

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/129874>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Model Misspecification in ABC: Consequences and Diagnostics.*

David T. Frazier[†], Christian P. Robert[‡] and Judith Rousseau[§]

July 11, 2019

Abstract

We analyze the behavior of approximate Bayesian computation (ABC) when the model generating the simulated data differs from the actual data generating process; i.e., when the data simulator in ABC is misspecified. We demonstrate both theoretically and in simple, but practically relevant, examples that when the model is misspecified different versions of ABC can yield substantially different results. Our theoretical results demonstrate that even though the model is misspecified, under regularity conditions, the accept/reject ABC approach concentrates posterior mass on an appropriately defined pseudo-true parameter value. However, under model misspecification the ABC posterior does not yield credible sets with valid frequentist coverage and has non-standard asymptotic behavior. In addition, we examine the theoretical behavior of the popular local regression adjustment to ABC under model misspecification and demonstrate that this approach concentrates posterior mass on a completely different pseudo-true value than accept/reject ABC. Using our theoretical results, we suggest two approaches to diagnose model misspecification in ABC. All theoretical results and diagnostics are illustrated in a simple running example.

Keywords: Likelihood-free methods, model misspecification, approximate Bayesian Computation (ABC), Asymptotics, regression adjustment ABC

*We would like to thank the Editor, David Dunson, an Associate Editor and two anonymous referees for their constructive comments that greatly improved the paper.

[†]Monash University, Melbourne Australia. Email: david.frazier@monash.edu.

[‡]Université Paris Dauphine PSL, CEREMADE CNRS, Paris, France. Email: xian@ceremade.dauphine.fr

[§]University of Oxford, Université Paris Dauphine PSL, CEREMADE CNRS, Paris, France. Email: rousseau@ceremade.dauphine.fr

1 Introduction

It is now routine in the astronomic, ecological and genetic sciences, as well as in economics and finance, that the models used to describe observed data are so complex that the likelihoods associated with these models can be computationally intractable. In a Bayesian inference paradigm, these situations have led to the rise of approximate Bayesian computation (ABC) methods that eschew calculation of the likelihood in favor of simulation; for reviews on ABC methods see, e.g., Marin et al. (2012), Robert (2016) and Sisson et al. (2018).

ABC is predicated on the belief that the observed data $\mathbf{y} := (y_1, y_2, \dots, y_n)^\top$ is drawn from the class of models $\{\theta \in \Theta : P_\theta^n\}$, where $\theta \in \Theta \subset \mathbb{R}^{k_\theta}$ is an unknown vector of parameters and where $\pi(\theta)$ describes our prior beliefs about θ . The goal of ABC is to conduct inference on the unknown θ by simulating pseudo-data \mathbf{z} , $\mathbf{z} := (z_1, \dots, z_n)^\top$, from P_θ^n and then ‘‘comparing’’ \mathbf{y} and \mathbf{z} . In most cases, this comparison is carried out using a vector of summary statistics $\eta(\cdot)$, and a metric $d(\cdot, \cdot)$. Generally speaking, in ABC values of θ are accepted, and used to build an approximation to the exact posterior, if they satisfy an acceptance rule that depends on the tolerance parameter ϵ .

Algorithm 1 ABC Algorithm

- 1: Simulate θ^i , $i = 1, 2, \dots, N$, from $\pi(\theta)$,
 - 2: Simulate $\mathbf{z}^i = (z_1^i, z_2^i, \dots, z_n^i)^\top$, $i = 1, 2, \dots, N$, from $P_{\theta^i}^n$;
 - 3: For each $i = 1, \dots, N$, accept θ^i if $d(\eta(\mathbf{z}^i), \eta(\mathbf{y})) \leq \epsilon$, where ϵ denotes an user chosen tolerance parameter ϵ .
-

Algorithm 1 details the common accept/reject implementation of ABC, which can be augmented with additional steps to increase sampling efficiency; see, e.g., the MCMC-ABC approach of Marjoram et al. (2003), or the SMC-ABC approach of Sisson et al. (2007). Post-processing of the simulated pairs $\{\theta^i, \eta(\mathbf{z}^i)\}$ has also been proposed as a means of obtaining more accurate posterior approximations (see, e.g., the local linear regression adjustment approach of Beaumont et al., 2002, the marginal adjustment approach of Nott et al., 2014, or the recalibration approach of Rodrigues et al., 2018).

While several post processing strategies exist, the most common approach is the so-called local linear regression adjustment (Beaumont et al., 2002), which involves post-processing the output from Algorithm 1 using a linear regression model to improve the resulting posterior approximation; we refer the interested reader to Blum (2018) for an overview of regression adjustment methods in ABC. For $\{\theta^i, \eta(\mathbf{z}^i)\}_{i \geq 1}$ denoting a sample from the ABC posterior based on Algorithm 1, the local linear regression adjustment uses the sample $\{\theta^i, \eta(\mathbf{z}^i)\}_{i \geq 1}$ to produce the adjusted posterior sample $\{\theta^i - \hat{\beta}^\top[\eta(\mathbf{z}^i) - \eta(\mathbf{y})]\}_{i \geq 1}$, where $\hat{\beta}$ is obtained from a regression of θ^i on $\{\eta(\mathbf{z}^i) - \eta(\mathbf{y})\}$.

Regardless of the ABC algorithm chosen, the very nature of ABC is such that the researcher must believe there are values of θ in the prior support that can yield simulated summaries $\eta(\mathbf{z})$ that are ‘close to’ the observed summaries $\eta(\mathbf{y})$. Therefore, in order for ABC to yield meaningful inference about θ there must exist values of $\theta \in \Theta$ such that $\eta(\mathbf{z})$ and $\eta(\mathbf{y})$ are similar.

While complex models allow us to explain many features of the observed data, it is unlikely that any researcher will be able to construct a model P_θ^n that perfectly reproduces all features of \mathbf{y} . In other words, by the very nature of the complex models to which ABC is applied, the class of

models $\{\theta \in \Theta : P_\theta^n\}$ used to simulate pseudo-data \mathbf{z} is likely misspecified. Even when accounting for the use of summary statistics that are not sufficient, and which might be compatible with several models, the value these summaries take for the observed data may well be incompatible, i.e., highly unlikely, with the realised values of these statistics in the assumed model.

Given the likelihood of model misspecification in empirical applications, understanding the behavior of popular ABC approaches under model misspecification, and the consequences of this behavior, is of paramount importance for practitioners. As the following example illustrates, a particular consequence of model misspecification is that different ABC approaches can yield significantly different results.

Example 1: Consider an artificially simple example where the assumed data generating process (DGP) is z_1, \dots, z_n iid as $\mathcal{N}(\theta, 1)$ but the actual DGP is y_1, \dots, y_n iid as $\mathcal{N}(\theta, \sigma^2)$. That is, for $\sigma^2 \neq 1$, the assumed DGP maintains an incorrect assumption about the variance of the observed data. We consider as the basis of our ABC analysis the following summary statistics:

- the sample mean $\eta_1(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i$,
- the sample variance $\eta_2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n (y_i - \eta_1(\mathbf{y}))^2$.

Consider conducting inference on θ based on two versions of ABC: the accept/reject approach (hereafter, ABC-AR), where we take $d(x, y) = \|x - y\|$ to be the Euclidean norm, and a local linear regression adjustment approach to ABC (hereafter, ABC-Reg). ABC-Reg adjusts the accepted draws from ABC-AR using a weighted linear regression of θ^i on $\{\eta(\mathbf{z}^i) - \eta(\mathbf{y})\}$, with weights $K_\epsilon(\|\eta(\mathbf{z}^i) - \eta(\mathbf{y})\|)$, where $K_\epsilon(\cdot)$ is a kernel function and where the tolerance ϵ operates as a bandwidth for the kernel function. Following Beaumont et al. (2002), we take as the kernel function, $K_\epsilon(t)$, the Epanechnikov kernel: $K_\epsilon(t) = \mathbb{1}_{t \leq \epsilon} \cdot c\epsilon^{-1} (1 - (t/\epsilon)^2)$, where $\mathbb{1}_{t \leq \epsilon}$ denotes the indicator function on the event $t \leq \epsilon$, and where c is a normalizing constant.

To determine how these two ABC approaches behave under varying levels of model misspecification, we fix $\theta = 1$ and simulate “observed data” \mathbf{y} according to different values of σ^2 . The sample size across the experiments is taken to be $n = 100$. We consider a sequence of simulated data sets for \mathbf{y} such that each corresponds to a different value of σ^2 , with σ^2 taking values from $\sigma^2 = .5$ to $\sigma^2 = 5$ with evenly spaced increments of 0.05. Across all the datasets we fix the random numbers used to generate the simulated observed data and only change the value of σ^2 to isolate the impact of model misspecification; i.e., we generate one common set of random numbers $\nu_i \sim \mathcal{N}(0, 1)$, $i = 1, \dots, 100$, then for a value of σ^2 we generate observed data according to $y_i = 1 + \nu_i \cdot \sigma$.

Our prior beliefs are given by $\theta \sim \mathcal{N}(0, 25)$. We implement ABC-AR using $N = 25,000$ simulated pseudo-datasets generated iid according to $z_i^j \sim \mathcal{N}(\theta^j, 1)$. For both ABC-AR and ABC-Reg, we set ϵ to be the 1% quantile of the simulated distances $\|\eta(\mathbf{y}) - \eta(\mathbf{z}^j)\|$. To further isolate the impact of randomness on this procedure, we use the same simulated data across the different observed datasets; i.e., both ABC procedures will use the same simulated data across the different values of σ^2 . By fixing the simulated data across the experiments, and by fixing the random numbers in the observed data, differences in the ABC output across the experiments can be attributed to the changing value of σ^2 .

Figure 6 compares the posterior means of ABC-AR, and ABC-Reg across the different values for σ^2 .¹ The results demonstrate that model misspecification induces dramatic differences between the two ABC approaches, even at a relatively small sample size.² We draw two specific conclusions from Figure 6: one, the posterior mean of ABC-AR remains relatively stable across the different levels of misspecification, but does shift away from the true value ($\theta = 1$) as the level of model misspecification increases; two, the posterior mean of ABC-Reg becomes unstable even at relatively small levels of misspecification. The performance of the local linear regression adjustment is particularly interesting given that this method can have theoretical advantages over ABC-AR, i.e., Algorithm 1, when the model is correctly specified (Li and Fearnhead, 2018b). We formally explore these issues in Sections two and three but note here that when $\sigma^2 \approx 1$ (i.e., correct model specification) both ABC approaches give similar results. \square




Figure 1: Comparison of posterior means, denoted by $\hat{\theta}$, for ABC-AR, and ABC-Reg across varying levels of model misspecification.

In the remainder of the paper, we elaborate on the above issues and rigorously characterizes the asymptotic behavior of accept/reject ABC and local regression adjustment ABC when the model generating the pseudo-data is misspecified. In Section two, we discuss model misspecification in the ABC context and demonstrate that under model misspecification, for a certain choice of the tolerance, the posterior associated with Algorithm 1 asymptotically concentrates on an appropriately defined pseudo-true value. In addition, we demonstrate that the asymptotic shape of the ABC posterior is non-standard under model misspecification, and can yield credible sets with arbitrary levels of coverage. In Section three, we provide a rigorous justification for the behavior observed in Figure 6: we demonstrate that under model misspecification the posteriors

¹The posterior densities for ABC-AR and ABC-Reg obtained from these experiments also display a similar pattern of behavior and are presented in the appendix.

²Even though the DGP for \mathbf{z} is misspecified, because of the nature of the model misspecification and the limiting behavior of $\eta(\mathbf{y})$, if one were to only use the first summary statistic (the sample mean) model misspecification would have little impact in this example. However, in general both the nature of the model misspecification and the precise limiting form of $\eta(\mathbf{y})$ are unknown. Therefore, choosing a set of summaries that can mitigate the impact of model misspecification will be difficult, if not impossible, in practical applications of ABC.

for local linear (and nonlinear) regression adjustment ABC can asymptotically concentrate onto a completely different region of the parameter space than accept/reject ABC (Algorithm 1). We then use these theoretical results to devise an alternative regression adjustment approach that performs well regardless of model specification. Motivated by our asymptotic results, in Section four we develop two model misspecification detection procedures: a graphical detection approach based on comparing acceptance probabilities from Algorithm 1 and an approach based on comparing the output from Algorithm 1 and its linear regression adjustment counterpart. Section five concludes with a brief overview and discussion of our results. Proofs of all theoretical results are contained in the supplementary appendix.

2 Model Misspecification in ABC

Before rigorously characterising the behavior of ABC under model misspecification, we first set the notation used throughout the remainder of the paper. For \mathbf{y} denoting the observed data, we let P_0^n denote the true distribution generating \mathbf{y} . The class of implied distributions used in ABC to simulate pseudo-data is denoted by $\mathcal{P} := \{\theta \in \Theta \subseteq \mathbb{R}^{k_\theta} : P_\theta^n\}$, while \mathbf{z} denotes pseudo-data with support \mathcal{Z} generated according to P_θ^n . The simulated summary statistics $\eta(\mathbf{z}) = (\eta_1(\mathbf{z}), \dots, \eta_{k_\eta}(\mathbf{z}))^\top$ are a k_η -dimensional random vector with support $\mathcal{B} := \{\eta(\mathbf{z}) : \mathbf{z} \in \mathcal{Z}\} \subseteq \mathbb{R}^{k_\eta}$. We let $d_1(\cdot, \cdot)$ denote a metric on Θ , and $d_2(\cdot, \cdot)$ a metric on \mathcal{B} . However, when no confusion will result we simply denote a generic metric by $d(\cdot, \cdot)$. $\Pi(\theta)$ denotes the prior measure and $\pi(\theta)$ its corresponding density. For economy of notation, in what follows, we disregard the dependence of P_0^n and P_θ^n on n , and simply represent these quantities as P_0 and P_θ , respectively.

2.1 On the Notion of Model Misspecification in ABC

Recall that, in likelihood-based inference, model misspecification is taken to mean that $P_0 \notin \mathcal{P}$. The result of this model misspecification is that the Kullback-Leibler divergence,

$$\mathcal{D}(P_0||P_\theta) = - \int \log \left\{ \frac{dP_\theta(\mathbf{y})}{dP_0(\mathbf{y})} \right\} dP_0(\mathbf{y}),$$

satisfies

$$\inf_{\theta \in \Theta} \mathcal{D}(P_0||P_\theta) > 0.$$

In this case, the parameter value

$$\theta^* = \arg \inf_{\theta \in \Theta} \mathcal{D}(P_0||P_\theta)$$

is referred to as the pseudo-true value. Even though the model is misspecified, under reasonable regularity conditions, Bayesian procedures predicated on the likelihood of P_θ yield posteriors that concentrate on θ^* ; see, e.g., Kleijn and van der Vaart (2012) and Muller (2013).

In this paper, we assume the researcher conducts posterior inference on θ via ABC when the observed sample \mathbf{y} is generated according to P_0 , and in the case where $P_0 \notin \mathcal{P}$. However, in contrast to likelihood-based procedures, ABC is not based on the full data \mathbf{y} but on two separate

approximations, the summary statistics $\eta(\mathbf{y})$ and the threshold ϵ . Therefore, even if $P_0 \notin \mathcal{P}$ the model class \mathcal{P} may still be capable of generating a simulated summary $\eta(\mathbf{z})$ that is compatible with the observed summary $\eta(\mathbf{y})$, or is within an ϵ -neighbourhood of $\eta(\mathbf{y})$. The approximate nature of ABC means that the notion of model misspecification in a likelihood-based setting, namely $\inf_{\theta} \mathcal{D}(P_0 || P_{\theta}) > 0$, is not necessarily a meaningful notion of model misspecification associated with the output of an ABC algorithm, or ABC posterior distributions.

Recalling that the ABC posterior measure is given by, for $A \subset \Theta$,

$$\Pi_{\epsilon}[A|\eta(\mathbf{y})] = \int_A P_{\theta} [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon] d\Pi(\theta) \Big/ \int_{\Theta} P_{\theta} [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon] d\Pi(\theta),$$

we see that misspecification in ABC will be driven by the behavior of $\eta(\mathbf{y}), \eta(\mathbf{z})$ and the set $\{\theta \in \Theta : d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon\}$. To rigorously formulate the notion of model misspecification in ABC, we must therefore study the limiting behaviour of the ABC likelihood $P_{\theta} [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon]$ as the amount of information in the data accumulates.

To this end, we follow the framework of Marin et al. (2014), Frazier et al. (2018) and Li and Fearnhead (2018a,b), where it is assumed that the summary statistics concentrate around some fixed value, namely, $\eta(\mathbf{y})$ to b_0 (under P_0) and $\eta(\mathbf{z})$ to $b(\theta)$ (under P_{θ}). In Marin et al. (2014), the authors study the case where $\epsilon = 0$, while Frazier et al. (2018) and Li and Fearnhead (2018a,b) study $\epsilon > 0$ but allow ϵ to vary with n and set $\epsilon = \epsilon_n$. In the latter papers the authors demonstrate that the amount of information ABC obtains about a given θ depends on: (1) the rate at which the observed (resp. simulated) summaries converge to a well-defined limit counterpart b_0 (resp., $b(\theta)$); (2) the rate at which the tolerance ϵ_n goes to zero; (3) the link between b_0 and $b(\theta)$. When $P_0 \in \mathcal{P}$, there exists some θ_0 such that $b(\theta_0) = b_0$ and the results of Frazier et al. (2018) completely characterize the asymptotic behaviour of the ABC posterior distribution. Furthermore, this analysis remains correct even if $P_0 \notin \mathcal{P}$, so long as there exists some $\theta_0 \in \Theta$ such that $b_0 = b(\theta_0)$.

Therefore, the meaningful concept of model misspecification in ABC is that there does not exist any $\theta_0 \in \Theta$ satisfying $b_0 = b(\theta_0)$, which is precisely the notion of model incompatibility defined in Marin et al. (2014). Throughout the remainder, we say that the model is (ABC) misspecified if

$$\epsilon^* = \inf_{\theta \in \Theta} d(b_0, b(\theta)) > 0 \tag{1}$$

and note here that this condition is more likely to occur when $k_{\theta} < k_{\eta}$.

Heuristically, the implication of misspecification in ABC is that, under concentration of $\eta(\mathbf{z})$ to $b(\theta)$ and $\eta(\mathbf{y})$ to b_0 , by the triangle inequality and the definition of ϵ^*

$$d(\eta(\mathbf{y}), \eta(\mathbf{z})) \geq d(b_0, b(\theta)) - o_{P_{\theta}}(1) - o_{P_0}(1) \geq \epsilon^* - o_p(1), \text{ for all } \epsilon_n = o(1),$$

and the event $\{\theta \in \Theta : d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n\}$ becomes extremely rare, and corresponds to the event

$$\{\theta \in \Theta : d(\eta(\mathbf{z}), b(\theta)) > \epsilon^* - o(1)\}.$$

Consequently, for a sequence of tolerances $\epsilon_n = o(1)$, once $\epsilon_n < \epsilon^* + o(1)$ hardly any draws of θ will be selected regardless of how many simulated samples from $\pi(\theta)$ we generate, and the ABC

posterior $\Pi_\epsilon[A|\eta(\mathbf{y})]$ will become ill-behaved as n increases.

While tolerance sequences $\epsilon_n = o(1)$ will eventually cause the posterior $\Pi_\epsilon[A|\eta(\mathbf{y})]$ to be ill-behaved, it is possible that other choices for ϵ_n will produce a well-behaved posterior. In the following section we show that (certain) tolerance sequences satisfying $\epsilon_n \rightarrow \epsilon^*$, as $n \rightarrow +\infty$, yield well-behaved ABC posteriors that concentrate posterior mass on an appropriately defined pseudo-true value.

2.2 ABC Posterior Concentration Under Misspecification

Building on the intuition in the previous section, in this and the following section we rigorously characterize the asymptotic behaviour of

$$\Pi_\epsilon[A|\eta(\mathbf{y})] = \int_A P_\theta [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n] d\Pi(\theta) \Big/ \int_\Theta P_\theta [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n] d\Pi(\theta)$$

when $P_0 \notin \mathcal{P}$ and $\epsilon^* > 0$. To do so, we first define the following additional notations: for sequences $\{a_n\}$ and $\{b_n\}$, real valued, $a_n \lesssim b_n$ denotes $a_n \leq Cb_n$ for some $C > 0$, $a_n \asymp b_n$ denotes equivalent order of magnitude, $a_n \gg b_n$ indicates a larger order of magnitude and the symbols $o_P(a_n)$, $O_P(b_n)$ have their usual meaning. Unless otherwise noted, all limits are taken as $n \rightarrow +\infty$.

We maintain the following assumptions.

[A0] There exist a unique b_0 such that $d(\eta(\mathbf{y}), b_0) = o_{P_0}(1)$ and a positive sequence $v_{0,n} \rightarrow +\infty$ such that

$$\liminf_{n \rightarrow +\infty} P_0 [d(\eta(\mathbf{y}), b_0) \geq v_{0,n}^{-1}] = 1.$$

[A1] There exist a continuous, injective map $b : \Theta \rightarrow \mathcal{B} \subset \mathbb{R}^{k_n}$ and a function $\rho_n(\cdot)$ satisfying: $\rho_n(u) \rightarrow 0$ as $n \rightarrow +\infty$ for all $u > 0$, and $\rho_n(u)$ monotone non-increasing in u (for any given n), such that, for all $\theta \in \Theta$,

$$P_\theta [d(\eta(\mathbf{z}), b(\theta)) > u] \leq c(\theta)\rho_n(u), \quad \int_\Theta c(\theta)d\Pi(\theta) < +\infty,$$

where $\mathbf{z} \sim P_\theta$, and we assume either of the following:

- (i) *Polynomial deviations:* There exist a positive sequence $v_n \rightarrow +\infty$ and $u_0, \kappa > 0$ such that $\rho_n(u) = v_n^{-\kappa}u^{-\kappa}$, for $u \leq u_0$.
- (ii) *Exponential deviations:* There exists $h_\theta(\cdot) > 0$ such that $P_\theta[d(\eta(\mathbf{z}), b(\theta)) > u] \leq c(\theta)e^{-h_\theta(uv_n)}$ and there exists $m, C > 0$ such that

$$\int_\Theta c(\theta)e^{-h_\theta(uv_n)} d\Pi(\theta) \leq Ce^{-m \cdot (uv_n)^\tau}, \quad \text{for } u \leq u_0.$$

[A2] There exist some $D > 0$ and $M_0, \delta_0 > 0$ such that, for all $\delta_0 \geq \delta > 0$ and $M \geq M_0$, there exists $S_\delta \subset \{\theta \in \Theta : d(b(\theta), b_0) - \epsilon^* \leq \delta\}$ for which

- (i) In case (i) of **[A1]**, $D < \kappa$ and $\int_{S_\delta} \left(1 - \frac{c(\theta)}{M}\right) d\Pi(\theta) \gtrsim \delta^D$.

(ii) In case (ii) of [A1], $\int_{S_\delta} (1 - c(\theta)e^{-h_\theta(M)}) d\Pi(\theta) \gtrsim \delta^D$.

The above assumptions are similar to those given in Frazier et al. (2018), and we refer the interested reader to Remarks 1 and 2, and Example 1 in that paper for a detailed discussion of these assumptions. Under the above assumptions, we have the following result.

Theorem 1. *Assume that the data generating process for \mathbf{y} satisfies [A0] and assume that equation (1) holds. Assume also that conditions [A1] and [A2] are satisfied and $\epsilon_n \downarrow \epsilon^*$ with*

$$\epsilon_n \geq \epsilon^* + Mv_n^{-1} + v_{0,n}^{-1},$$

for M large enough. Let M_n be any positive sequence going to infinity and $\delta_n \geq M_n(\epsilon_n - \epsilon^*)$, then

$$\Pi_\epsilon [d(b(\theta), b_0) \geq \epsilon^* + \delta_n |\eta(\mathbf{y})|] = o_{P_0}(1), \quad (2)$$

as soon as

$$\begin{aligned} \delta_n &\geq M_n v_n^{-1} u_n^{-D/\kappa} = o(1) \quad \text{in case (i) of assumption [A1]} \\ \delta_n &\geq M_n v_n^{-1} |\log(u_n)|^{1/\tau} = o(1) \quad \text{in case (ii) of assumption [A1]}. \end{aligned}$$

with $u_n = \epsilon_n - (\epsilon^* + Mv_n^{-1} + v_{0,n}^{-1}) \geq 0$.

We remind the reader that the proofs of all theoretical results are contained in the appendix.

Remark 1. Theorem 1 states that even though the model is misspecified, the ABC posterior concentrates onto

$$\arg \inf_{\theta \in \Theta} d(b(\theta), b_0),$$

under the assumption that ϵ_n is slightly larger than ϵ^* . Under the more precise framework of Theorem 2, where the asymptotic shape of the posterior distribution is studied, this condition can be refined to allow ϵ_n to be slightly smaller than ϵ^* . However, we demonstrate that if $\epsilon^* - \epsilon_n$ is bounded below by a positive constant, then the posterior distribution does not necessarily concentrate.

Using the posterior concentration in Theorem 1, we have the following result.

Corollary 1. *Assume the hypotheses of Theorem 1 are satisfied and define $\theta^* \in \Theta$ as*

$$\theta^* = \arg \inf_{\theta \in \Theta} d(b_0, b(\theta)),$$

then, for any $\delta > 0$,

$$\Pi_\epsilon [d(\theta, \theta^*) > \delta |\eta(\mathbf{y})|] = o_{P_0}(1).$$

Remark 2. Theorem 1 and Corollary 1 demonstrate that, under an identification condition, the ABC posterior $\Pi_\epsilon[\cdot | \eta(\mathbf{y})]$ concentrates on θ^* if the model is misspecified. Therefore, Theorem 1 is an extension of Theorem 1 in Frazier et al. (2018) to the case of misspecified models. In addition, we note that Theorem 1 above is similar to Theorem 4.3 in Bernton et al. (2019) for ABC inference based on the Wasserstein distance. The validity of each of these results requires

that the map $\theta \mapsto b(\theta)$ be injective. If this condition is not satisfied, there can exist a continuum of values under which $d(b(\theta), b_0) = \epsilon^*$. In this case, the ABC posterior will no longer converge to a point mass, but will concentrate onto the set $\{\theta \in \Theta : \epsilon^* = d(b(\theta), b_0)\}$.

Remark 3. It is crucial to note that the pseudo-true value θ^* depends on the choice of $d(\cdot, \cdot)$. This implies that ABC based on two different metrics $d(\cdot, \cdot)$ and $\tilde{d}(\cdot, \cdot)$ will produce two different pseudo-true values, unless if by happenstance $\inf\{\theta \in \Theta : d(b(\theta), b_0)\}$ and $\inf\{\theta \in \Theta : \tilde{d}(b(\theta), b_0)\}$ coincide. This lies in stark contrast to the posterior concentration result in Frazier et al. (2018), which demonstrated that under correct model specification the posterior $\Pi_\epsilon[\cdot | \eta(\mathbf{y})]$ concentrates on the same true value regardless of the choice of $d(\cdot, \cdot)$.

2.3 Shape of the Asymptotic Posterior Distribution

In this section, we analyse the asymptotic shape of the ABC posterior under model misspecification. For simplicity, we take the rate at which the simulated and observed summaries converge to their limit counterparts to be the same, i.e., we take $v_{0,n} = v_n$,³ and we consider as the distance $d(\eta(\mathbf{z}), \eta(\mathbf{y})) = \|\eta(\mathbf{z}) - \eta(\mathbf{y})\|$ where $\|\cdot\|$ denotes is the norm associated to a given scalar product $\langle \cdot, \cdot \rangle$. Denote by I_{k_η} the $(k_\eta \times k_\eta)$ dimensional identity matrix and let

$$\Phi(B) = \Pr [\mathcal{N}(0, I_{k_\eta}) \in B],$$

for any measurable subset B of \mathbb{R}^{k_η} .

The following conditions are needed to establish the results of this section.

[A0'] Assumption [A0] is satisfied, and $\epsilon^* = d(b(\theta^*), b_0) > 0$, where $\theta^* = \arg \inf_{\theta \in \Theta} d(b(\theta), b_0)$.

[A1'] Assumption [A1] holds and for some positive-definite matrix $\Sigma_n(\theta^*)$, $c_0 > 0$, $\kappa > 1$ and $\delta > 0$, for all $\|\theta - \theta^*\| \leq \delta$, $P_\theta [\|\Sigma_n(\theta^*)\{\eta(\mathbf{z}) - b(\theta)\}\| > u] \leq c_0 u^{-\kappa}$ for all $0 < u \leq \delta v_n$.

[A3] The map $\theta \mapsto b(\theta)$ is twice continuously differentiable at θ^* and the Jacobian $\nabla_\theta b(\theta^*)$ has full column rank k_θ . The Hessian of $\|b(\theta) - b_0\|^2$ evaluated at θ^* , and denoted by H^* , is a positive-definite matrix.

[A4] There exists a sequence of $(k_\eta \times k_\eta)$ positive-definite matrices $\Sigma_n(\theta)$ such that for all $M > 0$ there exists $u_0 > 0$ for which

$$\sup_{|x| \leq M} \sup_{\|\theta - \theta^*\| \leq u_0} |P_\theta (\langle Z_n, e \rangle \leq x) - \Phi(x)| = o(1),$$

where $Z_n = \Sigma_n(\theta)(\eta(\mathbf{z}) - b(\theta))$ and $e = (b(\theta^*) - b_0) / \|b(\theta^*) - b_0\|$.

[A5] There exist v_n going to infinity and $u_0 > 0$ such that for all $\|\theta - \theta^*\| \leq u_0$, the sequence of functions $\theta \mapsto \Sigma_n(\theta)v_n^{-1}$ converges to some positive-definite matrix $A(\theta)$ and is equicontinuous at θ^* .

[A6] $\pi(\theta)$, the density of the prior measure $\Pi(\theta)$, is continuous and positive at θ^* .

³Allowing $v_{0,n}$ and v_n to differ will not greatly alter the following result. The result presented will still be valid, but only for the slower of the two rates.

[A7] For $Z_n^0 = \Sigma_n(\theta^*)\{\eta(\mathbf{y}) - b_0\}$ and all M_n going to infinity

$$P_0(\|Z_n^0\| > M_n) = o(1).$$

The above assumptions are similar to those used in Frazier et al. (2018) to deduce the limiting shape of the ABC posterior under correct model specification, and we refer the interested reader to Remarks 3 and 4 in that paper for a detailed discussion of these assumptions. Under the above assumptions, we have the following result.

Theorem 2. *Assume [A0'], [A1'] (with $\kappa \geq k_\theta$), [A2] and [A3]-[A7] are satisfied. We then have the following results.*

(i) *If $\lim_n v_n(\epsilon_n - \epsilon^*) = 2c$, with $c \in \mathbb{R}$, then for $\|\cdot\|_{TV}$ the total-variation norm*

$$\|\Pi_{v_n^{1/2}, \epsilon} - Q_c\|_{TV} = o_{P_0}(1)$$

where $\Pi_{z_n, \epsilon}$ is the ABC posterior distribution of $z_n(\theta - \theta^)$ for any sequence $z_n > 0$ and Q_c has density q_c with respect to Lebesgue measure on \mathbb{R}^{k_θ} proportional to*

$$q_c(x) \propto \Phi\left(\frac{c - \langle Z_n^0, A(\theta^*)e \rangle \epsilon^*}{\|A(\theta^*)e\| \epsilon^*} - \frac{x^\top H^* x}{4\|A(\theta^*)e\| \epsilon^*}\right)$$

(ii) *If $\lim_n v_n(\epsilon_n - \epsilon^*) = +\infty$ with $u_n = \epsilon_n - \epsilon^* = o(1)$, for $\mathcal{U}_{\{\|x\| \leq M\}}$ the uniform measure over the set $\{\|x\| \leq M\}$,*

$$\|\Pi_{u_n^{-1}, \epsilon} - \mathcal{U}_{\{x^\top H^* x \leq 2\}}\|_{TV} = o_{P_0}(1),$$

Remark 4. As is true in the case where the model is correctly specified, if ϵ_n is too large, which here means that $(\epsilon_n - \epsilon^*) \gg 1/v_n$, then the asymptotic distribution of the ABC posterior is uniform with a radius that is of the order $\epsilon_n - \epsilon^*$. In contrast to the case of correct model specification, if $\epsilon^* > 0$ and if $v_n\{\epsilon_n - \epsilon^*\} \rightarrow 2c \in \mathbb{R}$, then the limiting distribution is no longer Gaussian. Moreover, this result maintains even if $c = 0$.

Remark 5. In likelihood-based Bayesian inference, credible sets are not generally valid confidence sets if the model is misspecified, however, the resulting posterior is still asymptotically normal (see, e.g., Kleijn and van der Vaart, 2012 and Muller, 2013). In the case of ABC, not only will credible sets not be valid confidence sets, but the asymptotic shape of the ABC posterior is not even Gaussian.

Remark 6. In practice ϵ^* is unknown, and it is therefore not possible to choose ϵ_n directly. However, we note that the application of ABC is most often implemented by accepting draws of θ within some pre-specified (and asymptotically shrinking) quantile threshold; i.e., one accepts a simulated draw θ^i if $d(\eta(\mathbf{z}^i), \eta(\mathbf{y}))$ is smaller than the α -th empirical quantile of the simulated values $d(\eta(\mathbf{z}^j), \eta(\mathbf{y}))$, $j \leq N$. However, as discussed in Frazier et al., 2018, the two representations of the ABC approach are dual in the sense that choosing a value of α on the order of $\delta v_n^{-k_\theta}$, with δ small, corresponds to choosing $|\epsilon_n - \epsilon^*| \lesssim \delta^{1/k_\theta} v_n$ and choosing $\alpha_n \gtrsim M v_n^{-k_\theta}$ corresponds to choosing $\epsilon_n - \epsilon^* \gtrsim M v_n$. We further elaborate on the equivalence between both approaches in Section 4.1.

Interestingly, the proof of Theorem 2 (provided in the appendix) demonstrates that if $v_n(\epsilon_n - \epsilon^*) \rightarrow -\infty$, in particular when $\epsilon_n = o(1)$ and $\epsilon^* > 0$, posterior concentration of $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ need not occur. We present an illustration of this phenomena in the following simple example.

Example 2: Consider the case where $k_\theta = 1$ and $k_\eta = 2$. Let $\tilde{Z}_y = \sqrt{n}(\eta(\mathbf{y}) - b_0)$ and $\tilde{Z}_n = \sqrt{n}(\eta(\mathbf{z}) - b(\theta))$, where $\tilde{Z}_n \sim \mathcal{N}(0, v_\theta^2 I_2)$, for v_θ some known function of θ , and $b(\theta) = (\theta, \theta)^\top$. In addition, assume that $b_0 = (\bar{b}_0, -\bar{b}_0)$, with $\bar{b}_0 \neq 0$. Under this setting, and when $\|\cdot\|$ is the Euclidean norm, it follows that the unique pseudo-true value is $\theta^* = 0$. However, depending on v_θ , the approximate posterior need not concentrate on $\theta^* = 0$. This is summarized in the following Proposition.

Proposition 1. *In the setup described above, if $v_\theta/v_{\theta^*} = \sigma(\theta)$, for v_{θ^*} some known function, such that σ is continuous and $\sigma(\bar{b}_0/2)^2 \geq 3$, and if the prior has positive and continuous density on $[-\bar{b}_0, \bar{b}_0]$, then*

$$\Pi_\epsilon \{|\theta - \theta^*| \leq \delta|\eta(\mathbf{y})\} = o(\Pi_\epsilon \{|\theta - \bar{b}_0/2| \leq \delta|\eta(\mathbf{y})\}) = o(1).$$

3 Local Regression Adjustment under Misspecification

3.1 Posterior Concentration

Local regression adjustments to ABC have found broad applicability with practitioners. However, we caution against the blind application of local regression adjustment when one is willing to entertain the idea of model misspecification. As demonstrated by the introductory example, the use of this particular adjustment can lead to point estimators that behave very differently to those obtained from Algorithm 1, even in small samples.

In this section, we first rigorously characterize posterior concentration of local linear regression adjustment ABC (ABC-Reg) under model misspecification. Using this initial result, we then extend the conclusion to local nonlinear regression adjustment approaches. For simplicity, we only consider the case of scalar θ , however, we allow $\eta(\mathbf{y})$ to be multi-dimensional.⁴

ABC-Reg first runs Algorithm 1, with tolerance ϵ_n , to obtain a set of accepted draws and summaries $\{\theta^i, \eta(\mathbf{z}^i)\}$, and then uses a linear regression model to adjust the accepted values of θ . In this way, the original accepted value θ^i is artificially related to $\eta(\mathbf{y})$ and $\eta(\mathbf{z})$ through the linear regression model

$$\theta^i = \mu + \beta^\top \{\eta(\mathbf{y}) - \eta(\mathbf{z}^i)\} + \nu_i,$$

where ν_i denotes the model residual. Define $\bar{\theta} = \sum_{i=1}^N \theta^i/N$ and $\bar{\eta} = \sum_{i=1}^N \eta(\mathbf{z}^i)/N$. Given θ^i , ABC-Reg then produces an adjusted parameter draw according to

$$\tilde{\theta}^i = \theta^i - \hat{\beta}^\top \{\eta(\mathbf{z}^i) - \eta(\mathbf{y})\},$$

$$\hat{\beta} = \left[\frac{1}{N} \sum_{i=1}^N (\eta(\mathbf{z}^i) - \bar{\eta}) (\eta(\mathbf{z}^i) - \bar{\eta})^\top \right]^{-1} \left[\frac{1}{N} \sum_{i=1}^N (\eta(\mathbf{z}^i) - \bar{\eta}) (\theta^i - \bar{\theta}) \right] = \widehat{\text{Var}}^{-1}(\eta(\mathbf{z}^i)) \widehat{\text{Cov}}(\eta(\mathbf{z}^i), \theta^i)$$

⁴This result can be extended at the cost of more complicated arguments but we refrain from this setting to simplify the interpretation of our results.

Therefore, for $\theta^i \sim \Pi_\epsilon[\cdot|\eta(\mathbf{y})]$, the posterior measure for $\tilde{\theta}^i$ is nothing but a scaled and shifted version of $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$. Consequently, the asymptotic behavior of the ABC-Reg posterior, denoted by $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$, is determined by the behavior of $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$, $\hat{\beta}$, and $\{\eta(\mathbf{y}) - \eta(\mathbf{z}^i)\}$.

The following result describes the asymptotic behavior of the ABC-Reg posterior $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$.

Corollary 2. *Assume that [A0'], [A1] and [A2] are satisfied and $\epsilon_n \downarrow \epsilon^*$ with*

$$\epsilon_n \geq \epsilon^* + Mv_n^{-1} + v_{0,n}^{-1},$$

for M large enough. Furthermore, assume that for some β_0 with $\|\beta_0\| > 0$, $\|\hat{\beta} - \beta_0\| = o_{P_\theta}(1)$. Define $\tilde{\theta}^* = \theta^* - \beta_0^\top(b(\theta^*) - b_0)$. Let M_n be any positive sequence going to infinity and $\delta_n \geq M_n(\epsilon_n - \epsilon^*)$, then

$$\tilde{\Pi}_\epsilon[|\theta - \tilde{\theta}^*| > \delta|\eta(\mathbf{y})] = o_{P_0}(1),$$

as soon as

$$\begin{aligned} \delta_n &\geq Mv_n^{-1}u_n^{-D/\kappa} = o(1) \quad \text{in case (i) of assumption [A1]} \\ \delta_n &\geq Mv_n^{-1}|\log(u_n)|^{1/\tau} = o(1) \quad \text{in case (ii) of assumption [A1]}. \end{aligned}$$

with $u_n = \epsilon_n - (\epsilon^* + Mv_n^{-1} + v_{0,n}^{-1}) \geq 0$.

Remark 7. An immediate consequence of Corollary 1 and 2 is that the ABC posterior $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ concentrates mass on

$$\theta^* = \arg \inf_{\theta \in \Theta} d(b(\theta), b_0),$$

while the ABC-Reg posterior $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$ concentrates mass on

$$\tilde{\theta}^* = \theta^* - \beta_0^\top(b(\theta^*) - b_0).$$

As a consequence, ABC-Reg takes draws of θ that are (asymptotically) optimal in terms of minimizing the chosen distance, $d(\cdot, \cdot)$, between observed and simulated summaries, and perturbs them in a (linear) manner that need not preserve the optimality of the original draws. Furthermore, for $\|\beta_0\|$ large, the pseudo-true value $\tilde{\theta}^*$, onto which ABC-Reg concentrates, can easily lie outside Θ . Therefore, if the model is misspecified, there is no guarantee that ABC-Reg returns draws that are optimal in terms of minimizing $d(\cdot, \cdot)$ and there is even no guarantee that ABC-Reg returns values in Θ .⁵

Remark 8. Crucially, the result of Corollary 2 and the phenomena discussed in Remark 7, are not restricted to local linear regression adjustment, and extend to the nonlinear variety of regression adjustment. For brevity, we only sketch the general idea here, and note that a rigorous proof follows along the same lines as Corollary 2, and is therefore omitted. Consider the nonlinear regression model

$$\theta = m(\eta(\mathbf{z})) + v,$$

for some unknown function $m(\cdot)$. Denote by $\hat{m}(\cdot)$ a nonparametric estimator of the unknown regression function that is constructed using the accepted draws $\{\theta^\ell, \eta(\mathbf{z}^\ell)\}_{\ell \leq L}$ obtained from

⁵In the appendix, we give a concrete example of this later behavior in the confines of the g -and- k distribution.

the ABC posterior $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$. A nonlinear regression post-processing approach transforms the accepted θ^ℓ into

$$\tilde{\theta}^\ell = \theta^\ell + \{\hat{m}(\eta(\mathbf{y})) - \hat{m}(\eta(\mathbf{z}^\ell))\}, \text{ for } \ell = 1, \dots, L. \quad (3)$$

Under the regularity conditions [A0'], [A1] and [A2], if the (nonparametric) estimator $\hat{m}(\cdot)$ converges to a function $m^*(\cdot)$, $\hat{m}(\eta) = m^*(\eta) + o_p(1)$, uniformly in a neighbourhood of b_0 and $b(\theta^*)$, then by the concentration of $\eta(\mathbf{y})$ and $\eta(\mathbf{z}^\ell)$, equation (3) becomes

$$\tilde{\theta}^\ell = \theta^\ell + \{m^*(b_0) - m^*(b(\theta^*))\} + o_p(1) = \theta^* + \{m^*(b_0) - m^*(b(\theta^*))\} + o_p(1).$$

The last equality follows from the posterior concentration of $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ toward θ^* , i.e., $\theta^\ell = \theta^* + o_p(1)$. Hence, as soon as $m^*(b_0) - m^*(b(\theta^*)) \neq 0$ the nonlinear regression post-processed ABC posterior concentrates onto a different value than $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$, which offers no particular justification. Moreover, given that $\hat{m}(\cdot)$ is constructed using the simulated data, there is no reason to suspect that $m^*(b_0) = m^*(b(\theta^*))$. Lastly, we note that the above computations are not significantly altered if instead we had considered a nonlinear conditional heteroskedastic regression model, as in Blum and François (2010), from the outset. As a consequence, the nonlinear conditional heteroskedastic regression adjustment will have similar asymptotic behavior to the nonlinear regression adjustment.

Remark 9. An additional consequence of Corollary 1 and 2 is that the ABC posterior $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ and the ABC-Reg posterior $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$ will yield different posterior expectations. This difference between expectations calculated under $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ and $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$ explains the divergence between the ABC-AR and ABC-Reg posterior means observed in Figure 6. In the following section, we use this behavior to derive a procedure for detecting model misspecification.

3.2 Adjusting Local Regression Adjustment

The difference between accept/reject ABC and ABC-Reg under model misspecification is related to the regression adjustments re-centering of the accepted draws θ^i by $\hat{\beta}^\top\{\eta(\mathbf{y}) - \eta(\mathbf{z})\}$. Whilst useful under correct model specification, when the model is misspecified the adjustment can force θ^i away from θ^* and towards $\tilde{\theta}^*$, which need not lie in Θ or be optimal in terms of minimizing $d(\cdot, \cdot)$.

The cause of this behavior is the inability of $\eta(\mathbf{z})$ to replicate the asymptotic behavior of $\eta(\mathbf{y})$, which in the terminology of Marin et al. (2014) means that the model is incompatible with the observed summaries. This incompatibility of the summary statistics ensures that the influence of the centering term $\hat{\beta}^\top\{\eta(\mathbf{y}) - \eta(\mathbf{z})\}$ can easily dominate that of the accepted draws θ^i , with the introductory example being just one example of this behavior.

In an attempt to maintain the broad applicability of local linear regression adjustment in ABC, and still ensure it gives sensible results under model misspecification, we propose a useful modification of the regression adjustment approach. To motivate this modification recall that, under correct model specification and regularity conditions, at first-order the linear regression

adjustment approach ensures (see Theorem 4 in Frazier et al., 2018):

$$\begin{aligned}
\tilde{\theta}^i &= \theta^i + \hat{\beta}^\top \{\eta(\mathbf{y}) - \eta(\mathbf{z})\} \\
&= \theta^i + \hat{\beta}^\top \{b_0 - b(\theta^i)\} + O_p(1/v_n) \\
&= \theta^i - [\nabla_\theta b(\theta^*)^\top V_0^{-1} \nabla_\theta b(\theta^*)]^{-1} \nabla_\theta b(\theta^*)^\top V_0^{-1} \nabla_\theta b(\bar{\theta})(\theta^i - \theta^*) + O_p(1/v_n), \tag{4}
\end{aligned}$$

where $b_0 = b(\theta^*)$ by correct model specification, $\bar{\theta}$ is an intermediate value satisfying $|\bar{\theta} - \theta^*| \leq |\theta^i - \theta^*|$, $V_0 = \lim_n \text{Var}[\sqrt{n}\{\eta(\mathbf{y}) - b_0\}]$, and the third line follows from a mean-value expansion and the definition of the local linear regression adjustment. Therefore, it follows from (4) that, even if $k_\eta > k_\theta$, the dimension of $\eta(\mathbf{y})$ will not impact the asymptotic variance of the ABC-Reg posterior mean. This result, at least in part, helps explain (from a technical standpoint) the popularity of the ABC-Reg approach as a dimension reduction method.

However, under model misspecification, $b_0 \neq b(\theta)$ for any $\theta \in \Theta$, and hence there does not exist an intermediate value $\bar{\theta}$ such that

$$b_0 - b(\theta^i) \neq \nabla_\theta b(\bar{\theta})(\theta^* - \theta^i).$$

As a consequence, equation (4) cannot be valid (in general) if the model is misspecified.

The behavior of ABC-Reg under correct and incorrect model specifications suggests that the methods poor behavior under the latter can be mitigated by replacing $\eta(\mathbf{y})$ with an alternative term. To this end, define $\hat{\theta} = \int \theta d\Pi_\epsilon[\theta|\eta(\mathbf{y})]$ to be the posterior mean of accept/reject ABC. Let $\hat{\mathbf{z}}^m$, $m = 1, \dots, M$, be a pseudo-data set of length n simulated under the assumed DGP and at the value $\hat{\theta}$, and define

$$\hat{\eta} = \sum_{m=1}^M \eta(\hat{\mathbf{z}}^m)/M.$$

Using $\hat{\eta}$, we can then implement the modified local linear regression adjustment

$$\check{\theta}^i = \theta^i + \hat{\beta}^\top \{\hat{\eta} - \eta(\mathbf{z}^i)\}.$$

The key to this modified approach is that under correct specification $\hat{\eta}$ behaves like $\eta(\mathbf{y})$, while under incorrect specification $\hat{\eta}$ behaves like $\eta(\mathbf{z})$.⁶ A direct consequence of this construction is that this approach avoids the incompatibility issue that arises under model misspecification. In addition, since this new regression adjustment approach uses a centering sequence calculated from the accept/reject ABC posterior mean, $\hat{\theta}$, the asymptotic behavior of this new approach is similar to accept/reject ABC.

Example 1 (Continued): Recall that the assumed DGP is z_1, \dots, z_n iid as $\mathcal{N}(\theta, 1)$ but the actual DGP is y_1, \dots, y_n iid as $\mathcal{N}(\theta, \sigma^2)$. ABC is conducted using the following summary statistics:

- the sample mean $\eta_1(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i$,
- the sample variance $\eta_2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n (y_i - \eta_1(\mathbf{y}))^2$.

⁶While the choice of M will not matter asymptotically, we argue that M should be chosen so that the variability of $\hat{\eta}$ is small relative to $\eta(\mathbf{z})$.

We consider three different DGPs corresponding to $\sigma^2 \in \{1, 2, 3\}$. For each of these cases we generate 1000 artificial samples for \mathbf{y} of length $n = 100$ and apply four different ABC approaches: the accept/reject ABC approach (ABC-AR), the local linear regression adjustment of Beaumont et al. (2002) (ABC-Reg), our new local linear regression adjustment (ABC-RegN), and the nonlinear regression adjustment of Blum and François (2010) that is fit using neural nets (ABC-NN). Each procedure relies on $N = 25,000$ pseudo-datasets generated according to $z_i^j \sim \mathcal{N}(\theta^j, 1)$, where the prior is again given by $\theta \sim \mathcal{N}(0, 25)$. For each of the procedures we set the tolerance ϵ to be the 1% quantile of the simulated distances $\|\eta(\mathbf{y}) - \eta(\mathbf{z}^i)\|$.

Figure 2 plots the posterior mean of each approach across the Monte Carlo replications and across all designs. The results demonstrate that the new regression adjustment maintains stable performance across both correct and incorrect model specification, while the point estimators obtained from ABC-Reg and ABC-NN are more varied. More specifically, for $\sigma^2 \in \{1, 2\}$ we see that all regression adjustments tend to give similar results. However, for $\sigma^2 = 3$, it is clear that the traditional linear and nonlinear adjustment approaches produce point estimators with more variability, across the repeated samples, than the other ABC approaches. This additional variability is a direct consequence of the fact that ABC-Reg and ABC-NN enforce a regression relationship between the accepted draws θ^i and $\eta(\mathbf{y}) - \eta(\mathbf{z})$ when one does not necessarily exist. As we have seen in Corollary 2, enforcing this additional (wrong) model to produce new values of θ^i , which need not respect the actual relationship between θ^i and $\eta(\mathbf{y}) - \eta(\mathbf{z})$, will (randomly) shift the adjusted draws away from their initial center of posterior mass, θ^* , and can yield point estimators with more variability in a repeated sampling context.




Figure 2: Posterior mean comparison of ABC-AR (AR), standard local linear regression adjustment (Reg), the new regression adjustment approach (RegN), and the local nonlinear regression adjustment (NN) across $\sigma^2 \in \{1, 2, 3\}$. Recall that $\sigma^2 = 1$ corresponds to correct model specification.

In addition to the results for the posterior means, in Table 1 we record for each method the posterior standard deviation, the length of the corresponding 95% credible set and the Monte Carlo coverage across the different designs. The values given in Table 1 represent the average values of these quantities across the Monte Carlo replications.

The results demonstrate that all of the local regression adjustments, ABC-Reg, ABC-RegN,

and ABC-NN, have much smaller posterior variability and much shorter credible sets than ABC-AR (on average). As a consequence, when the model is misspecified, this behavior gives researchers a false sense of precision, and leads to poor coverage rates (for the pseudo-true value) across all the local regression adjustment procedures (linear and nonlinear). Therefore, even though our new regression adjustment procedure gives stable performance under correct and incorrect model specification, it still suffers from the coverage issues alluded to in the remarks given after Theorem 2.

Furthermore, additional numerical experiments conducted in the supplemental appendix demonstrate that the use of the so-called heteroskedasticity correction for the local regression adjustment, as in Blum and François (2010), does not significantly alter these results. More specifically, the resulting heteroskedasticity corrected local regression adjustment approaches yield results that are very similar to those obtained without the heteroskedasticity correction. We refer the interested reader to the appendix for these details.

From this simple example we can conclude that under model misspecification local regression adjustment approaches, both linear and nonlinear, with or without a heteroskedasticity correction, can lead to significant overconfidence in the resulting point estimates obtained from these methods, and can result in poor coverage rates for the pseudo-true value.

Table 1: Monte Carlo coverage (Cov), credible set length (Len), and posterior standard deviation (Std) for the simple normal example under various levels of model misspecification. Cov is the percentage of times that the 95% credible set contained $\theta = 1$. Len is the average length of the credible set across the Monte Carlo trials. Std is the average posterior standard deviation across the Monte Carlo trials.

	ABC-AR			ABC-RegN		
	Cov	Len	Std	Cov	Len	Std
$\sigma^2 = 1$	0.9820	0.4666	0.1221	0.9380	0.3851	0.1001
$\sigma^2 = 2$	0.9610	0.6147	0.1576	0.8020	0.3837	0.0998
$\sigma^2 = 3$	0.9130	0.6164	0.1581	0.7070	0.3839	0.0997
	ABC-Reg			ABC-NN		
	Cov	Len	Std	Cov	Len	Std
$\sigma^2 = 1$	0.9410	0.3820	0.0997	0.9500	0.3853	0.1006
$\sigma^2 = 2$	0.7170	0.3826	0.0998	0.7290	0.4440	0.1228
$\sigma^2 = 3$	0.4600	0.3821	0.0997	0.4190	0.5043	0.1490

□

4 Detecting Misspecification

In this section we propose two methods to detect model misspecification in ABC. The first approach is based on the behavior of the acceptance probability under correct and incorrect model specification. The second approach is based on comparing posterior expectations calculated under $\Pi_\epsilon[\cdot|\eta(\mathbf{y})]$ (obtained from Algorithm 1) and $\tilde{\Pi}_\epsilon[\cdot|\eta(\mathbf{y})]$ (obtained using local linear regression adjustment).

4.1 A Simple Graphical Approach to Detecting Misspecification

From the results of Frazier et al. (2018), under regularity and correct model specification, the acceptance probability $\alpha_n = \Pr [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n]$ satisfies, for n large and $\epsilon_n \gg v_n^{-1}$,

$$\alpha_n = \Pr [d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n] \asymp \epsilon_n^{k_\theta}.$$

In this way, as $\epsilon_n \rightarrow 0$ the acceptance probability $\alpha_n \rightarrow 0$ in a manner that is approximately linear in $\epsilon_n^{k_\theta}$.

However, this relationship between α_n and ϵ_n does not extend to the case where the model is misspecified. In particular, if $\epsilon^* > 0$, once $\epsilon_n < \epsilon^*$ the acceptance probability α_n will be small or zero, even for a large number of simulations N .

The behavior of α_n under correct and incorrect model specifications means that one can potentially diagnose misspecification by comparing the behavior of α_n over a decreasing sequence of tolerance values. In particular, if we take a decreasing sequence of equally spaced tolerances $\epsilon_{1,n} < \epsilon_{2,n} < \dots < \epsilon_{J,n}$ we can construct and plot the resulting sequence $\{\alpha_{j,n}\}_j$ to determine if $\{\alpha_{j,n}\}_j$ decays in an (approximately) linear fashion as $\epsilon_{j,n}$ decreases.

While α_n is infeasible to obtain in practice, the same procedure can be applied with α_n replaced by the estimator $\hat{\alpha}_n = \sum_{i=1}^N \mathbb{1}_{[d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_n]} / N$. In this way, such a graphical check can easily be performed using the ABC reference table. The only difference is that, instead of considering a single tolerance ϵ , one would consider a sequence of tolerances $\{\epsilon_{j,n}\}_j$ and record, for each j ,

$$\hat{\alpha}_{j,n} = \sum_{i=1}^N \mathbb{1}_{[d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq \epsilon_{j,n}]} / N.$$

Once $\hat{\alpha}_{j,n}$ has been obtained, it can be plotted against $\epsilon_{j,n}^{k_\theta}$ (in some fashion) and the relationship can be analyzed to determine if deviations from linearity are in evidence.

To understand exactly how such a procedure can be implemented, we return to the simple normal example.

Example 1 (Continued): The assumed DGP is z_1, \dots, z_n iid as $\mathcal{N}(\theta, 1)$ but the actual DGP is y_1, \dots, y_n iid as $\mathcal{N}(\theta, \sigma^2)$. We again consider ABC using the following summary statistics:

- the sample mean $\eta_1(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i$,
- the sample variance $\eta_2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n (y_i - \eta_1(\mathbf{y}))^2$.

Taking $\sigma^2 \in \{1, 1 + 1/9, \dots, 1 + 8/9\}$, we generate observed samples of size $n = 100$ according to $y_i \sim \mathcal{N}(1, \sigma^2)$, iid, where, for each of the nine different simulated datasets, we keep the random numbers fixed and only change σ^2 . We consider $N = 25,000$ simulated datasets generated iid according to $z_i^j \sim \mathcal{N}(\theta^j, 1)$, with $\theta^j \sim \mathcal{N}(0, 25)$, and for $d(\cdot, \cdot)$ we take the Euclidean norm. For the sequence of $\epsilon_{j,n}$ values, we consider $J = 100$ evenly spaced increments, with $\epsilon_{J,n}$ chosen to correspond to the 10% quantile of the simulated distances, and where $\epsilon_{1,n}$ is taken to be the 0.1% quantile of the simulated distances.

In Figure 3, we plot the results over the nine different levels of misspecification. Each figure contains two different curves: the dashed curve represents the observed relationship between $\hat{\alpha}_{j,n}$

and $\epsilon_{j,n}$, while the solid curve plots a linear relationship between $\hat{\alpha}_{j,n}$ and $\epsilon_{j,n}$ that can be used to help visually diagnose departures from linearity.⁷




Figure 3: Graphical comparison of empirical acceptance probabilities $\hat{\alpha}_{j,n}$ (dashed line) and theoretical acceptance probabilities (solid line) for a decreasing sequence of tolerance values $\epsilon_{j,n}$.

We recall that, in this example correct specification warrants a linear relationship between $\hat{\alpha}_n$ and ϵ_n , since we are only conducting inference on a single parameter. More generally, under correct model specification, we would expect a linear relationship between $\hat{\alpha}_n$ and $\epsilon_n^{k_\theta}$.

Analyzing Figure 3, we see that the relationship is fairly linear for $\sigma^2 \leq 1 + 1/3$. However, for $\sigma^2 > 1 + 1/3$ the relationship between the acceptance probabilities $\hat{\alpha}_{j,n}$ and $\epsilon_{j,n}$ exhibits significant nonlinear behavior. Therefore, in this example the diagnostic would suggest that the model is misspecified once $\sigma^2 > 1 + 1/3$, which is evidenced by the nonlinear relationship between $\hat{\alpha}_{n,j}$ and $\epsilon_{j,n}$.

Clearly, obtaining broad conclusions about model misspecification from this graphical approach depends on many features of the underlying model, the dimension of θ , and the exact nature of misspecification. However, it is always possible to benchmark the results for the observed data against those obtained under the ABC reference table. That is, using the ABC reference table, we can easily apply the above diagnostic to one, or even several, “observed” se-

⁷Across the different values of σ^2 , the solid line is constructed using the endpoint pairs $(\hat{\alpha}_{J,n}, \epsilon_{J,n}) = (0.10, \epsilon_{J,n})$ and $(\hat{\alpha}_{1,n}, \epsilon_{1,n}) = (0.001, \epsilon_{1,n})$.

ries from the reference table. Then, if the relationship between $(\hat{\alpha}_{j,n}, \epsilon_{j,n})$ in the observed data deviates from that observed in the reference table, this is evidence that the model is misspecified. \square

4.2 Detecting Model Misspecification Using Regression Adjustment

Corollary 1 and 2 demonstrate that accept/reject ABC (ABC-AR) and local linear regression adjustment ABC (ABC-Reg), place posterior mass in different regions of the parameter space. Therefore, for $\theta \mapsto h(\theta)$ a known, smooth function, under model misspecification the posterior expectation of $h(\cdot)$ under ABC-AR and ABC-Reg,

$$\hat{h} = \int h(\theta) d\Pi_\epsilon[\theta|\eta(\mathbf{y})], \quad \tilde{h} = \int h(\theta) d\tilde{\Pi}_\epsilon[\theta|\eta(\mathbf{y})],$$

will converge in probability, as $n \rightarrow +\infty$ and $\epsilon_n \downarrow \epsilon^*$, to distinct values. However, if the model is correctly specified, it can be shown that \hat{h} and \tilde{h} will not differ, up to first-order, so long as $\epsilon_n = o(1/\sqrt{n})$ (this result follows from Theorem 4 in Frazier et al., 2018, or Theorem 3.1 in Li and Fearnhead, 2018a). Therefore, a useful approach for detecting model misspecification is to compare various posterior expectations, such as moments or quantiles, calculated under the two posteriors.

More specifically, if the model is correctly specified and if we use Algorithm 1 based on quantile thresholding with $\alpha_n = \delta n^{-k_\theta/2}$, for $\delta > 0$, then

$$\sqrt{n}\|\hat{h} - \tilde{h}\| = o_{P_0}(1).$$

However, if $\epsilon^* = \inf_{\theta \in \Theta} d(b_0, b(\theta)) > 0$, under regularity conditions, we can deduce that

$$\|\hat{h} - \tilde{h}\| = O_{P_0}(1).$$

Therefore, if $\sqrt{n}\|\hat{h} - \tilde{h}\|$ is large, this is meaningful evidence that the model may be misspecified.

Detecting misspecification by analyzing the magnitude of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ requires the specification of a cutoff value, denoted by t_n , such that if $\sqrt{n}\|\hat{h} - \tilde{h}\|$ is larger than t_n , we conclude that the model is likely misspecified. While there are several ways to choose the cutoff value t_n , we propose a simulation-based approach that uses the ABC reference table. Namely, we use the fact that for any fixed value of θ , which for instance could be drawn from the prior or from the ABC-AR posterior, we can always simulate “observed data” from the assumed model to determine what the magnitude of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ should be under correct specification. More specifically, for a given value of θ , we can always simulate an “observed data series”, and then run both ABC-AR and ABC-Reg on this simulated observed data to calculate $\sqrt{n}\|\hat{h} - \tilde{h}\|$ under correct model specification.

To operationalize this approach, we generate $b = 1, \dots, B$ such simulated “observed” datasets, all at the same value of θ , and for each of these datasets we calculate $\sqrt{n}\|\hat{h}_b - \tilde{h}_b\|$. The cutoff value t_n can then be defined as an empirical quantile of the simulated distances $\{\sqrt{n}\|\hat{h}_b - \tilde{h}_b\|\}_{b=1}^B$, which have all been calculated under correct specification. For instance, t_n could be defined as the 95% quantile of $\{\sqrt{n}\|\hat{h}_b - \tilde{h}_b\|\}_{b=1}^B$. The value of t_n obtained from this procedure can then be compared with the corresponding value of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ obtained from the actual observed data \mathbf{y} ,

and if $\sqrt{n}\|\hat{h} - \tilde{h}\| > t_n$ we conclude that the model is likely misspecified.

We demonstrate this approach to diagnosing model misspecification in our simple running example.

Example 1 (Continued): The assumed DGP is z_1, \dots, z_n iid as $\mathcal{N}(\theta, 1)$ but the actual DGP is y_1, \dots, y_n iid as $\mathcal{N}(\theta, \sigma^2)$. We again consider the following summary statistics:

- the sample mean $\eta_1(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i$,
- the sample variance $\eta_2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n (y_i - \eta_1(\mathbf{y}))^2$.

For the observed data series, \mathbf{y} , we simulate $n = 100$ observed data points from a normal random variable with mean $\theta = 1$, variance σ^2 , and take $\sigma^2 = \{2, 3\}$. For both ABC-AR and ABC-Reg, we again take $N = 25,000$ simulated datasets generated according to $z_i^j \sim \mathcal{N}(\theta^j, 1)$, with $\theta^j \sim \mathcal{N}(0, 25)$. For $d(\cdot, \cdot)$ we take the Euclidean norm and we again take the tolerance to be the 1% quantile of the simulated distances.

For the choice of the function $h(\theta)$, we consider $h(\theta) = (\theta^2, \theta^3)^\top$, so that \hat{h} and \tilde{h} represent the second and third posterior moments calculated under ABC-AR and ABC-Reg:

$$\hat{h} = \left(\int \theta^2 d\Pi_\epsilon[\theta|\eta(\mathbf{y})], \int \theta^3 d\Pi_\epsilon[\theta|\eta(\mathbf{y})] \right)^\top, \quad \tilde{h} = \left(\int \theta^2 d\tilde{\Pi}_\epsilon[\theta|\eta(\mathbf{y})], \int \theta^3 d\tilde{\Pi}_\epsilon[\theta|\eta(\mathbf{y})] \right)^\top.$$

Recall that under correct specification $\sqrt{n}\|\hat{h} - \tilde{h}\| = o_{P_0}(1)$. Thus, if the model was correctly specified, we would expect to realize a value of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ that is relatively small.

To gauge the magnitude of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ under correct specification, we first simulate $b = 1, \dots, B$ “observed data” series, where each series contains $n = 100$ observations generated iid from a normal random variable with mean $\theta = 1$ and $\sigma^2 = 1$. Then, using these series as our observed data, we run ABC-AR and ABC-Reg and calculate $\{\sqrt{n}\|\hat{h}_b - \tilde{h}_b\|\}_{b=1}^B$. For this experiment, we take $B = 100$ replications. The cutoff value t_n is then defined to be the 95% empirical quantile.⁸

Having obtained t_n , we now analyze the ability of this approach to detect misspecification. For the two different misspecified DGPs, which corresponds to $\sigma^2 \in \{2, 3\}$, we simulate 100 Monte Carlo replications and calculate $\sqrt{n}\|\hat{h} - \tilde{h}\|$ in each replication. The resulting sampling distributions of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ are given in Figure 4. For comparison purposes, in Panel A of Figure 4, we also give the sampling distribution of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ calculated under correct specification ($\sigma^2 = 1$). The results demonstrate that there are dramatic differences between the sampling distribution of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ under correct and incorrect specification.

Using the chosen cutoff value t_n , we detect misspecification 91% of the time when $\sigma^2 = 2$ and 97% when $\sigma^2 = 3$. Given these results, and the results in Figure 4, it is clear that this simple diagnostic allows us to detect model misspecification in this example.

It is important to realize that the behavior of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ observed in Figure 4 is not a function of the specific $h(\theta)$ we have chosen but is driven by the differences in the ABC-AR and ABC-Reg posteriors under model misspecification. In particular, the results of Corollary 1 and 2 imply that similar behavior will be observed for virtually any well-behaved function $\theta \mapsto h(\theta)$.

⁸The choice of θ used to simulate the pseudo “observed data” needed to find t_n has little impact on the magnitude of the cutoff value in large samples. Therefore, it is possible to randomly choose a value of θ from the ABC-AR posterior without significantly altering the results presented here.




Figure 4: Monte Carlo sampling distribution of $\sqrt{n}\|\hat{h}-\tilde{h}\|$ in the normal example for $\sigma^2 \in \{1, 2, 3\}$, where $h = (\theta^2, \theta^3)^\top$. Recall that $\sigma^2 = 1$ corresponds to correct specification.

To illustrate this fact, we repeat the example above using the alternative function $h(\theta) = \theta$. For this experiment, we use the exact same datasets as in the above example, and implement the same procedure to find the cutoff value t_n , but now in the case where $h(\theta) = \theta$. Across the Monte Carlo replications, this approach detects misspecification 96% of the time when $\sigma^2 = 2$ and 99% when $\sigma^2 = 3$. Similarly, Figure 5 plots the sampling distributions of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ when $h(\theta) = \theta$. While the scales are different to those in Figure 4, the results are qualitatively the same: there are significant differences between the sampling distribution of $\sqrt{n}\|\hat{h} - \tilde{h}\|$ under correct and incorrect specification. We reiterate that this result is not surprising since, under our results, we know that ABC-AR and ARC-Reg concentrate posterior mass on different values and therefore any sufficiently smooth function $h(\theta)$ will concentrate onto different values under ABC-AR and ABC-Reg. \square




Figure 5: Monte Carlo sampling distribution of $\sqrt{n}\|\hat{h}-\tilde{h}\|$ in the normal example for $\sigma^2 \in \{1, 2, 3\}$, where $h(\theta) = \theta$. Recall that $\sigma^2 = 1$ corresponds to correct specification.

Lastly, we note that the choice of θ used to construct t_n will not significantly alter the reported results. To demonstrate this, we rerun the diagnostic procedure at two additional values of θ used to construct t_n . If we use the value $\theta = 0$ to obtain t_n , in the case where $h(\theta) = (\theta^2, \theta^3)^\top$, the procedure would lead us to conclude in favor of misspecification 100% of the time for $\sigma^2 = 2$ and $\sigma^2 = 3$, while using the function $h(\theta) = \theta$ we would conclude in favor of misspecification 97% of the time under $\sigma^2 = 2$ and 100% under $\sigma^2 = 3$. Similarly, if we consider the value $\theta = 2$ and construct t_n , in the case where $h(\theta) = (\theta^2, \theta^3)^\top$, we would conclude in favor of misspecification 83% of the time for $\sigma^2 = 2$ and 92% under $\sigma^2 = 3$, while taking $h(\theta) = \theta$ we would conclude in favor of misspecification 95% of the time under $\sigma^2 = 2$ and 98% under $\sigma^2 = 3$.

5 Discussion

At first glance, ABC techniques seem less than appropriate in misspecified settings as they rely heavily on the assumed models capacity to reproduce the features of the observed dataset. Furthermore, they scale the proximity (or tolerance) in terms of the actual data as a percentile, henceforth providing a relative as opposed to absolute measure of proximity. This paper demonstrates that ABC is indeed open to poor performances in these settings, with the performance seemingly worse the more involved the version of ABC. This pattern is not immensely surprising in that more complex methods try to extract more information from the simulated data. What we find more exciting, and of particular interest, are the convergence results for the more basic ABC versions, provided an identifiability constraint on the summary statistic is satisfied. The role of the tolerance sequence also appears to be more crucial than in the well-specified case, in connection with the fact that there is a minimal non-zero limit for this tolerance. Furthermore, we demonstrate that post-processing the ABC output by local regression may lead to poor inference in misspecified models, and we propose an alternative approach that is less sensitive to the correctness of the model specification. At this stage, it is unclear if other post-processing ABC approaches, such as, for instance, the marginal adjustment approach (Nott et al., 2014) or the

recalibration approach (Rodrigues et al., 2018), will perform similarly to local regression post-processing under model misspecification, and more research on this front is necessary to obtain conclusive results.

In addition to the above, we demonstrate that, rather naturally, the discrepancies between the different ABC versions can be exploited to detect model misspecification, despite the inherent difficulty in separating Monte Carlo variability from a genuine difference. Both the examples in the paper and in the appendix illustrate that detection is achievable, since the reference tables produced by ABC can be used to automatically calibrate the expected differences between the different versions of ABC.

A potential direction of future research on misspecification not investigated in this paper would be to connect it with generative adversarial networks (GAN, Goodfellow et al., 2014) since this machine-learning technique aims both at learning a data generating mechanism (the generative network) and at separating the actual data from the generated one (the discriminating network). While the proximity between ABC and GANs has already been noted, there is however little known from a statistical viewpoint in this area about misspecification properties.

Concerning other versions of ABC based on a machine-learning approach to the intractable likelihood, for example, such as in Papamakarios and Murray (2016), our intuition is that, similar to ABC based on local nonlinear regression adjustment, the massive and unstructured dependence of the calibration on the model will similarly yield poor performance under misspecification. Apart from a GAN strategy, at this stage we fail to see a robust solution to this difficulty.

References

- Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. *Genetics*, 162(4):2025–2035.
- Bernton, E., Jacob, P. E., Gerber, M., and Robert, C. P. (2019). Approximate Bayesian computation with the wasserstein distance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Blum, M. and François, O. (2010). Non-linear regression models for approximate bayesian computation. *Statistics and Computing*, 20(1):63–73.
- Blum, M. G. (2018). Regression approaches for ABC. *Handbook of Approximate Bayesian Computation*, pages 71–85.
- Csilléry, K., François, O., and Blum, M. G. (2012). abc: an r package for approximate Bayesian computation (ABC). *Methods in ecology and evolution*, 3(3):475–479.
- Drovandi, C. C. and Pettitt, A. N. (2011). Likelihood-free bayesian estimation of multivariate quantile distributions. *Computational Statistics & Data Analysis*, 55(9):2541–2556.
- Fearnhead, P. and Prangle, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):419–474.

- Frazier, D. T., Martin, G. M., Robert, C. P., and Rousseau, J. (2018). Asymptotic properties of approximate Bayesian computation. *Biometrika*, 105(3):593–607.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Networks. *Proceedings of the International Conference on Neural Information Processing Systems (NIPS 2014)*, pages 2672–2680.
- Kleijn, B. and van der Vaart, A. (2012). The Bernstein-von-Mises theorem under misspecification. *Electron. J. Statist.*, 6:354–381.
- Li, W. and Fearnhead, P. (2018a). On the asymptotic efficiency of approximate Bayesian computation estimators. *Biometrika*, 105(2):285–299.
- Li, W. and Fearnhead, P. (2018b). Convergence of regression-adjusted approximate Bayesian computation. *Biometrika*, 105(2):301–318.
- Marin, J.-M., Pillai, N. S., Robert, C. P., and Rousseau, J. (2014). Relevant statistics for Bayesian model choice. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(5):833–859.
- Marin, J.-M., Pudlo, P., Robert, C. P., and Ryder, R. J. (2012). Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180.
- 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.
- Muller, U. K. (2013). Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix. *Econometrica*, 81(5):1805–1849.
- Nott, D. J., Fan, Y., Marshall, L., and Sisson, S. A. (2014). Approximate Bayesian computation and Bayes linear analysis: toward high-dimensional ABC. *Journal of Computational and Graphical Statistics*, 23(1):65–86.
- Papamakarios, G. and Murray, I. (2016). Fast ε -free inference of simulation models with bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*, pages 1028–1036.
- Robert, C. P. (2016). *Approximate Bayesian Computation: A Survey on Recent Results*, pages 185–205. Springer International Publishing.
- Rodrigues, G., Prangle, D., and Sisson, S. A. (2018). Recalibration: A post-processing method for approximate Bayesian computation. *Computational Statistics & Data Analysis*, 126:53–66.
- Sisson, S. A., Fan, Y., and Beaumont, M. (2018). *Handbook of Approximate Bayesian Computation*. Chapman and Hall/CRC, New York.
- Sisson, S. A., Fan, Y., and Tanaka, M. M. (2007). Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765.

A Proofs of Main Results

This section contains the proofs of the theoretical results given in the main text.

Proof of Theorem 1

This theorem is an adaptation of Frazier et al. (2018). Let $\delta_n \geq M_n(\epsilon_n - \epsilon^*) \geq 3M_nv_{0,n}^{-1}$, then $P_0(\Omega_d) = 1 + o(1)$ for $\Omega_d := \{\mathbf{y} : d(\eta(\mathbf{y}), b_0) \leq \delta_n/2\}$. Assume that $\mathbf{y} \in \Omega_d$. Consider the event

$$A_d(\delta_n) := \{(\mathbf{z}, \theta) : \{d(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_n\} \cap \{d(b(\theta), b_0) \geq \epsilon^* + \delta_n\}\}.$$

Note that, by definition $d(b(\theta), b_0) \geq \epsilon^*$, with $\epsilon^* > 0$. For all $(\mathbf{z}, \theta) \in A_d(\delta_n)$ and if $\mathbf{y} \in \Omega_d$,

$$\begin{aligned} \delta_n < d(b(\theta), b_0) - \epsilon^* &\leq d(b(\theta), \eta(\mathbf{z})) + d(\eta(\mathbf{z}), \eta(\mathbf{y})) + d(\eta(\mathbf{y}), b_0) - \epsilon^* \\ &\leq d(b(\theta), \eta(\mathbf{z})) + \epsilon_n - \epsilon^* + \delta_n/2 \end{aligned}$$

so that $\delta_n \leq 4d(b(\theta), \eta(\mathbf{z}))$. This implies in particular that

$$\begin{aligned} \Pr(A_d(\delta_n)) &= \int_{\{d(b(\theta), b_0) \geq \epsilon^* + \delta_n\}} P_\theta [d(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_n] d\Pi(\theta) \\ &\leq \int_{\Theta} P_\theta [d(b(\theta), \eta(\mathbf{z})) \geq \delta_n/4] d\Pi(\theta). \end{aligned} \tag{5}$$

In case (i) of polynomial tails,

$$\Pr(A_d(\delta_n)) \leq (v_n \delta_n)^{-\kappa} \int_{\Theta} c(\theta) d\Pi(\theta) = o(1) \tag{6}$$

as soon as $v_n \delta_n \rightarrow +\infty$, or in case (ii) of exponential tails

$$\Pr(A_d(\delta_n)) \leq C e^{-c(\delta_n v_n)^\tau}. \tag{7}$$

Moreover, we can bound from below

$$\alpha_n = \int_{\Theta} P_\theta [d(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_n] d\Pi(\theta)$$

Note that on $\{d(\eta(\mathbf{z}), b(\theta)) \leq Mv_n^{-1}/2\} \cap \Omega_d$

$$d(\eta(\mathbf{y}), \eta(\mathbf{z})) \leq d(\eta(\mathbf{z}), b(\theta)) + d(\eta(\mathbf{y}), b_0) + d(b(\theta), b_0) \leq v_{0,T}^{-1} + Mv_n^{-1}/2 + d(b(\theta), b_0) \leq \epsilon_n$$

as soon as $\epsilon^* \leq d(b(\theta), b_0) \leq \epsilon_n - v_{0,T}^{-1} + Mv_n^{-1}/2$. Since $\epsilon_n - \epsilon^* \geq v_{0,T}^{-1} + Mv_n^{-1}$, on Ω_d ,

$$\begin{aligned} \int_{\Theta} P_{\theta}(d(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_n) d\Pi(\theta) &\geq \int_{d(b(\theta), b_0) \leq (\epsilon_n - \epsilon^*)/4 \vee v_n^{-1} M/2} (1 - P_{\theta}(d(\eta(\mathbf{z}), b(\theta)) \geq Mv_n^{-1}/2)) d\Pi(\theta) \\ &\geq \int_{d(b(\theta), b_0) \leq (\epsilon_n - \epsilon^*)/4 \vee v_n^{-1} M/2} \left(1 - \frac{c(\theta)2^{\kappa}}{M^{\kappa}}\right) d\Pi(\theta) \\ &\gtrsim (\epsilon_n - \epsilon^*)^D \vee v_n^{-D} \gtrsim (\epsilon_n - \epsilon^*)^D \end{aligned}$$

in case (i) of [A1], under [A2]. If case (ii) of [A1] holds, under [A2], we have

$$\begin{aligned} \int_{\Theta} P_{\theta}(d(\eta(\mathbf{z}), \eta(\mathbf{y})) \leq \epsilon_n) d\Pi(\theta) &\geq \int_{d(b(\theta), b_0) \leq (\epsilon_n - \epsilon^*)/4 \vee v_n^{-1} M/2} (1 - c(\theta)e^{-h_{\theta}(M/2)}) d\Pi(\theta) \\ &\gtrsim (\epsilon_n - \epsilon^*)^D \end{aligned}$$

Combining these two inequality with the upper bounds (6) or (7) leads to

$$\Pi_{\epsilon}[d(b(\theta), b_0) \geq \epsilon^* + \delta_n |\eta(\mathbf{y})|] \lesssim (\epsilon_n - \epsilon^*)^{-D} (v_n \delta_n)^{-\kappa},$$

in case (i) and

$$\Pi_{\epsilon}[d(b(\theta), b_0) \geq \epsilon^* + \delta_n |\eta(\mathbf{y})|] \lesssim (\epsilon_n - \epsilon^*)^{-D} e^{-c(\delta_n v_n)^{\tau}},$$

in case (ii). These are of order $o(1)$ if $\delta_n \geq M_n v_n^{-1} (\epsilon_n - \epsilon^*)^{-D/\kappa}$ in case (i), or if $\delta_n \geq M_n v_n^{-1} |\log(\epsilon_n - \epsilon^*)|^{1/\tau}$ in case (ii).

Proof of Corollary 1

Proof. Define $Q(\theta) = |d(b(\theta), b_0) - d(b(\theta^*), b_0)|$. From the continuity of $\theta \mapsto b(\theta)$ and the definition of θ^* , for any $\delta > 0$ there exists a $\gamma(\delta) > 0$ such that $\inf_{\theta: d\{\theta, \theta^*\} > \delta} Q(\theta) \geq \gamma(\delta) > 0$. Then,

$$\begin{aligned} \Pi_{\epsilon}[d(\theta, \theta^*) > \delta |\eta(\mathbf{y})|] &\leq \Pi_{\epsilon}[|Q(\theta) - Q(\theta^*)| > \gamma(\delta) |\eta(\mathbf{y})|] = \Pi_{\epsilon}[|d(b(\theta), b_0) - d(b(\theta^*), b_0)| > \gamma(\delta) |\eta(\mathbf{y})|] \\ &= \Pi_{\epsilon}[d(b(\theta), b_0) > \epsilon^* + \gamma(\delta) |\eta(\mathbf{y})|]. \end{aligned}$$

The result follows if $\Pi_{\epsilon}[|d(b(\theta), b_0) - d(b(\theta^*), b_0)| > \gamma(\delta) |\eta(\mathbf{y})|] = o_{P_0}(1)$. For $\delta_n > 0$ and $\delta_n = o(1)$ as defined in Theorem 1, by the conclusion of Theorem 1, the result follows once $\gamma(\delta) \geq \delta_n$. \square

Proof of Theorem 2

Proof. For the sake of simplicity and without loss of generality we write $v_n = \sqrt{n}$, $\tilde{Z}_n = \sqrt{n}(\eta(\mathbf{z}) - b(\theta))$ and $\tilde{Z}_y = \sqrt{n}(\eta(\mathbf{y}) - b_0)$. Denote by $B_n(K) = \{\|\theta - \theta^*\| \leq K\}$. Throughout the proof C

denotes a generic constant which may vary from line to line. We have for all θ

$$\begin{aligned}
P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) &= P_\theta \left(\|\tilde{Z}_n - \tilde{Z}_y + \sqrt{n}(b(\theta) - b_0)\|^2 \leq n\epsilon_n^2 \right) \\
&= P_\theta \left(\|\tilde{Z}_n\|^2 + 2\langle Z_n, \sqrt{n}(b(\theta) - b_0) - \tilde{Z}_y \rangle \leq n[\epsilon_n^2 - \|b(\theta) - b_0 - \tilde{Z}_y/\sqrt{n}\|^2] \right) \\
&= P_\theta \left(\langle \tilde{Z}_n, b(\theta) - b_0 \rangle \leq \frac{\sqrt{n}[\epsilon_n^2 - \|b(\theta) - b_0 - \tilde{Z}_y/\sqrt{n}\|^2]}{2} - \frac{\|\tilde{Z}_n\|^2 - 2\langle \tilde{Z}_n, \tilde{Z}_y \rangle}{2\sqrt{n}} \right)
\end{aligned}$$

Now on $\Omega_n = \{\|\tilde{Z}_y\| \leq M_n/2\}$ with M_n a sequence going to infinity arbitrarily slowly and such that $M_n = o(n^{1/4})$,

$$\begin{aligned}
\sqrt{n}\|b(\theta) - b_0 - \tilde{Z}_y/\sqrt{n}\|^2 &= \sqrt{n}\|b(\theta^*) - b_0\|^2 + \frac{\sqrt{n}(\theta - \theta^*)^\top H^*(\theta - \theta^*)}{2} - 2\langle b(\theta^*) - b_0, \tilde{Z}_y \rangle \\
&\quad + O(M_n^2/\sqrt{n}) + O(\sqrt{n}\|\theta - \theta^*\|^3 + \|\theta - \theta^*\|M_n)
\end{aligned}$$

where H^* is the second derivative of $\theta \mapsto \|b(\theta) - b_0\|^2$ at θ^* , noting that the first derivative is equal to 0 at θ^* . Let $\epsilon^* = \|b(\theta^*) - b_0\|$, $e' = (b(\theta^*) - b_0)$, and $\epsilon > 0$. If $\|\theta - \theta^*\| \leq \epsilon$ and on the event $\Omega_n = \{\|Z_y\| \leq M_n\}$ where $M_n^2 = o(\sqrt{n})$,

$$\begin{aligned}
&P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \\
&\leq P_\theta \left(\langle \tilde{Z}_n, e' \rangle \leq \frac{\sqrt{n}[\epsilon_n^2 - (\epsilon^*)^2 - (1 + C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \langle \tilde{Z}_y, e' \rangle + C\epsilon + \|\theta - \theta^*\|M_n \right) \\
&\quad + P_\theta (\|Z_n\|^2 > \epsilon^2\sqrt{n}/4) \\
&\geq P_\theta \left(\langle \tilde{Z}_n, e' \rangle \leq \frac{\sqrt{n}[\epsilon_n^2 - (\epsilon^*)^2 - (1 - C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \langle \tilde{Z}_y, e' \rangle - C\epsilon - \|\theta - \theta^*\|M_n \right) \\
&\quad - P_\theta (\|\tilde{Z}_n\|^2 > \epsilon^2\sqrt{n}/4) \tag{8}
\end{aligned}$$

Consider the case where $\sqrt{n}(\epsilon_n^2 - (\epsilon^*)^2) \rightarrow 2c \in \mathbb{R}$. We split Θ into $\{\|\theta - \theta^*\|^2\sqrt{n} \leq M\}$, $\{\epsilon > \|\theta - \theta^*\| > M/n^{1/4}\}$ and $\{\epsilon \leq \|\theta - \theta^*\|\}$, where ϵ is arbitrarily small.

First if $\|\theta - \theta^*\|^2\sqrt{n} \leq M$

$$\begin{aligned}
&\frac{\sqrt{n}[\epsilon_n^2 - (\epsilon^*)^2 - (1 - C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \langle \tilde{Z}_y, e' \rangle + \frac{M_n^2}{\sqrt{n}} + \|\theta - \theta^*\|M_n \\
&\leq c + \langle \tilde{Z}_y, e' \rangle - \frac{(1 - C\epsilon)\sqrt{n}(\theta - \theta^*)^\top H^*(\theta - \theta^*)}{4} + \epsilon
\end{aligned}$$

and

$$\begin{aligned}
&\frac{\sqrt{n}[\epsilon_n^2 - (\epsilon^*)^2 - (1 + C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \frac{M^2}{\sqrt{n}} + \|\theta - \theta^*\|M_n \\
&\geq c + \langle \tilde{Z}_y, e' \rangle - \frac{(1 + C\epsilon)\sqrt{n}(\theta - \theta^*)^\top H^*(\theta - \theta^*)}{4} - \epsilon
\end{aligned}$$

Moreover, using assumption **[A5]**,

$$\langle \tilde{Z}_n, e' \rangle = \sqrt{n} \langle \Sigma_n(\theta)^{-1} Z_n, e' \rangle = \langle Z_n, \sqrt{n} \Sigma_n(\theta)^{-1} e' \rangle = \langle Z_n, A(\theta^*) e' \rangle + o(\|Z_n\|).$$

We then have with $c' = c + \langle \tilde{Z}_y, e' \rangle$, $x = n^{1/4}(1 - \epsilon)^{1/2}(\theta - \theta^*)$ and $\|\theta - \theta^*\| \leq M/n^{1/4} \leq u_0$ if n is large enough

$$\begin{aligned} P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) &\leq P_\theta \left(\langle Z_n, A(\theta^*) e' \rangle \leq c' - \frac{x^\top H^* x}{4} + \epsilon \right) + \epsilon \\ &\leq \Phi \left(\frac{c' + M\epsilon}{\|A(\theta^*) e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*) e'\|} \right) + \epsilon \\ &+ \sup_{\|\theta' - \theta^*\| \leq u} \left| P_{\theta'} \left(\langle Z_n, A(\theta^*) e' \rangle \leq c' - \frac{x^\top H^* x}{4} + \epsilon \right) - \Phi \left(\frac{c' + \epsilon}{\|A(\theta^*) e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*) e'\|} \right) \right| \\ &\leq \Phi \left(\frac{c' + \epsilon}{\|A(\theta^*) e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*) e'\|} \right) + C\epsilon. \end{aligned}$$

Similarly with $y = n^{1/4}(1 + \epsilon)^{1/2}(\theta - \theta^*)$

$$\begin{aligned} P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) &\geq \Phi \left(\frac{c' - \epsilon}{\|A(\theta^*) e'\|} - \frac{y^\top H^* y}{4\|A(\theta^*) e'\|} \right) - \epsilon \\ &- \sup_{\|\theta - \theta^*\| \leq u} \left| P_\theta \left(\langle Z_n, A(\theta^*) e' \rangle \leq c' - \frac{y^\top H^* y}{4} - \epsilon \right) - \Phi \left(\frac{c' - \epsilon}{\|A(\theta^*) e'\|} - \frac{y^\top H^* y}{4\|A(\theta^*) e'\|} \right) \right| \\ &\geq \Phi \left(\frac{c' - M\epsilon}{\|A(\theta^*) e'\|} - \frac{y^\top H^* y}{4\|A(\theta^*) e'\|} \right) + C\epsilon \end{aligned}$$

Moreover for all $t \in \mathbb{R}$, writing $x(\theta)$ to emphasize its dependence in θ ,

$$\begin{aligned} \Delta_1 &= \int_{B_n(\frac{M}{n^{1/4}})} \sup_{\|\theta' - \theta^*\| \leq u_0} \left| P_{\theta'} \left(\langle Z_n, A(\theta^*) e' \rangle \leq t - \frac{x(\theta)^\top H^* x(\theta)}{4} \right) - \Phi \left(\frac{t}{\|A(\theta^*) e'\|} - \frac{x(\theta)^\top H^* x(\theta)}{4\|A(\theta^*) e'\|} \right) \right| d\theta \\ &= n^{-k_\theta/4} \int_{\|x\| \leq M} \sup_{\|\theta' - \theta^*\| \leq u_0} \left| P_{\theta'} \left(\langle Z_n, A(\theta^*) e' \rangle \leq t - \frac{x^\top H^* x}{4} \right) - \Phi \left(\frac{t}{\|A(\theta^*) e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*) e'\|} \right) \right| dx \\ &= n^{-k_\theta/4} M o(1) \end{aligned}$$

where the last equality follows from **[A4]** We then have

$$\begin{aligned} &\int_{B_n(M/n^{1/4})} P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta \\ &\leq \pi(\theta^*) (1 + o(1)) n^{-k_\theta/4} \int_{\|x\| \leq M} \Phi \left(\frac{c' + M\epsilon}{\|A(\theta^*) e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*) e'\|} \right) dx + o(n^{-k_\theta/4} M) \end{aligned} \quad (9)$$

Similarly

$$\begin{aligned}
& \int_{\Theta} P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta \\
& \geq \int_{\|\theta - \theta^*\| \leq n^{-1/4} M} P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta \\
& \geq \pi(\theta^*) (1 + o(1)) n^{-k_{\theta}/4} \int_{\|x\| \leq M} \Phi \left(\frac{c' - M\epsilon}{\|A(\theta^*)e'\|} - \frac{x^{\top} H^* x}{4\|A(\theta^*)e'\|} \right) dx (1 + o_{P_0}(1)).
\end{aligned} \tag{10}$$

Also if $\epsilon > \|\theta - \theta^*\| > M/n^{1/4}$, since there exists $a > 0$ such that $z^{\top} H^* z \geq a\|z\|^2$ for all z , if M is large enough,

$$P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \leq P_{\theta} \left(\langle Z_n, A(\theta^*)e' \rangle \leq -\frac{aM^2}{8} \right) \leq P_{\theta} \left(\|Z_n\| > \frac{aM^2}{8\|A(\theta^*)e'\|} \right) \lesssim M^{-2\kappa}$$

Now let $j \geq 0$ and set $M_j = 2^j M$. On $M_j n^{-1/4} \leq \|\theta - \theta^*\| \leq M_{j+1} n^{-1/4}$

$$P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \leq P_{\theta} \left(\langle Z_n, A(\theta^*)e' \rangle \leq -\frac{aM_j^2}{8} \right) \leq P_{\theta} \left(\|Z_n\| > \frac{aM_j^2}{8\|A(\theta^*)e'\|} \right) \lesssim M_j^{-2\kappa}$$

so that

$$\begin{aligned}
\int_{Mn^{-1/4} \leq \|\theta - \theta^*\| \leq \epsilon} P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta & \lesssim n^{-k_{\theta}/4} \sum_{j=0}^{J_n} M_j^{-2\kappa} (M_{j+1} - M_j) \\
& \lesssim n^{-k_{\theta}/4} \sum_{j=0}^{J_n} M_j^{-2\kappa+1} \lesssim n^{-k_{\theta}/4} M^{-2\kappa+1}.
\end{aligned} \tag{11}$$

Finally if $\|\theta - \theta^*\| > \epsilon$, $\|b(\theta) - b_0 - Z_y/\sqrt{n}\|^2 - (\epsilon^*)^2 \geq C\epsilon$ on Ω_n and when n is large enough $\sqrt{n}(\epsilon_n^2 - (\epsilon^*)^2) \leq c + C\epsilon$ so that

$$\begin{aligned}
P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) & \leq P_{\theta} \left(\langle \tilde{Z}_n, b(\theta) - b_0 \rangle \leq -\sqrt{n}C\epsilon + c' + C\epsilon \right) + O(n^{-\kappa/4}) \\
& \leq P_{\theta} (\|Z_n\| > Cn^{1/2} \|b(\theta) - b_0\|^{-1}/2) + O(n^{-\kappa/4}) \\
& \leq c(\theta) C^{-\kappa} n^{-\kappa/2} \|b(\theta) - b_0\|^{\kappa} + O(n^{-\kappa/4}).
\end{aligned}$$

Therefore

$$\begin{aligned}
\int_{\|\theta - \theta^*\| \geq \epsilon} P_{\theta} (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta & \leq C^{-\kappa} n^{-\kappa/2} \int_{\|\theta - \theta^*\| \geq \epsilon} c(\theta) \|b(\theta) - b_0\|^{\kappa} \pi(\theta) d\theta \\
& + Cn^{-\kappa/4} \int_{\|\theta - \theta^*\| \geq \epsilon} c(\theta) \pi(\theta) d\theta
\end{aligned} \tag{12}$$

Finally combining (9), (11), (12) and (10) we obtain that if $\kappa > k_\theta$

$$\int_{\Theta} P_\theta (\|\eta(\mathbf{z}) - \eta(\mathbf{y})\|^2 \leq \epsilon_n^2) \pi(\theta) d\theta = n^{-k_\theta/4} \int_{\mathbb{R}} \Phi \left(\frac{c'}{\|A(\theta^*)e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*)e'\|} \right) dx + o(n^{-k_\theta/4})$$

and for all $x = n^{1/4}(\theta - \theta^*) \in \mathbb{R}^{k_\theta}$ fixed, writing $\pi_{z_n, \epsilon}(\cdot)$ the density of $\Pi_{z_n, \epsilon}$, the ABC posterior distribution of $z_n(\theta - \theta^*)$,

$$\pi_{n^{1/4}, \epsilon}(x) = \frac{\Phi \left(\frac{c'}{\|A(\theta^*)e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*)e'\|} \right)}{\int_{\mathbb{R}} \Phi \left(\frac{c'}{\|A(\theta^*)e'\|} - \frac{x^\top H^* x}{4\|A(\theta^*)e'\|} \right) dx} + o(1) := q_c(x) + o(1)$$

so that $\|\pi_{n^{1/4}, \epsilon} - q_c\|_1 = o(1)$.

We now study the case where $\sqrt{n}(\epsilon_n^2 - (\epsilon^*)^2) := \sqrt{n}u_n^2 \rightarrow +\infty$ with $u_n = o(1)$ and we show that the limiting distribution is uniform. Using (8), we have that if $B_{0,n} = \{(\theta - \theta^*)^\top H^*(\theta - \theta^*) \leq 2u_n^2 - 4M_n/\sqrt{n}\}$, with $M_n < u_n^2\sqrt{n}\epsilon$ going to infinity

$$\begin{aligned} P_\theta(\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n) &\leq 1 \\ &\geq P_\theta \left(\langle \tilde{Z}_n, e' \rangle \leq 2M_n + \langle \tilde{Z}_y, e' \rangle - \epsilon - \|\theta - \theta^*\|M_n \right) \\ &\geq P_\theta \left(\langle \tilde{Z}_n, e' \rangle \leq M_n/2 \right) \geq 1 - \frac{c_1}{M_n^{-\kappa}} \end{aligned}$$

for some $c_1 > 0$ on the event $\{|\langle \tilde{Z}_y, e' \rangle| \leq M_n/2\}$, which has probability going to 1.

This implies in particular that

$$\begin{aligned} \int_{B_{0,n}} P_\theta(\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n) \pi(\theta) d\theta &\leq \pi(\theta^*)(1 + o(1)) \text{Vol}(B_{0,n}) \\ &\geq \pi(\theta^*)(1 + o(1)) \text{Vol}(B_{0,n}) \left(1 - \frac{c_1}{M_n^{-\kappa}}\right) \end{aligned} \quad (13)$$

Also

$$\text{Vol}(B_{0,n}) \asymp u_n^{k_\theta} \quad (14)$$

Let $K_n \geq 4$, if $K_n u_n^2 \geq (\theta - \theta^*)^\top H^*(\theta - \theta^*) > 2u_n^2(1 - C\epsilon)^{-1} + 4M_n/\sqrt{n}$, then there exists $C' > 0$ such that

$$\begin{aligned} &\frac{\sqrt{n}[u_n^2 - (1 - C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \langle \tilde{Z}_y, e' \rangle - \epsilon - \|\theta - \theta^*\|M_n \\ &\leq -2M_n + \langle \tilde{Z}_y, e' \rangle - \epsilon - C'\epsilon^2 M_n \leq -M_n \end{aligned}$$

on the event $\{|\langle \tilde{Z}_y, e' \rangle| \leq M_n/2\}$. Therefore writing $B_{1,n} = \{K_n u_n^2 \geq (\theta - \theta^*)^\top H^*(\theta - \theta^*) > 2u_n^2(1 - C\epsilon)^{-1} + 4M_n/\sqrt{n}\}$

$$\int_{B_{1,n}} P_\theta(\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n) \pi(\theta) d\theta \lesssim M_n^{-\kappa} \text{Vol}(B_{1,n}) \lesssim M_n^{-\kappa} K_n^{k_\theta/2} \text{Vol}(B_{0,n}) \quad (15)$$

Moreover

$$\text{Vol}(\{(\theta - \theta^*)^\top H^*(\theta - \theta^*) \leq 2u_n^2(1 - C\epsilon)^{-1} + 4M_n/\sqrt{n}\}) - \text{Vol}(B_{0,n}) \lesssim \epsilon \text{Vol}(B_{0,n}). \quad (16)$$

If $K_n u_n^2 \leq (\theta - \theta^*)^\top H^*(\theta - \theta^*) \leq \epsilon^2$, then

$$\begin{aligned} & \frac{\sqrt{n}[u_n^2 - (1 - C\epsilon)(\theta - \theta^*)^\top H^*(\theta - \theta^*)/2]}{2} + \langle \tilde{Z}_y, e' \rangle - \epsilon - \|\theta - \theta^*\| M_n \\ & \leq -\frac{\sqrt{n}(\theta - \theta^*)^\top H^*(\theta - \theta^*)}{8} \end{aligned}$$

when n is large enough and there exists $b > 0$ such that

$$\begin{aligned} \int_{B_{1,n}^c} \mathbb{1}_{(\theta - \theta^*)^\top H^*(\theta - \theta^*) \leq \epsilon^2} P_\theta(\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n) \pi(\theta) d\theta & \lesssim n^{-\kappa/2} \int_{B_n(A\epsilon^2)} \mathbb{1}_{\|\theta - \theta^*\| \geq b\sqrt{K_n}u_n} \|\theta - \theta^*\|^{-2\kappa} d\theta \\ & \lesssim n^{-\kappa/2} \int_{b\sqrt{K_n}u_n}^{\sqrt{A}\epsilon} r^{k_\theta - 2\kappa - 1} dr \end{aligned} \quad (17)$$

Since $k_\theta < 2\kappa$ then the above term is of order

$$K_n^{(k_\theta - 2\kappa)/2} (\sqrt{n}u_n^2)^{-\kappa/2} u_n^{k_\theta} \asymp K_n^{(k_\theta - 2\kappa)/2} (\sqrt{n}u_n^2)^{-\kappa/2} \text{Vol}(B_{0,n}) = o(\text{Vol}(B_{0,n}))$$

Finally if $\|\theta - \theta^*\| \geq \epsilon$, similarly to the case where $\sqrt{n}u_n^2 \rightarrow c \in \mathbb{R}$, we obtain (12) and this term is $o(\text{Vol}(B_{0,n}))$ as soon as $n^{-\kappa/4} = o(u_n^{k_\theta})$. Since $n^{-1/4} = o(u_n)$ the latter is true as soon as $\kappa \geq k_\theta$. Combining (13), (15), (17), (16) and (12), we obtain that

$$\left\| \int_{\Theta} P_\theta(\|\eta(Z) - \eta(\mathbf{y})\| \leq \epsilon_n) \pi(\theta) d\theta - \pi(\theta^*) \text{Vol}(\tilde{B}_{0,n}) \right\| = o_p(1) \quad (18)$$

where $\tilde{B}_{0,n} = \{(\theta - \theta^*)^\top H^*(\theta - \theta^*) \leq 2u_n^2\}$. Let $x = u_n^{-1}(\theta - \theta^*)$ be fixed and $x^\top H^* x < 2$, then for n large enough $x^\top H^* x \leq 2 - 4M_n^2/(\sqrt{n}u_n^2)$ and using

$$\pi_{u_n^{-1}, \epsilon}(x) = \pi_\epsilon(\theta^* + u_n x | \mathbf{y}) u_n^{k_\theta}$$

then

$$\pi_{u_n^{-1}, \epsilon}(x) = 1 + o_p(1).$$

If $x^\top H^* x > 2$, then if $\epsilon > 0$ is small enough and n is large enough $x^\top H^* x \geq 2(1 - C\epsilon)^{-1} + 4M_n^2/(\sqrt{n}u_n^2)$ and

$$\pi_{u_n^{-1}, \epsilon}(x) = o_p(1).$$

This implies that the ABC posterior distribution of $u_n^{-1}(\theta - \theta^*)$ converges to the Uniform distribution over the ellipsoid $\{x^\top H^* x \leq 2\}$ in total variation. \square

Proof of Proposition 1

Proof. To prove Proposition 1, we prove that the approximate likelihood

$$P_\theta \left(\left\| \tilde{Z}_n - Z_y + \sqrt{n}(b(\theta) - b_0) \right\|^2 \leq \epsilon_n^2 \right)$$

is highly peaked around $\theta \neq 0$, and, as such, concentration around $\theta^* = 0$ can not result.

As in the proof of Theorem 2, writing $Z_n = Z = (Z_1, Z_2)^\top$, we can define $W = Z/\|Z\|$ and $R = \|Z\|/v_\theta$ and we have that W and R are independent and that their distribution does not depend on θ . In particular $R^2 \sim \chi^2(2)$. Now, set $h = b(\theta) - b_0 - Z_y/\sqrt{n}$, so that

$$\|Z - Z_y + \sqrt{n}(b(\theta) - b_0)\|^2 - n\epsilon_n^2 = v_\theta^2 R^2 + 2\sqrt{n}Rv_\theta \langle W, h \rangle + n(\|h\|^2 - \epsilon_n^2) \leq 0 \quad (19)$$

if and only if

$$\Delta(W) = v_\theta^2 n (\langle W, h \rangle^2 - \|h\|^2 + \epsilon_n^2) = v_\theta^2 n \tilde{\Delta}(W) \geq 0, \quad R \in (r_1(W), r_2(W)) \cap \mathbb{R}_+,$$

where

$$r_1(W) = \frac{\sqrt{n}}{v_\theta} \left[-\langle W, h \rangle - \sqrt{\tilde{\Delta}(W)} \right], \quad r_2(W) = \frac{\sqrt{n}}{v_\theta} \left[-\langle W, h \rangle + \sqrt{\tilde{\Delta}(W)} \right].$$

Note that $\tilde{\Delta}(W) \leq \epsilon_n^2$ so that if $\tilde{\Delta}(W) \geq 0$ then $|\langle W, h \rangle| = \|h\|(1 + O(\epsilon_n))$, given that $\|h\| \asymp 1$ on the event $\|Z_y\| \leq M$ for some arbitrarily large M . Therefore if $\langle -W, h \rangle \leq 0$, then $\langle -W, h \rangle \asymp -\|h\|$ and there is no solution for R in (19). Hence (19) holds if and only if $\langle -W, h \rangle \geq 0$, $\tilde{\Delta}(W) \geq 0$ and $R \in (r_1(W), r_2(W))$.

By symmetry we can set $W = -W$ and, on the set $\langle W, h \rangle \geq 0$, using the fact that $R^2 \sim \chi^2(2)$,

$$P_\theta \left(\|Z - Z_y + \sqrt{n}(b(\theta) - b_0)\|^2 \leq n\epsilon_n^2 \mid W \right) = e^{-r_1(W)^2/2} \left(1 - e^{-\frac{n\tilde{\Delta}(W)}{v_\theta^2}} \right) \\ r_1(W) \in (\sqrt{n}\|h\|/v_\theta \{1 - 2\epsilon_n\}, \sqrt{n}\|h\|/v_\theta)$$

To derive an approximation of $P_\theta \left(\|Z - Z_y + \sqrt{n}(b(\theta) - b_0)\|^2 \leq n\epsilon_n^2 \right)$ we study more precisely $r_1(W)$. For the sake of simplicity we assume that $\sqrt{n}\epsilon_n = o(1)$, since the case where $\sqrt{n}\epsilon_n = O(1)$ can be treated similarly. Then

$$P_\theta \left(\|Z - Z_y + \sqrt{n}(b(\theta) - b_0)\|^2 \leq n\epsilon_n^2 \right) \leq e^{-\frac{n\|h\|^2}{2v_\theta^2}(1-2\epsilon_n)^2} P_\theta \left(\tilde{\Delta}(W) \geq 0 \right) \\ \geq e^{-\frac{n\|h\|^2}{2v_\theta^2}} P_\theta \left(\tilde{\Delta}(W) \geq 0 \right)$$

Consider $\theta = r\bar{b}_0$ with $r \in [-1, 1]$ so that $r = 0$ corresponds to $\theta = \theta^* \equiv 0$, then $h =$

$\bar{b}_0(r-1, r+1)^\top + O_{P_0}(1/\sqrt{n})$ and

$$\begin{aligned} P_\theta \left(\tilde{\Delta}(W) \geq 0 \right) &\leq 2P_\theta \left(0 \leq \langle W - h/\|h\|, h \rangle \leq \epsilon_n^2/(2\|h\|^2)(1 + \epsilon_n^2/\|h\|^2) \right) \\ &\geq 2P_\theta \left(0 \leq \langle W - h/\|h\|, h \rangle \leq \epsilon_n^2/(2\|h\|^2) \right) \\ &= \frac{\epsilon_n^2}{(1+r^2)\bar{b}_0^2} g \left\{ \frac{(1-r)}{\sqrt{2(1+r^2)}} \right\} \left\{ 1 + O(\epsilon_n^2 \vee 1/\sqrt{n}) \right\} \end{aligned}$$

where $g(\cdot)$ is the density of W_1 , with $W = (W_1, W_2)$. We thus obtain that for n large enough

$$\begin{aligned} P_\theta \left(\left\| Z - Z_y + \sqrt{n}(b(\theta) - b_0) \right\|^2 \leq n\epsilon_n^2 \right) &\leq \frac{\epsilon_n^2}{(1+r^2)\bar{b}_0^2} e^{-\frac{n\bar{b}_0^2(1+r^2)}{2v_\theta^2}(1-3\epsilon_n)^2} \times \\ &g \left\{ \frac{\bar{b}_0(1-r)}{\sqrt{2(1+r^2)}} \right\} \left\{ 1 + O(\epsilon_n^2 \vee 1/\sqrt{n}) \right\} \\ &\geq \frac{\epsilon_n^2}{(1+r^2)\bar{b}_0^2} e^{-\frac{n\bar{b}_0^2(1+r^2)}{2v_\theta^2}(1+\epsilon_n)^2} \times \\ &g \left\{ \frac{\bar{b}_0(1-r)}{\sqrt{2(1+r^2)}} \right\} \left\{ 1 + O(\epsilon_n^2 \vee 1/\sqrt{n}) \right\} \end{aligned}$$

Take $v_{r\bar{b}_0} = \bar{b}_0 v(r)$ such that $(1+1/4)/v(1/2)^2 \leq 1/(2v(0)^2)$, then for $\delta > 0$ small enough,

$$\Pi_\epsilon (|\theta - \theta^*| \leq \delta | \eta(\mathbf{y})) = o\left(\Pi_\epsilon \left\{ |\theta - \bar{b}_0/2| \leq \delta | \eta(\mathbf{y}) \right\}\right)$$

since there exists $c > 0$ such that

$$\int_{|\theta| \leq \delta} e^{-\frac{n(\bar{b}_0^2 + \theta^2)}{2v_\theta^2}} \pi(\theta) d\theta \leq e^{-nc} \int_{|\theta - \bar{b}_0/2| \leq \delta} e^{-\frac{n(\bar{b}_0^2 + \theta^2)}{2v_\theta^2}} \pi(\theta) d\theta.$$

□

Proof of Corollary 2

Proof. The proof is a consequence of Theorem 1 and the structure of $\tilde{\theta} = \theta - \hat{\beta}^\top \{\eta(\mathbf{z}) - \eta(\mathbf{y})\}$, and $\tilde{\theta}^* = \theta^* - \beta_0^\top \{b(\theta^*) - b_0\}$. Therefore, we only sketch the idea here.

Take $\delta_n \geq M_n(\epsilon_n - \epsilon^*) \geq M_n v_n^{-1}$. By assumption $\epsilon^* > 0$ and $\|\beta_0\| > 0$. Define $\Omega_d = \{\mathbf{y} : \|\eta(\mathbf{y}) - b_0\| \leq \delta_n/u_0\}$ for some $u_0 \geq 2(1 + \|\beta_0\|)$. By the result of Theorem 1 we have that

$$\begin{aligned} \tilde{\Pi}_\epsilon \left[|\tilde{\theta} - \tilde{\theta}^*| > \delta_n | \mathbf{y} \right] &= \Pi_\epsilon \left[\{\theta : |\tilde{\theta} - \tilde{\theta}^*| > \delta_n\} \cap \{\theta : |\theta - \theta^*| \leq \delta_n/u_0\} | \mathbf{y} \right] + o_{P_0}(1) \\ &= \frac{\int_{|\theta - \theta^*| \leq \delta_n/u_0} \mathbb{1} \left[|\tilde{\theta} - \tilde{\theta}^*| \geq \delta_n \right] P_\theta \left[\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n \right] d\Pi(\theta)}{\int_{|\theta - \theta^*| \leq \delta_n/u_0} P_\theta \left[\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n \right] d\Pi(\theta)} + o_{P_0}(1), \end{aligned}$$

where both equalities follow by posterior concentration of $|\theta - \theta^*|$ at rate $\delta_n \gg v_{0,n}^{-1}$. Similar steps to that of Theorem 1 yield

$$D_n = \int_{|\theta - \theta^*| \leq \delta_n/u_0} P_\theta [\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n] d\Pi(\theta) \gtrsim \delta_n^D,$$

under case (i) or case (ii) of **[A1]**. Define the event

$$S(\delta_n) = \left\{ (\mathbf{z}, \theta) : \{\tilde{\theta} - \tilde{\theta}^* > \delta_n\} \cap \{\theta : |\theta - \theta^*| \leq \delta_n/u_0\} \cap \{\mathbf{z} : \|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n\} \right\}$$

Note that

$$\begin{aligned} \tilde{\theta} - \tilde{\theta}^* &= \theta - \theta^* + [\hat{\beta} - \beta_0]^\top \{b(\theta) - b_0\} + [\hat{\beta} - \beta_0]^\top \{b_0 - \eta(\mathbf{y})\} + [\hat{\beta} - \beta_0]^\top \{\eta(\mathbf{z}) - b(\theta)\} \\ &\quad + [\hat{\beta} - \beta_0]^\top \{b(\theta) - b(\theta^*)\} + \beta_0^\top \{b_0 - \eta(\mathbf{y})\} + \beta_0^\top \{\eta(\mathbf{z}) - b(\theta)\} + \beta_0^\top \{b(\theta) - b(\theta^*)\} \end{aligned}$$

For $\mathbf{y} \in \Omega_d$, we have

$$\begin{aligned} \delta_n < |\tilde{\theta} - \tilde{\theta}^*| &\leq |\theta - \theta^*| + \|\hat{\beta} - \beta_0\| \|b(\theta) - b_0\| + \|\hat{\beta} - \beta_0\| \|b_0 - \eta(\mathbf{y})\| + \|\hat{\beta} - \beta_0\| \|\eta(\mathbf{z}) - b(\theta)\| \\ &\quad + \|\hat{\beta} - \beta_0\| \|b(\theta) - b(\theta^*)\| + \|\beta_0\| \|b_0 - \eta(\mathbf{y})\| + \|\beta_0\| \|b(\theta) - b(\theta^*)\| + \|\beta_0\| \|\eta(\mathbf{z}) - b(\theta)\| \\ &\leq \delta_n/u_0 + \|\beta_0\| \delta_n/u_0 + o(\delta_n) + (O(\delta_n) + \|\beta_0\|) \|\eta(\mathbf{z}) - b(\theta)\| + o_{P_\theta}(1) \end{aligned}$$

where the last inequality follows from $\|\hat{\beta} - \beta_0\| = o_{P_\theta}(1)$ and concentration of $|\theta - \theta^*|$ at rate $\delta_n \gg v_n^{-1}$. Therefore, take $u_0 \geq 2(1 + \|\beta_0\|)$ and rearrange the above to obtain

$$0 < \frac{\delta_n}{2(O(\delta_n) + \|\beta_0\|)} < \|\eta(\mathbf{z}) - b(\theta)\| + o(\delta_n).$$

This then implies that

$$\begin{aligned} \Pr[S(\delta_n)] &= \int_{\{\theta : |\theta - \theta^*| \leq \delta_n/u_0\}} \mathbb{1} \left[|\tilde{\theta} - \tilde{\theta}^*| > \delta_n \right] P_\theta [\|\eta(\mathbf{z}) - \eta(\mathbf{y})\| \leq \epsilon_n] d\Pi(\theta) \\ &\leq \int_{\Theta} P_\theta [\|\eta(\mathbf{z}) - b(\theta)\| > c \cdot \delta_n] d\Pi(\theta) \\ &\lesssim (v_n \delta_n)^{-\kappa} \text{ under case (i) of [A1]} \\ &\lesssim \exp(-c v_n^T \delta_n^T) \text{ under case (ii) of [A1]} \end{aligned}$$

Recalling that $D_n \gtrsim \delta^D$ and using the above, the result follows similarly to Theorem 1. \square

B Additional Computations for Example 1

In this section, we consider several additional aspects of the simulation exercise for Example 1 in the main text. For completeness, we now recall the general features of the example. The assumed DGP is z_1, \dots, z_n iid as $\mathcal{N}(\theta, 1)$ but the actual DGP is y_1, \dots, y_n iid as $\mathcal{N}(\theta, \sigma^2)$. We conduct ABC using the following summary statistics:

- the sample mean $\eta_1(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i$,
- the sample variance $\eta_2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n (y_i - \eta_1(\mathbf{y}))^2$.

Our prior beliefs are given by $\theta \sim \mathcal{N}(0, 25)$. For accept/reject ABC (ABC-AR), we use $N = 25,000$ simulated pseudo datasets generated iid according to $z_i^j \sim \mathcal{N}(\theta^j, 1)$. For both ABC-AR and the local linear regression adjustment (ABC-Reg), we set the tolerance ϵ to be the 1% quantile of the simulated distances $\|\eta(\mathbf{y}) - \eta(\mathbf{z}^j)\|$.

B.1 Example 1: ABC-AR and ABC-Reg Posterior Comparison

In this section, we use the exact same data generated for the experiments conducted in Example 1 in Section 1 of the main text, however, we now analyze the posteriors for ABC-AR and ABC-Reg. For clarity, we briefly recall the details of the Monte Carlo design used in this example. A sequence of “observed” datasets for \mathbf{y} are generated, each corresponding to a different value of σ^2 . To isolate the impact of model misspecification, which occurs if $\sigma^2 \neq 1$, across the experiments, the same set of random numbers are used to generate the observed data. The sample size across the experiments is taken to be $n = 100$. To further isolate the impact of misspecification, we implement ABC using the same pseudo-data across all experiments.

The example in Section 1 clearly demonstrated that the ABC-AR and ABC-Reg posterior means behave very differently depending on the level of model misspecification. Unsurprisingly, the ABC-AR and ABC-Reg posteriors themselves displaying very different patterns of behavior under model misspecification.

In Figure 6 we plot the posterior densities obtained from ABC-AR (Panel A) and ABC-Reg (Panel B) across a subset of the different datasets used in the experiment that correspond to $\sigma^2 \in \{1, 1.05, 1.10, \dots, 5\}$.

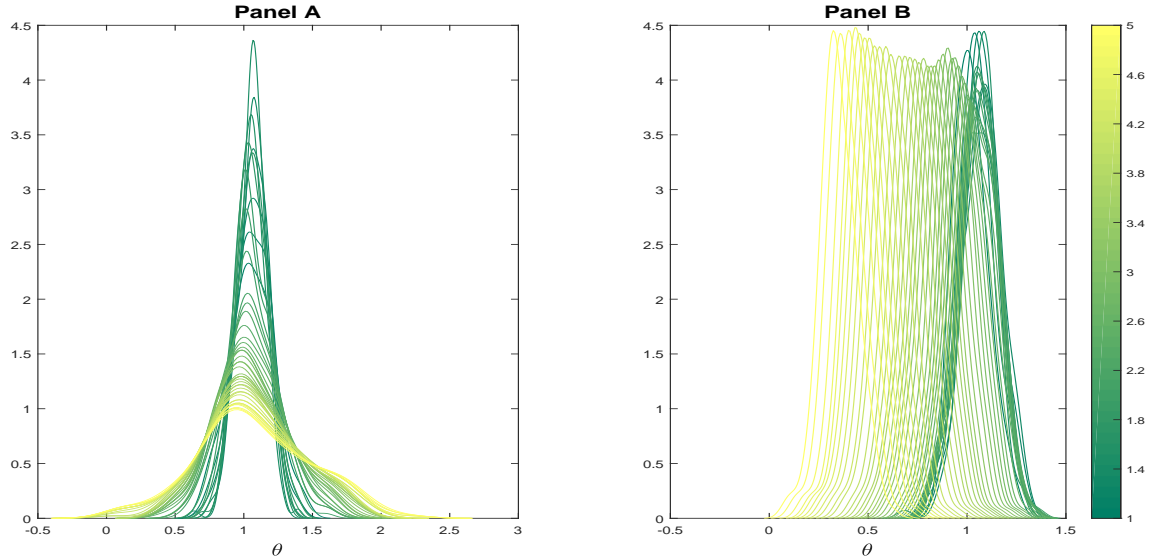


Figure 6: Comparison of posterior densities for ABC-AR, and ABC-Reg across varying levels of model misspecification. The true value generating the data is $\theta = 1$. The posteriors are colored so that darker colors represent less model misspecification (σ^2 closer to unity), and lighter colors representing larger degrees of model misspecification (larger values of σ^2).

Analyzing Panel A in Figure 6, we see that the ABC-AR posteriors remain roughly centered near $\theta = 1$, regardless of the value of σ^2 . However, as σ^2 increases the ABC-AR posteriors noticeably flatten to accommodate the additional variability in the observed data. In contrast to ABC-AR, the mean of the ABC-Reg posterior (Panel B of Figure 6) shifts significantly as σ^2 increases and the variability of the ABC-Reg posterior remains roughly constant across all levels of misspecification. Therefore, ABC-Reg completely neglects the fact that the variance of the observed data is increasing as σ^2 increases.

This finding is further visually confirmed by the results in Figure 7, which plots the corresponding 95% credible intervals (HPD intervals) for ABC-AR and ABC-Reg across the different values of σ^2 used in the experiments. The results in Figure 7 demonstrate that as σ^2 increases the credible regions for ABC-AR widen to accommodate the increased variance in the observed data, whereas the credible intervals for ABC-Reg maintains roughly the same length as σ^2 increases.

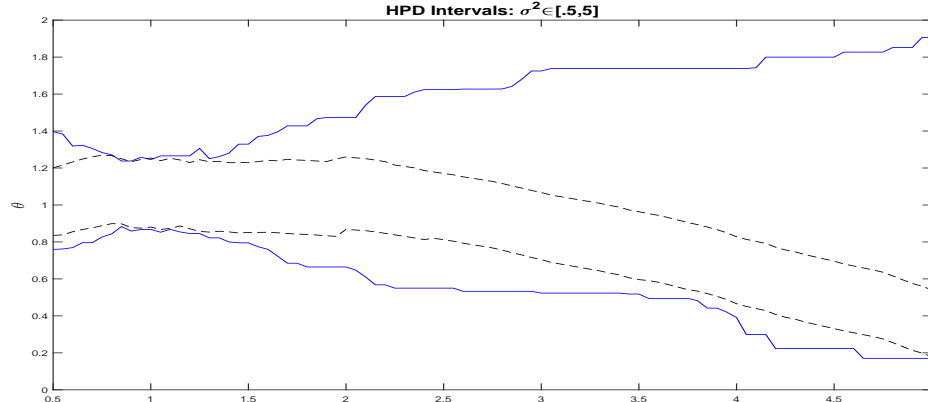


Figure 7: Comparison of credible intervals (HPD² intervals) for ABC-AR, and ABC-Reg across varying levels of model misspecification. The true value generating the data is $\theta = 1$. The solid lines represent the HPD intervals for ABC-AR, while the dashed lines are the HPD intervals for ABC-Reg.

B.2 Example 1: Alternative Regression Adjustments

In this section, we further explore the behavior of local regression adjustment approaches under misspecification. In particular, in the confines of Example 1, we will compare the repeated sampling behavior of ABC-AR and three different regression adjustment approaches: the local linear regression adjustment (ABC-Reg), the proposed local linear adjustment approach described in Section 3.2 (ABC-RegN), and the local nonlinear regression adjustment of Blum and François (2010) (ABC-NN). For each regression adjustment approach, we consider two versions: one version with the heteroskedasticity adjustment of Blum and François (2010) and one version without. All told, we will compare six different regression adjustment approaches across various levels of model misspecification. All adjustment methods are carried out using the R package `abc` (Csilléry et al., 2012).

We follow the Monte Carlo design considered in Section 3.2 of the main paper: for each Monte Carlo replication we simulate observed data $y_i \sim \mathcal{N}(1, \sigma^2)$, iid, and consider three different values of σ^2 corresponding to $\sigma^2 \in \{1, 2, 3\}$. For each value of σ^2 we generate 1,000 artificial observed data sets each of length $n = 100$. Every ABC procedure relies on $N = 25,000$ pseudo-data sets generated according to $z_i \sim \mathcal{N}(\theta, 1)$, iid, and for each method the tolerance is chosen to the 1% quantile of the simulated distances $\|\eta(\mathbf{y}) - \eta(\mathbf{z})\|$. As in the main text, the local regression adjustment procedures use the Epanechnikov kernel.

In Figure 8, we plot the posterior means of the different procedures, across the values of σ^2 and across the Monte Carlo replications, without the heteroskedasticity correction, and the results in Figure 9 consider the case with the heteroskedasticity correction. While the results of Figure 8 have already been presented in the main text, we reproduce them here to simplify the comparison.

The results of this experiment demonstrate that the point estimators obtained from the ABC regression adjustment procedures can display larger variability, in this repeated sampling context, than those obtained from ABC-AR. Moreover, there seems to be no real benefit from the local nonlinear regression adjustment. In addition, the heteroskedasticity correction does not significantly alter the behavior of any local regression adjustment approach. Indeed, the two sets of

results are visually very similar.




Figure 8: Posterior mean comparison of ABC-AR (AR), local linear regression adjustment (Reg), the proposed local linear regression adjustment (RegN), and the local nonlinear regression adjustment (NN) across three values of σ^2 : $\sigma^2 \in \{1, 2, 3\}$. Recall that $\sigma^2 = 1$ corresponds to correct model specification. All plots are presented without outliers.




Figure 9: Posterior mean comparison of ABC-AR (AR), local linear regression adjustment with heteroskedasticity correction (RegC), the proposed local linear regression adjustment with heteroskedasticity correction (RegNC), and the local nonlinear regression adjustment with heteroskedasticity correction (NNC) across three separate values of σ^2 : $\sigma^2 \in \{1, 2, 3\}$. Recall that $\sigma^2 = 1$ corresponds to correct model specification. All results are presented without outliers.

Across the Monte Carlo replications and for each ABC procedure, in Table 2 we present the median posterior standard deviation, the Monte Carlo coverage, the median length of the 95% posterior credible interval, as well as the 2.5% and 97.5% posterior quantiles (represented as the median value over the replications). In this experiment, we compare medians instead of means due to the fact that the heteroskedasticity corrected regression adjustment approaches returned several large outliers that could otherwise skew the comparison.

The results in Table 2 demonstrate that, in terms of posterior variability, as measured by posterior standard deviation, ABC-AR displays the largest variability across all designs. The local linear regression adjustments without heteroskedasticity correction (ABC-Reg, ABC-RegN) have posterior standard deviations that are virtually unchanged across the designs, while the standard deviation of the heteroskedasticity corrected local linear approach (ABC-RegC) is actually decreasing as σ^2 increases, i.e., as the variability in the data increase. In contrast, the posterior standard deviation for the nonlinear regression adjustment without heteroskedasticity correction (ABC-NN) increases as σ^2 increases, however, the posterior standard deviation of the corrected version (ABC-NNC) is stable across the designs.

As discussed in the main text, the fact that the local regression adjustment has small posterior variability, relative to ABC-AR, results in small credible sets and a false sense of precision. As a direct consequence, all the adjustment procedures, with and without the heteroskedasticity correction, have poor Monte Carlo coverage when $\sigma^2 = 2, 3$.

Overall, these results suggest that, at least in this context, there is no meaningful difference between the results obtained by local linear or local nonlinear regression adjustments. In addition, the heteroskedasticity correction does not improve the behavior of the regression adjustment, and can potentially exacerbate the coverage issues observed in these procedure (see, e.g., the results of ABC-RegC in Table 2). Lastly, these results suggest that our proposed local linear regression adjustment (ABC-RegN) performs well relative to the other local regression adjustments (with and without heteroskedasticity correction).

Table 2: Monte Carlo coverage (Cov), credible set length (Len), and posterior standard deviation (Std) for the normal example under various levels of model misspecification. Cov is the percentage of times that the 95% credible set contained $\theta = 1$. Len is the median length of the 95% credible set, across the Monte Carlo trials. Std is the median posterior standard deviation across the Monte Carlo trials. Q025 represents the median (over the replications) of the 2.5% posterior quantiles across the different methods, while Q975 represents the corresponding 97.5% quantile.

$\sigma^2 = 1$	Cov	Len	Std	Q025	Q975
ABC-AR	0.982	0.461	0.120	0.770	1.234
ABC-RegN	0.938	0.385	0.100	0.802	1.190
ABC-Reg	0.941	0.382	0.100	0.811	1.192
ABC-NN	0.950	0.385	0.100	0.811	1.194
ABC-RegNC	0.938	0.389	0.100	0.802	1.193
ABC-RegC	0.945	0.386	0.100	0.809	1.197
ABC-NNC	0.941	0.384	0.101	0.808	1.200
$\sigma^2 = 2$	Cov	Len	Std	Q025	Q975
ABC-AR	0.961	0.614	0.158	0.698	1.310
ABC-RegN	0.802	0.383	0.100	0.809	1.194
ABC-Reg	0.717	0.382	0.100	0.818	1.200
ABC-NN	0.729	0.428	0.113	0.766	1.204
ABC-RegNC	0.809	0.387	0.100	0.805	1.194
ABC-RegC	0.645	0.361	0.094	0.791	1.223
ABC-NNC	0.636	0.383	0.100	0.773	1.201
$\sigma^2 = 3$	Cov	Len	Std	Q025	Q975
ABC-AR	0.913	0.613	0.157	0.694	1.311
ABC-RegN	0.707	0.383	0.099	0.816	1.195
ABC-Reg	0.460	0.381	0.100	0.791	1.180
ABC-NN	0.419	0.479	0.128	0.632	1.137
ABC-RegNC	0.708	0.386	0.100	0.811	1.199
ABC-RegC	0.462	0.261	0.069	0.700	1.249
ABC-NNC	0.409	0.383	0.101	0.624	1.171

C Additional Example: Misspecified g -and- k Model

To further demonstrate the behavior of different ABC approaches under model misspecification, we consider an additional example based on the g -and- k distribution, which is an oft-used example in the ABC literature to compare the behavior of different ABC approaches (see, e.g., Drovandi and Pettitt, 2011, Fearnhead and Prangle, 2012, and Bernton et al., 2019). The g -and- k model is most commonly stated through its quantile function:

$$q \in (0, 1) \mapsto a + b \left(1 + 0.8 \frac{1 - \exp(-gz(q))}{1 + \exp(-gz(q))} \right) (1 + z(q)^2)^k z(q),$$

where $z(q)$ refers to the q -th quantile of the standard normal distribution. The four parameters of the g -and- k distributions have specific interpretations. The parameter a represents the location, b the scale, while g and k control the skewness and kurtosis, respectively. Following Drovandi and Pettitt (2011), we consider the following priors on the parameters

$$a \sim \mathcal{U}[0, 10], \quad b \sim \mathcal{U}[0, 10], \quad g \sim \mathcal{U}[0, 10] \quad b \sim \mathcal{U}[0, 10],$$

where $\mathcal{U}[0, 1]$ denotes the uniform distribution on $[0, 1]$.

ABC-based inference in the g -and- k model is usually conducted using the quantiles of the simulated and observed data. Therefore, in what follows we take as our summary statistics for ABC the octiles of the data:

- $\eta_j(\mathbf{y}) = O_j(\mathbf{y})$, for $1, \dots, 7$, where (O_1, \dots, O_7) partitions the data into eight equal parts.

The g -and- k distribution is a highly-flexible class of distributions that is capable of modeling data with a complex unconditional distribution. While highly flexible, g -and- k distributions are unimodal and are not capable of capturing multi-modality that may exist in the data.

In this section, we compare the behavior of different ABC-based procedures when the underlying assumption is that the data comes from the g -and- k distribution, but when the observed data is actually generated from a distribution with minor bi-modality. In particular, we generate observed data iid from the Gaussian mixture

$$y_i \sim w \cdot \mathcal{N}(\mu_1, \sigma_1^2) + (1 - w) \cdot \mathcal{N}(\mu_2, \sigma_2^2). \quad (20)$$

In what follows, we fix the parameters in equation (20) to be

$$(\mu_1, \sigma_1^2)^\top = (1, 2)^\top, \quad (\mu_2, \sigma_2^2)^\top = (7, 2)^\top, \quad \text{and } w = 0.9.$$

This DGP produces observed data that exhibits positive skewness and excess kurtosis, owing to the fact that the density exhibits a minor ‘‘hump’’ in the right tail of the data. This specification was chosen particularly to generate minor bi-modality in the observed data, which is a feature that the g -and- k model can not capture. For illustration, the kernel density of a representative data set of size $n = 100$ simulated from this mixture model is given in Figure 10.




Figure 10: Kernel density of data simulated from the DGP in equation (20).

C.1 Monte Carlo Experiments

Similar to Example 1 in the main paper, we now compare the behavior of ABC-AR and various local ABC regression adjustments. In particular, we consider the following local regression adjustments: ABC-Reg (the standard weighted local-linear adjustment, Beaumont et al., 2002), ABC-RegN (our weighted local-linear adjustment) and ABC-NN (a local nonlinear regression adjustment using neural nets, Blum and François, 2010). From the DGP in equation (20), we simulate $n = 100$ observations for the observed data \mathbf{y} . Just as in Example 1 in the main text, each ABC procedure is based on $N = 25,000$ simulated draws, where the tolerance ϵ_n is chosen to be to the 1% quantile of the overall simulated distances. For ABC-AR we use the Euclidean norm. Figure 11 plots the posteriors from the different ABC methods across the four parameters of the g -and- k distribution for a representative experiment.




Figure 11: Posterior comparison of various ABC procedures in the g -and- k model when the true DGP is given by equation (20).

Analyzing Figure 11, three features are immediately in evidence. Firstly, for most of the ABC procedures, and across virtually all the parameters, the posteriors are decidedly non-Gaussian (as suggested by the results of Theorem 2 in the main text). Secondly, across the parameters, the different ABC procedures produce posteriors with very different behavior, which mirrors the results observed in Example 1. As mentioned earlier, this behavior is a direct consequence of the local regression adjustments transformation of the accepted ABC-AR draws. Thirdly, for the g parameter the local regression posteriors place significant mass outside of the original prior space, which was $\mathcal{U}[0, 10]$. Similarly, for the parameter k these procedures yield some negative values for k .

This latter feature is troubling given the specific nature of the mismatch between the actual

DGP and the assumed g -and- k distribution. In particular, the theoretical moments in the actual DGP imply that population skewness is greater than 1.5 and that excess kurtosis is nearly 3, respectively. Recalling that skewness and kurtosis in the g -and- k distribution are controlled by g and k , respectively, we know that $g > 0$ is associated with positive skewness, and $k < 0$ is associated with kurtosis that is less than that of the normal distribution. Given this, the results of Figure 11 demonstrate that the local regression adjustment procedures can place significant posterior mass on $g < 0$, even though the observed data is always positively skewed. Similarly, these procedures can place some positive posterior mass on $k < 0$, even though the observed data exhibits positive excess kurtosis.

These behaviors observed in the local regression adjustments are entirely due to the fact that these approaches disregard the nature of the parameter space when adjusting the accepted draws, and, hence, can transport posterior mass outside the original parameter space. As a consequence, in this example the local regression adjustment approaches can place significant posterior mass on values of g (and to a lesser extent k) that are incompatible with the observed data.

Similar to the conclusions obtained in Example 1, this behavior of the local regression adjustment is not a feature of any particular dataset but is persistent across different datasets. To demonstrate this fact, we simulate 1000 replications from the true DGP in equation (20), and rerun each ABC procedure on these observed datasets. Across the replications, for each of the different ABC procedures, we record the posterior mean, standard deviation, as well as the length of the 95% credible regions for each of the parameters in the g -and- k distribution, which we calculate via the corresponding 2.5 and 97.5 quantiles of the individual parameter posteriors. In Table 3, we present the averages of these results across the replications.

To help interpret the findings in Table 3, before discussing the results, we first calculate the pseudo-true value

$$\theta^* = \arg \min_{\theta \in \Theta} \|b_0 - b(\theta)\|,$$

where we reminded the reader that b_0 is the probability limit of $\eta(\mathbf{y})$, given by the octiles of the observed data, and $b(\theta)$ is the probability limit of $\eta(\mathbf{z})$, which corresponds to the octiles of the simulated data. To obtain this pseudo-true value, we must first calculate b_0 , the population octiles of the Gaussian mixture distribution. The quantile function of the Gaussian mixture has no closed-form but the quantiles, and hence the value of b_0 , can be obtained by numerically inverting the corresponding CDF of the Gaussian mixture. Give the value of b_0 , and the fact that the quantiles of the g -and- k distribution have an analytical form, in terms of the standard normal quantile function, we can numerically solve for the value of $\theta = (a, b, g, k)^\top$ that minimizes $\|b_0 - b(\theta)\|$. Using this approach, the pseudo-true value under this particular Monte Carlo design is given by

$$\theta^* = (a^*, b^*, g^*, k^*)^\top = (1.17, 1.50, 0.41, 0.23)^\top. \tag{21}$$

This value of θ^* corresponds to a g -and- k distribution with positive skewness and kurtosis that is larger than that of the normal distribution. Therefore, as one would hope, the pseudo-true value reflects the actual features of the true DGP, namely, positive skewness and excess kurtosis. Using this pseudo-true value, we also calculate the Monte Carlo coverage of the different ABC procedures, which we also display in Table 3.

Table 3: Posterior summaries for accept/reject ABC (AR), local linear regression adjusted ABC (Reg), see Beaumont et al. (2002), our proposed local linear regression adjusted ABC (RegN), and a local nonlinear regression adjustment (NN) based on neural nets, see Blum and François (2010). Mean is the average posterior mean, Std is the average posterior standard deviation, Len is the average length of the 95% credible set, Cov is the Monte Carlo coverage calculated using the pseudo-true value defined in (21), and Q025 and Q975 refer to the average (across the Monte Carlo replications) 2.5% and 97.5% quantiles of the posterior.

a	AR	RegN	Reg	NN	b	AR	RegN	Reg	NN
Mean	1.1190	1.1568	1.2060	1.1862	Mean	0.8961	1.0143	1.6342	1.6685
Std	0.6564	0.0027	0.0031	0.1333	Std	0.7114	0.0104	0.0113	0.0994
Len	2.2517	0.0112	0.0125	0.4893	Len	2.5380	0.0414	0.0439	0.3914
Cov	1.0000	0.0180	0.0140	0.5110	Cov	0.9990	0.1730	0.5220	0.5600
Q025	0.0831	1.1509	1.1993	0.9419	Q025	0.0478	0.9940	1.6119	1.5326
Q975	2.3349	1.1621	1.2118	1.4312	Q975	2.5858	1.0354	1.6558	1.9240

g	AR	RegN	Reg	NN	k	AR	RegN	Reg	NN
Mean	4.7758	5.3310	0.8639	0.3707	Mean	1.5321	2.1790	0.7973	0.3889
Std	2.9404	1.5265	1.4632	0.7212	Std	1.4145	0.9426	0.7810	0.3866
Len	9.4511	6.1010	5.8798	2.9902	Len	5.1661	3.8949	3.2841	1.5546
Cov	0.8900	0.3380	0.6320	0.6870	Cov	1.0000	0.5760	0.8420	0.8060
Q025	0.2484	2.4728	-1.7755	-0.9119	Q025	0.0510	0.3648	-0.6637	-0.3487
Q975	9.6995	8.5738	4.1043	2.0783	Q975	5.2171	4.2597	2.6204	1.2059

Analyzing the results in Table 3, we see that all procedures give relatively accurate point estimates of the location parameter a . However, the Monte Carlo coverage for this parameter varies drastically across the different ABC procedures, which reflects the extremely small posterior standard deviations of the local regression adjustment approaches. For the scale parameter, b , a similar story is in evidence. Namely, all procedures give point estimators that are not too far from the pseudo-true value, $b^* = 1.5$, however, the small posterior standard deviations for the local adjustment procedures leads to significant under-coverage.

For the parameter g , the different ABC procedures have very different posterior behavior. The most striking feature is that, while the average posterior means for the local linear adjustment (Reg) and the nonlinear adjustment (NN) are not too far from the pseudo-true value, $g^* = 0.41$, both procedures place a significant amount of posterior mass on $g < 0$.⁹ We recall that the true DGP is such that the observed data always has positive skewness.

The results for the parameter k , which governs the kurtosis of the g -and- k distribution, are similar to those obtained for the parameter g . Namely, while the local linear and nonlinear regression adjustments generally give point estimators that are closer to the pseudo-true value of $k^* \approx 0.23$ than the other ABC procedures, both approaches place a non-negligible amount of posterior mass on $k < 0$. This behavior is clearly at odds with the observed data: values of $k < 0$ imply that the observed data exhibits kurtosis that is less than that of the normal.

As demonstrated by Corollary 2 in the main paper, this behavior is a direct consequence of model misspecification and the nature of the local regression adjustment. In essence, under

⁹This point can be seen by analyzing the 2.5% and 97.5% quantiles from these procedures.

model misspecification, these adjustment procedures can be viewed as taking accepted draws, which asymptotically yield the smallest distance between observed and simulated summaries, and perturbing them according to a criterion that does not respect the original optimality of these draws.

Given the behavior of the local regression adjustment in this example, as well as the results of Example 1 in the main text, we suggest researchers treat the output of local regression adjustment ABC procedures with a healthy level of skepticism in settings where model misspecification is a possibility.