



HAL
open science

Approximate Bayesian Computation: a non-parametric perspective

Michael Blum

► **To cite this version:**

Michael Blum. Approximate Bayesian Computation: a non-parametric perspective. 2009. hal-00373301v4

HAL Id: hal-00373301

<https://hal.science/hal-00373301v4>

Preprint submitted on 8 Sep 2009 (v4), last revised 8 Apr 2010 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Approximate Bayesian Computation: a non-parametric perspective

Michael G.B. Blum

CNRS, Laboratoire TIMC-IMAG, Faculté de médecine, 38706 La Tronche

Université Joseph Fourier, Grenoble, France

email: `michael.blum@imag.fr`

Abstract

Approximate Bayesian Computation is a family of likelihood-free inference techniques that are well-suited to models defined in terms of a stochastic generating mechanism. In a nutshell, Approximate Bayesian Computation proceeds by computing summary statistics \mathbf{s}_{obs} from the data and simulating synthetic summary statistics for different values of the parameter Θ . The posterior distribution is then approximated by an estimator of the conditional density $g(\Theta|\mathbf{s}_{obs})$. In this paper, we derive the asymptotic bias and variance of the standard estimators of the posterior distribution which are based on rejection sampling and linear adjustment. Additionally, we introduce an original estimator of the posterior distribution based on quadratic adjustment and we show that its bias contains a smaller number of terms than the estimator with linear adjustment. Although we find that the estimators with adjustment are not universally superior to the estimator based on rejection sampling, we find that they can achieve better performance when there is a nearly homoscedastic relationship between the summary statistics and the parameter of interest. Last, we present model selection in Approximate Bayesian Computation and provide asymptotic properties of two estimators of the model probabilities. As for parameter estimation, the asymptotic results raise the importance of the curse of dimensionality in Approximate Bayesian Computation. Performing numerical simulations in a simple normal model confirms that the estimators may be less efficient as the number of summary statistics increases. Supplemental materials containing the details of the proofs are available online.

KEYWORDS: Conditional density estimation, implicit statistical models, kernel regression, curse of dimensionality, local polynomials

1 Introduction

Inference in Bayesian statistics relies on the *full posterior distribution* defined as

$$g(\Theta|D) = \frac{p(D|\Theta)\pi(\Theta)}{p(D)} \quad (1)$$

where Θ denotes the vector of parameters and D denotes the observed data. The expression given in (1) depends on the *prior distribution* $\pi(\Theta)$, the *likelihood function* $p(D|\Theta)$ and the marginal probability of the data $p(D) = \int_{\Theta} p(D|\Theta)\pi(\Theta) d\Theta$. However, when the statistical model is defined in term of a stochastic generating mechanism, the likelihood can be intractable. Methods of inference in the context of these so-called *implicit statistical models* have been proposed by Diggle & Gratton (1984) in a frequentist setting. Implicit statistical models can be thought of as a computer generating mechanism that mimics data generation. In the past ten years, interests in implicit statistical models have reappeared in population genetics where Beaumont et al. (2002) gave the name of Approximate Bayesian Computation (ABC) to a family of likelihood-free inference methods.

Since its original developments in population genetics (Fu & Li 1997, Tavaré et al. 1997, Pritchard et al. 1999, Beaumont et al. 2002), ABC has successfully been applied in a large range of scientific fields such as archaeological science (Wilkinson & Tavaré 2009), ecology (François et al. 2008, Jabot & Chave 2009), epidemiology (Tanaka et al. 2006, Blum & Tran 2009), stereology (Bortot et al. 2007) or in the context of protein networks (Ratmann et al. 2007). Despite the increasing number of ABC applications, theoretical results concerning its properties are still lacking and the present paper contributes to filling this gap.

In ABC, inference is no more based on the full posterior distribution $g(\Theta|D)$ but on the *partial* posterior distribution $g(\Theta|\mathbf{s}_{obs})$ where \mathbf{s}_{obs} denotes a vector of d -dimensional summary statistics computed from the data D . The partial posterior distribution (Doksum and Lo 1990) is defined as

$$g(\Theta|\mathbf{s}_{obs}) = \frac{p(\mathbf{s}_{obs}|\Theta)\pi(\Theta)}{p(\mathbf{s}_{obs})}. \quad (2)$$

Of course, the partial and the full posterior distributions are the same if the summary statistics are sufficient with respect to the parameter Θ .

To generate a sample from the partial posterior distribution $g(\Theta|\mathbf{s}_{obs})$, ABC proceeds by simulating n values $\Theta_i, i = 1, \dots, n$ from the prior distribu-

tion π , and then simulating summary statistics \mathbf{s}_i according to $p(\mathbf{s}|\Theta_i)$. Once the couples (Θ_i, \mathbf{s}_i) , $i = 1 \dots n$, have been obtained, the estimation of the partial posterior distribution is a problem of conditional density estimation. Here we will derive the asymptotic bias and variance of a Nadaraya-Watson type estimator (Nadaraya 1964, Watson 1964), of an estimator with linear adjustment proposed by Beaumont et al. (2002) in the context of ABC, and of an original estimator with quadratic adjustment that we propose.

Although replacing the full posterior by the partial one is an approximation crucial in ABC, we will not investigate its consequences here. The reader is referred to Cam (1964), and Abril (1994) for theoretical works on the concept of approximate sufficiency; and to Joyce & Marjoram (2008) for a practical method that selects informative summary statistics in ABC. Here, we concentrate on the second type of approximation arising from the discrepancy between the estimated partial posterior distribution and the true partial posterior distribution.

In addition to parameter inference, Bayesian model selection can also be handled within ABC. For sake of simplicity, we assume here that there are two competitive models \mathcal{M}_1 and \mathcal{M}_2 that are *a priori* equally likely. The extent to which the data support \mathcal{M}_1 over \mathcal{M}_2 is measured by the partial Bayes factor defined as

$$\text{BF} = \frac{p_1(\mathbf{s}_{obs})}{p_2(\mathbf{s}_{obs})},$$

in which $p_1(\mathbf{s}_{obs})$ and $p_2(\mathbf{s}_{obs})$ denote the marginal probability of the summary statistics in each model. A related criteria for model selection is the posterior (partial) probability that \mathcal{M}_1 is the correct model given either \mathcal{M}_1 or \mathcal{M}_2

$$p(\mathcal{M}_1|\mathbf{s}_{obs}) = \frac{p_1(\mathbf{s}_{obs})}{p_1(\mathbf{s}_{obs}) + p_2(\mathbf{s}_{obs})}. \quad (3)$$

As for parameter inference, we do not study here the error arising from the difference between the partial posterior model probability and the *full* one $p(\mathcal{M}_1|D) = p_1(D)/(p_1(D) + p_2(D))$ but we focus on the error arising from the estimation of $p(\mathcal{M}_1|\mathbf{s}_{obs})$.

In this paper, we investigate the asymptotic bias and variance of the estimators of the posterior distribution and the model probabilities. Parameter inference within ABC is introduced in Section 2 and model selection is introduced in Section 3. Section 4 presents the main theorem concerning the asymptotic bias and variance of the estimators of the partial posterior distribution. It ends with the asymptotic theory for the estimation of the model

probability $p(\mathcal{M}_1|\mathbf{s}_{obs})$. In Section 5, we perform a numerical study in which the properties of the ABC estimators are investigated in a simple normal model. The proofs of the theorems are given in the Appendix.

2 Parameter inference

2.1 Smooth rejection

In the context of ABC, the Nadaraya-Watson estimator of the partial posterior mean $E[\Theta|\mathbf{s}_{obs}]$ can be written as

$$m_0 = \frac{\sum_{i=1}^n \Theta_i K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs})} \quad (4)$$

where $K_B(\mathbf{u}) = |B|^{-1}K(B^{-1}\mathbf{u})$, B is the *bandwidth matrix* that is assumed to be non-singular, K is a d-variate kernel such that $\int K(\mathbf{u}) d\mathbf{u} = 1$, and $|B|$ denotes the determinant of B . To estimate the partial posterior distribution $g(\theta|\mathbf{s}_{obs})$ ($\theta \in \mathbb{R}$) of a one-dimensional coordinate of Θ , we introduce a kernel \tilde{K} that is a symmetric density function on \mathbb{R} . Here we will restrict our analysis to univariate density estimation but multivariate density estimation can also be implemented in the same vein. The bandwidth corresponding to \tilde{K} is denoted b' ($b' > 0$) and we use the notation $\tilde{K}_{b'}(\cdot) = \tilde{K}(\cdot/b')/b'$. As the bandwidth b' goes to 0, a simple Taylor expansion shows that

$$E_{\theta'}[\tilde{K}_{b'}(\theta' - \theta)|\mathbf{s}_{obs}] \approx g(\theta|\mathbf{s}_{obs}).$$

The estimation of the partial posterior distribution $g(\theta|\mathbf{s}_{obs})$ can thus be viewed as a problem of nonparametric regression. After substituting Θ_i by $\tilde{K}_{b'}(\theta_i - \theta)$ in equation (4), we obtain the following estimator of $g(\theta|\mathbf{s}_{obs})$ (Rosenblatt 1969)

$$\hat{g}_0(\theta|\mathbf{s}_{obs}) = \frac{\sum_{i=1}^n \tilde{K}_{b'}(\theta_i - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs})}. \quad (5)$$

The initial rejection-based ABC estimator consisted of using a kernel K that took 0 or 1 values (Pritchard et al. 1999). Estimation with smooth kernels K was proposed by Beaumont et al. (2002).

2.2 Regression adjustment

Besides introducing smoothing in the ABC algorithm, Beaumont et al. (2002) proposed additionally to adjust the θ_i 's to weaken the effect of the discrepancy between \mathbf{s}_i and \mathbf{s}_{obs} . In the neighborhood of \mathbf{s}_{obs} , they proposed to approximate the conditional expectation of θ given \mathbf{s}_{obs} by \hat{m}_1 where

$$\hat{m}_1(\mathbf{s}) = \hat{\alpha} + (\mathbf{s} - \mathbf{s}_{obs})^t \hat{\beta} \text{ for } \mathbf{s} \text{ such that } K_B(\mathbf{s} - \mathbf{s}_{obs}) > 0. \quad (6)$$

The estimates $\hat{\alpha}$ ($\alpha \in \mathbb{R}$) and $\hat{\beta}$ ($\beta \in \mathbb{R}^d$) are found by minimizing the weighted least squares criterion

$$\sum_{i=1}^n \{\theta_i - (\alpha + (\mathbf{s}_i - \mathbf{s}_{obs})^t \beta)\} K_B(\mathbf{s}_i - \mathbf{s}_{obs}). \quad (7)$$

The solution to (7) is given by (Ruppert & Wand 1994, Härdle et al. 2004)

$$(\hat{\alpha}, \hat{\beta}) = (X^t W X)^{-1} X^t W \theta, \quad (8)$$

where W is a diagonal matrix whose i^{th} element is $K_B(\mathbf{s}_i - \mathbf{s}_{obs})$, and

$$X = \begin{pmatrix} 1 & s_1^1 - s_{obs}^1 & \cdots & s_1^d - s_{obs}^d \\ \vdots & \cdots & \ddots & \vdots \\ 1 & s_n^1 - s_{obs}^1 & \cdots & s_n^d - s_{obs}^d \end{pmatrix}, \quad \theta = \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_n \end{pmatrix}.$$

The principle of regression adjustment consists of forming the empirical residuals $\epsilon_i = \theta_i - \hat{m}_1(\mathbf{s}_i)$, and to adjust the θ_i by computing

$$\theta_i^* = \hat{m}_1(\mathbf{s}_{obs}) + \epsilon_i, \quad i = 1, \dots, n. \quad (9)$$

Estimation of $g(\theta|\mathbf{s}_{obs})$ is obtained with the estimator of equation (5) after replacing the θ_i 's by the θ_i^* 's. This leads to the estimator proposed by Beaumont et al. (2002, eq. (9))

$$\hat{g}_1(\theta|\mathbf{s}_{obs}) = \frac{\sum_{i=1}^n \tilde{K}_{b'}(\theta_i^* - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs})}. \quad (10)$$

To improve the estimation of the conditional mean, we suggest a slight modification to $\hat{g}_1(\theta|\mathbf{s}_{obs})$ using a quadratic rather than a linear adjustment. The conditional expectation of θ given \mathbf{s} is now approximated by \hat{m}_2 where

$$\hat{m}_2(\mathbf{s}) = \tilde{\alpha} + (\mathbf{s} - \mathbf{s}_{obs})^t \tilde{\beta} + \frac{1}{2} (\mathbf{s} - \mathbf{s}_{obs})^t \tilde{\gamma} (\mathbf{s} - \mathbf{s}_{obs}) \text{ for } \mathbf{s} \text{ such that } K_B(\mathbf{s} - \mathbf{s}_{obs}) > 0. \quad (11)$$

The three parameters $(\alpha, \beta, \gamma) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d^2}$ are found by minimizing the quadratic extension of the least square criterion given in (7). Because γ is a symmetric matrix, the inference of γ only requires the lower triangular part and the diagonal of the matrix to be estimated. The solution to this new minimization problem is given by (8) where the design matrix X is now equal to

$$X = \begin{pmatrix} 1 & s_1^1 - s_{obs}^1 & \cdots & s_1^d - s_{obs}^d & \frac{(s_1^1 - s_{obs}^1)^2}{2} & (s_1^1 - s_{obs}^1)(s_1^2 - s_{obs}^2) & \cdots & \frac{(s_1^d - s_{obs}^d)^2}{2} \\ \vdots & \cdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & s_n^1 - s_{obs}^1 & \cdots & s_n^d - s_{obs}^d & \frac{(s_n^1 - s_{obs}^1)^2}{2} & (s_n^1 - s_{obs}^1)(s_n^2 - s_{obs}^2) & \cdots & \frac{(s_n^d - s_{obs}^d)^2}{2} \end{pmatrix},$$

Letting $\theta_i^{**} = \hat{m}_2(\mathbf{s}_{obs}) + (\theta_i - \hat{m}_2(\mathbf{s}_i))$, the new estimator of the partial posterior distribution is given by

$$\hat{g}_2(\theta | \mathbf{s}_{obs}) = \frac{\sum_{i=1}^n \tilde{K}_{b'}(\theta_i^{**} - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs})}. \quad (12)$$

Estimators with regression adjustment in the same vein as those proposed in equations (10) and (12) have already been proposed by Hyndman et al. (1996) and Hansen (2004) for performing conditional density estimation when $d = 1$.

3 Model selection

3.1 Smooth rejection

We assume here that either $n/2$ simulations have been performed in each model \mathcal{M}_1 and \mathcal{M}_2 or that the two generative models were chosen with equal probability for each simulation. We denote by Y_i , $i = 1, \dots, n$ an indicator variable equal to 1 if the i^{th} simulation was performed using the generative model of \mathcal{M}_1 and 0 otherwise. An estimator of $p(\mathcal{M}_1 | \mathbf{s}_{obs})$ is obtained using the following Nadaraya-Watson estimator

$$\hat{p}_0(\mathcal{M}_1 | \mathbf{s}_{obs}) = \frac{\sum_{i=1}^n Y_i K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs})}. \quad (13)$$

Using equation (13) to compute the partial Bayes factor, we get

$$\hat{\text{BF}}_0 = \frac{\sum_{i=1}^n Y_i K_B(\mathbf{s}_i - \mathbf{s}_{obs})}{\sum_{i=1}^n (1 - Y_i) K_B(\mathbf{s}_i - \mathbf{s}_{obs})}. \quad (14)$$

When K takes 0 or 1 values, the partial Bayes factor is simply estimated as the ratio of the acceptance rates in each model (Pritchard et al. 1999).

3.2 Local logistic regression

An alternative method has been proposed by Beaumont (2008) to estimate the model probabilities in ABC. Viewing the estimation problem in a regression setting in which \mathbf{s} is the predictive variable and Y is the indicator variable to predict, Beaumont (2008) proposed to use local logistic regression to estimate $E[Y|\mathbf{s}_{obs}] = p(\mathcal{M}_1|\mathbf{s}_{obs})$. In local logistic regression, the log-odds of the model probabilities are approximated by a linear function so that $\log[p(\mathcal{M}_1|\mathbf{s})/(1 - p(\mathcal{M}_1|\mathbf{s}))] = \delta_0 + (\mathbf{s} - \mathbf{s}_{obs})^t \delta_1$. The log of the weighted likelihood \mathcal{L} can be written as

$$\mathcal{L}(\delta_0, \delta_1) = \sum_{i=1}^n \{Y_i \log(g_{\delta_0, \delta_1}(\mathbf{s}_i)) + (1 - Y_i) \log(1 - g_{\delta_0, \delta_1}(\mathbf{s}_i))\} K_B(\mathbf{s}_i - \mathbf{s}_{obs}), \quad (15)$$

where $g_{\delta_0, \delta_1}(\mathbf{s}) = e^{\delta_0 + (\mathbf{s} - \mathbf{s}_{obs})^t \delta_1} / (1 + e^{\delta_0 + (\mathbf{s} - \mathbf{s}_{obs})^t \delta_1})$. Denoting by $\hat{\delta}_0$ and $\hat{\delta}_1$, the parameters that maximize the log-likelihood given in equation (15), the probability of model \mathcal{M}_1 is estimated as

$$\hat{p}_1(\mathcal{M}_1|\mathbf{s}_{obs}) = \frac{e^{\hat{\delta}_0}}{1 + e^{\hat{\delta}_0}}. \quad (16)$$

The optimization of equation (15) has no explicit solution and iterative algorithms such as iteratively reweighted least squares shall be considered (McCullagh & Nelder 1990).

4 Asymptotic bias and variance in ABC

4.1 Parameter inference

To study the asymptotic bias and variance of the three estimators of the partial posterior distribution $\hat{g}_j(\cdot|\mathbf{s}_{obs})$, $j = 0, 1, 2$, we assume that the bandwidth matrix is diagonal $B = bD$. A more general result for non-singular

matrix B is given in the Appendix. In practice, the bandwidth matrix B may depend on the simulations, but we will assume in this Section that it has been fixed independently of the simulations. This assumption facilitates the computations and is classical when investigating the asymptotic bias and variance of non-parametric estimators (Ruppert & Wand 1994).

The first (resp. second) derivative of a function f with respect the variable x is denoted f_x (resp. f_{xx}). When the derivative is taken with respect to a vector \mathbf{x} , $f_{\mathbf{x}}$ denotes the gradient of f and $f_{\mathbf{xx}}$ denotes the Hessian of f . The variance-covariance matrix of K is assumed to be diagonal and equal to $\mu_2(K)I_d$. We additionally introduce the following notations $\mu_2(\tilde{K}) = \int_u u^2 \tilde{K}(u) du$, $R(K) = \int_{\mathbf{u}} K^2(\mathbf{u}) d\mathbf{u}$, and $R(\tilde{K}) = \int_u \tilde{K}^2(u) du$.

Theorem 1 *Assume that $B = bD$ and assume that conditions (A1):(A5) of the Appendix hold. The bias and the variance of the estimators $\hat{g}_j(\cdot|\mathbf{s}_{obs})$, $j = 0, 1, 2$, are given by*

$$E[\hat{g}_j(\theta|\mathbf{s}_{obs}) - g(\theta|\mathbf{s}_{obs})] = C_1 b'^2 + C_{2,j} b^2 + O_P((b^2 + b'^2)^2) + O_P\left(\frac{1}{n|B|}\right), \quad j = 0, 1, 2, \quad (17)$$

$$\text{Var}[\hat{g}_j(\theta|\mathbf{s}_{obs})] = \frac{C_3}{nb^d b'} (1 + o_P(1)), \quad j = 0, 1, 2, \quad (18)$$

with

$$C_1 = \frac{\mu_2(\tilde{K}) g_{\theta\theta}(\theta|\mathbf{s}_{obs})}{2},$$

$$C_{2,0} = \mu_2(K) \left(\frac{g_{\mathbf{s}}(\theta|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} D^2 p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(D^2 g_{\mathbf{ss}}(\theta|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} \right), \quad (19)$$

$$C_{2,1} = \mu_2(K) \left(\frac{h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} D^2 p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(D^2 h_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} - \frac{h_{\epsilon}(\epsilon|\mathbf{s}_{obs}) \text{tr}(D^2 m_{\mathbf{ss}}(\mathbf{s}_{obs}))}{2} \right), \quad (20)$$

$$C_{2,2} = \mu_2(K) \left(\frac{h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} D^2 p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(D^2 h_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} \right), \quad (21)$$

and

$$C_3 = \frac{R(K) R(\tilde{K}) g(\theta|\mathbf{s}_{obs})}{|D| p(\mathbf{s}_{obs})}. \quad (22)$$

Remark 1. Curse of dimensionality The mean square error (MSE) of an estimator is equal to the sum of its squared bias and its variance. With standard algebra, we find that the MSEs are minimized when both b and b' are of the order of $n^{-1/(d+5)}$. This implies that the minimal MSEs are of the order of $n^{-4/(d+5)}$. Thus, the rate at which the minimal MSEs converge to 0 decreases importantly as the dimension d of \mathbf{s}_{obs} increases. This phenomenon known as the *curse of dimensionality* is a particular acute issue for the three estimators of $g(\theta|\mathbf{s}_{obs})$.

Remark 2. Effective local size and effect of design As shown by equations (18) and (22), the variance of the estimators can be expressed, up to a constant, as $\frac{1}{\tilde{n}} \frac{g(\theta|\mathbf{s}_{obs})}{b'}$, where the effective local size is $\tilde{n} = n|D|p(\mathbf{s}_{obs})b^d$. The effective local size is an approximation of the expected number of simulations that fall within the ellipsoid of radii equal to the diagonal elements of D times b . Thus equations (18) and (22) reflect that the variance is penalized by sparser simulations around \mathbf{s}_{obs} (Ruppert & Wand 1994). Sequential Monte Carlo samplers (Sisson et al. 2007, Beaumont et al. 2009, Toni et al. 2009) precisely aim at adapting the sampling distribution of the parameters, a.k.a the design, to increase the probability of targeting close to \mathbf{s}_{obs} .

Remark 3. A closer look at the bias There are two terms in the bias of $\hat{g}_0(\cdot|\mathbf{s}_{obs})$ (equation (19)) that are related to the smoothing in the space of the summary statistics. The first term in equation (19) corresponds to the effect of the design and is large when the gradient of $Dp(\cdot)$ is collinear to the gradient of $Dg(\theta|\cdot)$. This term reflects that, in the neighborhood of \mathbf{s}_{obs} , there will be an excess of points in the direction of $Dp_{\mathbf{s}}(\mathbf{s}_{obs})$. Up to a constant, the second term in equation (19) is proportional to $\text{tr}(D^2 g_{\text{ss}}(\theta|\mathbf{s})|_{\mathbf{s}=\mathbf{s}_{obs}})$ which is simply the sum of the elementwise product of D and the Hessian $g_{\text{ss}}(\theta|\mathbf{s})|_{\mathbf{s}=\mathbf{s}_{obs}}$.

For the estimator $\hat{g}_2(\cdot|\mathbf{s}_{obs})$ with quadratic adjustment, the asymptotic bias is the same as the bias of an estimator for which the conditional mean would be known exactly. Results of the same nature were found, for $d = 1$, by Fan & Yao (1998) when estimating the conditional variance and Hansen (2004) when estimating the conditional density using an estimator with adjustment. Compared to the bias of $\hat{g}_2(\cdot|\mathbf{s}_{obs})$, the bias of the estimator with linear adjustment $\hat{g}_1(\cdot|\mathbf{s}_{obs})$ contains an additional term depending on the curvature of the conditional mean.

Remark 4. Bias comparison between the estimators with and without adjustment To investigate the differences between the three es-

timators, we first assume that the partial posterior distribution of θ can be written as $h(\theta - m(\mathbf{s}))$ in which the function h does not depend on \mathbf{s} . This amounts at assuming an homoscedastic model in which the conditional distribution of θ given \mathbf{s} depends on \mathbf{s} only through the conditional mean $m(\mathbf{s})$. If the conditional mean m is linear in \mathbf{s} , both $C_{2,1}$ and $C_{2,2}$ are null involving that both estimators with regression adjustment have a smaller bias than $\hat{g}_0(\cdot|\mathbf{s}_{obs})$. For such ideal models, the bandwidth b of the estimators with regression adjustment can be taken infinitely large so that the variance will be inversely proportional to the total number of simulations n . Still assuming that $g(\theta|\mathbf{s}) = h(\theta - m(\mathbf{s}))$, but with a non-linear m , the constant $C_{2,2}$ is null so that the estimator $\hat{g}_2(\cdot|\mathbf{s}_{obs})$ has the smallest asymptotic MSE. However, for general partial posterior distributions, it is not possible to rank the three different biases. Consequently, when using the estimators with adjustment, the parameterization of the model shall be guided toward making the distributions $g(\theta|\mathbf{s})$ as homoscedastic as possible. This partly explains why the log and the logit transformations are usually considered in ABC before making regression adjustment (Box & Cox 1964). Comparison of asymptotic biases for conditional density estimators with and without adjustment can also be found in Hansen (2004) for $d = 1$.

4.2 Model selection

In the following, we give the main theorem concerning the bias and variance of the estimators $\hat{p}_j(\mathcal{M}_1|\mathbf{s}_{obs})$, $j = 0, 1$. We assume here that the bandwidth matrix is diagonal $B = bD$ but a more general theorem for non singular bandwidth matrix B could also be obtained from Fan et al. (1995)

Theorem 2 *Assume that conditions (A1) and (A5') of the Appendix hold and that the condition (A3) holds for both \mathcal{M}_1 and \mathcal{M}_2 . The bias and the variance of the estimators $\hat{p}_0(\mathcal{M}_1|\mathbf{s}_{obs})$ and $\hat{p}_1(\mathcal{M}_1|\mathbf{s}_{obs})$ are given by*

$$E[\hat{p}_j(\mathcal{M}_1|\mathbf{s}_{obs}) - p(\mathcal{M}_1|\mathbf{s}_{obs})] = \left(\mu_2(K) \frac{\text{tr}(D^2 p_{\mathbf{s}\mathbf{s}}(\mathcal{M}_1|\mathbf{s})|_{\mathbf{s}=\mathbf{s}_{obs}})}{2} + E_j + o_P(1) \right) b^2, \quad (23)$$

$$\text{Var}[\hat{p}_j] = \frac{1}{nb^d} \frac{\mu_2(K) p(\mathcal{M}_1|\mathbf{s}_{obs})(1 - p(\mathcal{M}_1|\mathbf{s}_{obs}))}{|D| p(\mathbf{s}_{obs})} + o_P\left(\frac{1}{nb^d}\right), \quad j = 0, 1,$$

in which

$$E_0 = \mu_2(K) \frac{p_{\mathbf{s}}(\mathcal{M}_1|\mathbf{s})|_{\mathbf{s}=\mathbf{s}_{obs}}^t D^2 p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})},$$

$$E_1 = \frac{\mu_2(K)}{2} \left(p_{\mathbf{s}}(\mathcal{M}_1|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} D^2 p_{\mathbf{s}}(\mathcal{M}_1|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}} \left(\frac{1}{1 - p(\mathcal{M}_1|\mathbf{s}_{obs})} - \frac{1}{p(\mathcal{M}_1|\mathbf{s}_{obs})} \right) \right),$$

and $p(\mathbf{s}) = \frac{p_1(\mathbf{s})}{2} + \frac{p_2(\mathbf{s})}{2}$.

Remark 1 Curse of dimensionality Standard algebra shows that the optimal bandwidths are found when $b \propto n^{-1/(d+4)}$ for which the MSE is of the order of $n^{-4/(d+4)}$.

Remark 2 Effect of design As for parameter estimation, it is not possible to give an universal ranking of the two different estimators. However, we find that the estimator based on local logistic regression is design-adaptive meaning that its bias does not depend on the design $p(\mathbf{s}_{obs}) = \frac{p_1(\mathbf{s}_{obs})}{2} + \frac{p_2(\mathbf{s}_{obs})}{2}$. By contrast, the constant E_0 involved in the bias of the Nadaraya-Watson estimator is inversely proportional to $p(\mathbf{s}_{obs})$.

Remark 3 Bayes factor The estimators of the logarithm of the partial Bayes factor are simply obtained as $\log \hat{\text{BF}}_j = \phi(\hat{p}_j)$, for $j = 0, 1$, where $\phi(x) = \log(x/(1-x))$. Using a Taylor expansion, we find that the bias and variance of $\log \hat{\text{BF}}_j$ have the same asymptotic behavior as those obtained in Theorem 2 except that the bias is multiplied by $\phi'(\hat{p}_j) = 1/(\hat{p}_j(1-\hat{p}_j))$ and that the variance is multiplied by $\phi'(\hat{p}_j)^2$. When estimating the log of the Bayes factor using local logistic regression, the bias takes a simple form as (Fan et al. 1995, Theorem 3)

$$E[\log \hat{\text{BF}}_1 - \log \text{BF}] = \frac{\mu_2(K)}{2} \text{tr}(D^2 \psi_{\text{ss}}(\mathbf{s}_{obs})) b^2,$$

where $\psi(\mathbf{s}) = \phi(\hat{p}_1(\mathbf{s}))$.

5 Numerical comparisons between the estimators

In this section, we consider a simple normal model to illustrate the curse of dimensionality in ABC. We assume that the d -dimensional data are drawn according to a normal distribution with mean $\mu = (\mu_1, \dots, \mu_d)$ and a variance-covariance matrix equal to the $d \times d$ identity matrix I_d . The prior for μ is a d -dimensional Gaussian distribution of mean μ_0 and variance-covariance matrix I_d . Here, we study the properties of the estimators of the distribution of

e^{μ_1} based on the summary statistics consisting of the d -dimensional empirical mean $\bar{\mathbf{x}}$ of the data. Note that the empirical mean of the first component of the data \bar{x}_1 is a sufficient statistic with respect to e^{μ_1} so that the partial posterior is the same as the full posterior in this example. The $d - 1$ additional empirical means convey no information for estimating e^{μ_1} and are added here to show that incorporating useless summary statistics can have a dramatic effect on the estimation of the posterior. We assume in the following that $\bar{x}_i = 0$, for $i = 1, \dots, d$ and that the sample size is $M = 10$. In this simple model, the posterior distribution is known and is a log-normal distribution with mean and variance (on a log scale) $\mu_0/(M + 1)$, and $1/(M + 1)$. We choose a spherically symmetric kernel (Härdle et al. 2004, p. 68) for K so that $K(\mathbf{u}) = K_1(\|\mathbf{u}\|)$ where $\|\cdot\|$ denotes the Euclidean norm and we consider the same kernel K_1 for \hat{K} where K_1 denotes the Epanechnikov kernel.

5.1 How many simulations are required to reach a given level of accuracy

Here we compute the minimum number n_{\min} of simulations that are required to reach a given level of accuracy when estimating the parameter e^{μ_1} . The number n_{\min} is defined as the smallest number of simulations so that the relative squared error is less than 10% when estimating the posterior distribution at 0. Similar computations were performed by Silverman (1986) to illustrate the curse of dimensionality for density estimation. The mean square errors are computed using equations (17) and (18) in which the constants can be analytically derived in this simple example. We assume that $B = bI_d$ and the optimal bandwidths for b and b' are found by numerical minimization of the asymptotic mean square errors. To simplify the computations, we assume that the prior for μ is a Gaussian distribution of mean $\mu_0 = 0$. As displayed by Figure 1, the regression-based estimators require a significantly smaller number of simulations for $d \leq 7$ but the improvement becomes negligible for $d > 7$. To give a quantitative flavor of the importance of the curse of dimensionality, we note that a minimum number of approximately one million of simulations is required to have a relative mean square error smaller than 10% when $d = 6$ and this number increases to more than a thousand of billions when $d = 10$. Compared to the estimator with linear adjustment, the estimator with quadratic adjustment requires less simulation when $d = 1, 3, 4$ but this is not true for $d = 2$. When $d \geq 5$, there is no more significant

difference (on a log scale) between the two estimators with adjustment.




Figure 1: Minimum number of simulations that are required to reach a given level of accuracy as a function of the dimension of the summary statistics. The minimum number of simulations is defined as the smallest number of simulations that are required so that the relative mean square error is less than 10%

Note that estimating e^{μ_1} rather μ_1 is a really loose parameterization here because the regression-based estimators would manage to cope with the curse of dimensionality when estimating μ_1 . Indeed the model $\mu_1 = m(\bar{\mathbf{x}}) + \epsilon$ is linear and homoscedastic so that the term involving b^2 in the bias (equation (17)) would be null for the regression-based estimators. The (loose) parameterization that has been chosen here illustrates 1) the importance of parameterization in ABC and 2) that the regression-based estimators will

typically cope with the curse of dimensionality for intermediate values of d but will be inefficient for large values of d .

5.2 Monte-Carlo comparison for parameter inference

To further investigate the differences between the three estimators, we compare the three ABC estimators \hat{g}_j , $j = 0, 1, 2$, to the true posterior distribution of e^{μ_1} . Here we set $\mu_0 = -1$. To compare the different estimators, we compute the mean integrated square error (MISE) defined as

$$\text{MISE} = E \left[\int_{e^{\mu_1} \in \mathbb{R}^+} \{g(e^{\mu_1} | \mathbf{s}_{obs}) - \hat{g}_j(e^{\mu_1} | \mathbf{s}_{obs})\}^2 de^{\mu_1} \right], \quad j = 0, 1, 2.$$

The integrated square errors are averaged using a total of 500 replicates, each consisting of performing $n = 10,000$ simulations. We choose a diagonal bandwidth matrix $B = bD$ where the diagonal elements of D contain the mean absolute deviations of each component of the summary statistics. For choosing the bandwidth b , we follow a common practice in the ABC literature consisting of choosing independently of the simulations the percentage of accepted simulations, i.e. the percentage of simulations for which $K_B(\mathbf{s}_i - \mathbf{s}_{obs})$ is different from 0. Here the percentage of accepted simulations is set to 5%. Note that many alternative methods have been considered for the choice of the bandwidth b when performing conditional density estimation or estimation of the conditional distribution function (Hall et al. 1999, Fan & Yim 2004, Hall et al. 2004). The bandwidth b' is computed using the Silverman's rule of thumb (Silverman 1986, page 48, eq. (3.31)). The integration required to compute the MISE is performed using a simple trapezoidal integration rule with 512 points equally spaced points between 0 and 3.

As displayed by Figure 2 and in accordance with Theorem 1, the MISE increases as the dimension of the summary statistics increases. For $d = 1$, the MISEs of the three estimators are the same and the improvement achieved by the estimators with regression-adjustment appears when $d \geq 2$. When $3 \leq d \leq 8$, the estimators with quadratic regression adjustment achieves the lowest MISE but the situation reverses for $d \geq 9$ where the lowest MISE is achieved by the estimator with linear adjustment. This can be explained by the variance of the estimator with quadratic adjustment that can be important for large values of d . Indeed, the estimator with quadratic adjustment requires the inference of $d(d+3)/2 + 1$ parameters whereas the estimator with linear adjustment requires the inference of only

$d + 1$ parameters. Note that the potential differences between the variances of the different estimators is not captured by the derivation of the asymptotic variances obtained in Theorem 1.

5.3 Monte-Carlo comparison for model selection

Here we assume that μ_1 is equal to 0 and the vector $(\mu_2, \dots, \mu_d) \rightsquigarrow \mathcal{N}(0, I_{d-1})$ in model \mathcal{M}_1 whereas model \mathcal{M}_2 assumes that $(\mu_1, \dots, \mu_d) \rightsquigarrow \mathcal{N}(0, I_d)$. This simple example amounts at testing $\mu_1 = 0$ against $\mu_1 \neq 0$. Standard computations lead to $p(\mathcal{M}_1|\bar{\mathbf{x}} = (0, \dots, 0)) = \sqrt{M+1}/(1 + \sqrt{M+1})$. Similarly to the example of parameter estimation, the summary statistics $(\bar{\mathbf{x}}_2, \dots, \bar{\mathbf{x}}_d)$ convey no information and are added here to illustrate the curse of dimensionality. We consider both the Nadaraya-Watson estimator $\hat{p}_0(\mathcal{M}_1|\bar{\mathbf{x}})$ and the estimator with local logistic regression $\hat{p}_1(\mathcal{M}_1|\bar{\mathbf{x}})$. The maximization of the weighted likelihood (equation (15)) is performed using the R routine `glm` (R Development Core Team 2008).

Figure 3 displays the mean squared error (MSE) of the two estimators of $p(\mathcal{M}_1|\bar{\mathbf{x}})$ using a total of $n = 10,000$ simulations and setting the percentage of accepted simulations to 5%. The local logistic regression provides a smaller MSE than the Nadaraya-Watson estimator for $d \geq 3$ and both estimators have comparable properties for $d = 1, 2$. The curse of dimensionality is once again displayed by this example since the MSEs increase with the dimension of the summary statistics. We note however that both estimators infer the probability of $p(\mathcal{M}_1|\bar{\mathbf{x}})$ accurately even when $d = 10$. Indeed, the mean squared errors divided by the square of the true value of $p(\mathcal{M}_1|\bar{\mathbf{x}})$ are equal to 0.65% (Nadaraya-Watson estimator) and 0.55% (local logistic regression).

6 Discussion

In this paper, we presented Approximate Bayesian Computation as a technique of inference that relies on stochastic simulations and non-parametric statistics. We have introduced an estimator of $g(\theta|\mathbf{s}_{obs})$ based on quadratic adjustment for which the asymptotic bias involves less terms than the asymptotic bias of the estimator with linear adjustment proposed by Beaumont et al. (2002). More generally, we have shown that the gain obtained with the estimators based on regression adjustment (equation (10) and (12)) is all the more important that the distribution of the residual ϵ in the model

$\theta(\mathbf{s}) = m(\mathbf{s}) + \epsilon$ is independent of \mathbf{s} . This observation emphasizes the importance of model parameterization when considering estimators based on regression adjustment.

The crucial point raised by the asymptotic results given in Theorem 1 and 2 concerns the curse of dimensionality when performing Approximate Bayesian Computation. For both parameter estimation and model selection, Theorems 1 and 2 show that the properties of the ABC estimators may seriously deteriorate as the dimension of the summary statistics increases. This is a particularly acute issue for ABC since it is tempting to use as many summary statistics as possible so that not much information is lost when summarizing the data. To increase the probability of targeting close to the observed summary statistics and consequently improve the properties of ABC estimators, two types of alternatives have been proposed. The first alternative consists of reducing the dimension of the summary statistics in the regression framework. Different techniques of dimension reduction have already been proposed in the context of ABC. Blum & François (2009) gave an estimator of the posterior distribution based on neural network regression and Leuenberger et al. (2009) proposed to reduce the number of summary statistics using principal component analysis or partial least-squares regression. The second type of alternative aims at performing simulations of the generative model in a parameter region for which the partial posterior distribution is substantial. Such adaptive ABC algorithms encompass ABC-MCMC algorithms (Marjoram et al. 2003, Bortot et al. 2007) and ABC sequential Monte Carlo samplers (Sisson et al. 2007, Beaumont et al. 2009, Toni et al. 2009).

APPENDIX: PROOFS OF THE THEOREMS

APPENDIX A. Hypotheses of Theorem 1 and 2

- A1)** The kernel K has a finite second order moment such that $\int \mathbf{u}\mathbf{u}^T K(\mathbf{u}) d\mathbf{u} = \mu_2(K)I_d$ where $\mu_2(K) \neq 0$. We also require that all first-order moments of K vanish, that is, $\int \mathbf{u}_i K(\mathbf{u}) d\mathbf{u} = 0$ for $i = 1, \dots, d$. As noted by Ruppert & Wand (1994), this condition is fulfilled by spherically symmetric kernels and product kernels based on symmetric univariate kernels.

A2) The kernel \tilde{K} is a symmetric univariate kernel with finite second order moment $\mu_2(\tilde{K})$.

A3) The observed summary statistics \mathbf{s}_{obs} lie in the interior of the support of p . At \mathbf{s}_{obs} , all the second order derivatives of the function p exist and are continuous.

A4) The point θ is in the support of the partial posterior distribution. At the point $(\theta, \mathbf{s}_{obs})$, all the second order derivatives of the partial posterior g exist and are continuous. The conditional mean of θ , $m(\mathbf{s})$, exists in a neighborhood of \mathbf{s}_{obs} and is finite. All its second order derivatives exist and are continuous.

For the theorem 1, we assume that

A5) The sequence of non-singular bandwidth matrices B and bandwidths b' is such that $1/(n|B|b')$, each entry of $B^t B$, and b' tend to 0 as $n \rightarrow \infty$.

For the theorem 2, we assume that

A5') The sequence of non-singular bandwidth matrices B is such that $1/(n|B|)$, and each entry of $B^t B$ tend to 0 as $n \rightarrow \infty$.

APPENDIX B. Proof of Theorem 1

The three estimators of the partial posterior distribution $\hat{g}_j(\cdot | \mathbf{s}_{obs})$, $j = 0, 1, 2$, are all of the Nadaraya-Watson type. The difficulty in the computation of the bias and the variance of the Nadaraya-Watson estimator comes from the fact that it is a ratio of two random variables. Following Pagan & Ullah (1999, p. 98) or Scott (1992), we linearize the estimators in order to compute their biases and their variances. We write the estimators of the partial posterior distribution \hat{g}_j , $j = 0, 1, 2$, as

$$\hat{g}_j(\theta | \mathbf{s}_{obs}) = \frac{\hat{g}_{j,N}}{\hat{g}_D}, \quad j = 0, 1, 2,$$

where

$$\hat{g}_{0,N} = \frac{1}{n} \sum_{i=1}^n \tilde{K}_{b'}(\theta_i - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs}),$$

$$\hat{g}_{1,N} = \frac{1}{n} \sum_{i=1}^n \tilde{K}_{b'}(\theta_i^* - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs}),$$

$$\hat{g}_{2,N} = \frac{1}{n} \sum_{i=1}^n \tilde{K}_{b'}(\theta_i^{**} - \theta) K_B(\mathbf{s}_i - \mathbf{s}_{obs}),$$

and

$$\hat{g}_D = \sum_{i=1}^n K_B(\mathbf{s}_i - \mathbf{s}_{obs}).$$

To compute the asymptotic expansions of the moments of the three estimators, we use the following lemma

Lemma 1 For $j = 0, 1, 2$, we have

$$\begin{aligned} \hat{g}_j(\theta|\mathbf{s}_{obs}) &= \frac{E[\hat{g}_{j,N}]}{E[\hat{g}_D]} + \frac{\hat{g}_{j,N} - E[\hat{g}_{j,N}]}{E[\hat{g}_D]} - \frac{E[\hat{g}_{j,N}](\hat{g}_D - E[\hat{g}_D])}{E[\hat{g}_D]^2} \\ &\quad + O_P(\text{Cov}(\hat{g}_{j,N}, \hat{g}_D) + \text{Var}[\hat{g}_D]) \end{aligned} \quad (24)$$

Proof. Lemma 1 is a simple consequence of the Taylor expansion for the function $(x, y) \rightarrow x/y$ in the neighborhood of the point $(E[\hat{g}_{j,N}], E[\hat{g}_D])$ (see Pagan & Ullah (1999) for another proof). The order of the reminder follows from the weak law of large numbers. □

The following Lemma gives an asymptotic expansion of all the expressions involved in equation (24).

Lemma 2 Suppose assumption (A1)-(A5) hold, denote $\epsilon = \theta - m(\mathbf{s}_{obs})$, then we have

$$E[\hat{g}_D] = p(\mathbf{s}_{obs}) + \frac{1}{2} \mu_2(K) \text{tr}(BB^t p_{ss}(\mathbf{s}_{obs})) + o(\text{tr}(B^t B)), \quad (25)$$

$$\begin{aligned} E[\hat{g}_{0,N}] &= p(\mathbf{s}_{obs}) g(\theta|\mathbf{s}_{obs}) + \frac{1}{2} b^2 \mu_2(\tilde{K}) g_{\theta\theta}(\theta|\mathbf{s}_{obs}) p(\mathbf{s}_{obs}) \\ &\quad + \mu_2(K) [g_{\mathbf{s}}(\theta|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} BB^t p_{\mathbf{s}}(\mathbf{s}_{obs}) + \frac{1}{2} g(\theta|\mathbf{s}_{obs}) \text{tr}(BB^t p_{\mathbf{ss}}(\mathbf{s}_{obs})) \\ &\quad + \frac{1}{2} p(\mathbf{s}_{obs}) \text{tr}(BB^t g_{\mathbf{ss}}(\theta|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})] + o(b^2) + o(\text{tr}(B^t B)), \end{aligned} \quad (26)$$

$$\begin{aligned}
E[\hat{g}_{1,N}] &= p(\mathbf{s}_{obs})h(\epsilon|\mathbf{s}_{obs}) + \frac{1}{2}b'^2\mu_2(\tilde{K})h_{\epsilon\epsilon}(\epsilon|\mathbf{s}_{obs})p(\mathbf{s}_{obs}) \\
&+ \mu_2(K)[h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}}BB^tp_{\mathbf{s}}(\mathbf{s}_{obs}) + \frac{1}{2}h(\epsilon|\mathbf{s}_{obs})\text{tr}(BB^tp_{\mathbf{ss}}(\mathbf{s}_{obs})) \\
&+ \frac{1}{2}p(\mathbf{s}_{obs})\text{tr}(BB^th_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}}) - \frac{h_{\epsilon}(\epsilon|\mathbf{s}_{obs})}{2}\text{tr}(BB^tm_{\mathbf{ss}}(\mathbf{s}_{obs}))] \\
&+ o(b'^2) + o(\text{tr}(B^tB)), \tag{27}
\end{aligned}$$

$$\begin{aligned}
E[\hat{g}_{2,N}] &= p(\mathbf{s}_{obs})h(\epsilon|\mathbf{s}_{obs}) + \frac{1}{2}b'^2\mu_2(\tilde{K})h_{\epsilon\epsilon}(\epsilon|\mathbf{s}_{obs})p(\mathbf{s}_{obs}) \\
&+ \mu_2(K)[h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}}BB^tp_{\mathbf{s}}(\mathbf{s}_{obs}) + \frac{1}{2}h(\epsilon|\mathbf{s}_{obs})\text{tr}(BB^tp_{\mathbf{ss}}(\mathbf{s}_{obs})) \\
&+ \frac{1}{2}p(\mathbf{s}_{obs})\text{tr}(BB^th_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}}) + o(b'^2) + o(\text{tr}(B^tB)), \tag{28}
\end{aligned}$$

$$\text{Var}[\hat{g}_D] = \frac{R(K)p(\mathbf{s}_{obs})}{n|B|} + O\left(\frac{1}{n}\right) + O\left(\frac{\text{tr}(BB^t)}{n|B|}\right), \tag{29}$$

$$\text{Var}[\hat{g}_{j,N}] = \frac{R(K)R(\tilde{K})g(\theta|\mathbf{s}_{obs})p(\mathbf{s}_{obs})}{nb'|B|} + O\left(\frac{1}{n}\right) + O\left(\frac{\text{tr}(BB^t)}{nb'|B|}\right) + O\left(\frac{b'}{n|B|}\right), \quad j = 0, 1, 2, \tag{30}$$

$$\text{Cov}[\hat{g}_{j,N}, \hat{g}_D] = \frac{R(K)p(\mathbf{s}_{obs})g(\theta|\mathbf{s}_{obs})}{n|B|} + O\left(\frac{1}{n}\right), \quad j = 0, 1, 2. \tag{31}$$

Proof. See the Supplemental Material available online \square

Theorem 1 is a particular case of the following theorem that gives the bias and variance of the three estimators of the partial posterior distribution for a general nonsingular bandwidth matrix B .

Theorem 3 *Assume that B is a non-singular bandwidth matrix and assume that conditions (A1)-(A5) holds, then the bias of \hat{g}_j , $j = 0, 1, 2$, is given by*

$$E[\hat{g}_j(\theta|\mathbf{s}_{obs}) - g(\theta|\mathbf{s}_{obs})] = D_1b'^2 + D_{2,j} + O_P((\text{tr}(B^tB) + b'^2)^2) + O_P\left(\frac{1}{n|B|}\right), \quad j = 0, 1, 2, \tag{32}$$

with

$$\begin{aligned}
D_1 = C_1 &= \frac{\mu_2(\tilde{K})g_{\theta\theta}(\theta|\mathbf{s}_{obs})}{2}, \\
D_{2,0} &= \mu_2(K) \left(\frac{g_{\mathbf{s}}(\theta|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}}BB^tp_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(BB^tg_{\mathbf{ss}}(\theta|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} \right),
\end{aligned}$$

$$D_{2,1} = \mu_2(K) \left(\frac{h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} BB^t p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(BB^t h_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} - \frac{h_{\epsilon}(\epsilon|\mathbf{s}_{obs}) \text{tr}(BB^t m_{\mathbf{ss}})}{2} \right),$$

and

$$D_{2,2} = \mu_2(K) \left(\frac{h_{\mathbf{s}}(\epsilon|\mathbf{s})^t_{|\mathbf{s}=\mathbf{s}_{obs}} BB^t p_{\mathbf{s}}(\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})} + \frac{\text{tr}(BB^t h_{\mathbf{ss}}(\epsilon|\mathbf{s})_{|\mathbf{s}=\mathbf{s}_{obs}})}{2} \right),$$

The variance of the estimators \hat{g}_j , $j = 0, 1, 2$, is given by

$$\text{Var}[\hat{g}_j(\theta|\mathbf{s}_{obs})] = \frac{R(K)R(\tilde{K})g(\theta|\mathbf{s}_{obs})}{p(\mathbf{s}_{obs})n|B|b'}(1 + o_P(1)), \quad j = 0, 1, 2, \quad (33)$$

Proof.

Theorem 3 is a consequence of Lemma 1 and 2. Taking expectations on both sides of equation (24), we find that

$$E[\hat{g}_j(\theta|\mathbf{s}_{obs})] = \frac{E[\hat{g}_{j,N}]}{E[\hat{g}_D]} + O_P[\text{Cov}(g_{j,N}, \hat{g}_D) + \text{Var}(\hat{g}_D)]. \quad (34)$$

Using a Taylor expansion, and the equations (25)-(28), (29), and (31) given in Lemma 2, we find the bias of the estimators given in equation (32).

For the computation of the variance, we find from equation (24) and (34) that

$$\hat{g}_j(\theta|\mathbf{s}_{obs}) - E[\hat{g}_j(\theta|\mathbf{s}_{obs})] = \frac{\hat{g}_{j,N} - E[\hat{g}_{j,N}]}{E[\hat{g}_D]} - \frac{E[\hat{g}_{j,N}](\hat{g}_D - E[\hat{g}_D])}{E[\hat{g}_D]^2} + O_P\left(\frac{1}{n|B|}\right). \quad (35)$$

The order of the reminder follows from equations (29) and (31). Taking the expectation of the square of equation (35), we now find

$$\text{Var}[\hat{g}_j(\theta|\mathbf{s}_{obs})] = \frac{\text{Var}[\hat{g}_{j,N}]}{E[\hat{g}_D]} + \frac{E[\hat{g}_{j,N}]^2 \text{Var}[\hat{g}_D]}{E[\hat{g}_D]^4} - 2\text{Cov}(\hat{g}_D, \hat{g}_{j,N}) \frac{E[\hat{g}_{j,N}]}{E[\hat{g}_D]^3} + o_P\left(\frac{1}{n|B|}\right). \quad (36)$$

The variance of the estimators given in equation (33) follows from a Taylor expansion that makes use of equations (25)-(31) given in Lemma 2.

□

APPENDIX C. Proof of Theorem 2

The estimator $\hat{p}_0(\mathcal{M}_1|\mathbf{s}_{obs})$ is a Nadaraya-Watson estimator and its asymptotics are given by the standard asymptotic bias and variance of the Nadaraya-Watson estimator (Härdle et al. 2004, p. 131). The asymptotic bias and variance of the estimator $\hat{p}_1(\mathcal{M}_1|\mathbf{s}_{obs})$ given in equation (16) can be obtained from the Theorem 3 of Fan et al. (1995) dedicated to multivariate local regression for generalized linear models.

APPENDIX D. Supplemental Materials.

- Proof of Lemma 2. See Supplemental Materials at

http://membres-timc.imag.fr/Michael.Blum/my_publications.html

References

- Abril, J. C. (1994), ‘On the concept of approximate sufficiency’, *Pakistan Journal of Statistics* **10**, 171–177.
- Beaumont, M. A. (2008), Joint determination of topology, divergence time, and immigration in population trees, *in* R. C. Matsumura S, Forster P, ed., ‘Simulation, Genetics and Human Prehistory’, McDonald Institute Monographs, McDonald Institute Monographs, UK, pp. 134–1541.
- Beaumont, M. A., Marin, J.-M., Cornuet, J.-M. & Robert, C. P. (2009), ‘Adaptivity for ABC algorithms: the ABC-PMC scheme’, *Biometrika* **to appear**.
- Beaumont, M. A., Zhang, W. & Balding, D. J. (2002), ‘Approximate Bayesian computation in population genetics’, *Genetics* **162**, 2025–2035.
- Blum, M. G. B. & François, O. (2009), ‘Non-linear regression models for Approximate Bayesian Computation’, *Statistics and Computing*, *to appear* **arXiv:0809.4178**.
- Blum, M. G. B. & Tran, V.-C. (2009), ‘Hiv with contact-tracing: a case study in approximate bayesian computation’, **ArXiv:0810.0896**.

- Bortot, P., Coles, S. G. & Sisson, S. A. (2007), ‘Inference for stereological extremes’, *Journal of the American Statistical Association* **102**, 84–92.
- Box, G. E. P. & Cox, D. R. (1964), ‘An analysis of transformations’, *Journal of the Royal Statistical Society: Series B* **26**, 211–246.
- Cam, L. L. (1964), ‘Sufficiency and approximate sufficiency’, *The Annals of Mathematical Statistics* **35**, 1419–1455.
- Diggle, P. J. & Gratton, R. J. (1984), ‘Monte Carlo methods of inference for implicit statistical models’, *Journal of the Royal Society: Series B* **46**, 193–227.
- Fan, J., Heckman, N. E. & Wand, M. P. (1995), ‘Local polynomial kernel regression for generalized linear models and quasi-likelihood functions’, *Journal of the American Statistical Association* **90**, 141–150.
- Fan, J. & Yao, Q. (1998), ‘Efficient estimation of conditional variance functions in stochastic regression’, *Biometrika* **85**, 645–660.
- Fan, J. & Yim, T. H. (2004), ‘A data-driven method for estimating conditional densities’, *Biometrika* **91**, 819–834.
- François, O., Blum, M. G. B., Jakobsson, M. & Rosenberg, N. A. (2008), ‘Demographic history of european populations of arabidopsis thaliana.’, *PLoS genetics* **4**(5).
- Fu, Y. X. & Li, W. H. (1997), ‘Estimating the age of the common ancestor of a sample of DNA sequences’, *Molecular Biology and Evolution* **14**, 195–199.
- Hall, P., Racine, J. & Li, Q. (2004), ‘Cross-validation and the estimation of conditional probability densities’, *Journal of the American Statistical Association* **99**, 1015–1026.
- Hall, P., Wolff, R. C. L. & Yao, Q. (1999), ‘Methods for estimating a conditional distribution function’, *Journal of the American Statistical Association* **94**, 154–163.
- Hansen, B. E. (2004), Nonparametric conditional density estimation. Working paper available at <http://www.ssc.wisc.edu/~bhansen/papers/ncde.pdf>.

- Härdle, W., Müller, M., Sperlich, S. & Werwatz, A. (2004), *Nonparametric and Semiparametric Models*, Springer, Berlin.
- Hyndman, R. J., Bashtannyk, D. M. & Grunwald, G. K. (1996), ‘Estimating and visualizing conditional densities’, *Journal of Computing and Graphical Statistics* **5**(4), 315–336.
- Jabot, F. & Chave, J. (2009), ‘Inferring the parameters of the neutral theory of biodiversity using phylogenetic information and implications for tropical forests’, *Ecology Letters* **12**(3), 239–248.
- Joyce, P. & Marjoram, P. (2008), ‘Approximately sufficient statistics and Bayesian computation’, *Statistical Applications in Genetics and Molecular Biology* **7**(1). Article 26.
- Leuenberger, C., Wegmann, D. & Excoffier, L. (2009), ‘Bayesian computation and model selection in population genetics’, [ArXiv/0901.2231](https://arxiv.org/abs/0901.2231).
- Marjoram, P., Molitor, J., Plagnol, V. & Tavaré, S. (2003), ‘Markov chain monte carlo without likelihoods.’, *Proceedings of the National Academy of Sciences of the United States of America* **100**(26), 15324–15328.
- McCullagh, P. & Nelder, J. A. (1990), *Generalized Linear Models*, 2nd edn, CRC Press, London.
- Nadaraya, E. (1964), ‘On estimating regression’, *Theory of Probability and Applications* **9**, 141–142.
- Pagan, A. & Ullah, A. (1999), *Nonparametric econometrics / Adrian Pagan, Aman Ullah*, Cambridge University Press, Cambridge ; New York ;, Cambridge, UK.
- Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A. & Feldman, M. W. (1999), ‘Population growth of human Y chromosomes: a study of y chromosome microsatellites’, *Molecular Biology and Evolution* **16**, 1791–1798.
- R Development Core Team (2008), *R: A Language and Environment for Statistical Computing*. ISBN 3-900051-07-0.

- Ratmann, O., Jørgensen, O., Hinkley, T., Stumpf, M., Richardson, S. & Wiuf, C. (2007), ‘Using likelihood-free inference to compare evolutionary dynamics of the protein networks of *h. pylori* and *p. falciparum*’, *PLoS Computational Biology* **3**(11), e230+.
- Rosenblatt, M. (1969), Conditional probability density and regression estimates, *in* ‘Multivariate Analysis II’, Academic Press, New York, pp. 25–31.
- Ruppert, D. & Wand, M. P. (1994), ‘Multivariate locally weighted least squares regression’, *Annals of Statistics* **22**, 1346–1370.
- Scott, D. W. (1992), *Multivariate density estimation*, Wiley, New York.
- Silverman, B. W. (1986), *Density estimation: for statistics and data analysis*, Chapman and Hall, London.
- Sisson, S. A., Fan, Y. & Tanaka, M. (2007), ‘Sequential Monte Carlo without likelihoods’, *Proceedings of the National Academy of Sciences of the United States of America* **104**, 1760–1765.
- Tanaka, M., Francis, A., Luciani, F. & Sisson, S. (2006), ‘Estimating tuberculosis transmission parameters from genotype data using approximate Bayesian computation’, *Genetics* **173**, 1511–1520.
- Tavaré, S., Balding, D. J., Griffiths, R. C. & Donnelly, P. (1997), ‘Inferring coalescence times from DNA sequence data’, *Genetics* **145**, 505–518.
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A. & Stumpf, M. P. (2009), ‘Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems’, *Journal of The Royal Society Interface* **6**, 187–202.
- Watson, G. S. (1964), ‘Smooth regression analysis’, *Shankya Series A* **26**, 359–372.
- Wilkinson, R. D. & Tavaré, S. (2009), ‘Estimating primate divergence times by using conditioned birth-and-death processes’, *Theoretical Population Biology* **75**(4), 278–285.




Figure 2: Mean integrated square error as a function of the dimension of the summary statistics. The parameter to infer is the exponential of the location parameter e^{μ_1} of the first component of a Gaussian sample. The total number of simulations was set to $n = 10,000$ and the percentage of accepted simulation was set to 5%




Figure 3: Mean square error of the probability of model \mathcal{M}_1 as a function of the dimension of the summary statistics. The total number of simulations was set to $n = 10,000$ and the percentage of accepted simulation was set to 5%.