begin{aligned}\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big) &= -\frac{1}{n}\sum_{i=1}^{n} h\Big(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\Big) \\ &= -\frac{1}{n}\sum_{i=1}^{n}\Bigg(h\Big(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\Big) + \Big(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\Big)^T \nabla_{\boldsymbol{\theta}} h\Big(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\Big) \\ &\quad + \frac{1}{2}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\Big)^T \nabla^2_{\boldsymbol{\theta}} h\Big(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\Big)\Big(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\Big)\Bigg) \\ &\quad + R_{5,n} \end{aligned} \tag{S.40}$$

where $R_{5,n}$ is the remainder of the expansion that involves $\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^3$ and the third-order derivative of h. By assumption (B4), this quantity has an expectation $\mathbb{E}(|R_{5,n}|) = O(\mathbb{E}\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^3)$. The zeroth-order term is

$\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)$ by definition. To re-express the first-order term, we use the fact

$$\nabla_{\boldsymbol{\theta}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) = 0$$

or equivalently

$$\begin{aligned} 0 &= \frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\right) \\ &= \frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) + \frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right) \\ &\quad + R_{6,n} \end{aligned} \tag{S.41}$$

where $R_{6,n}$ is another remainder of the expansion that involves $\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^2$ and the third-order derivative of h. By assumption (B4), this quantity has an expectation $\mathbb{E}(|R_{6,n}|) = O(\mathbb{E}\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\|^2)$.

Rearranging the equation in (S.41) yields

$$-\frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) = \frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right) + R_{6,n}.$$

Plugging this to the equation in (S.40) results in

$$\begin{aligned}\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= -\frac{1}{n}\sum_{i=1}^{n}\left(h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)+\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T\nabla_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\right.\\ &\quad\left.+\frac{1}{2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T\nabla^2_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)\right)+R_{5,n}\\ &= \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)-\frac{1}{n}\boldsymbol{\delta}_n^T\left(-\frac{1}{n}\sum_{i=1}^{n}\nabla^2_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\right)\boldsymbol{\delta}_n\\ &\quad+\frac{1}{2n}\boldsymbol{\delta}_n^T\left(-\frac{1}{n}\sum_{i=1}^{n}\nabla^2_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\right)\boldsymbol{\delta}_n+R_{5,n}+\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T R_{6,n}.\end{aligned}$$

Let $\boldsymbol{\Gamma}_n\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right) := -\frac{1}{n}\sum_{i=1}^{n}\nabla^2_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)$. We can simplify the above equation as:

$$\begin{aligned}\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)-\frac{1}{n}\boldsymbol{\delta}_n^T\boldsymbol{\Gamma}_n\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\boldsymbol{\delta}_n+\frac{1}{2n}\boldsymbol{\delta}_n^T\boldsymbol{\Gamma}_n\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\boldsymbol{\delta}_n\\ &\quad+R_{5,n}+\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T R_{6,n}\\ &= \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)-\frac{1}{2n}\boldsymbol{\delta}_n^T\boldsymbol{\Gamma}_n\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)\boldsymbol{\delta}_n+R_{5,n}+\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T R_{6,n}\end{aligned}\tag{S.42}$$

Note that the first term $\mathbb{E}\left[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)\right]=\mathcal{C}(\boldsymbol{\theta}^*)$ and by equation (S.25),

$$\boldsymbol{\Gamma}_n\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}^*\right)=\boldsymbol{\Gamma}+\Delta_{n,\boldsymbol{\Gamma}}\left(\widehat{\boldsymbol{\theta}}_n^{(t-1)}\right),$$

where the later analysis in the proof of Lemma 3 showed that

$$\mathbb{E}\Big(\Big\|\Delta_{n,\boldsymbol{\Gamma}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t-1)}\Big)\Big\|\Big) = O\Big(\mathbb{E}\Big(\Big\|\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\Big\|\Big)\Big) = O(1/\sqrt{n}).$$

Thus, equation (S.38) yields that $\mathbb{E}\Big(\boldsymbol{\delta}_n^T \boldsymbol{\Gamma}_n\Big(\widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\Big)\,\boldsymbol{\delta}_n\Big) = \rho d + o(1)$. So the expectation of equation (S.42) becomes

$$\mathbb{E}\Big[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)\Big] = \mathcal{C}(\boldsymbol{\theta}^*) - \frac{\rho d}{2n} + o\Big(\frac{1}{n}\Big). \tag{S.43}$$

Therefore, combining (S.39) and (S.43), the bias of interest is

$$\begin{aligned}\mathbb{E}\Big[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big) - \mathcal{C}\Big(\widehat{\boldsymbol{\theta}}_n^{(t)}\Big)\Big] &= \mathcal{C}(\boldsymbol{\theta}^*) - \frac{\rho d}{2n} - \Big(\mathcal{C}(\boldsymbol{\theta}^*) + \frac{\rho d}{2n}\Big) + o\Big(\frac{1}{n}\Big) \\ &= -\frac{\rho d}{n} + o\Big(\frac{1}{n}\Big).\end{aligned}$$

□

Appendix B: EM Algorithm for Minimizing the Cross-Entropy

This appendix describes our version of the expectation-maximization (EM) algorithm to minimize the cross-entropy from a Gaussian mixture model (GMM) to a target distribution.

The GMM density is expressed as

$$q(\mathbf{x};\boldsymbol{\theta}) = \sum_{j=1}^{k} \alpha_j \, q_j\big(\mathbf{x};\boldsymbol{\mu}_j,\boldsymbol{\Sigma}_j\big)\,, \tag{S.44}$$

where the component weights, $\alpha_j > 0, j = 1,\ldots,k$, satisfy $\sum_{j=1}^{k}\alpha_j = 1$. The jth Gaussian component density q_j is parametrized by the mean $\boldsymbol{\mu}_j$ and the covariance $\boldsymbol{\Sigma}_j$. Thus, the model parameter $\boldsymbol{\theta}$ denotes $(\alpha_1,\ldots,\alpha_k,\boldsymbol{\mu}_1,\ldots,\boldsymbol{\mu}_k,\boldsymbol{\Sigma}_1,\ldots,\boldsymbol{\Sigma}_k)$.

To find the MCE $\widehat{\boldsymbol{\theta}}^{(t)}$ of $\boldsymbol{\theta}$, we want to minimize $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.18) and thus set its gradient to zero:

$$-\frac{1}{\sum_{s=0}^{t-1} n_s}\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\nabla_{\boldsymbol{\theta}} h\Big(\mathbf{X}_i^{(s)},\widehat{\boldsymbol{\theta}}^{(s)},\boldsymbol{\theta}\Big) = 0.$$

This leads to the following updating equations for our version of the EM

algorithm:

$$\alpha_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}}, \tag{S.45}$$

$$\boldsymbol{\mu}_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}\mathbf{X}_i^{(s)}}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}, \tag{S.46}$$

$$\boldsymbol{\Sigma}_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}(\mathbf{X}_i^{(s)}-\boldsymbol{\mu}_j)(\mathbf{X}_i^{(s)}-\boldsymbol{\mu}_j)^T}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}, \tag{S.47}$$

where

$$\gamma_{ij}^{(s)} = \frac{\alpha_j\, q_j\left(\mathbf{X}_i^{(s)};\boldsymbol{\mu}_j,\boldsymbol{\Sigma}_j\right)}{\sum_{j'=1}^{k}\alpha_{j'}\, q_{j'}\left(\mathbf{X}_i^{(s)};\boldsymbol{\mu}_{j'},\boldsymbol{\Sigma}_{j'}\right)}. \tag{S.48}$$

The right-hand sides of the updating equations in (S.45), (S.46), and (S.47) involve

$$\boldsymbol{\theta} = (\alpha_1,\ldots,\alpha_k,\boldsymbol{\mu}_1,\ldots,\boldsymbol{\mu}_k,\boldsymbol{\Sigma}_1,\ldots,\boldsymbol{\Sigma}_k)$$

either explicitly or implicitly through $\gamma_{ij}^{(s)}$. Thus, the equations cannot be analytically solved for $\boldsymbol{\theta}$. Instead, starting with an initial guess of $\boldsymbol{\theta}$, our version of the EM algorithm alternates between the expectation step (computing $\gamma_{ij}^{(s)}$) and the maximization step (updating $\boldsymbol{\theta}$) based on the

updating equations until convergence is reached.

Prior studies [Botev et al., 2013, Wang and Zhou, 2015, Kurtz and Song, 2013] using mixture models for the cross-entropy method for importance sampling (recall that this method is a special case of the procedure in Figure 1 when r is proportional to the optimal importance sampling density) do not iterate their updating equations; instead, they solve them only once when new data are gathered. This paper uses the aforementioned EM algorithm (i.e., iterating the updating equations until convergence) within the t^{th} iteration to minimize $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.18).

Appendix C: Implementation Details of the Numerical Example

This appendix describes the implementation details of the numerical example in Section 4.3 to ensure the reproducibility. For the numerical experiment (with 500 repetitions), we randomly determined $\widehat{\boldsymbol{\theta}}^{(0)} = \boldsymbol{\eta}$ by drawing $\mu_1, \ldots, \mu_{30}$ from a standard multivariate Gaussian and setting all $\Sigma_1, \ldots, \Sigma_{30}$ as $3I_{p\times p}$.

For the implementation of the EM algorithm, we used multiple random initial values of $\boldsymbol{\theta}$ and chose the best minimizer of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.18) to reduce the impact of initial guess of $\boldsymbol{\theta}$ on the algorithm's performance and avoid getting stuck with a local minimizer [Figueiredo and Jain, 2002]. In

the t^{th} iteration for $t \geq 1$, we randomly selected $\mu_1, \ldots, \mu_k$ from $\{\mathbf{X}_i^{(s-1)} : h\Big(\mathbf{X}_i^{(s-1)}, \widehat{\boldsymbol{\theta}}_n^{(s-1)}, \boldsymbol{\theta}\Big) > 0, i = 1, \ldots, n; s = 1, \ldots, t\}$ without replacement. However, if the set's cardinality was smaller than k, we randomly selected any elements in $\{\mathbf{X}_i^{(s-1)} : h\Big(\mathbf{X}_i^{(s-1)}, \widehat{\boldsymbol{\theta}}_n^{(s-1)}, \boldsymbol{\theta}\Big) = 0, i = 1, \ldots, n; s = 1, \ldots, t\}$ for the remaining parameters. We set $\Sigma_1, \ldots, \Sigma_k$ as $(3/p)\,\text{trace}\big(cov(\bar{\mathbf{X}})\big)\, I_{p\times p}$, where $\bar{\mathbf{X}}$ is the data matrix created by augmenting $\{\mathbf{X}_i^{(s-1)} : i = 1, \ldots, n; s = 1, \ldots, t\}$, and *cov* is the sample covariance. We used equal component weights for the initialization, $\alpha_j = 1/k, j = 1, \ldots, k$.

When the number of components, k, became large enough to cause an overfitting issue within the EM algorithm, we caught it by monitoring the condition numbers of the Gaussian components' covariances [Figueiredo and Jain, 2002]. We aborted the EM algorithm when the condition number of any covariance exceeded 10^5. If we needed to abort most of the EM algorithms that started with different initial parameter guesses (we used the threshold of 5 aborted out of 10), it indicated that k is already too large for the given sample size.

To check the convergence of the EM algorithm, we checked the reduction of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (3.18). We stopped iterating updating equations in the EM algorithm if the reduction of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ was less than 1% or a specified maximum number of EM iterations, 10, is reached.

We computed the CIC for $k = k_{\min}, k_{\min}+1, \ldots, k_{\max}$ for the grid search of the minimizer $k^{*(t)}$ in the t^{th} iteration. We set $k_{\min}$ as one for $t = 1$ and $\max\bigl(1, k^{*(t-1)} - 3\bigr)$ for $t \geq 2$. At $k = k_{\min}$, if all random initializations failed to converge, then we reduced $k_{\min}$ by one. In practice, it is generally unnecessary to increase k up to $k_{\max}$, which is upper bounded by a known function of the sample size, because the overfitting is detected within the EM algorithm. To reduce the grid search time, we computed the moving average (with the window size of four) of the CIC and stopped increasing k when the moving average started to increase.