

Approximate Confidence Distribution Computing

SUZANNE THORNTON*, WENTAO LI, AND MINGE XIE

Abstract

Approximate confidence distribution computing (ACDC) offers a new take on the rapidly developing field of likelihood-free inference from within a frequentist framework. The appeal of this computational method for statistical inference hinges upon the concept of a *confidence distribution*, a special type of estimator which is defined with respect to the repeated sampling principle. An ACDC method provides frequentist validation for computational inference in problems with unknown or intractable likelihoods. The main theoretical contribution of this work is the identification of a matching condition necessary for frequentist validity of inference from this method. In addition to providing an example of how a modern understanding of confidence distribution theory can be used to connect Bayesian and frequentist inferential paradigms, we present a case to expand the current scope of so-called approximate Bayesian inference to include non-Bayesian inference by targeting a confidence distribution rather than a posterior. The main practical contribution of this work is the development of a data-driven approach to drive ACDC in both Bayesian or frequentist contexts. The ACDC algorithm is data-driven by the selection of a data-dependent proposal function, the structure of which is quite general and adaptable to many settings. We explore three numerical examples that both verify the theoretical arguments in the development of ACDC and suggest instances in which ACDC outperform approximate Bayesian computing methods computationally.

KEYWORDS AND PHRASES: Approximate Bayesian inference, Confidence distribution, Computational inference.

1. INTRODUCTION

1.1 Approximate Confidence Distribution Computing

Approximate confidence distribution computing (ACDC) is a new take on likelihood-free inference within a frequentist setting. The development of this computational method for statistical inference hinges upon the modern notion of a *confidence distribution*, a special type of estimator which will be defined shortly. Through targeting this special distribution estimator rather than a specific likelihood or posterior distribution as in variational inference and approximate Bayesian inference, respectively, ACDC provides frequentist validation for inference in complicated settings with an unknown or intractable likelihood where dimension-reducing sufficient summary statistics may not even exist. This work demonstrates another example where confidence distribution estimators connect Bayesian and frequent inference, in the surprising context of computational methods for likelihood-free inference [23, 21].

Let $x_{\text{obs}} = \{x_1, \dots, x_n\}$ be an observed sample originating from a data-generating model that belongs to some complex parametric family M_θ . Suppose the likelihood function is intractable (either analytically or computationally), but that this model is generative, i.e. given any $\theta \in \mathcal{P}$, we

can simulate artificial data from M_θ . Let $S_n(\cdot)$ be a summary statistic that maps the sample space into a smaller dimensional space and $r_n(\theta)$ be a data-dependent function on the parameter space. The simplest version of ACDC is the rejection algorithm labeled Algorithm 1 below, where $K_\varepsilon(u) = \varepsilon^{-1}K(u/\varepsilon)$ for a kernel density $K(\cdot)$ satisfying $\max K(u) = 1$ [8] and ε is a small positive value, referred to as the *tolerance level*.

Algorithm 1: Accept-reject approximate confidence distribution computing (ACDC).

1. Simulate $\theta_1, \dots, \theta_N \sim r_n(\theta)$;
 For each $i = 1, \dots, N$,
2. Simulate $x^{(i)} = \{x_1^{(i)}, \dots, x_n^{(i)}\}$ from M_{θ_i} ;
3. Accept θ_i with probability $K_\varepsilon(\|s^{(i)} - s_{\text{obs}}\|)$,
 where $s_{\text{obs}} = S_n(x_{\text{obs}})$ and $s^{(i)} = S_n(x^{(i)})$.

The output of many iterations of Algorithm 1 are potential parameter values, and these potential parameter values are draws from the probability density

$$q_\varepsilon(\theta \mid s_{\text{obs}}) = \frac{\int_{\mathcal{S}} r_n(\theta) f_n(s \mid \theta) K_\varepsilon(\|s - s_{\text{obs}}\|) ds}{\int_{\mathcal{P} \times \mathcal{S}} r_n(\theta) f_n(s \mid \theta) K_\varepsilon(\|s - s_{\text{obs}}\|) ds d\theta}. \quad (1.1)$$

Here, $f_n(s \mid \theta)$ denotes the likelihood of the summary statistic, implied by the intractable likelihood of the data. There-

*Corresponding author.

fore $f_n(s | \theta)$ is typically also intractable. We will refer to $f_n(s | \theta)$ as an *s-likelihood* to emphasize that this is distinct from a traditional likelihood function $f_n(x_{obs} | \theta)$. We denote the cumulative distribution function corresponding to $q_\varepsilon(\theta | s_{obs})$ by $Q_\varepsilon(\theta | s_{obs})$.

The main contribution of this paper is the establishment of a matching condition under which $Q_\varepsilon(\theta | s_{obs})$ is an *approximate confidence distribution* for θ and can be used to derive various types of frequentist inferences. These conditions depend on the choice of $r_n(\theta)$ but are rather general and we present a strategy for choosing an appropriate data-dependent function in Section 3. Practically, this new perspective allows the data to drive the algorithm in a way that can make it more computationally effective than other existing likelihood-free approaches. Theoretically, this perspective establishes frequentist validation for inference from ACDC based upon general conditions that do not depend on the sufficiency of s_{obs} .

Our justification for these practical and theoretical advantages of ACDC relies on the frequentist notion of a confidence distribution. Some background information on confidence distributions is presented next. To motivate the concept, let us consider parameter estimation within a frequentist paradigm. We often desire that our estimators, whether point estimators or interval estimators, have certain properties such as unbiasedness or similar performance under repeated, randomly sampling. A confidence distribution is an extension of this tradition in that it is a distribution estimator (i.e., it is a sample-dependent distribution function) that satisfies certain desirable properties. Following [23] and [18], we define a confidence distribution as follows:

A sample-dependent function on the parameter space is a CONFIDENCE DISTRIBUTION (CD) for a parameter θ if 1) For each given sample, the function is a distribution function on the parameter space; 2) The function can provide valid confidence sets of all levels for θ .

A confidence distribution has a similar appeal to a Bayesian posterior in that it is a distribution function carrying much information about the parameter. A confidence distribution however, is a frequentist notion which treats the parameter as a fixed, unknown quantity and the sampling of data as the random event. A confidence distribution is a sample-dependent function that can be used to estimate the parameter of interest to quantify the uncertainty of the estimation. If $H_n(\cdot)$ is a CD for some parameter θ , then one can simulate $\xi_{CD} \sim H_n(\cdot)$, conditioned upon the observed data. We will refer to the random estimator $\xi_{CD} \sim H_n(\cdot)$ as a CD-RANDOM VARIABLE.

From a Bayesian perspective, the function $r_n(\theta)$ can be viewed as if it is a data-dependent prior. From a frequentist perspective, the data-dependent function $r_n(\theta)$ acts as an initial distribution estimate for θ and Algorithm 1 is a way to update this estimate in search of a better-preforming distribution estimate. This is analogous to any updating algorithm in point estimation requiring an initial estimate that

is updated in search for a better-performing one (e.g., say, a Newton-Raphson algorithm or an expectation-maximization algorithm). Of critical concern in this perspective is avoiding double use of the data for inference. An appropriate choice of the initial distribution estimate, $r_n(\theta)$, addresses this concern and a general strategy for choosing $r_n(\theta)$ is proposed later in Section 3. Therefore, we assert that $Q_\varepsilon(\theta | s_{obs})$ can be used for valid frequentist inference on θ (e.g., deriving confidence sets, p -values, etc.) even if it may (sometimes) not be the most efficient estimator (i.e., may not produce the tightest confidence sets for all $\alpha \in (0, 1)$ levels).

1.2 Related Work on Approximate Bayesian Computing (ABC)

Approximate Bayesian computation (ABC) refers to a family of computing algorithms to approximate posterior densities of θ by bypassing direct likelihood evaluations [cf. 6, 4, 17]. The target of an ABC algorithm is the posterior distribution rather than a confidence distribution. A simple rejection sampling ABC method proceeds in the same manner as Algorithm 1, but it replaces $\theta_1, \dots, \theta_N \sim r_n(\theta)$ with $\theta_1, \dots, \theta_N \sim \pi(\theta)$, a pre-specified prior distribution for θ , in Step 1. In this article, we view ABC as a special case of ACDC where $r_n(\theta) = \pi(\theta)$. The simple rejection sampling ABC is computationally inefficient. Some advanced computing techniques have been used to improve upon the simple ABC approach. One such improvement is the importance sampling version of ABC detailed in Algorithm 2 below. This algorithm can be thought of as a version of ACDC where $r_n(\theta)$ is treated as a proposal distribution. However, as in the rejection sampling version of ABC, IS-ABC still requires a choice of prior which can result in a loss of computational efficiency as we will see in a later section.

Algorithm 2: Importance sampling ABC (IS-ABC).

1. Simulate $\theta_1, \dots, \theta_N \sim r_n(\theta)$.
For each $i = 1, \dots, N$,
2. Simulate $x^{(i)} = \{x_1^{(i)}, \dots, x_n^{(i)}\}$ from M_θ .
3. Accept θ_i with probability $K_\varepsilon(\|s^{(i)} - s_{obs}\|)$,
where $s_{obs} = S_n(x_{obs})$ and $s^{(i)} = S_n(x^{(i)})$, and assign
importance weights $w(\theta_i) = \pi(\theta_i)/r_n(\theta_i)$.

The theoretical argument behind an approximate Bayesian inference (either using the simple rejection sampling ABC or IS-ABC) depends upon $q_\varepsilon(\theta | s_{obs})$ converging to the posterior, $p(\theta | x_{obs}) = \pi(\theta)f_n(x_{obs} | \theta)/\int \pi(\theta)f_n(x_{obs} | \theta)d\theta$, as the tolerance level approaches zero; c.f., e.g. [13] and [2]. However, it is well-known that the quality of this approximation depends not only on the size of ε (and choice of prior) but, also importantly, upon the choice of summary statistic. The choice of $K(\cdot)$ is not essential and does not have significant effect on

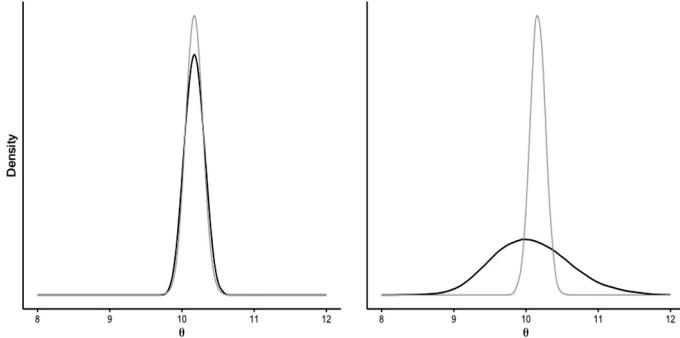


Figure 1: The gray curves below represent the target posterior distribution (gray lines), $p(\theta | x)$, for an $n = 100$ IID sample from $\text{Cauchy}(\theta = 10, 0.55)$. The curves in black represent $q_\varepsilon(\theta | s_{\text{obs}})$ for two different summary statistics, $S_{n1} = \text{Median}(x)$ (left) and $S_{n2} = \bar{x}$ (right). In each case $\varepsilon = 0.005$ and $r_n(\theta) \propto 1$.

the accuracy of ABC estimators when the tolerance level is reasonably small [8, 11]. Common choices of $K(\cdot)$ include normal and uniform kernels. If s_{obs} is not sufficient (as is generally the case in applications of ABC), then the s-likelihood $f_n(s_{\text{obs}} | \theta)$ can be very different from the likelihood of the data $f_n(x_{\text{obs}} | \theta)$ and thus $q_\varepsilon(\theta | s_{\text{obs}})$ can be a very poor approximation to $p(\theta | x_{\text{obs}})$, even as ε approaches zero and $n \rightarrow \infty$.

For example, consider using Algorithm 1 with two different choices of summary statistic, the sample mean or median, for estimating the location parameter of a random sample ($n = 100$) from $\text{Cauchy}(10, 0.55)$. If we suppose $r_n(\theta) \propto 1$, then our algorithm is not data-driven and $r_n(\theta)$ instead acts as an uninformative prior. Hence this example corresponds to an accept-reject version of ABC algorithm. Figure 1 shows $q_\varepsilon(\theta | s_{\text{obs}})$ for each choice of summary statistic (black lines) where $\varepsilon = 0.005$. The posterior distribution (gray lines) does not match well with either approximate ABC posterior distribution $q_\varepsilon(\theta | s_{\text{obs}})$ because only the entire data vector itself is sufficient in this example. This example demonstrates how a strictly Bayesian approach that targets $p(\theta | x_{\text{obs}})$ can produce inconsistent results and even misleading inferential conclusions.

To quote [15]: “the choice of the summary statistic is essential to ensure ABC produces a reliable approximation to the true posterior distribution.” Much of the current literature on ABC methods is appropriately oriented towards the selection and evaluation of the summary statistic. The theoretical justification for inference from ACDC on the other hand, does not require an optimal selection of S_n . Although less informative summary statistics may lead to less efficient CDs, the validity of the inferential conclusions can remain intact even for less informative (and non-sufficient) summary statistics. (See Sections 2.1 and 4.)

The large sample theoretical results presented in Section 3 specify conditions under which Algorithm 1 produces

an asymptotically normal confidence distribution. These results are similar to those in [10] but our work is distinct because we do not target an approximation to a posterior distribution. Instead, the theoretical results in this section focus on the properties and performance of ACDC inherited through its connection to CDs. Additionally, in Section 3 we propose a regression-adjustment technique based on that of [10] and [3]. This post-processing step for ACDC is applied to Algorithms 1 and 2 in Section 4 to improve the accuracy of the CDs.

Computationally, the numerical studies in Section 4 also suggest that ACDC can be more stable than IS-ABC even when both approaches utilize the same data-driven function $r_n(\theta)$. This difference in performance is due to the fact that the importance weights, $w(\theta) = \pi(\theta)/r_n(\theta)$, in IS-ABC can fluctuate greatly, depending on the prior, resulting in numerical instability of the generated parameter values. The steep computing cost associated with the generative model is an expansive area of current research on likelihood-free methods including adaptations that decrease the computing cost of ABC methods such as MCMC methods [14] and sequential Monte Carlo techniques [20]. Although an exploration of these adaptations is beyond the scope of this paper, we expect that many of these approaches can be readily applied to improve the computational performance of ACDC as well. The numerical examples in Section 4 demonstrate how accept-reject ACDC accepts more simulations than IS-ABC suggesting that merely incorporating $r_n(\theta)$ as a data-dependent proposal function is not always computationally preferable.

1.3 Notation and Outline of Topics

In addition to the notation from the introduction, throughout the remainder of paper we will use the following notation. The observed data is $x_{\text{obs}} \in \mathcal{X} \subset \mathbb{R}^n$, the summary statistic is a mapping $S_n : \mathcal{X} \rightarrow \mathcal{S} \subset \mathbb{R}^d$ and the observed summary statistic is $s_{\text{obs}} = S_n(x_{\text{obs}})$. The parameter of interest is $\theta \in \mathcal{P} \subset \mathbb{R}^p$ with $p \leq d < n$; i.e. the number of unknown parameters is no greater than the number of summary statistics and the dimension of the summary statistic is smaller than the dimension of the data. If some function of S_n is an estimator for θ , we will denote this function by $\hat{\theta}_S$.

The next section presents the core theoretical result of this paper which establishes a necessary condition for ACDC methods to produce a valid CD, thereby establishing ACDC as a likelihood-free method that provides valid frequentist inference. Section 3 presents general large sample conditions for ACDC that produce asymptotic CDs and establishes precise conditions for an appropriate choice of the data-dependent function $r_n(\theta)$. Section 4 contains three numerical examples that verify the inferential conclusions of ACDC and illustrate the computational advantages of this data-driven algorithm. Section 5 concludes with a brief discussion. All proofs for Sections 2 and 3 are contained in the Supplementary Material and the Appendix.

2. ESTABLISHING FREQUENTIST GUARANTEES

2.1 General Conditions

In this section, we formally establish conditions under which ACDC can be used to produce confidence regions with guaranteed frequentist coverage for any significance level. To motivate our main theoretical result, first consider the simple case of a scalar parameter and a function $\hat{\theta}_S = \hat{\theta}(S_n)$ which maps the summary statistic, $S_n \in \mathcal{S}$, into the parameter space \mathcal{P} .

Claim. *If*

$$\text{pr}^*(\theta - \hat{\theta}_S \leq t \mid S_n = s_{\text{obs}}) = \text{pr}(\hat{\theta}_S - \theta \leq t \mid \theta = \theta_0), \quad (2.1)$$

then $H_n(t) \stackrel{\text{def}}{=} 1 - Q_\varepsilon(2\hat{\theta}_S - t \mid s_{\text{obs}})$ is a CD for θ .

In the claim, $\text{pr}^*(\cdot \mid S_n = s_{\text{obs}})$ refers to the probability measure on the simulation, conditional on the observed summary statistic, and $\text{pr}(\cdot \mid \theta = \theta_0)$ is the probability measure on the data before it is observed. The proof of this claim (provided in Appendix A) involves showing that $H_n(\theta_0)$ follows a uniform distribution. Once this is established, any $(1 - \alpha)100\%$ level confidence interval for θ can be found by inverting the confidence distribution, $H_n(t)$.

This claim is conceptually similar to the bootstrap central limit theorem which states conditions under which the variability of the bootstrap estimator matches the variability induced by the random sampling procedure. Equation (2.1) instead matches the variability induced by the Monte-Carlo sampling to the random sampling variability. On the left hand side, $\hat{\theta}_S$ is fixed given s_{obs} and the (conditional) probability measure is defined with respect to the Monte-Carlo copies of θ . Thus θ on the left hand side of this equation plays the role of a CD random variable. On the right hand side, the probability measure is defined with respect to the sampling variability where θ is the true parameter value. See also [21] for more discussions of similar matching that link Monte-Carlo randomness with sample randomness across Bayesian, fiducial and frequentist paradigms.

The main condition necessary for valid frequentist inference from ACDC methods is a generalization of the claim above for vector θ .

Condition 1. *For \mathfrak{B} a Borel set on \mathbb{R}^k ,*

$$\sup_{A \in \mathfrak{B}} \left| \text{pr}^*\{V(\theta, S_n) \in A \mid S_n = s_{\text{obs}}\} - \text{pr}\{W(\theta, S_n) \in A \mid \theta = \theta_0\} \right| = o_p(\delta_{n,\varepsilon}),$$

where $\text{pr}^*(\cdot \mid S_n = s_{\text{obs}})$ refers to the probability measure on the simulation, conditional on the observed summary statistic, $\text{pr}(\cdot \mid \theta = \theta_0)$ is the probability measure on the data before it is observed, and $\delta_{n,\varepsilon}$ is a positive rate of convergence that depends on n and ε .

Rather than consider only the linear functions $(\theta - \hat{\theta}_S)$ and $(\hat{\theta}_S - \theta)$, Condition 1 considers any functions $V(\theta, S_n)$ and $W(\theta, S_n)$. (For example, the claim above is a special case of Condition 1 where $V(t_1, t_2) = -W(t_1, t_2) = t_1 - \hat{\theta}_S(t_2)$.) We use the notation $\text{pr}^*\{\cdot \mid S_n = s_{\text{obs}}\}$ because this probability measure is defined over a transformation of the $\theta \sim Q_\varepsilon(\cdot \mid s_{\text{obs}})$.

Furthermore, Condition 1 permits the parameter space and the sample space of the summary statistic to be different from each other. In short, a matching condition on the relationship between two general, multi-dimensional mappings, $V, W : \mathcal{P} \times \mathcal{S} \rightarrow \mathbb{R}^k$ is the key to establishing when ACDC can be used to produce a confidence distribution.

For a given s_{obs} and $\alpha \in (0, 1)$, we can define a set $A_{1-\alpha} \subset \mathbb{R}^k$ such that,

$$\text{pr}^*\{V(\theta, S_n) \in A_{1-\alpha} \mid S_n = s_{\text{obs}}\} = (1 - \alpha) + o(\delta'), \quad (2.2)$$

where $\delta' > 0$ is a pre-selected small, positive precision number, designed to control the Monte-Carlo approximation error. If Condition 1 holds, then

$$\Gamma_{1-\alpha}(s_{\text{obs}}) \stackrel{\text{def}}{=} \{\theta : W(\theta, s_{\text{obs}}) \in A_{1-\alpha}\} \subset \mathcal{P} \quad (2.3)$$

is a level $(1 - \alpha)100\%$ confidence set for θ_0 . We summarize this in the following lemma which is proved in Appendix B.

Lemma 1. *Suppose there exist mappings V and $W : \mathcal{P} \times \mathcal{S} \rightarrow \mathbb{R}^k$ such that Condition 1 holds. Then, $\text{pr}\{\theta_0 \in \Gamma_{1-\alpha}(S_n) \mid \theta = \theta_0\} = (1 - \alpha) + o_p(\delta)$, where $\delta = \max\{\delta_{n,\varepsilon}, \delta'\}$. If Condition 1 holds almost surely, then $\text{pr}\{\theta_0 \in \Gamma_{1-\alpha}(S_n) \mid \theta = \theta_0\} \stackrel{\text{a.s.}}{=} (1 - \alpha) + o(\delta)$.*

Nowhere in Lemma 1 is the sufficiency (or near sufficiency) of S_n required. Of course, if the selected summary statistic happens to be sufficient, then inference from the CD will be equivalent to maximum likelihood inference. Remarkably, Lemma 1 may hold for finite n provided Condition 1 does not require $n \rightarrow \infty$, i.e. provided $\delta_{n,\varepsilon}$ only depends on ε . Later in this section we will consider a special case of Lemma 1 that may be independent of sample-size.

In the next sections we explore some specific situations in which Condition 1 holds. First however, we relate equation (2.2) to a random sample from $Q_\varepsilon(\cdot \mid s_{\text{obs}})$ for vector θ . Suppose θ'_i , $i = 1, \dots, m$, are m draws from $Q_\varepsilon(\cdot \mid s_{\text{obs}})$ and let $v_i = V(\theta'_i, s_{\text{obs}})$. The set $A_{1-\alpha}$ may be a $(1 - \alpha)100\%$ contour set of $\{v_1, \dots, v_m\}$ such that $o(\delta') = o(m^{-1/2})$. For example, we can directly use $\{v_1, \dots, v_m\}$ to construct a $100(1 - \alpha)\%$ depth contour as $A_{1-\alpha} = \{\theta : (1/m) \sum_{i=1}^m \mathbb{I}\{\hat{D}(v_i) < \hat{D}(\theta)\} \geq \alpha\}$, where $\hat{D}(\cdot)$ is an empirical depth function on \mathcal{P} computed from the empirical distribution of $\{v_1, \dots, v_m\}$. See, e.g., [19] and [12] for more on the development of data depth and depth contours in nonparametric multivariate analysis.

2.2 Finite Sample Size Case

We now explore a special case of Lemma 1 where the mappings V and W correspond an approximate pivot statistic. We call a mapping $T = T(\theta, S_n)$ from $\mathcal{P} \times \mathcal{S} \rightarrow \mathbb{R}^d$ an *approximate pivot statistic*, if

$$\text{pr}\{T(\theta, S_n) \in A \mid \theta = \theta_0\} = \int_{t \in A} g(t) dt \{1 + o(\delta'')\}, \quad (2.4)$$

where $g(t)$ is a density free of θ , $A \subset \mathbb{R}^d$ is any Borel set, and δ'' is either zero or a small number (tending to zero) that may or may not depend on the sample size n . For example, suppose $S_n | \theta = \lambda \sim \text{Poisson}(\lambda)$. Then, $T(\lambda, S_n) = (S_n - \lambda)/\sqrt{\lambda}$ is an approximate pivot when λ is large, and the density function is $\phi(t)\{1 + o(\lambda^{-1})\}$, where $\phi(t)$ the density function of the standard normal distribution [5]. The usual pivotal cases are special examples of approximate pivots that may not rely on large sample theory. Examples of approximate pivots where δ'' is a function of n are discussed later in Section 3.

Theorem 1. Suppose $T = T(\theta, S_n)$ is an approximate pivot statistic that is differentiable with respect to the summary statistic and, for given t and θ , let $s_{t,\theta}$ denote a solution to the equation $t = T(\theta, s)$. If

$$\int_{\mathcal{P}} r_n(\theta) K_{\varepsilon}(s_{t,\theta} - s_{\text{obs}}) d\theta = C, \quad (2.5)$$

where C is a constant free of t , then, Condition 1 holds almost surely, for $V(\theta, S_n) = W(\theta, S_n) = T(\theta, S_n)$. (Proof in Appendix C.)

A direct implication of Theorem 1 is that $\Gamma_{1-\alpha}(s_{\text{obs}})$, as defined in (2.3), is a level $(1 - \alpha)100\%$ confidence region where $\text{pr}\{\theta \in \Gamma_{1-\alpha}(S_n) \mid \theta_0\} \stackrel{\text{a.s.}}{=} (1 - \alpha) + o(\delta)$.

The assumption in equation (2.5) needs to be verified on a case-by-case basis. Location and scale families contain natural pivot statistics and satisfy these conditions. This is formally stated in the following corollary.

Corollary 1. (a) Suppose S_n is a point estimator for μ such that $S_n \sim g_1(S_n - \mu)$ and suppose $r_n(\mu) \propto 1$. Then, for any u , $|\text{pr}^*(\mu - S_n \leq u \mid S_n = s_{\text{obs}}) - \text{pr}(S_n - \mu \leq u \mid \mu = \mu_0)| \stackrel{\text{a.s.}}{=} o(1)$.

(b) Suppose S_n is a point estimator for σ such that $S_n \sim g_2(S_n/\sigma)/\sigma$ and suppose $r_n(\sigma) \propto 1/\sigma$. then, for any $v > 0$, $|\text{pr}^*\left(\frac{\sigma}{S_n} \leq v \mid S_n = s_{\text{obs}}\right) - \text{pr}\left(\frac{S_n}{\sigma} \leq v \mid \sigma = \sigma_0\right)| \stackrel{\text{a.s.}}{=} o(1)$.

(c) If $S_{n,1}$ and $S_{n,2}$ are point estimators for μ and σ , respectively, where $S_{n,1} \sim g_1\{(S_{n,1} - \mu)/\sigma\}/\sigma$ and $S_{n,2} \sim g_2(S_{n,2}/\sigma)/\sigma$ are independent and if $r_n(\mu, \sigma) \propto 1/\sigma$, then, for any u and any $v > 0$,

$$\begin{aligned} & \left| \text{pr}^* \left[\left(\begin{array}{l} \mu - S_{n,1} \leq u \\ \frac{\sigma}{S_{n,2}} \leq v \end{array} \right) \mid \begin{pmatrix} S_{n,1} \\ S_{n,2} \end{pmatrix} = \begin{pmatrix} s_{1,\text{obs}} \\ s_{2,\text{obs}} \end{pmatrix} \right] \right. \\ & \quad \left. - \text{pr} \left[\left(\begin{array}{l} S_{n,1} - \mu \leq u \\ \frac{S_{n,2}}{\sigma} \leq v \end{array} \right) \mid \begin{pmatrix} \mu \\ \sigma \end{pmatrix} = \begin{pmatrix} \mu_0 \\ \sigma_0 \end{pmatrix} \right] \right| \stackrel{\text{a.s.}}{=} o(1). \end{aligned}$$

Consequently, $H_1(S_{n,1}, x) = \int_{-\infty}^x g_1(S_{n,1} - u) du$ is a CD for μ and $H_2(S_{n,2}^2, x) = 1 - \int_0^x g_2(\hat{\sigma}_S/u) du$, is a CD for σ .

Note that Theorem 1 and Corollary 1 cover some finite sample size scenarios, including the Cauchy example discussed in Section 1. For this example, Corollary 1 (part (a)) asserts that the different posterior approximations obtained by approximate Bayesian computing with either $S_n = \text{Median}(x)$ or $S_n = \bar{x}$ are both CDs. That is, both densities in black in Figure 2 are densities for confidence distributions of θ , obtained by Algorithm 2. These distribution estimators lead to valid frequentist inference even though neither summary statistic is sufficient. This development represents a departure from the typical asymptotic arguments for likelihood-free computational inference. Note, a practical issue with applying Corollary 1 to Algorithm 1 is that the user can not actually simulate from the required $r_n \propto 1$. In such cases, we suggest using the minibatch scheme, introduced in Section 3.2, to approximate r_n as we did for the Cauchy example in Section 4.1. The comparison of this approach with Algorithm 2 is given in Table 1.

This section has considered the case in which the tolerance level, ε , does not necessarily depend on the sample size n . In the next section, the tolerance may depend on the sample size so we adopt the notation ε_n to reflect this.

3. LARGE SAMPLE THEORY

3.1 A Bernstein-von Mises Theorem for ACDC

In the Bayesian ABC framework, Condition 1 holds as $n \rightarrow \infty$ by selecting a ε_n that decreases to zero at a certain rate. [11]. We now verify Condition 1 holds more generally for ACDC methods that use a data-dependent $r_n(\theta)$, in a large sample setting. The results presented here are generalizations of results in [10] and [11] and extend them from the Bayesian framework to the confidence distribution framework. We give sufficient conditions with which allowing $r_n(\theta)$ to depend on the data does not lead to the overestimation of statistical efficiency, ie. the ‘double use’ of data which is of concern for calibrated inferential methods. Roughly speaking, the next theorem establishes that the distribution of a centered random draw from $Q_{\varepsilon}(\theta \mid s_{\text{obs}})$ and the distribution of its centered expectation (before the data is observed), i.e. $\int \theta dQ_{\varepsilon}(\theta \mid S_n)$, are asymptotically the same.

The next condition concerning the asymptotic behavior of the summary statistic is crucial for the proofs of the theorems in this section (see Appendix F).

Condition 2. There exists a sequence $\{a_n\}$, satisfying $a_n \rightarrow \infty$ as $n \rightarrow \infty$, a d -dimensional vector $s(\theta)$, a $d \times d$ matrix $A(\theta)$, and some $\delta_0 > 0$ such that for $S_n \sim f_n(\cdot \mid \theta)$ and all $\theta \in \mathcal{P}_0 \stackrel{\text{def}}{=} \{\theta : \|\theta - \theta_0\| < \delta_0\} \subset \mathcal{P}$,

$$a_n \{S_n - s(\theta)\} \xrightarrow{d} N\{0, A(\theta)\}, \text{ as } n \rightarrow \infty,$$

and $s_{\text{obs}} \xrightarrow{P} s(\theta_0)$. Furthermore, assume that

- (i) $s(\theta), A(\theta) \in C^1(\mathcal{P}_0)$, $A(\theta)$ is positive definite for all θ ;
- (ii) for any $\delta > 0$ there exists a $\delta' > 0$ such that $\|s(\theta) - s(\theta_0)\| > \delta'$ for all θ such that $\|\theta - \theta_0\| > \delta$; and
- (iii) $I(\theta) \stackrel{\text{def}}{=} \left\{ \frac{\partial}{\partial \theta} s(\theta) \right\}^T A(\theta)^{-1} \left\{ \frac{\partial}{\partial \theta} s(\theta) \right\}$ has full rank at $\theta = \theta_0$.

This is a standard condition but, notably, does not depend on the sufficiency of this statistic. Because of this, we refrain from discussing this condition further so we may instead focus on our main contribution, the development of the following regulatory conditions on $r_n(\theta)$.

Condition 3. For all $\theta \in \mathcal{P}_0$, $r_n(\theta) \in C^2(\mathcal{P}_0)$ and $r_n(\theta_0) > 0$.

Condition 4. There exists a sequence $\{\tau_n\}$ such that $\tau_n = o(a_n)$ and $\sup_{\theta \in \mathcal{P}_0} \tau_n^{-p} r_n(\theta) = O_p(1)$.

Condition 5. There exists constants m, M such that $0 < m < |\tau_n^{-p} r_n(\theta_0)| < M < \infty$.

Condition 6. It holds that $\sup_{\theta \in \mathbb{R}^p} \tau_n^{-1} D\{\tau_n^{-p} r_n(\theta)\} = O_p(1)$.

Condition 3 is a general assumption regarding the differentiability of $r_n(\theta)$ within an open neighborhood of the true parameter value. Condition 4 and 5 essentially require $r_n(\theta)$ to be more dispersed than the s-likelihood within a compact set containing θ_0 . They require $r_n(\theta)$ converge to a point mass more slowly than $f_n(\theta | s_{\text{obs}})$. Condition 6 requires the gradient of the standardized version of $r_n(\theta)$ to converge with rate τ_n . Condition 3–6 are relatively weak conditions and can be satisfied with locally asymptotic $r_n(\theta)$, for example. They are also satisfied by the data-independent prior used in approximate Bayesian computation.

The proofs of the theorems in this section also require additional conditions (Conditions 7–10 of the supplementary material) that are typical of BvM-type theorems. These additional conditions are not presented here for readability reasons and because they do not directly relate to $r_n(\theta)$ which is our emphasis.

Theorem 2. Let $\hat{\theta}_S = \hat{\theta}(S_n) = \int \theta dQ_\varepsilon(\theta | S_n)$. Assume Condition 2. Assume $r_n(\theta)$ satisfies Condition 3–6 above and also Conditions 7–10 in the supplementary material. If $\varepsilon_n = o(a_n^{-1})$ as $n \rightarrow \infty$, then Condition 1 is satisfied with $V(\theta, S_n) = a_n(\theta - \hat{\theta}_{s_{\text{obs}}})$ and $W(\theta, S_n) = a_n(\hat{\theta}_S - \theta)$. (Proof in Appendix 8.)

Theorem 2 says when $\varepsilon_n = o(a_n^{-1})$, the coverage of $\Gamma_{1-\alpha}(s_{\text{obs}})$ is asymptotically correct as n and the number of accepted parameter values increase to infinity. In practice, $\hat{\theta}_S$ typically will not have a closed form. To construct $\Gamma_{1-\alpha}(s_{\text{obs}})$, the value of $\hat{\theta}$ at $S_n = s_{\text{obs}}$ can be estimated using the accepted parameter values from ACDC. Here Condition 1 is satisfied by generalizing the limit distributions of the approximate posterior in [10] so they hold also for $Q_\varepsilon(\theta | s_{\text{obs}})$, when $\varepsilon_n = o(a_n^{-1})$. Specifically, for A defined

as in equation (2.2),

$$\sup_{A \in \mathfrak{B}^p} \left| \int_{\{\theta: a_n(\theta - \hat{\theta}) \in A\}} dQ_\varepsilon(\theta | s_{\text{obs}}) - \int_A N\{t; 0, I(\theta_0)^{-1}\} dt \right| \xrightarrow{P} 0 \quad (3.1)$$

and

$$a_n(\hat{\theta} - \theta_0) \xrightarrow{d} N\{0, I(\theta_0)^{-1}\}, \quad (3.2)$$

as $n \rightarrow \infty$, where $I(\theta_0)$ is a non-singular matrix defined in Condition 2. Thus inference based on $Q_\varepsilon(\theta | s_{\text{obs}})$ is valid for $n \rightarrow \infty$ regardless of whether or not $r_n(\theta)$ depends on the data. For the same tolerance level, Theorem 2 asserts that the limiting distribution of $Q_\varepsilon(\theta | S_n)$ matches the limiting distribution of the approximate posterior from [10] which is the output distribution of the accept-reject version of ABC. In comparison however, ACDC has a better acceptance rate since the data-dependent $r_n(\theta)$ concentrates more probability mass around θ_0 than a typical prior.

Although inference from ACDC is validated with $\varepsilon_n = o(a_n^{-1})$, a well-known issue in approximate Bayesian literature is that this tolerance level is too small in practice, causing the degeneration of the acceptance rate as $n \rightarrow \infty$ for any proposal distribution [11]. Obviously ACDC methods will suffer from this same issue. (For an example with Normal data, see Appendix E.)

One remedy that relaxes the restriction on ε_n is to post-process the sample from $Q_\varepsilon(\theta | s_{\text{obs}})$ with a regression adjustment [1]. When the data-generating model is correctly specified, the regression adjusted sample correctly quantifies the CD uncertainty and yields an accurate point estimate with ε_n decaying at a rate of $o(a_n^{-3/5})$ [10].

Let $\theta^* = \theta - \beta_\varepsilon(s - s_{\text{obs}})$ be the post-processed sample from $Q_\varepsilon(\theta | s_{\text{obs}})$, where β_ε is the minimizer from

$$\begin{pmatrix} \alpha_\varepsilon \\ \beta_\varepsilon \end{pmatrix} = \arg \min_{\alpha \in \mathbb{R}^p, \beta \in \mathbb{R}^{d \times p}} E_\varepsilon \{ \|\theta - \alpha - \beta(s - s_{\text{obs}})\|^2 | s_{\text{obs}} \}$$

for expectation under the joint distribution of accepted θ values and corresponding summary statistics.

Theorem 3. Under the conditions of Theorem 2, if $\varepsilon_n = o(a_n^{-3/5})$ as $n \rightarrow \infty$, Condition 1 holds with $V(\theta, S_n) = a_n(\theta^* - \hat{\theta}_{s_{\text{obs}}})$ and $W(\theta, S_n) = a_n(\hat{\theta}_S^* - \theta)$, where $\hat{\theta}_S^*$ is the expectation of the post-processed observations of the CD random variable. (Proof in Appendix F.)

Here, Condition 1 is implied by the following convergence results (where A defined as in equation (2.2)),

$$\sup_{A \in \mathfrak{B}^p} \left| \int_{\{\theta: a_n(\theta - \hat{\theta}^*) \in A\}} dQ_\varepsilon^*(\theta | s_{\text{obs}}) - \int_A N\{t; 0, I(\theta_0)^{-1}\} dt \right| \xrightarrow{P} 0,$$

and

$$a_n(\hat{\theta}_S^* - \theta_0) \xrightarrow{d} N\{0, I(\theta_0)^{-1}\},$$

as $n \rightarrow \infty$. The limiting distributions above are the same as those in (3.1) and (3.2), therefore $\Gamma_{1-\alpha}(\mathbf{s}_{\text{obs}})$ constructed using the post-processed sample achieves the same efficiency as that constructed with the original ACDC sample of θ values. The benefit of permitting larger tolerance levels is a huge improvement in the computing costs associated with ACDC.

3.2 Designing r_n

Condition 4 implies that in practice, one must take care to choose $r_n(\theta)$ so that its growth with respect to the sample size is slower than the growth of the s-likelihood. In this section we propose a generic algorithm to construct such $r_n(\theta)$ based on subsets of the observed data.

Notably, there is a trade-off in ACDC inference between faster computations and guaranteed coverage of the approximate CD-based confidence sets. When $r_n(\theta)$ grows at a similar rate as the s-likelihood for $n \rightarrow \infty$, the computing time for ACDC methods may be reduced but this risks violating Conditions 4–6. If these assumptions are violated, the distribution of the resulting simulations is not necessarily a CD and consequently, inference may not be valid in terms of producing confidence sets with guaranteed coverage. Therefore, $r_n(\theta)$ should be designed such that its convergence rate is bounded away from that of the s-likelihood. The minibatch scheme presented below is one way to ensure $r_n(\theta)$ is appropriately bounded.

Assume that a point estimator $\hat{\theta}_S(z)$ of θ can be computed for a dataset, z , of any size.

Minibatch scheme

1. Choose k subsets of the observations, each with size n^ν for some $0 < \nu < 1$.
2. For each subset z_i of x_{obs} , compute the point estimate $\hat{\theta}_{S,i} = \hat{\theta}_S(z_i)$, for $i = 1, \dots, k$.
3. Let $r_n(\theta) = (1/kh) \sum_{i=1}^k K \left\{ h^{-1} \|\theta - \hat{\theta}_{S,i}\| \right\}$, where $h > 0$ is the bandwidth of the kernel density estimate using $\{\hat{\theta}_{S,1}, \dots, \hat{\theta}_{S,k}\}$ and kernel function K .

The choice of $K(\cdot)$ follows that of the standard multivariate kernel density estimation. If $\hat{\theta}_S$ is consistent, then for $\nu < 3/5$, $r_n(\theta)$ as obtained by this minibatch procedure will satisfy Conditions 4–6. Based on our experience, if n is large one may simply choose $\nu = 1/2$ to partition the data. For small n , say $n < 100$, it is better to select $\nu > 1/2$ and to overlap the subsets (or “mini” batches of the observed data) so that each subset contains a reasonable number of observations. For a given summary statistic, there are many methods to construct this type of point estimator including: a minimum distance-based optimizer [9, 16], the synthetic likelihood method and its variants [22, 7], or accept-reject ACDC with $\hat{\theta}_S = E\{\theta | S_n(z_i)\}$, the s-likelihood-based expectation over a subset of the observed data. The choice of $\hat{\theta}_S$ does not need to be an accurate estimator since it is only used to construct the initial rough distribution estimator

for θ . But, a heavily biased $\hat{\theta}_S$ causes bias in confidence sets derived from the CD, since $r_n(\theta)$ does not cover parameter values resulting in high values of $f_n(s | \theta)$ very well. In practice, the computing cost will depend on which particular optimization scheme is followed. However, a full study on the selection of $\hat{\theta}_S$ is beyond the scope of this paper.

The computational cost associated with implementing the minibatch scheme is comparable to the cost of constructing a proposal distribution for IS-ABC methods. Multiple runs to compute $\hat{\theta}_{S,i}$ values can be parallelized easily and any procedure to obtain a proposal distribution for IS-ABC can be applied on the mini batches of data to yield a point estimate for θ . For example, for each subset z_i , the conditional mean $E\{\theta | S_n(z_i)\}$ can be estimated by population Monte Carlo ABC on $S_n(z_i)$. This is not any more computationally expensive than computing the same estimate on the full data. This, together with the fact that accept-reject ACDC accepts more simulations than IS-ABC, make ACDC the favorable choice in terms of overall computational performance. The numerical examples in Section 4 support this conclusion.

At this point, our reader may wonder if $\hat{\theta}_S$ can be computed, why not simply use a bootstrap method to construct confidence sets? Although it requires no likelihood evaluation, this method has two significant drawbacks. First, the bootstrap method is heavily affected by the quality of $\hat{\theta}_S$. For example, a bootstrapped confidence interval for θ is based on quantiles of $\hat{\theta}_S$ from simulated data. A poor estimator typically leads to poor performing confidence sets. In contrast, in ACDC methods, $\hat{\theta}_S$ is only used to construct the initial distribution estimate which is then updated by the data. Second, when it is more computationally expensive to obtain $\hat{\theta}_S$ than the summary statistic, the bootstrap will be much more costly than ACDC methods since $\hat{\theta}_S$ must be calculated for each pseudo data set. Example 4.3 in the next section illustrates such an example.

4. NUMERICAL EXAMPLES

4.1 Location and Scale Parameters for Cauchy Data

In the Cauchy example presented in Figure 1 we saw how the lack of a sufficient summary statistic can change the validity of inferential conclusions from an ABC approach. Through a CD perspective however, the inferential conclusions from ACDC are valid under the frequentist criterion even if the summary statistic is not sufficient. Provided Condition 1 is satisfied, different summary statistics produce different CDs. Here we present a continuation of this Cauchy example where random data ($n = 400$) is drawn from a $\text{Cauchy}(\theta, \tau)$ distribution with data-generating parameter values $(\theta_0, \tau_0) = (10, 0.55)$. We investigate the performance of 500 independent 95% confidence intervals for θ alone (settings one and two) and τ alone (setting three) and 500 independent 95% confidence regions when both parameters are

unknown (settings four and five). For settings two and four, the summary statistics are less informative than those in the other settings. Here the information of summary statistic is measured by a sandwich-type of information matrix, $I(\theta_0)$, defined in Condition 2. They are chosen in order to demonstrate the performance in the scenario of Section 2.2 and show that ACDC can provide valid confidence regions when it is difficult to construct a valid confidence region with the summary-based posterior distribution.

In each setting, the confidence intervals (regions) are generated for both ACDC and IS-ABC utilizing the same minibatch scheme to construct r_n with the median and/or the median absolute deviation (MAD) as point estimators and $v = 1/2$. The main difference in the two algorithms is the use of r_n . In ACDC r_n is a data-driven initial CD estimate whereas the r_n in IS-ABC represents a Bayesian approach that assumes the uninformative prior in Corollary 1 and employs $r_n(\theta)$ as the proposal distribution for the importance sampling updates. Both algorithms are improved by adapting the regression adjustments mentioned in Section 3.1, so the output for every run of each algorithm is post-processed in this manner. The Monte Carlo sample size is chosen so that the tolerance level ε is small enough and the number of accepted simulations is reasonable. Therefore, the reported results show the effect of the importance sampling weights.

Table 1 compares frequentist coverage proportions of confidence regions from both algorithms. The acceptance proportion determines how many simulated parameter values are kept and thus is directly related to the tolerance level. Most coverage rates are close to the nominal levels when the acceptance proportion is small, which is expected from the asymptotic theory in Section 3. Overall the coverage performance is similar for both algorithms. For settings one, three, and five with the more informative summary statistic(s), both algorithms give similar confidence regions which undercover a bit in the finite-sample regime. For settings two and four with less informative summary statistics, ACDC is preferable because it produces tighter confidence bounds while still attaining at least the nominal 95% coverage level).

The main reason for the favorable performance of ACDC is related to the skewed importance weights for IS-ABC. This can be seen in the sizes of the confidence sets of the two algorithms in Table 1 but is even more clear when comparing the CD densities of each method as in Figure 2. Figure 2 shows the impact of importance weights in IS-ABC on the variances of point estimators and CDs. For settings 1, 3 and 5 where an informative summary statistic is used, the importance weights do not much affect either the point estimator or resulting CDs. In these cases, $r_n(\theta)$ is a good proposal distribution according to the criteria in [11]. For settings 2 and 4 however, where the summary statistic is less informative, Figure 2 shows how the importance weights inflate both point estimate and CD variances with Monte Carlo variation in IS-ABC. One reason for the severe skewedness in the importance weights $\pi(\theta)/r_n(\theta)$ is that the high variance of

S_n means more parameter simulations are accepted in the tails of $r_n(\theta)$. This results in broader confidence regions for IS-ABC than ACDC. An experiment with a smaller sample size ($n = 100$) was also carried out and we found that the performance of both algorithms and their comparison is similar to that of the case were $n = 400$. The only difference is that for setting 1, the coverage of ACDC is 0.9 and that of IS-ABC is 0.92, reflecting slight undercoverage likely due to the smaller sample size shifting away from the asymptotic regime.

This numerical study validates inference for both ACDC and IS-ABC even in the case where typical asymptotic arguments do not apply (settings 2 and 4). Furthermore, this example demonstrates two valid but distinct uses of the minibatch scheme for constructing a data driven distribution estimator. In Algorithm 1, r_n drives the search for a distribution estimator or acts as a data-dependent prior within a Bayesian context. In Algorithm 2, r_n acts as a proposal distribution for ABC with a flat prior. The computational differences in the performance of confidence regions in this example suggest that the former application of r_n is preferable to the latter if the summary statistic is not very informative. Interestingly, even though the IS-ABC algorithm fails to produce Bayesian posterior distributions, it can still provide us with valid frequentist inference.

4.2 Linear Regression with Cauchy Errors

To study the impact of dimensionality on ACDC, we consider the linear model $y_i = \sum_{j=1}^d x_j \beta_j + e_i$, where x_j and β_j are scalars, $i = 1, \dots, n$ and e_i are identically and independently distributed from $\text{Cauchy}(0, 1)$, and the parameter of interest is $\beta = (\beta_1, \dots, \beta_d)$. We examine synthetic data generated from the model with covariates identically and independently following the standard normal distribution and the true coefficients $\beta_j = j$. For a data set Y , the least squares estimator is chosen as the summary statistic, denoted by $\hat{\beta}_Y$. Similar to setting 2 in Example 4.1, it is unbiased but does not have finite variance. The least squares estimator is also used as the point estimator to construct r_n in the minibatch scheme. When constructing r_n , we use $v = 3/5$, e.g. subsets of size 24 which is a reasonable data size for estimating the 5-dimensional linear coefficients. To ensure that the proposed parameter values cover the parameter space with high likelihood values reasonable well, we bootstrap the data to obtain a total of 60 subsets. Since $\hat{\beta}_{y_s}$ follows the multivariate Cauchy distribution with the mean vector β and the scale matrix $(X_s^T X_s)^{-1/2}$, where X_s is the design matrix of the subset y_s , r_n is chosen to be the equally weighted mixture of the 60 Cauchy distributions centred at $\hat{\beta}_{y_{1:n}}$ and has the scale matrix $(X_s^T X_s)^{-1/2}$. To obtain a confidence region for β , since $\hat{\beta}_{y_{1:n}}$ follows a location family, the confidence region can be obtained via the pivot $(X_{1:n}^T X_{1:n})^{-1/2}(\beta - \hat{\beta}_{y_{1:n}})$. More specifically, the function $(\beta - \hat{\beta}_{y_{1:n}})^T (X_{1:n}^T X_{1:n})^{-1}(\beta - \hat{\beta}_{y_{1:n}})$ is used as the data depth function.

Table 1. Coverage proportions of confidence sets from ACDC applied to Cauchy data under five different settings. Coverage is calculated over 500 independent runs that draw a $n = 400$ IID sample from a $\text{Cauchy}(\theta = 10, \tau = 0.55)$ distribution. The Monte Carlo sample size for both algorithms is 50,000 and the nominal coverage level in every setting is 95%. The last column displays the median and standard deviation (in bracket) of size ratios of confidence sets from ACDC divided by those from IS-ABC.

	Acceptance proportion	ACDC Coverage	IS-ABC Coverage	Ratio of Widths/Volumes
Setting 1: θ unknown				
$S_n = \text{Median}(x)$	0.005 0.05 0.10	0.93 0.94 0.93	0.94 0.94 0.94	0.94(0.042) 0.94(0.03) 0.94(0.029)
Setting 2: θ unknown				
$S_n = \bar{x}$	0.005 0.05 0.10	0.97 0.97 0.97	0.98 0.98 0.98	0.65(0.28) 0.60(0.20) 0.56(0.19)
Setting 3: τ unknown				
$S_n = \text{MAD}(x)$	0.005 0.05 0.10	0.93 0.92 0.93	0.94 0.94 0.94	1.00(0.075) 1.00(0.044) 1.00(0.038)
Setting 4: $(\theta, \tau)'$ both unknown				
$S_n = \begin{pmatrix} \bar{x} \\ \text{SD}(x) \end{pmatrix}$	0.005 0.05 0.10	0.96 0.99 0.99	0.96 0.98 0.98	0.58(0.41) 0.48(0.59) 0.47(0.58)
Setting 5: $(\theta, \tau)'$ both unknown				
$S_n = \begin{pmatrix} \text{Median}(x) \\ \text{MAD}(x) \end{pmatrix}$	0.005 0.05 0.10	0.91 0.94 0.94	0.93 0.96 0.97	0.98(0.060) 1.00(0.052) 1.00(0.048)

In Table 2, we see that ACDC has coverage closer to the nominal level than IS-ABC and produces significantly smaller confidence regions in all settings, up to more than an order of magnitude. This is because the skewness of the importance weights of IS-ABC is more severe when more simulations are accepted in the tails of the proposal distribution, which is the case when the tolerance level is larger. In ACDC, a larger tolerance level corresponds to a greater reduction in the Monte Carlo variance because it avoids these importance weights. The coverage of ACDC is better because of the smaller Monte Carlo variance. Since the ratio of the volumes of the confidence regions is an exponential function of the number of unknown parameters, the improvement is more significant in higher-dimensions. This example shows the practical benefit of Corollary 1 in a moderately high dimensional setting.

4.3 Inference for a Ricker Model

A Ricker map is a non-linear dynamical system, often used in Ecology, that describes how a population changes over time. The population, N_t , is noisily observed and is described by the following model,

$$y_t \sim \text{Pois}(\phi N_t), \\ N_t = r N_{t-1} e^{-N_{t-1} + e_t}, e_t \sim N(0, \sigma^2),$$

where $t = 1, \dots, T$ and parameters r , ϕ and σ are positive constants, interpreted as the intrinsic growth rate of the population, a scale parameter, and the environmental noise, respectively. This model is computationally challenging since its likelihood function is intractable for $\sigma > 0$ and is highly irregular in certain regions of the parameter space.

We investigate the performance of confidence sets for each parameter marginally and two pairs of parameters jointly. We adopt the setting and choice of summary statistic from [22]. In [22], the summary statistic is crafted based on domain knowledge to provides accurate inference. It satisfies a central limit theorem, so according to Section 3.1, both the ABC posterior distribution and the CD produced by ACDC have the same limit distributions. This means that without considering the Monte Carlo error, the inference of both achieves the nominal coverage. The experiments here illustrates the effect of additional Monte Carlo variance from the importance weights using by Algorithm 2. The output of both algorithms are post-processed using the regression adjustment.

In the minibatch scheme, for the point estimator we use $E\{\theta \mid S_n(z_i)\}$ estimated by the population Monte Carlo version of IS-ABC. The maximum synthetic likelihood estimator proposed in [22] was also tried, but the estimates obtained this way over-concentrated in a certain area of the parameter space. The corresponding $r_n(\theta)$ did not cover the

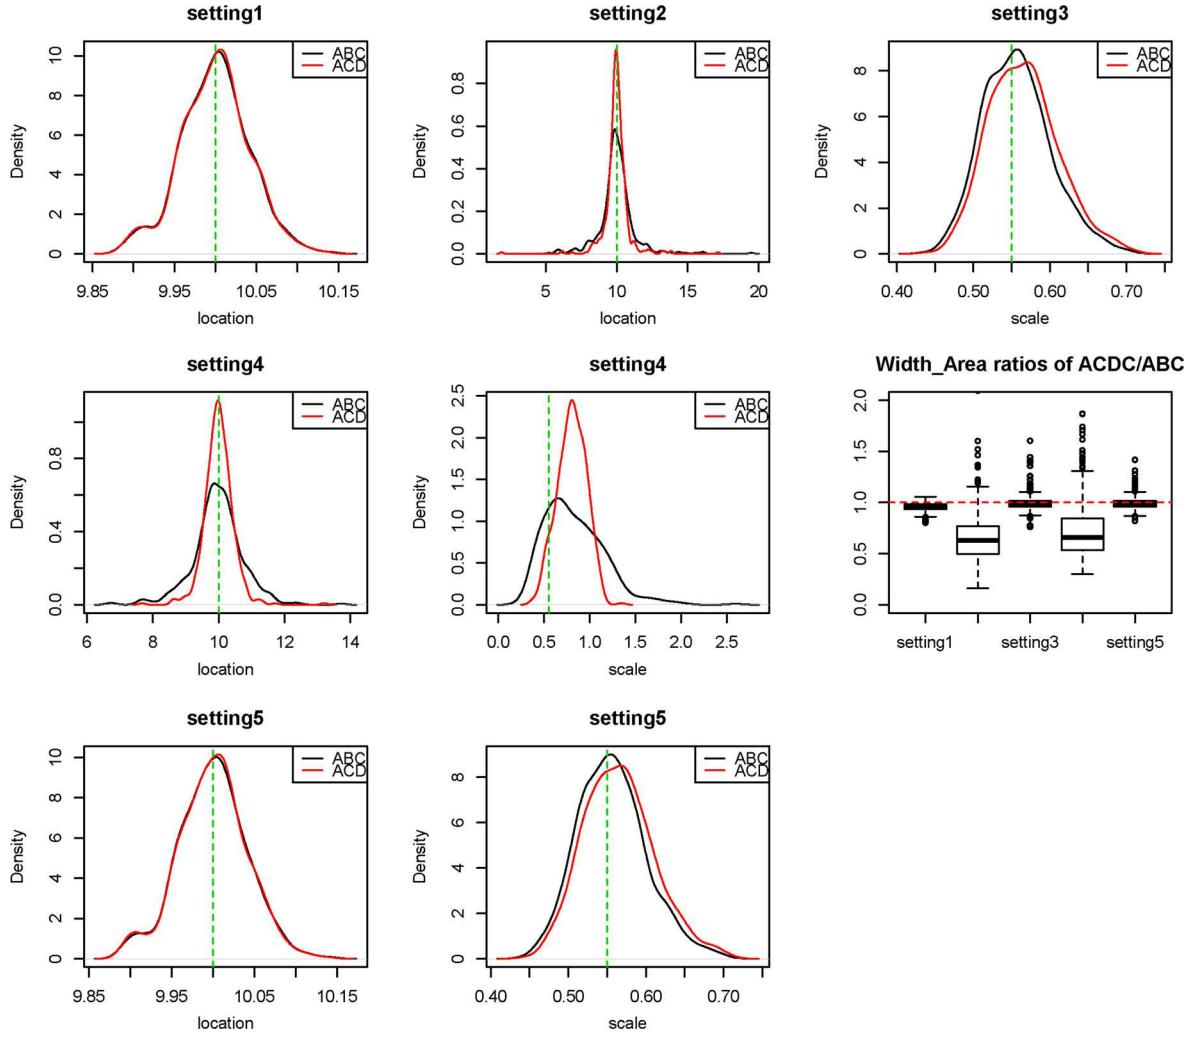


Figure 2: These are densities of point estimators from ACDC (red) and IS-ABC (black) for the 500 independent data sets for each of the five settings in Table 1. Additionally, this figure shows a box plot of the relative sizes of the 500 confidence sets, that is, the length (or volume) of regions produced by ACDC divided by those of IS-ABC.

target mass very well biasing the coverage levels. Instead, $r_n(\theta)$ is used to initialise the population Monte Carlo iterations. Since the sample size is small in this example, overlapping minibatches are chosen with a total number of 40 where each minibatch contains an observed series of length 10. In this example, the bootstrap method is not feasible because it is too computationally expensive to use the simulation-based methods in obtaining the point estimates.

In Table 3, when the acceptance proportion is small, most coverage rates of ACDC are close to the nominal level. In contrast, the confidence bounds from IS-ABC display more over-coverage, indicating an even smaller ε is needed to reduce the variance inflation. Furthermore, all ACDC confidence regions are tighter with a size reduction up to 51%. The box plots in Figure 3 show that the CD variances from IS-ABC are inflated substantially by the importance

weights, resulting in broader confidence regions as observed in the last column of Table 3.

As in Section 4.1, this numerical example validates the inferential conclusions from both algorithms but here we see ACDC consistently producing tighter confidence regions than IS-ABC. The summary statistics in this example were carefully selected to be informative based on domain knowledge. Nevertheless, ACDC still avoids the excessive Monte Carlo variation that impedes IS-ABC.

5. DISCUSSION

In this article, we propose ACDC as a new inference-based approach to likelihood-free methods. ACDC can provide valid frequentist inference for target parameters from data without a tractable likelihood. Although ACDC can be viewed as an extension of ABC, a crucial difference is that ACDC does not require any prior assumptions nor does

Table 2. Coverage proportions of confidence sets from ACDC applied to Cauchy regression to estimate the linear coefficients under different data size n and number of covariates d . Coverage is calculated over 200 independent runs. The Monte Carlo sample size is 10^6 and the nominal coverage level in every setting is 95%. The last column displays the median and median absolute deviation (in bracket) of size ratios of confidence sets from ACDC divided by those from IS-ABC.

Acceptance proportion	ACDC Coverage	IS-ABC Coverage	Ratio of Widths/Volumes
Setting 1: $n = 100, d = 2$			
0.005	0.95	0.96	0.23(0.13)
0.05	0.94	0.92	0.13(0.064)
0.1	0.94	0.88	0.11(0.06)
Setting 2: $n = 100, d = 5$			
0.005	0.96	0.88	0.13(0.11)
0.05	0.88	0.75	0.047(0.031)
0.1	0.84	0.73	0.04(0.027)
Setting 3: $n = 200, d = 2$			
0.005	0.94	0.98	0.23(0.14)
0.05	0.95	0.94	0.15(0.075)
0.1	0.94	0.90	0.12(0.063)
Setting 4: $n = 200, d = 5$			
0.005	0.96	0.89	0.15(0.13)
0.05	0.88	0.74	0.056(0.046)
0.1	0.86	0.74	0.048(0.037)

the validity of inferential conclusions depend upon the near-sufficiency of the summary statistic. Computationally, an ACDC approach is preferable when compared to the corresponding IS-ABC method which suffers from skewed importance weights. The most costly step of likelihood-free methods in general is typically the generation of artificial data (i.e. Step 2 in Algorithm 1 and 2). The ACDC approach with the minibatch scheme is more efficient than other likelihood-free approaches that are not able to orient the data-generating step around a data-driven choice of parameter value. However, the main advantage of ACDC over other likelihood-free approaches, is the ability to draw proper, calibrated inferential conclusions about the unknown parameter.

The main theoretical contribution of this work is the identification of a matching condition (Condition 1) necessary for valid frequentist inference from ACDC methods. This condition is similar to the theoretical support for bootstrap estimation and is met in cases that rely on typical asymptotic arguments (e.g. reference citations in Section 3) but also applies to certain small-sample cases. Additionally, a key practical contribution of this work is the general minibatch method for initializing ACDC estimators. This approach guides the search for a well-behaved distribution estimator using a data-dependent distribution $r_n(\theta)$. This can result in improved computational performance even compared to an IS-ABC method that is similarly data-driven. In cases where $r_n(\theta)$ does not yield reasonable acceptance

Table 3. Coverage proportions of marginal confidence intervals (or joint confidence regions) for ACDC and IS-ABC applied to Ricker data. Coverage is calculated over 150 independent runs that produce observations from $t = 51$ to 100 for data generated by a Ricker model with $(r, \sigma, \phi) = (e^{3.8}, 0.3, 10)$. The Monte Carlo sample size for both algorithms is 50,000 and the nominal coverage level in every setting is 95%. The last column displays the median ratio of the sizes of confidence sets from ACDC divided by those from IS-ABC.

Acceptance proportion	ACDC Coverage	IS-ABC Coverage	Ratio of Widths/Volumes
Setting 1: $\log(r)$ unknown			
0.005	0.953	0.980	0.793
0.05	0.960	0.987	0.734
0.10	0.967	0.987	0.707
Setting 2: $\log(\sigma)$ unknown			
0.005	0.967	0.987	0.782
0.05	0.987	0.993	0.732
0.10	0.987	0.993	0.717
Setting 3: $\log(\phi)$ unknown			
0.005	0.953	0.967	0.828
0.05	0.947	0.993	0.762
0.10	0.960	0.987	0.734
Setting 4: $(\log(r), \log(\sigma))'$ unknown			
0.005	0.960	0.947	0.611
0.05	0.960	0.973	0.519
0.10	0.960	0.947	0.484
Setting 5: $(\log(r), \log(\phi))'$ unknown			
0.005	0.973	0.987	0.749
0.05	1.0	1.0	0.619
0.10	1.0	1.0	0.557

probabilities we expect that many of the established techniques used in ABC can be readily adapted to ACDC to further improve its computational performance without sacrificing the frequentist inferential guarantees.

An ACDC approach quantifies the uncertainty in estimation by drawing upon a direct connection to confidence distribution estimators. Different choices of summary statistic yield different approximate CDs, some producing tighter confidence sets than others. However, inference from ACDC is validated, regardless of the sufficiency of S_n , provided Condition 1 can be established. Within a Bayesian framework, there is no clear way to choose among different posterior approximations associated with different summary statistics. By pivoting to a frequentist perspective, different summary statistics produce different (CD) estimators but all of these estimators are well-behaved in the long run, yielding valid inferential statements about θ . Supported by the theoretical developments and examples in this paper, it appears as though ACDC provides a more parsimonious solution to validating likelihood-free inference than attempts to reconcile differences among posteriors and their various approximations.

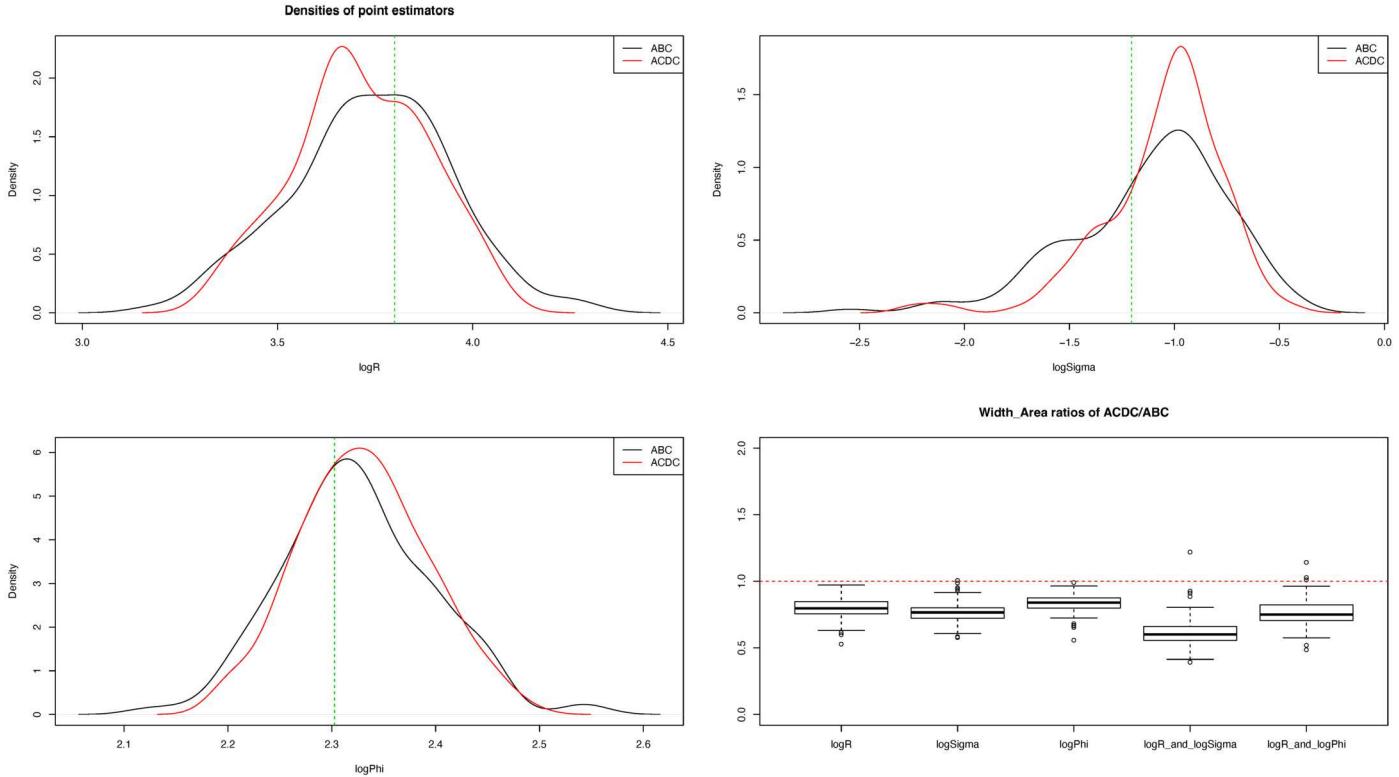


Figure 3: These are densities of point estimators from ACDC (red) and IS-ABC (black) for the 150 independent data sets produced by the Ricker model. Additionally, this figure shows a box plot of the ratio of the sizes of the 150 confidence sets, that is, the length (or volume) of regions produced by ACDC divided by those of IS-ABC.

SUPPLEMENTARY MATERIAL

The Appendices and Supplementary Material are available online. Besides detailed proofs, this material contains additional conditions and provides a few additional remarks as noted earlier in this paper.

FUNDING

The research is supported in part by research grants from the US National Science Foundation (DMS1812048, DMS2015373 and DMS2027855). This research stems from a chapter of the first author's PhD Thesis. The first author also acknowledges the generous graduate support from Rutgers University.

Accepted 24 April 2023

REFERENCES

- [1] BEAUMONT, M. A., ZHANG, W. and BALDING, D. J. (2002). Approximate Bayesian Computation in Population Genetics. *Genetics* **162**(4) 2025–2035.
- [2] BEAUMONT, M. A., CORNUET, J.-M., MARIN, J.-M. and ROBERT, C. P. (2009). Adaptive approximate Bayesian computation. *Biometrika* **96**(10) 1–9. <https://doi.org/10.1093/biomet/asp052>. MR2767283
- [3] BLUM, M. G. (2010). Approximate Bayesian computation: A non-parametric perspective. *Journal of the American Statistical Association* **105**(491) 1178–1187. <https://doi.org/10.1198/jasa.2010.tm09448>. MR2752613
- [4] CAMERON, E. and PETTITT, A. N. (2012). Approximate Bayesian Computation for astronomical model analysis: A case study in galaxy demographics and morphological transformation at high redshift. *Monthly Notices of the Royal Astronomical Society* **425**(1) 44–65.
- [5] CHENG, T. T. (1949). The normal approximation to the Poisson distribution and a proof of a conjecture of Ramanujan. *Bulletin of the American Mathematical Society* **55**(4) 396–401. <https://doi.org/10.1090/S0002-9904-1949-09223-6>. MR0029487
- [6] CSILLÉRY, K., BLUM, M. G. B., GAGGIOTTI, O. E. and FRANÇOIS, O. (2010). Approximate Bayesian Computation (ABC) in practice. *Trends in Ecology and Evolution* **25**(7) 410–418.
- [7] FASIOLO, M., WOOD, S. N., HARTIG, F. and BRAVINGTON, M. V. (2018). An extended empirical saddlepoint approximation for intractable likelihoods. *Electronic Journal of Statistics* **12**(1) 1544–1578. <https://doi.org/10.1214/18-ejs1433>. MR3806432
- [8] FEARNHEAD, P. and PRANGLE, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation (with discussion). *Journal of the Royal Statistical Society, Series B* **74** 419–474. <https://doi.org/10.1111/j.1467-9868.2011.01010.x>. MR2925370
- [9] GOURIEROUX, C., MONFORT, A. and RENAULT, E. (1993). Indirect Inference. *Journal of Applied Econometrics* **8**(S1) 85–118.
- [10] LI, W. and FEARNHEAD, P. (2018). Convergence of regression-adjusted approximate Bayesian computation. *Biometrika* **105**(2) 301–318. <https://doi.org/10.1093/biomet/asx081>. MR3804404

- [11] LI, W. and FEARNHEAD, P. (2018). On the Asymptotic Efficiency of approximate Bayesian computation estimators. *Biometrika* **105**(2) 286–299. <https://doi.org/10.1093/biomet/asx078>. MR3804403
- [12] LIU, R. Y., PARELIUS, J. M. and SINGH, K. (1999). Multivariate analysis by data depth: descriptive statistics, graphics and inference (with discussion). *Annals of Statistics* **27** 783–858. <https://doi.org/10.1214/aos/1018031260>. MR1724033
- [13] MARIN, J.-M., PUDLO, P., ROBERT, C. P. and RYDER, R. J. (2011). Approximate Bayesian Computational methods. *Statistics and Computing* **22**(6) 1167–1180. <https://doi.org/10.1007/s11222-011-9288-2>. MR2992292
- [14] MARJORAM, P., MOLITOR, J., PLAGNOL, V. and TAVARÉ, S. (2003). Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences* **100**(26) 15324–15328.
- [15] MARQUETTE, J.-B. (2018) *Statistics for Astrophysics: Bayesian Methodology*.
- [16] MCFADDEN, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society* 995–1026. <https://doi.org/10.2307/1913621>. MR1014539
- [17] PETERS, G. W., FAN, Y. and SISSON, S. A. (2012). On sequential Monte Carlo partial rejection control approximate Bayesian computation. *Statistical Computing* **22** 1209–1222. <https://doi.org/10.1007/s11222-012-9315-y>. MR2992295
- [18] SCHWEDER, T. and HJORT, N. L. (2016) *Confidence, Likelihood, Probability*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139046671>. MR3558738
- [19] SERFLING, R. (2002). Quantile functions for multivariate analysis: approaches and applications. *Statistica Neerlandica* **56**(2) 214–232. <https://doi.org/10.1111/1467-9574.00195>. MR1916321
- [20] SISSON, S., FAN, Y. and TANAKA, M. (2007). Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Science USA* **104** 1760–1765. <https://doi.org/10.1073/pnas.0607208104>. MR2301870
- [21] THORNTON, S. and XIE, M. (2022). Bridging Bayesian, frequentist and fiducial (BFF) inferences using confidence distribution. In *Handbook on Bayesian, Fiducial and Frequentist (BFF) Inferences* (J. O. Berger, X.-L. Meng, N. Reid and M. Xie, eds.) Chapman & Hall. To appear. <https://doi.org/10.1007/s11425-017-9325-9>. MR4119542
- [22] WOOD, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature* **466**(7310) 1102.
- [23] XIE, M. and SINGH, K. (2013). Confidence distribution, the frequentist distribution estimator of a parameter: A review. *International Statistical Review* **81**(1) 3–39. <https://doi.org/10.1111/insr.12000>. MR3047496

Suzanne Thornton. Swarthmore College, U.S.A. E-mail address: s thornt1@swarthmore.edu

Wentao Li. The University of Manchester, U.K. E-mail address: wentao.li@manchester.ac.uk

Minge Xie. Rutgers, The State University of New Jersey, U.S.A. E-mail address: mxie@stat.rutgers.edu