

# Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation

Paul Fearnhead and Dennis Prangle

*Lancaster University, UK*

[*Read before The Royal Statistical Society at a meeting organized by the Research Section on Wednesday, December 14th, 2011, Professor D. M. Titterton in the Chair*]

**Summary.** Many modern statistical applications involve inference for complex stochastic models, where it is easy to simulate from the models, but impossible to calculate likelihoods. Approximate Bayesian computation (ABC) is a method of inference for such models. It replaces calculation of the likelihood by a step which involves simulating artificial data for different parameter values, and comparing summary statistics of the simulated data with summary statistics of the observed data. Here we show how to construct appropriate summary statistics for ABC in a semi-automatic manner. We aim for summary statistics which will enable inference about certain parameters of interest to be as accurate as possible. Theoretical results show that optimal summary statistics are the posterior means of the parameters. Although these cannot be calculated analytically, we use an extra stage of simulation to estimate how the posterior means vary as a function of the data; and we then use these estimates of our summary statistics within ABC. Empirical results show that our approach is a robust method for choosing summary statistics that can result in substantially more accurate ABC analyses than the *ad hoc* choices of summary statistics that have been proposed in the literature. We also demonstrate advantages over two alternative methods of simulation-based inference.

**Keywords:** Indirect inference; Likelihood-free inference; Markov chain Monte Carlo methods; Simulation; Stochastic kinetic networks

## 1. Introduction

### 1.1. Background

Many modern statistical applications involve inference for stochastic models given partial observations. Often it is easy to simulate from the models but calculating the likelihood of the data, even by using computationally intensive methods, is impracticable. In these cases a natural approach to inference is to use simulations from the model for different parameter values, and to compare the simulated data sets with the observed data. Loosely, the idea is to estimate the likelihood of a given parameter value from the proportion of data sets, simulated using that parameter value, that are ‘similar to’ the observed data. This idea dates back at least as far as Diggle and Gratton (1984).

If we replace ‘similar to’ with ‘the same as’ (see for example Tavaré *et al.* (1997)), then this approach would give an unbiased estimate of the likelihood; and asymptotically as we increase the amount of simulation we obtain a consistent estimate. However, in most applications the probability of an exact match of the simulated data with the observed data is negligible, or 0,

*Address for correspondence:* Paul Fearnhead, Department of Mathematics and Statistics, Fylde College, Lancaster University, Lancaster, LA1 4YF, UK.  
E-mail: [p.fearnhead@lancaster.ac.uk](mailto:p.fearnhead@lancaster.ac.uk)

so we cannot consider such exact matches. The focus of this paper is how to define ‘similar to’ for these cases.

In this paper we focus on a particular approach: approximate Bayesian computation (ABC). This approach combines an estimate of the likelihood with a prior to produce an approximate posterior, which we shall refer to as the ABC posterior. The use of ABC initially became popular within population genetics, where simulation from a range of population genetic models is possible by using the coalescent (Kingman, 1982), but where calculating likelihoods is impracticable for realistic sized data sets. The first use of ABC was by Pritchard *et al.* (1999), who looked at inference about human demographic history. Further applications include inference for recombination rates (Padhukasahasram *et al.*, 2006), evolution of pathogens (Wilson *et al.*, 2009) and evolution of protein networks (Ratmann *et al.*, 2009). Its increasing importance can be seen by the current range of application of ABC, which has recently been applied within epidemiology (McKinley *et al.*, 2009; Tanaka *et al.*, 2006), inference for extremes (Bortot *et al.*, 2007), dynamical systems (Toni *et al.*, 2009) and Gibbs random fields (Grelaud *et al.*, 2009) among many others. Part of the appeal of ABC is its flexibility; it can easily be applied to any model for which forward simulation is possible. For example Wegmann *et al.* (2009) stated that ABC

‘should allow evolutionary geneticists to reasonably estimate the parameters they are really interested in, rather than require them to shift their interest to problems for which full-likelihood solutions are available’.

Recently software has been developed to help to implement ABC within population genetics (Cornuet *et al.*, 2008; Lopes *et al.*, 2009) and systems biology (Liepe *et al.*, 2010).

### 1.2. Approximate Bayesian computation algorithms and approximations

Consider analysing  $n$ -dimensional data  $\mathbf{y}_{\text{obs}}$ . We have a model for the data, which depends on an unknown  $p$ -dimensional parameter  $\theta$ . Denote the probability density of the data given a specific parameter value by  $\pi(\mathbf{y}|\theta)$ , and denote our prior by  $\pi(\theta)$ . We assume that it is simple to simulate  $\mathbf{Y}$  from  $\pi(\mathbf{y}|\theta)$  for any  $\theta$ , but that we do not have an analytic form for  $\pi(\mathbf{y}|\theta)$ .

We define the ABC posterior in terms of

- (a) a function  $S(\cdot)$  which maps the  $n$ -dimensional data onto a  $d$ -dimensional summary statistic,
- (b) a density kernel  $K(\mathbf{x})$  for a  $d$ -dimensional vector  $\mathbf{x}$ , which integrates to 1, and
- (c) a bandwidth  $h > 0$ .

Let  $\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}})$ . If we now define an approximation to the likelihood as

$$p(\theta|\mathbf{s}_{\text{obs}}) = \int \pi(\mathbf{y}|\theta) K[\{S(\mathbf{y}) - \mathbf{s}_{\text{obs}}\}/h] d\mathbf{y},$$

then the ABC posterior can be defined as

$$\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) \propto \pi(\theta) p(\theta|\mathbf{s}_{\text{obs}}). \quad (1)$$

The idea of ABC is that the ABC posterior will approximate, in some way, the true posterior for  $\theta$  and can be used for inference about  $\theta$ . The form of approximation will depend on the choice of  $S(\cdot)$ ,  $K(\cdot)$  and  $h$ . For example, if  $S(\cdot)$  is the identity function then we can view  $p(\theta|\mathbf{s}_{\text{obs}})$  as a kernel density approximation to the likelihood. If also  $h \rightarrow 0$ , then the ABC posterior will converge to the true posterior. For other choices of  $S(\cdot)$ , the kernel is measuring the closeness of  $\mathbf{y}$  to  $\mathbf{y}_{\text{obs}}$  just via the closeness of  $S(\mathbf{y})$  to  $\mathbf{s}_{\text{obs}}$ . The reason for considering ABC is that we can

**Table 1.** Algorithm 1: importance (and rejection) sampling implementation of ABC

<p><i>Input</i>—a set of data <math>\mathbf{y}_{\text{obs}}</math>, and a function <math>S(\cdot)</math>;  a density kernel <math>K(\cdot)</math>, with <math>\max\{K(\mathbf{x})\} = 1</math> and a bandwidth <math>h &gt; 0</math>;  a proposal density <math>g(\theta)</math>, with <math>g(\theta) &gt; 0</math> if <math>\pi(\theta) &gt; 0</math>;  an integer <math>N &gt; 0</math></p> <p><i>Initialize</i>—define <math>\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}})</math></p> <p><i>Iterate</i>—for <math>i = 1, \dots, N</math>:  step 1, simulate <math>\theta_i</math> from <math>g(\theta)</math>;  step 2, simulate <math>\mathbf{y}_{\text{sim}}</math> from <math>\pi(\mathbf{y} \theta_i)</math>, and calculate <math>\mathbf{s} = S(\mathbf{y}_{\text{sim}})</math>.  step 3, with probability <math>K\{(\mathbf{s} - \mathbf{s}_{\text{obs}})/h\}</math> set <math>w_i = \pi(\theta_i)/g(\theta_i)</math>; otherwise set <math>w_i = 0</math></p> <p><i>Output</i>—a set of parameter values <math>\{\theta_i\}_{i=1}^N</math> and corresponding weights <math>\{w_i\}_{i=1}^N</math></p>
--

construct Monte Carlo algorithms which approximate the ABC posterior and only require the ability to simulate from  $\pi(\mathbf{y}|\theta)$ .

For simplicity of the future exposition, we shall assume that  $\max\{K(\mathbf{x})\} = 1$ . This assumption imposes no restriction, since, if  $K_0(\mathbf{x})$  is a density kernel, then so is  $K(\mathbf{x}) = h_0^{-d} K_0(\mathbf{x}/h_0)$  for any  $h_0 > 0$ . Thus we can choose  $h_0$  so that  $\max\{K(\mathbf{x})\} = 1$ . Note that a bandwidth  $\lambda$  for kernel  $K_0(\mathbf{x})$  is equivalent to a bandwidth  $h = \lambda h_0$  for  $K(\mathbf{x})$ , so the value of  $h_0$  just redefines the units of the bandwidth  $h$ .

One algorithm for approximating the ABC posterior (1), based on importance sampling, is given by algorithm 1 (Table 1). For alternative importance sampling approaches, based on sequential Monte Carlo sampling, see Beaumont *et al.* (2009), Sisson *et al.* (2007) and Toni *et al.* (2009). A standard importance sampler for approximating the ABC posterior would repeatedly simulate a parameter value  $\theta'$  from a proposal  $g(\theta)$  and then assign it an importance sampling weight,

$$\frac{\pi(\theta')}{g(\theta')} \int \pi(\mathbf{y}|\theta') \frac{K[\{S(\mathbf{y}) - \mathbf{s}_{\text{obs}}\}]}{h} d\mathbf{y}. \quad (2)$$

This is not possible here, as  $\pi(\mathbf{y}|\theta)$  is intractable, so algorithm 1 introduces an extra Monte Carlo step. This step first simulates  $\mathbf{y}_{\text{sim}}$  from  $\pi(\mathbf{y}|\theta')$  and then accepts this value with probability  $K[\{S(\mathbf{y}_{\text{sim}}) - \mathbf{s}_{\text{obs}}\}/h]$ . Accepted values are assigned weights  $\pi(\theta')/g(\theta')$ . The key thing is that the expected value of such weights is just expression (2), which is all that is required for this to be a valid importance sampling algorithm targeting expression (1).

The output of algorithm 1 is a weighted sample of  $\theta$ -values, which approximates the ABC posterior of  $\theta$ . Many of the weights will be 0, and in practice we would remove the corresponding  $\theta$ -values from the sample. A specific case of algorithm 1 occurs when  $g(\theta) = \pi(\theta)$ , and then we have a rejection sampling algorithm. The most common implementation of ABC has a deterministic accept–reject decision in step 3, and this corresponds to  $K(\cdot)$  being the density of a uniform random variable—the support of the uniform random variable defining the values of  $\mathbf{s} - \mathbf{s}_{\text{obs}}$  which are accepted.

An alternative Monte Carlo procedure for implementing ABC is based on Markov chain Monte Carlo (MCMC) sampling (Marjoram *et al.*, 2003; Bortot *et al.*, 2007) and given in algorithm 2 (Table 2). Both this and algorithm 1 target the same ABC posterior (for a proof of the validity of this algorithm see Sisson *et al.* (2010)).

Good implementation of ABC requires a trade-off between the approximation error between the ABC posterior and the true posterior, and the Monte Carlo approximation error of the ABC posterior. The latter of these will also be affected by the algorithm that is used to implement

**Table 2.** Algorithm 2: MCMC sampling implementation of ABC

*Input*—a set of data  $\mathbf{y}_{\text{obs}}$ , and a function  $S(\cdot)$ ;  
 a density kernel  $K(\cdot)$ , with  $\max\{K(\mathbf{x})\} = 1$  and a bandwidth  $h > 0$ ;  
 a transition kernel  $g(\cdot)$ ;  
 an integer  $N > 0$   
*Initialize*—define  $\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}})$ , and choose or simulate  $\theta_0$  and  $\mathbf{s}_0$   
*Iterate*—for  $i = 1, \dots, N$ :  
 step 1, simulate  $\theta$  from  $g(\theta|\theta_{i-1})$ ;  
 step 2, simulate  $\mathbf{y}_{\text{sim}}$  from  $\pi(\mathbf{y}|\theta)$ , and calculate  $\mathbf{s} = S(\mathbf{y}_{\text{sim}})$ ;  
 step 3, with probability

$$\min\left[1, \frac{K\{(\mathbf{s} - \mathbf{s}_{\text{obs}})/h\}}{K\{(\mathbf{s}_{i-1} - \mathbf{s}_{\text{obs}})/h\}} \frac{\pi(\theta) g(\theta_{i-1}|\theta)}{\pi(\theta_{i-1}) g(\theta|\theta_{i-1})}\right]$$

accept  $\theta$  and set  $\theta_i = \theta$  and  $\mathbf{s}_i = \mathbf{s}$ ; otherwise set  $\theta_i = \theta_{i-1}$  and  $\mathbf{s}_i = \mathbf{s}_{i-1}$   
*Output*—a set of parameter values  $\{\theta_i\}_{i=1}^N$ .

ABC, and the specific implementation of that algorithm. We shall consider both of these approximations in turn.

To understand the former approximation, consider the following alternative expression of the ABC posterior as a continuous mixture of posteriors. Consider the random variable  $\mathbf{S} = S(\mathbf{Y})$ , which is just the summary statistic of the data. Denote the posterior for  $\theta$  given  $\mathbf{S} = \mathbf{s}$  by  $\pi(\theta|\mathbf{s})$ , and the marginal density for  $\mathbf{S}$  by  $\pi(\mathbf{s}) := \int \pi(\mathbf{s}|\theta) \pi(\theta) d\theta$ . Then

$$\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) = \int \beta(\mathbf{s}) \pi(\theta|\mathbf{s}) d\mathbf{s}, \quad \beta(\mathbf{s}) = \frac{\pi(\mathbf{s}) K\{(\mathbf{s} - \mathbf{s}_{\text{obs}})/h\}}{\int \pi(\mathbf{s}) K\{(\mathbf{s} - \mathbf{s}_{\text{obs}})/h\} d\mathbf{s}}. \quad (3)$$

The mixing weight for a given value of  $\mathbf{S}$  is just the conditional density for such a value of  $\mathbf{S}$  given acceptance in step 3 of algorithm 1. If  $h$  is small then  $\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) \approx \pi(\theta|\mathbf{s}_{\text{obs}})$ . Thus we can split the approximation of the posterior by ABC into the approximation of  $\pi(\theta|\mathbf{y}_{\text{obs}})$  by  $\pi(\theta|\mathbf{s}_{\text{obs}})$ , and the further approximation of  $\pi(\theta|\mathbf{s}_{\text{obs}})$  by  $\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}})$ . The former is controlled by the choice of  $S(\cdot)$ ; the latter by  $K(\cdot)$  and  $h$ .

Now consider the Monte Carlo approximation within importance sampling ABC. Consider a scalar function  $a(\theta)$ , and define the ABC posterior mean of  $a(\theta)$  as

$$E_{\text{ABC}}\{a(\theta)|\mathbf{s}_{\text{obs}}\} = \int a(\theta) \pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) d\theta.$$

Assuming that this exists then, by the law of large numbers, as  $N \rightarrow \infty$

$$\frac{\sum_{i=1}^N w_i a(\theta_i)}{\sum_{i=1}^N w_i} \rightarrow E_{\text{ABC}}\{a(\theta)|\mathbf{s}_{\text{obs}}\}, \quad (4)$$

where the  $w_i$ s are the importance sampling weights, and convergence is in probability. For rejection sampling, for large  $N$  the variance of this estimator is just

$$\text{var}_{\text{ABC}}\{a(\theta)|\mathbf{s}_{\text{obs}}\} / N_{\text{acc}}, \quad (5)$$

where the numerator is the ABC posterior variance of  $a(\theta)$ , and the denominator is  $N_{\text{acc}} = N \int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) d\theta$ , the expected number of acceptances.

Similar calculations for both importance sampling in general and for the MCMC version of ABC are given in Appendix A. In all cases the Monte Carlo error depends on  $\int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) d\theta$ , the average acceptance probability of the rejection sampling algorithm. The following lemma

characterizes this acceptance probability for small  $h$  in this case of continuous summary statistics (similar results can be shown for discrete summary statistics).

*Lemma 1.* Assume that either

- (a) the marginal distribution for the summary statistics  $\pi(\mathbf{s})$  is continuous at  $\mathbf{s} = \mathbf{s}_{\text{obs}}$  and that the kernel  $K(\cdot)$  has finite support or
- (b)  $\pi(\mathbf{s})$  is continuously differentiable,  $|\partial\pi(\mathbf{s})/\partial s_i|$  is bounded above for all  $i$  and  $\int |\mathbf{x}_i| K(\mathbf{x}) \, d\mathbf{x}$  is bounded for all  $i$ . Then, in the limit as  $h \rightarrow 0$ ,

$$\int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) \, d\theta = \pi(\mathbf{s}_{\text{obs}})h^d + o(h^d), \quad (6)$$

where  $d$  is the dimension of  $\mathbf{s}_{\text{obs}}$ .

For a proof of lemma 1, see Appendix B.

This result gives insight into how  $S(\cdot)$  and  $h$  affect the Monte Carlo error. To minimize Monte Carlo error, we need  $h^d$  to be not too small. Thus ideally we want  $S(\cdot)$  to be a low dimensional summary of the data that is sufficiently informative about  $\theta$  that  $\pi(\theta|\mathbf{s}_{\text{obs}})$  is close, in some sense, to  $\pi(\theta|\mathbf{y}_{\text{obs}})$ . The choice of  $h$  affects the accuracy of  $\pi_{\text{ABC}}$  in approximating  $\pi(\theta|\mathbf{s}_{\text{obs}})$ , but also the average acceptance probability (6), and hence the Monte Carlo error.

### 1.3. Approach and outline

Previous justification for ABC has, at least informally, been based around  $\pi_{\text{ABC}}$  approximating  $\pi(\theta|\mathbf{y}_{\text{obs}})$  globally. This is possible if  $\mathbf{y}_{\text{obs}}$  is low dimensional, and  $S(\cdot)$  is the identity function, or if we have low dimensional sufficient statistics for the data. In these cases we can control the error in the ABC approximation by choosing  $h$  sufficiently small. In general applications this is not so. Arguments have been made about trying to choose approximate sufficient statistics (see for example Joyce and Marjoram (2008)). However, the definition of such statistics is not clear, and more importantly it is difficult or impossible to construct a general method for finding such statistics.

We take a different approach and weaken the requirement for  $\pi_{\text{ABC}}$  to be a good approximation to  $\pi(\theta|\mathbf{y}_{\text{obs}})$ . We argue for  $\pi_{\text{ABC}}$  to be a good approximation solely in terms of the *accuracy of certain estimates* of the parameters. We also would like to know how to interpret the ABC posterior and probability statements that are derived from it. To do this we consider a property that we call *calibration*, which we define formally below. If  $\pi_{\text{ABC}}$  is calibrated, then this means that probability statements that are derived from it are appropriate, and in particular that we can use  $\pi_{\text{ABC}}$  to quantify uncertainty in estimates. We argue that using such criteria we can construct ABC posteriors which have both good inferential properties, and which can be estimated well by using Monte Carlo methods such as in algorithm 1.

In Section 2 we define formally what we mean by calibration and accuracy. Standard ABC is not calibrated, but a simple modification, which we call *noisy ABC*, is. We further show that if the bandwidth  $h$  is small then standard ABC is approximately calibrated. Theoretical results are used to produce recommendations about when standard ABC and when noisy ABC should be used. Then, for a certain definition of accuracy we show that the optimal choice of summary statistics  $S(\mathbf{Y})$  are the true posterior means of the parameters. Although these are unknown, simulation can be used to estimate the summary statistics. Our approach to doing this is described in Section 3, together with results that show the advantage of ABC over directly using the summary statistics to estimate parameters. Section 4 gives examples comparing our implementation of ABC with previous methods. The paper ends with a discussion.

## 2. Calibration and accuracy of approximate Bayesian calibration

First we define what we mean by calibration and introduce a version of ABC that is calibrated. We show how the idea of calibration can be particularly important when analysing multiple data sets. We then discuss our definition of accuracy of the ABC posterior, and how this can be used to guide the choice of summary statistic that we use.

### 2.1. Calibration and noisy approximate Bayesian calibration

Consider a subset of the parameter space  $\mathcal{A}$ . For given data  $\mathbf{y}_{\text{obs}}$ , the ABC posterior will assign a probability to the event  $\theta \in \mathcal{A}$

$$\Pr_{\text{ABC}}(\theta \in \mathcal{A} | \mathbf{s}_{\text{obs}}) = \int_{\mathcal{A}} \pi_{\text{ABC}}(\theta | \mathbf{s}_{\text{obs}}) d\theta.$$

For a given probability  $q$ , consider the event  $E_q(\mathcal{A})$  that  $\Pr_{\text{ABC}}(\theta \in \mathcal{A} | \mathbf{s}_{\text{obs}}) = q$ . Then the ABC posterior is calibrated if

$$\Pr\{\theta \in \mathcal{A} | E_q(\mathcal{A})\} = q. \quad (7)$$

The probability is then defined in terms of the density on parameters and data given by our prior and likelihood model,  $\pi(\theta) \pi(\mathbf{y} | \theta)$ , but ignoring any Monte Carlo randomness. Statement (7) states that, under repeated sampling from the prior, data and summary statistics, events assigned probability  $q$  by the ABC posterior will occur with probability  $q$ . A consequence of calibration is that the ABC posterior will appropriately represent uncertainty in the parameters: for example we can construct appropriate credible intervals, i.e. calibration means that we can use the ABC posterior as we would any standard posterior distribution.

Standard ABC posteriors (3) are not calibrated in general. Instead we introduce the idea of *noisy ABC* which is calibrated. Noisy ABC involves defining summary statistics which are random. A noisy ABC importance sampling algorithm is obtained by changing the initialization within algorithm 1 or algorithm 2 to ‘simulate  $\mathbf{x}$  from  $K(\mathbf{x})$  and define  $\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}}) + h(\mathbf{x})$ ’, giving algorithm 3. The resulting ABC posterior for a given  $\mathbf{y}_{\text{obs}}$  is random, and the definition of the probability in equation (7) needs to account for this extra randomness.

*Theorem 1.* Algorithm 1 produces an ABC posterior that is calibrated.

*Proof.* The ABC posterior that is derived by algorithm 3 is  $\pi_{\text{ABC}}(\theta | \mathbf{s}_{\text{obs}})$  where  $\mathbf{s}_{\text{obs}}$  is related to the data  $\mathbf{y}_{\text{obs}}$  by

$$\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}}) + h\mathbf{x}, \quad (8)$$

and  $\mathbf{x}$  is the realization of a random variable with density  $K(\mathbf{x})$ . However, the definition of  $\pi_{\text{ABC}}(\theta | \mathbf{s}_{\text{obs}})$  is just that of the true posterior for  $\theta$  given data  $\mathbf{s}_{\text{obs}}$  generated by equation (8). It immediately follows that this density is calibrated.  $\square$

A related idea was used by Wilkinson (2008) who showed that the ABC posterior is equivalent to the true posterior under an assumption of appropriate model error. In the limit as  $h \rightarrow 0$  noisy ABC is equivalent to standard ABC; we discuss the links between the two in more detail in Section 2.4.

### 2.2. Inference from multiple sources of data

Consider combining data from  $m$  independent sources,  $\mathbf{y}_{\text{obs}}^{(1)}, \dots, \mathbf{y}_{\text{obs}}^{(m)}$ . It is possible to use

individual ABC analyses for each data set, and then to combine these inferences. One sequential approach is to use the ABC posterior after analysing the  $i$ th data set as a prior for analysing the  $(i + 1)$ th data set. Algorithms for such an approach are a special case of those discussed in Wilkinson (2011).

One consequence of calibration is that such inferences will be well behaved in the limit as  $m$  becomes large. To see this define the ABC approximation to the  $i$ th data set as  $p(\theta|\mathbf{s}_{\text{obs}}^{(i)})$ , and note that the above approach is targeting the following ABC posterior:

$$\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}^{(1)}, \dots, \mathbf{s}_{\text{obs}}^{(m)}) \propto \pi(\theta) \prod_{i=1}^m p(\theta|\mathbf{s}_{\text{obs}}^{(i)}).$$

If we use noisy ABC, where  $\mathbf{s}_{\text{obs}}^{(i)}$  is randomly centred on  $S(\mathbf{y}_{\text{obs}}^{(i)})$ , then by the same argument as for theorem 1 this ABC posterior will be calibrated. Furthermore, we have the following result.

*Theorem 2.* Let  $\theta_0$  be the true parameter value. Consider noisy ABC, where  $\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}}) + h\mathbf{x}$ , where  $\mathbf{x}$  is drawn from  $K(\cdot)$ . Then the expected noisy ABC log-likelihood,

$$E[\log\{p(\theta|\mathbf{S}_{\text{obs}})\}] = \iint \log[p\{\theta|\mathbf{S}(\mathbf{y}) + h\mathbf{x}\}] \pi(\mathbf{y}|\theta_0) K(\mathbf{x}) d\mathbf{y} d\mathbf{x},$$

has its maximum at  $\theta = \theta_0$ .

*Proof.* Make the change of variable  $\mathbf{s}_{\text{obs}} = \mathbf{S}(\mathbf{y}) + h\mathbf{x}$ . Then

$$E[\log\{p(\theta|\mathbf{S}_{\text{obs}})\}] = \frac{1}{h} \iint \log\{p(\theta|\mathbf{s}_{\text{obs}})\} \pi(\mathbf{y}|\theta_0) \frac{K[\{\mathbf{s}_{\text{obs}} - \mathbf{S}(\mathbf{y})\}]}{h} d\mathbf{y} d\mathbf{s}_{\text{obs}}.$$

Now by definition  $p(\theta|\mathbf{s}_{\text{obs}}) = \int \pi(\mathbf{y}|\theta) K[\{\mathbf{s}_{\text{obs}} - \mathbf{S}(\mathbf{y})\}/h] d\mathbf{y}$ , and thus we obtain

$$E[\log\{p(\theta|\mathbf{S}_{\text{obs}})\}] = \frac{1}{h} \iint \log\{p(\theta|\mathbf{s}_{\text{obs}})\} p(\theta|\mathbf{s}_{\text{obs}}) d\mathbf{s}_{\text{obs}}.$$

By Jensen's inequality this has its maximum at  $\theta = \theta_0$ . □

The importance of this result is that, under the standard regularity conditions (Bernardo and Smith, 1994), the noisy ABC posterior will converge onto a point mass on the true parameter value as  $m \rightarrow \infty$ . By comparison, if we use standard ABC, then we have no such guarantee. As a simple example assume that  $Y^{(i)}$  are independent and identically distributed from a normal distribution with mean 0 and variance  $\sigma^2$ , and our kernel is chosen to be normal with variance  $\tau < \sigma^2$ . If we use standard ABC, the ABC posterior given  $y^{(1)}, \dots, y^{(m)}$  will converge to a point mass on  $\sigma^2 - \tau$  as  $m \rightarrow \infty$ .

We look at this issue empirically in Section 4.3.

### 2.3. Accuracy and choice of summary statistics

Calibration itself is not sufficient to define a sensible ABC posterior. For example the prior distribution is always calibrated but will not give accurate estimates of parameters. Thus we also want to maximize the accuracy of estimates based on the ABC posterior. We shall define accuracy in terms of a loss function for estimating the parameters. A natural choice of loss function is quadratic loss. Let  $\theta_0$  be the true parameter values, and  $\hat{\theta}$  an estimate. Then we shall consider the class of loss functions, which is defined in terms of a  $p \times p$  positive definite matrix  $A$ ,

$$L(\theta_0, \hat{\theta}; A) = (\theta_0 - \hat{\theta})^T A (\theta_0 - \hat{\theta}). \quad (9)$$

We now consider implementing ABC to minimize this quadratic error loss of estimating the parameters. We consider the limit of  $h \rightarrow 0$ . This will give results that define the optimal choice of summary statistics,  $S(\cdot)$ .

For any choice of weight matrix  $A$  that is of full rank, the following theorem shows that the optimal choice of summary statistics is  $S(\mathbf{y}_{\text{obs}}) = E(\theta|\mathbf{y}_{\text{obs}})$ , which is the true posterior mean.

*Theorem 3.* Consider a  $p \times p$  positive definite matrix  $A$  of full rank. Given observation  $\mathbf{y}_{\text{obs}}$ , let  $\Sigma$  be the true posterior variance for  $\theta$ .

- (a) The minimal possible quadratic error loss  $E\{L(\theta, \hat{\theta}; A)|\mathbf{y}_{\text{obs}}\}$  occurs when  $\hat{\theta} = E(\theta|\mathbf{y}_{\text{obs}})$  and is  $\text{tr}(A\Sigma)$ .
- (b) If  $S(\mathbf{y}_{\text{obs}}) = E(\theta|\mathbf{y}_{\text{obs}})$  then in the limit as  $h \rightarrow 0$  the minimum loss, based on inference using the ABC posterior, is achieved by  $\hat{\theta} = E_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}})$ . The resulting expected loss is  $\text{tr}(A\Sigma)$ .

*Proof.* Part (a) is a standard result of Bayesian statistics (Bernardo and Smith, 1994). For (b) we just need to show that, in the limit as  $h \rightarrow 0$ ,  $E_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) = E(\theta|\mathbf{y}_{\text{obs}})$ . By definition, in the limit as  $h \rightarrow 0$ ,  $\mathbf{s}_{\text{obs}} = S(\mathbf{y}_{\text{obs}})$  with probability 1, and  $\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) = \pi(\theta|\mathbf{s}_{\text{obs}})$ . Furthermore

$$\begin{aligned} E_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) &= \int \theta \pi(\theta|\mathbf{s}_{\text{obs}}) d\theta, \\ &= \iint \theta \pi(\theta|\mathbf{y}) \pi(\mathbf{y}|\mathbf{s}_{\text{obs}}) d\mathbf{y} d\theta, \end{aligned}$$

where  $\pi(\mathbf{y}|\mathbf{s}_{\text{obs}})$  is the conditional distribution of the data  $\mathbf{y}$  given the summary statistic  $\mathbf{s}_{\text{obs}}$ . Finally by definition, all  $\mathbf{y}$  that are consistent with  $\mathbf{s}_{\text{obs}}$  satisfy  $\int \theta \pi(\theta|\mathbf{y}) d\theta = E(\theta|\mathbf{y}_{\text{obs}})$ , and hence the result follows.  $\square$

Use of squared error loss leads to ABC approximations that attempt to have the same posterior mean as the true posterior. Using alternative loss functions would mean matching other features of the posterior: for example absolute error loss would result in matching the posterior medians. It is possible to choose other summary statistics that also achieve the minimum expected loss. However, any such statistic with dimension  $d > p$  will cause larger Monte Carlo error (see lemma 1 and the discussion below).

#### 2.4. Comparison of standard approximate Bayesian computation and noisy approximate Bayesian computation

Standard and noisy ABC are equivalent in the limit as  $h \rightarrow 0$ , with the ABC posteriors converging to  $E\{\theta|\mathbf{S}(\mathbf{y}_{\text{obs}})\}$ . We can further quantify the accuracy of estimates based on standard or noisy ABC for  $h \approx 0$ . For noisy ABC we have the following result.

*Theorem 4.* Assume condition (a) of lemma 1, that  $\pi\{E(\theta|\mathbf{y}_{\text{obs}})\} > 0$ , and the kernel  $K(\cdot)$  corresponds to a random variable with mean 0. If  $\mathbf{S}(\mathbf{y}_{\text{obs}}) = E(\theta|\mathbf{y}_{\text{obs}})$  then for small  $h$  the expected quadratic loss that is associated with  $\hat{\theta} = E_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}})$  is

$$E\{L(\theta, \hat{\theta}; A)|\mathbf{y}_{\text{obs}}\} = \text{tr}(A\Sigma) + h^2 \int \mathbf{x}^T A \mathbf{x} K(\mathbf{x}) d\mathbf{x} + o(h^2).$$

*Proof.* The idea is that  $E_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) = \mathbf{s}_{\text{obs}} + o(h)$ , and the squared error loss based on  $\hat{\theta} = \mathbf{s}_{\text{obs}}$  is just  $\text{tr}(A\Sigma) + h^2 \int \mathbf{x}^T A \mathbf{x} K(\mathbf{x}) d\mathbf{x}$ . See Appendix C.  $\square$

A similar result exists for standard ABC (Prangle, 2011), which shows that in this case

$$E\{L(\theta, \hat{\theta}; A) | \mathbf{y}_{\text{obs}}\} = \text{tr}(A\Sigma) + O(h^4).$$

Extensions of these results can be used to give guidance on the choice of kernel. For noisy ABC they suggest a uniform kernel on the ellipse  $\mathbf{x}^T A \mathbf{x} < c$ , for some  $c$ ; for standard ABC they also suggest a uniform kernel on an ellipse, but the form of the ellipse is difficult to calculate in practice. We do not give these results in more detail, as in practice we have found that the choice of kernel has relatively little effect on the accuracy of either ABC algorithm.

These two results also give an insight into the overall accuracy of a Monte Carlo ABC algorithm (see also Blum (2010)) by using the following informal argument. For simplicity consider a rejection sampling algorithm. The Monte Carlo variance is inversely proportional to the acceptance probability. Thus using lemma 1 we have that the Monte Carlo variance is  $O(N^{-1}h^{-d})$ , where  $N$  is the number of proposals. The expected quadratic loss based on estimates from the ABC importance sampler will be increased by this amount. Thus for noisy ABC we want to choose  $h$  to minimize

$$\text{tr}(A\Sigma) + h^2 \int \mathbf{x}^T A \mathbf{x} K(\mathbf{x}) d\mathbf{x} + \frac{C_0}{Nh^d},$$

for some constant  $C_0$ . This gives that we want  $h = O(N^{-1/(2+d)})$ , and the overall expected loss above  $\text{tr}(A\Sigma)$  would then decay as  $N^{-2/(2+d)}$ . For standard ABC a similar argument gives  $h = O(N^{-1/(4+d)})$ , and the overall expected loss would then decay as  $N^{-4/(4+d)}$ .

Thus we can see that the choice between using standard ABC and noisy ABC is a choice of a trade-off between accuracy and calibration. Noisy ABC is calibrated, but for small  $h$  will give less accurate estimates. As such for the analysis of a single data set where the number of summary statistics is not too large, and hence  $h$  is small, we recommend the use of standard ABC. If we wish to combine inferences from ABC analyses of multiple data sets, then, in the light of the discussion in Section 2, we recommend noisy ABC. For all the examples in Section 4 we found that this approach worked well in practice.

Possibly the best approach is to use noisy ABC, but to use Rao–Blackwellization ideas to average out the noise that is added to the summary statistics. Such an approach would have the guarantee that the resulting expected quadratic loss for estimating any function of the parameters would be smaller than that from noisy ABC. However, implementing such a Rao–Blackwellization scheme efficiently appears non-trivial, with the only simple approach being to run noisy ABC independently on the same data set, and then to average the estimates across each of these runs.

### 3. Semi-automatic approximate Bayesian computation

The above theory suggests that we wish to choose summary statistics that are equal to posterior means. Although we cannot use this result directly, as we cannot calculate the posterior means, we can use simulation to estimate appropriate summary statistics.

Our approach is

- (a) to use a pilot run of ABC to determine a region of non-negligible posterior mass,
- (b) to simulate sets of parameter values and data,
- (c) to use the simulated sets of parameter values and data to estimate the summary statistics and
- (d) to run ABC with this choice of summary statistics.

Step (a) of this algorithm is optional. Its aim is to help to define an appropriate training region

of parameter space from which we should simulate parameter values. In applications where our prior distributions are relatively informative, this step should be avoided as we can simulate parameter values from the prior in step (b). However, it is important if we have uninformative priors, particularly if they are improper.

If we implement step (a), we assume that we have arbitrarily chosen some summary statistics to use within ABC. In our implementation below we choose our training region as a hypercube, with the range for each parameter being the range of that parameter observed within our pilot run. Then in step (b) we simulate parameter values from the prior truncated to this training region, and for each choice of parameter value we simulate an artificial data set. We repeat this  $M$  times, so that we have  $M$  sets of parameter values, each with a corresponding simulated data set.

There are various approaches that we can take for step (c). In practice we found that using linear regression, with appropriate functions of the data as predictors, is both simple and worked well. We also considered using the lasso (Hastie *et al.*, 2001) and canonical correlation analysis (Mardia *et al.*, 1979) but in general neither of these performed better than linear regression (though the lasso may be appropriate if we wish to use a large number of explanatory variables within the linear model).

Our linear regression approach involved considering each parameter in turn. First we introduce a vector-valued function  $f(\cdot)$ , so that  $f(\mathbf{y})$  is a vector of, possibly non-linear, transformations of the data. The simplest choice is  $f(\mathbf{y}) = \mathbf{y}$ , but in practice including other or different transformations as well may be beneficial. For example, in one application below we found  $f(\mathbf{y}) = (\mathbf{y}, \mathbf{y}^2, \mathbf{y}^3, \mathbf{y}^4)$ , i.e. a vector of length  $4n$  that consists of the data plus all second, third and fourth powers of individual data points, produced a better set of summary statistics.

For the  $i$ th summary statistic the simulated values of the  $i$ th parameter,  $\theta_i^{(1)}, \dots, \theta_i^{(M)}$ , are used as the responses; and the transformations of the simulated data,  $f(\mathbf{y}^{(1)}), \dots, f(\mathbf{y}^{(M)})$ , are used as the explanatory variables. We then fit the model

$$\theta_i = E(\theta_i | \mathbf{y}) + \varepsilon_i = \beta_0^{(i)} + \beta^{(i)} f(\mathbf{y}) + \varepsilon_i,$$

where  $\varepsilon_i$  is some zero-mean noise, using least squares. The fitted function  $\hat{\beta}_0^{(i)} + \hat{\beta}^{(i)} f(\mathbf{y})$  is then an estimate of  $E(\theta_i | \mathbf{y})$ . The constant terms can be neglected in practice as ABC uses only the difference in summary statistics. Thus the  $i$ th summary statistic for ABC is just  $\hat{\beta}^{(i)} f(\mathbf{y})$ .

Our approach of using a training region means that our models for the posterior means are based on only parameter values that are simulated within this region. We therefore suggest adapting the ABC run in step (d) so that the prior is truncated to lie within this training region (a similar idea was used in Blum and François (2010)). This can be viewed as using, weakly, the information that we have from the pilot ABC run within the final ABC run and has links with composite likelihood methods (Lindsay, 1988). More importantly it makes the overall algorithm robust to problems where  $E\{\hat{\beta}^{(i)} f(\mathbf{Y}) | \theta_i\}$  is similar for two dissimilar values of  $\theta_i$ : one inside the training region and one outside.

In practice below we use roughly a quarter of our total central processor unit time on steps (a) and (b) and half on step (d), with step (c) having negligible central processor unit cost. We call this semi-automatic ABC as the choice of summary statistics is now based on simulation, but there are still choices by the user in terms of fitting the linear model in step (c). This input is in terms of the choice of  $f(\mathbf{y})$  to be used. Step (c) is now a familiar statistical problem, and standard model checks can be used to decide whether that choice of  $f(\mathbf{y})$  is appropriate and, if not, how it could be improved. Also, repeating step (c) with a different choice of  $f(\mathbf{y})$  can be done without any further simulation of data and thus is quick in terms of the central processor unit cost. Furthermore standard model comparison procedures (e.g. using the Bayesian

information criterion BIC) can be used to choose between summary statistics that are obtained from linear regressions using different explanatory variables.

A natural question is whether this approach is better than the current approach to ABC, where summary statistics are chosen arbitrarily. In our implementation we still need to choose summary statistics for step (a), and we also need to choose the set of explanatory variables for the linear model. Thus it could be argued that all we have done is to replace one arbitrary choice with another. However, we believe that our approach is more robust to these choices than standard ABC is. First the choice of summary statistics in step (a) is purely to make step (b) more efficient, and as such the final results depend little on this choice. Secondly, when we choose the explanatory variables we can choose many such variables (of the order of hundreds). As such we are much more likely to include among these some variables which are informative about the parameters of interest than standard ABC is where generally a few summary statistics are used. If many summary statistics are used in ABC, then this will require a large value of  $h$  and will often be inefficient because the accept–reject decision is based not only on the informative summary statistics, but also those which are less informative. These issues are demonstrated empirically in the examples that we consider.

Our approach has similarities to that of Beaumont *et al.* (2002) (see also Blum and François (2010)), who used linear regression to correct the output from ABC. The key difference is that our approach uses linear regression to construct the summary statistics, whereas Beaumont *et al.* (2002) used linear regression to reduce the error between  $\pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}})$  and  $\pi(\theta|\mathbf{s}_{\text{obs}})$ . In particular the method of Beaumont *et al.* (2002) assumes that appropriate low dimensional summary statistics have already been chosen. We look at differences between our approach and that of Beaumont *et al.* (2002) empirically in the examples.

### 3.1. Why use approximate Bayesian computation?

Our approach involves using simulation to find estimates of the posterior mean of each parameter. A natural question is why not use these estimates directly? We think that using ABC has two important advantages over just using these estimates directly. The first is that ABC gives you a posterior distribution, and thus you can quantify uncertainty in the parameters as well as obtain point estimates.

Moreover we have the following result.

*Theorem 5.* Let  $\tilde{\theta} = E\{\theta|S(\mathbf{y}_{\text{obs}})\}$ . Then, for any function  $g$ ,

$$E\{L(\theta, \tilde{\theta}; A)|S(\mathbf{y}_{\text{obs}})\} \leq E\{L[\theta, g\{S(\mathbf{y}_{\text{obs}})\}; A]|S(\mathbf{y}_{\text{obs}})\}.$$

Furthermore, asymptotically as  $h \rightarrow 0$  the ABC posterior mean estimate of  $\theta$  is optimal among estimates based on  $S(\mathbf{y}_{\text{obs}})$ .

*Proof.* The proof of the first part is the standard argument that the mean is the optimal estimator under quadratic loss (Bernardo and Smith, 1994). The second part follows because as  $h \rightarrow 0$  the ABC posterior mean tends to  $\tilde{\theta}$ .  $\square$

Note that this result states that, in the limit as  $h \rightarrow 0$ , ABC gives estimates that are at least as accurate as or more accurate than any other estimators based on the same summary statistics.

#### 3.1.1. Comparison with indirect inference

Indirect inference (Gourieroux and Ronchetti, 1993) is a method that is similar to ABC in that it uses simulation from a model to produce estimates of the model's parameters. The general

1 procedure involves first analysing the data under an approximating model and estimating the  
 2 parameters, called auxiliary parameters, for this model. Then data are simulated for a range  
 3 of parameter values, and for each simulated data set we obtain an estimate of the auxiliary  
 4 parameters. Finally we estimate the true parameters on the basis of which parameter values  
 5 produced estimates of the auxiliary parameters that are closest to those estimated from the  
 6 true data. (In practice we simulate multiple data sets for each parameter value and obtain an  
 7 estimate of the auxiliary parameters on the basis of these multiple data sets.) The link to ABC  
 8 is that the auxiliary parameters in indirect inference are equivalent to the summary statistics  
 9 in ABC. Both methods then use (different) simulation approaches to produce estimates of the  
 10 true parameters from the values of the auxiliary parameters (or summary statistics) for the real  
 11 data.

12 For many approximating models, the auxiliary parameters depend on a small set of summary  
 13 statistics of the data; these were called auxiliary statistics in Heggland Frigessi (2004). In these  
 14 cases indirect inference is performing inference based on these auxiliary statistics. The above  
 15 result shows that, in the limit as  $h \rightarrow 0$ , ABC will be more accurate than an indirect inference  
 16 method whose auxiliary statistics are the same as the summary statistic that is used for ABC.  
 17 We investigate this empirically in Section 4.

## 18 4. Examples

19 The performance of semi-automatic ABC was investigated in a range of examples: independent  
 20 draws from a complex distribution (Section 4.2), a stochastic kinetic network for biochemical  
 21 reactions (Section 4.3), a partially observed  $M/G/1$ -queue (Section 4.5), an ecological popula-  
 22 tion size model (Section 4.4) and a model for the transmission of tuberculosis (Section 4.6).  
 23 Section 4.1 describes implementation details that are common to all examples. Section 4.3 con-  
 24 cerns a data set for which ABC and related methods have not previously been used and highlights  
 25 the use of noisy ABC in the sequential approach of Section 2.2. The other examples have previous  
 26 analyses in the literature, and we show that semi-automatic ABC compares favourably against  
 27 existing methods including indirect inference, the synthetic likelihood method of Wood (2010)  
 28 and ABC with *ad hoc* summary statistics, with or without the regression correction method of  
 29 Beaumont *et al.* (2002). We also show that direct use of the linear predictors that are created  
 30 during semi-automatic ABC can be inaccurate (e.g. Section 4.6). Apart from in Section 4.3 noisy  
 31 ABC runs are not shown; they are similar to non-noisy semi-automatic ABC but slightly less  
 32 accurate. The practical details of implementing our method are also explored: in particular how  
 33 the choice of explanatory variables  $f(\mathbf{y})$  is made.

### 34 4.1. Implementation details

35 Apart from the sequential implementation of ABC in Section 4.3, all ABC analyses were per-  
 36 formed by using algorithm 2 with a normal transition kernel. The density kernel was uniform  
 37 on an ellipsoid  $\mathbf{x}^T \mathbf{A} \mathbf{x} < c$ . This is a common choice in the literature and is close to optimal  
 38 for runs using semi-automatic ABC summary statistics as discussed in Section 2.4. For sum-  
 39 mary statistics that are not generated by our method we generally used  $A = I$ . In Section 4.4 a  
 40 different choice was necessary and is discussed there. For ABC using summary statistics from  
 41 our method, recall from Section 2.3 that  $A$  defines the relative weighting of the parameters  
 42 in our loss function. In Sections 4.2 and 4.6 the parameters are on similar scales so we used  
 43  $A = I$ . Elsewhere, marginal parameter variances were calculated for the output of each pilot run  
 44 and the means of these ( $s_1^2, s_2^2, \dots$ ) taken. A diagonal  $A$ -matrix was formed with  $i$ th diagonal  
 45 entry  $s_i^{-2}$ .

Other tuning details that are required by algorithm 2 are the choice of  $h$ , the variance matrix of the transition kernel and the starting values of the chain. Where possible, these were chosen by manual experimentation or based on previous analyses (e.g. from pilot runs). Otherwise they were based on a very short ABC rejection sampling analysis. Except where noted otherwise,  $h$  was tuned to give an acceptance rate of roughly 1% as this gave reasonable results in the applications that were considered. An alternative would be to use computational methods that try to choose  $h$  for each run; see Bortot *et al.* (2007) and Ratmann *et al.* (2007).

In the following examples our method is compared with an ABC analysis with summary statistics based on the existing literature, which is referred to as the ‘comparison’ analysis. To allow a fair comparison, this uses the same number of simulations as the entire semi-automatic method and a lower acceptance rate: roughly 0.5%.

For simulation studies on multiple data sets, the accuracies of the various analyses were compared as follows. The point estimate for each data set was calculated, and the quadratic loss (9) of each parameter estimate relative to the true parameter value was calculated. We present the mean quadratic losses of the individual parameters. In tables of results we highlight the smaller quadratic losses (all within 10% of the smallest values) by italicizing.

#### 4.2. Inference for $g$ -and- $k$ -distribution

The  $g$ -and- $k$ -distribution is a flexibly shaped distribution that is used to model non-standard data through a small number of parameters (Haynes, 1998). It is defined by its inverse distribution function (10), below, but has no closed form density. Likelihoods can be evaluated numerically but this is costly (Rayner and MacGillivray, 2002; Drovandi and Pettitt, 2009). ABC methods are attractive because simulation is straightforward by the inversion method. Here we use the fact that we can calculate the maximum likelihood estimate numerically to compare ABC with a full-likelihood analysis also.

The distribution is defined by

$$F^{-1}(x; A, B, c, g, k) = A + B \left[ 1 + c \frac{1 - \exp\{-gz(x)\}}{1 + \exp\{-gz(x)\}} \right] \{1 + z(x)^2\}^k z(x) \quad (10)$$

where  $z(x)$  is the  $x$ th standard normal quantile,  $A$  and  $B$  are location and scale parameters and  $g$  and  $k$  are related to skewness and kurtosis. The final parameter  $c$  is typically fixed as 0.8, and this is assumed throughout, leaving unknown parameters  $\theta = (A, B, g, k)$ . The only parameter restrictions are  $B > 0$  and  $k > -\frac{1}{2}$  (Rayner and MacGillivray, 2002).

Allingham *et al.* (2009) used ABC to analyse a simulated data set of  $n = 10^4$  independent draws from the  $g$ -and- $k$ -distribution with parameters  $\theta_0 = (3, 1, 2, 0.5)$ . A uniform prior on  $[0, 10]^4$  was used and the summary statistics were the full set of order statistics. We studied multiple data sets of a similar form as detailed below. Our aim is first to show how we can implement semi-automatic ABC in a situation where there are large numbers of possible explanatory variables (just using the order statistics gives  $10^4$  explanatory variables), and to see how the accuracy of semi-automatic ABC compares with the use of arbitrarily chosen summary statistics in Allingham *et al.* (2009). We also aim to look at comparing semi-automatic ABC with the linear regression correction of Beaumont *et al.* (2002) and with indirect inference.

##### 4.2.1. Comparison of approximate Bayesian computation methods

The natural choice of explanatory variables for this problem is based on the order statistics and also powers of the order statistics. Considering up to the fourth power seems appropriate as informally the four parameters are linked to location, scale, skewness and kurtosis. However,

fitting the linear model with the resulting  $4 \times 10^4$  explanatory variables is impracticable. As a result we considered using a subset of  $m$  evenly spaced order statistics, together with up to  $l$  powers of this subset. To choose appropriate values for  $m$  and  $l$  we fitted linear models with  $m$  ranging over a grid of values between 60 and 140, and  $l$  ranging between 1 and 4, and we used BIC (averaged across the models for the four parameter values) to choose an appropriate value for  $m$  and  $l$ . We then used the summary statistics that are obtained from the linear model with this value of  $m$  and  $l$  in the final run of ABC. For simplicity we did this for the first data set and kept the same value of  $m$  and  $l$  for analysing all subsequent data sets.

Using subsets of order statistics has computational advantages, as these can be generated efficiently by simulating corresponding standard uniform order statistics by using the exponential spacings method of Ripley (1987) (page 98) and performing inversion by substituting these in equation (10). The cost is linear in the number of order statistics required. Our pilot ABC run used the summary statistics from Allingham *et al.* (2009). Fitting the different linear models added little to the overall computational cost, which is dominated by the simulation of the data sets at the various stages of the procedure.

Our semi-automatic ABC procedure chose  $m = 100$  order statistics and  $l = 4$ . The accuracy of the resulting parameter estimates, measured by squared error loss across implementation of semi-automatic ABC on 50 data sets, is shown in Table 3. For comparison we show results of the ABC method of Allingham *et al.* (2009), implemented to have the same overall computational cost. We also show the accuracy of estimates that were obtained by post-processing the results of Allingham *et al.* (2009) by using the regression correction of Beaumont *et al.* (2002). We could use this regression correction on only 48 of the 50 data sets, as, for the remaining two, there were too few acceptances in the ABC run for the regression correction to be stable (on those two data sets the resulting loss after performing the regression correction was orders of magnitude greater than the original ABC estimates).

Although the analysis of Allingham *et al.* (2009) performed poorly, and hence produced a poor pilot region for semi-automatic ABC, semi-automatic ABC appears to perform well. It has losses that are between a factor of 2 and 100 smaller than for the method of Allingham *et al.* (2009) with or without the regression correction. Using the regression correction does

**Table 3.** Mean quadratic losses of various ABC analyses of 50  $g$ -and- $k$  data sets with parameters (3,1,2,0.5)<sup>†</sup>

<i>Method</i>	<i>A</i>	<i>B</i>	<i>g</i>	<i>k</i>
Allingham <i>et al.</i> (2009)	0.0059	0.0013	3.85	0.00063
Allingham + regression	0.00040	0.0017	0.28	0.00051
Semi-automatic ABC	<i>0.00016</i>	<i>0.00056</i>	0.044	0.00023
Comparison	0.00025	0.00063	0.0061	0.00041
Comparison + regression	<i>0.00016</i>	<i>0.00055</i>	<i>0.0014</i>	<i>0.00015</i>
Semi-automatic ABC	<i>0.00015</i>	<i>0.00053</i>	<i>0.0014</i>	<i>0.00015</i>
Maximum likelihood estimation	0.00016	0.00055	0.0013	0.00014

<sup>†</sup>The first three rows are based on using the summary statistics of Allingham *et al.* (2009) in ABC, and in the ABC pilot run for semi-automatic ABC. The next three rows use just 100 evenly spaced order statistics. Results based on using the regression correction of Beaumont *et al.* (2002) are denoted ‘regression’. For ‘Allingham + regression’ we give the mean loss for just 48 of the 50 data sets. For the remaining two data sets the number of ABC acceptances was low (about 200), and the regression correction was unstable. For comparison we give the mean quadratic loss of the true maximum likelihood estimates. Mean quadratic losses within 10% of the minimum value for that parameter are italicized.

improve the accuracy of the estimates for three of the four parameters, but to a lesser extent than semi-automatic ABC. For comparison, using the predictors from the linear regression to estimate the parameters directly had similar accuracy for all parameters except  $g$ , where the linear predictor's average error was greater by about a third.

To investigate the effect of the pilot run on semi-automatic ABC, and the choice of summary statistics on the regression correction, we repeated this analysis by implementing ABC with 100 order statistics. This greatly improved the performance of ABC, showing the importance of the choice of summary statistics. The improved pilot run also improves semi-automatic ABC but to a much lesser extent. In this case semi-automatic ABC has similar accuracy to that of the comparison ABC run with the regression correction. For a further comparison we also calculated the maximum likelihood estimates for each data set numerically. The two best ABC runs have mean quadratic loss that is almost identical to that of the maximum likelihood estimates.

#### 4.2.2. Comparison with indirect inference

Theorem 5 shows that asymptotically ABC is at least as accurate as other estimators based on the same summary statistics. We tested this with a comparison against indirect inference (see Section 3.1). The semi-automatic ABC analysis was repeated, and, to give a direct comparison, its summary statistics were used as the indirect inference auxiliary statistics.

Initial analysis showed that which method is more accurate depends on the true value of  $\theta$  and in particular the parameter  $g$ ; this is illustrated by Fig. 1. Therefore we studied data sets that were produced from varying  $g$ -values; we drew 50  $g$ -values from its prior and for each simulated data sets of  $n$   $g$ -and- $k$ -draws conditional on  $\theta = (3, 1, g, 0.5)$  for  $n = 10^2, 10^3, 10^4$ .

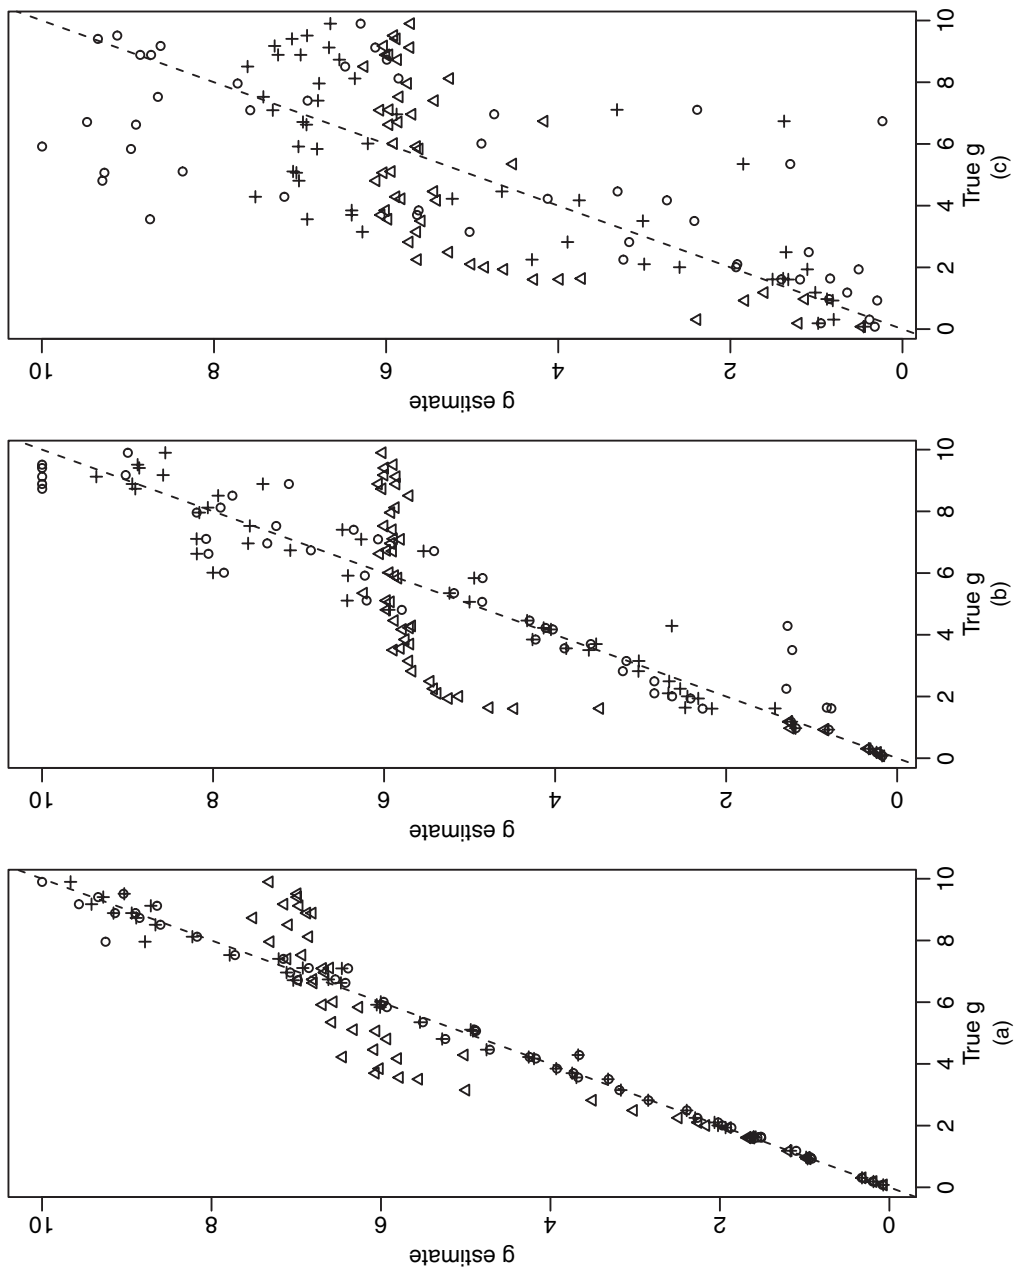
Each semi-automatic ABC analysis used a total of  $3.1 \times 10^6$  simulated data sets. Indirect inference was roughly tuned so that the total number of simulations equalled this. Similar results can be obtained from indirect inference using many fewer simulations, and indirect inference is thus a computationally quicker algorithm. Mean losses are given in Table 4, showing that, although the methods perform similarly for  $n = 10^4$ , ABC is more accurate for smaller  $n$ .

More detail is given in Fig. 1 which plots the true against estimated  $g$ -values. Of particular interest is the case  $n = 100$  where the  $g$ -parameter is very difficult to identify for  $g > 3$ . It is over this range that ABC outperforms indirect inference most clearly, with estimates from indirect inference being substantially more variable than those for ABC.

**Table 4.** Mean quadratic losses of semi-automatic ABC and indirect inference analyses of 50  $g$ -and- $k$  data sets with variable  $g$ -parameters<sup>†</sup>

$n$	Method	$A$	$B$	$g$	$k$
$10^4$	Pilot	0.0003	0.0008	1.7	0.0004
	Indirect inference	0.0003	0.0022	0.082	0.0063
	Semi-automatic ABC	<i>0.0001</i>	<i>0.0005</i>	<i>0.059</i>	<i>0.0002</i>
$10^3$	Pilot	0.0031	0.014	4.6	0.0073
	Indirect inference	0.0066	0.014	0.83	0.0053
	Semi-automatic ABC	<i>0.0012</i>	<i>0.0094</i>	<i>0.51</i>	<i>0.0042</i>
$10^2$	Pilot	0.0089	0.039	4.8	0.057
	Indirect inference	0.018	0.059	5.5	0.067
	Semi-automatic ABC	<i>0.0075</i>	<i>0.046</i>	<i>3.5</i>	<i>0.040</i>

<sup>†</sup>The smallest losses for each parameter and sample size are italicized.



**Fig. 1.** Estimated  $g$ -values from indirect inference (○) and semi-automatic ABC pilot (△) and final (×) analyses, plotted against true  $g$  values for 50  $g$ -and- $k$  data sets of sample size  $n$ : (a)  $n = 10^4$ ; (b)  $n = 10^3$ ; (c)  $n = 10^2$

### 4.3. Inference for stochastic kinetic networks

Stochastic kinetic networks are used to model biochemical networks. The state of the network is determined by the number of each of a discrete set of molecules and evolves stochastically through reactions between these molecules. See Wilkinson (2009) and references therein for further background.

Inference for these models is challenging, as the transition density of the model is intractable. However, simulation from the models is possible, e.g. by using the algorithm of Gillespie (Gillespie, 1977). As such they are natural applications for ABC methods. Here we focus on a simple example of a stochastic kinetic network: the Lotka–Volterra model of Boys *et al.* (2008). Although the model is simple, Boys *et al.* (2008) show the difficulty of full likelihood inference.

This model has a state which consists of two molecules. Denote the state at time  $t$  by  $\mathbf{Y}_t = (Y_t^{(1)}, Y_t^{(2)})$ . There are three types of reaction: the birth of a molecule of type 1, the death of a molecule of type 2 and a reaction between the two molecules which removes a type 1 molecule and adds a type 2 molecule. (The network is called the Lotka–Volterra model, because of its link with predator–prey models, type 1 molecules being prey and type 2 being predators.)

The dynamics for the model are Markov, and can be specified in terms of transition over a small time interval  $\delta t$ . For positive parameters  $\theta_1$ ,  $\theta_2$  and  $\theta_3$ , we have

$$\Pr\{\mathbf{Y}_{t+\delta t} = (z_1, z_2) | \mathbf{Y}_t = (y_1, y_2)\} = \begin{cases} 1 - (\theta_1 y_1 + \theta_2 y_1 y_2 + \theta_3 y_2) \delta t + o(\delta t) & \text{if } z_1 = y_1 \text{ and } z_2 = y_2, \\ \theta_1 y_1 \delta t + o(\delta t) & \text{if } z_1 = y_1 + 1 \text{ and } z_2 = y_2, \\ \theta_2 y_1 y_2 \delta t + o(\delta t) & \text{if } z_1 = y_1 - 1 \text{ and } z_2 = y_2 + 1, \\ \theta_3 y_2 \delta t + o(\delta t) & \text{if } z_1 = y_1 \text{ and } z_2 = y_2 - 1, \\ o(\delta t) & \text{otherwise.} \end{cases}$$

We shall focus on the case of the network being fully observed at discrete time points, and also on just observing the type 2 molecules initially, together with the type 1 molecules at all observation points. All simulations use the parameters from Boys *et al.* (2008), with  $\theta_1 = 0.5$ ,  $\theta_2 = 0.0025$  and  $\theta_3 = 0.3$ , and evenly sampled data collected at time intervals of length  $\tau$ . We shall perform analysis conditional on the known initial state of the system.

#### 4.3.1. Sequential approximate Bayesian computation analysis

Wilkinson (2011) considers simulation-based approaches for analysing stochastic kinetic

**Table 5.** Algorithm 4: a sequential ABC sampler for the Lotka–Volterra model

<p><i>Input</i>—a set of times <math>t_0, \dots, t_n</math> and data values <math>\mathbf{y}_{t_0}, \dots, \mathbf{y}_{t_n}</math>;  a number of particles, <math>N</math>, a kernel <math>K(\cdot)</math> and a bandwidth <math>h</math></p> <p><i>Initialize</i>—for <math>i = 1, \dots, N</math> sample <math>\theta^{(i)} = (\theta_1^{(i)}, \theta_2^{(i)}, \theta_3^{(i)})</math> from the prior distribution for the parameters</p> <p><i>Iterate</i>—for <math>j = 1, 2, \dots, n</math>:</p> <p style="padding-left: 2em;">step 1, for <math>i = 1, \dots, N</math>, sample a value for the state at time <math>t_j</math>, <math>\mathbf{y}_{t_j}^{(i)}</math>, given its value at time <math>t_{j-1}</math>, <math>\mathbf{y}_{t_{j-1}}</math>, and <math>\theta^{(i)}</math> by using the Gillespie algorithm;</p> <p style="padding-left: 2em;">step 2, for <math>i = 1, \dots, N</math>, calculate weights</p> $w^{(i)} = K[\{\mathbf{y}_{t_j}^{(i)} - \mathbf{y}_{t_j}\}/h];$ <p style="padding-left: 2em;">step 3, sample <math>N</math> times from a kernel density approximation to a weighted sample of <math>\theta</math>-values, <math>\{\theta^{(i)}, w_i\}_{i=1}^N</math> (see for example Liu and West (2001)); denote this sample <math>\theta^{(1)}, \dots, \theta^{(N)}</math></p> <p><i>Output</i>—a sample of <math>\theta</math>-values.</p>
---

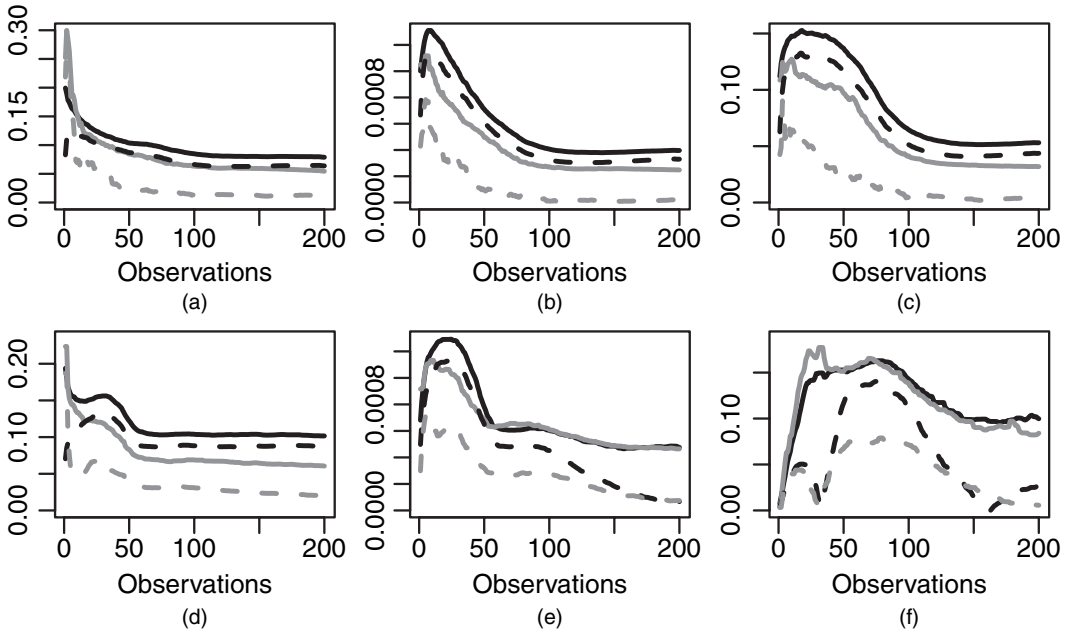
networks which are based on sequential Monte Carlo methods (see Doucet *et al.* (2000) for an introduction). In some applications, to get these methods to work for reasonable computational cost, Wilkinson (2011) suggests using ABC. A version of such an algorithm (though based on importance sampling rather than MCMC sampling for the sequential update) for the Lotka–Volterra model is given in algorithm 4 (Table 5). There are many approaches to improve the computational efficiency of this algorithm; see for example Doucet *et al.* (2000, 2001) for details.

The discussion following theorem 2 shows that an algorithm like algorithm 4 may give inconsistent parameter estimates, even when ignoring the Monte Carlo error. A simple remedy to this is to implement noisy ABC within this algorithm. This can be done by adding noise to the observed values. The noisy sequential ABC algorithm is given by algorithm 5, which differs from algorithm 4 by replacing step 2 with

step 2, simulate  $\mathbf{x}_j$  from  $K(\mathbf{x})$ ; for  $i = 1, \dots, N$ , calculate weights

$$w^{(i)} = K[\{\mathbf{y}_{t_j}^{(i)} - \mathbf{y}_{t_j} - h\mathbf{x}_j\}/h]$$

To evaluate the relative merits of the two sequential ABC algorithms, we analysed 100 simulated data sets for the Lotka–Volterra model. For stability of the sequential Monte Carlo algorithm we chose  $K(\cdot)$  to be the density function of a bivariate normal random variable, as a uniform kernel can lead to iterations where all weights are 0. We analysed two data scenarios, with  $\tau = 0.1$ : one with full observations, and one where only the number of type 1 molecules is observed. The sequential ABC algorithms were implemented with  $N = 5000$  and  $h = \sqrt{\tau}$ . The latter was chosen to be a small value for which the sequential algorithms still performed adequately in terms of Monte Carlo performance (as measured by variability of the weights after each iteration).



**Fig. 2.** Plots of root-square error loss (—) and absolute bias (---) for the standard sequential ABC algorithm (—, ---) and the noisy ABC version (—, ---) as a function of the number of observations ( $\tau = 0.1$ ), (a)–(c) observing the number of both molecules and (d)–(f) for observing only the number of type 1 molecules: (a), (d)  $\theta_1$ ; (b), (e)  $\theta_2$ ; (c), (f)  $\theta_3$

Results are shown in Fig. 2. These show both the absolute bias and the root mean quadratic loss for estimating each parameter for both ABC algorithms as the number of observations analysed varies between 1 and 200. For the full observation case, we can see evidence of bias in all three parameters for the standard sequential ABC algorithm. Noisy ABC shows evidence of being asymptotically unbiased as the number of observations increases. Overall, noisy ABC appears more accurate, except perhaps if only a handful of observations are made. When just type 1 molecules are observed, the picture is slightly more complex. We observed evidence of bias for the standard ABC algorithm for  $\theta_1$ , and for this parameter noisy ABC is more accurate. For the other parameters, there appears little difference between the accuracy and bias of the two ABC algorithms. The difference between the parameters is likely to be because  $\theta_1$  affects only the dynamics of the observed molecule.

#### 4.4. Inference for Ricker model

The Ricker map is an example of an ecological model with complex dynamics. It updates a population size  $N_t$  over a time step by

$$N_{t+1} = rN_t \exp(-N_t + e_t)$$

where the  $e_t$  are independent  $N(0, \sigma_e^2)$  noise terms. Wood (2010) studied a model in which Poisson observations  $y_t$  are made with means  $\phi N_t$ . The parameters of interest are  $\theta = (\log(r), \sigma_e, \phi)$ . The initial state is  $N_0 = 1$  and the observations are  $y_{51}, y_{52}, \dots, y_{100}$ . This is a complex inference scenario in which there is no obvious natural choice of summary statistics. Furthermore Wood (2010) argued that estimating parameters by maximum likelihood is difficult for this model owing to the chaotic nature of the system.

Wood (2010) analysed this model by using a Metropolis–Hastings algorithm to explore a ‘synthetic likelihood’,  $L_s(\theta)$ . A summary statistic function must be specified as in ABC, and  $L_s(\theta)$  is then the density of  $\mathbf{s}_{\text{obs}}$  under an  $N(\mu_\theta, \Sigma_\theta)$  model where  $\mu_\theta$  and  $\Sigma_\theta$  are an approximation of the mean and variance of  $\mathbf{s}_{\text{obs}}$  for the given parameter value obtained through simulation (see Wood (2010) for more details).

We simulated 50 data sets with  $\log(r) = 3.8$ ,  $\phi = 10$  and  $\log(\sigma_e)$ -values drawn from a uniform distribution on  $[\log(0.1), 0]$ . (Initial analysis showed that the true  $\log(\sigma_e)$ -value affected the performance of the various methods.) These were analysed under our approach and that of Wood (2010) (using the code in that paper’s supplementary material). The distribution just mentioned was used as prior for  $\log(\sigma_e)$ , and improper uniform priors were placed on  $\log(r) \geq 0$  and  $\phi$ . All parameters were assumed independent under the prior. As each parameter had a non-negativity constraint, the MCMC algorithms used log-transforms of  $\theta$  as the state. Each semi-automatic ABC analysis used  $10^6$  simulated data sets.

The summary statistics that were used by Wood (2010) were the autocovariances to lag 5, the coefficients of a cubic regression of the ordered differences  $\Delta_t = y_t - y_{t-1}$  on those of the observed data, least squares estimates for the model  $y_{t+1}^{0.3} = \beta_1 y_t^{0.3} + \beta_2 y_t^{0.6} + \varepsilon_t$ , the mean observation  $\bar{y}$  and  $\sum_{t=51}^{100} \mathbb{1}(y_t = 0)$  (the number of zero observations). We denote this set by  $E_0$  and use it in the semi-automatic ABC pilot runs. The pilot runs used acceptance kernel  $A = \Sigma^{-1}$  where  $\Sigma$  is the sample variance matrix of 500 simulated summary statistics vectors for a representative fixed parameter value.

In the regression stage of semi-automatic ABC, training data sets mostly consisting of 0s were fitted poorly. Since the observed data sets had at most 31 0s, any data sets with 45 or more 0s were discarded from our training data, and automatically rejected in the subsequent ABC runs. Summary statistics were constructed for two nested sets of explanatory variables. The smaller

**Table 6.** Mean quadratic losses for various analyses of 50 simulated Ricker data sets†

<i>Method</i>	$\log(r)$	$\sigma_e$	$\phi$
Synthetic likelihood	0.050	<i>0.032</i>	0.66
Comparison	<i>0.039</i>	0.038	0.54
Comparison + regression	0.046	0.041	0.78
Semi-automatic ABC	<i>0.039</i>	<i>0.032</i>	<i>0.36</i>

†Losses within 10% of the smallest values for that parameter are italicized.

set, E1, included E0 and additionally  $\sum_{t=51}^{100} \mathbb{1}(y_t = j)$  for  $1 \leq j \leq 4$ ,  $\log(\bar{y})$ , the logarithm of the sample variance,  $\log(\sum_{i=51}^{100} y_i^j)$  for  $2 \leq j \leq 6$  and auto-correlations up to lag 5. The larger set, E2, also added  $(y_t)_{51 \leq t \leq 100}$  (time-ordered observations),  $(y_{(t)})_{1 \leq t \leq 50}$  (magnitude-ordered observations),  $(y_t^2)_{51 \leq t \leq 100}$ ,  $(y_{(t)}^2)_{1 \leq t \leq 50}$ ,  $\{\log(1 + y_t)\}_{51 \leq t \leq 100}$ ,  $\{\log(1 + y_{(t)})\}_{1 \leq t \leq 50}$ ,  $(\Delta_t^2)_{52 \leq t \leq 100}$  and  $\{\Delta_{(t)}^2\}_{1 \leq t \leq 49}$ .

Using set E2 instead of E1 reduced BIC in each linear regression by the order of thousands for all except one data set, suggesting that E2 gives better predictors. Thus we used the summary statistics based on set E2 within the ABC analysis (Table 6). The results for these ABC analyses show an improvement over the synthetic likelihood for estimating  $\log(r)$  and  $\phi$ , and identical performance for estimating  $\sigma_e$ . The semi-automatic ABC analysis also does better than the comparison ABC analysis (based on summary statistics  $E_0$ ). For this application the linear regression adjustment of Beaumont *et al.* (2002) actually produces worse results than using the raw output of the comparison ABC analyses.

Finally we looked at 95% credible intervals that were constructed from the semi-automatic ABC and synthetic likelihood method. The coverage frequencies of these intervals were 0.86, 0.70 and 0.96 for synthetic likelihood and 0.98, 0.92 and 1 for our method. Whereas the synthetic likelihood intervals appear to have coverage frequencies that are too low for two of the parameters, those from ABC are consistent with 0.95 coverage given a sample size of 50 data sets.

#### 4.5. Inference for *M/G/1*-queue

Queuing models are an example of stochastic models which are easy to simulate from but often have intractable likelihoods. It has been suggested to analyse such models by using simulation-based procedures, and we shall look at a specific *M/G/1*-queue that has been analysed by both ABC (Blum and François, 2010) and indirect inference (Heggland and Frigessi, 2004) before. In this model, the service times are uniformly distributed in the interval  $[\theta_1, \theta_2]$  and inter-arrival times are exponentially distributed with rate  $\theta_3$ . The queue is initially empty and only the interdeparture times  $y_1, y_2, \dots, y_{50}$  are observed.

We analysed 50 simulated data sets from this model. The true parameters were drawn from the prior under which  $(\theta_1, \theta_2 - \theta_1, \theta_3)$  are uniformly distributed on  $[0, 10]^2 \times [0, \frac{1}{3}]$ . This choice gives arrival and service times of similar magnitudes, avoiding the less interesting situation where all  $y_i$ -values are independent draws from a single distribution.

The analysis of Blum and François (2010) used as summary statistics evenly spaced quantiles of the interdeparture times, including the minimum and maximum. Our semi-automatic ABC pilot analyses replicate this choice, using 20 quantiles. The explanatory variables  $f(\mathbf{y})$  that we used to construct summary statistics were the ordered interdeparture times. Adding powers of

**Table 7.** Mean quadratic losses for various analyses of 50  $M/G/1$  data sets<sup>†</sup>

<i>Method</i>	$\theta_1$	$\theta_2$	$\theta_3$
Comparison	1.1	2.2	0.0013
Comparison + regression	<i>0.020</i>	1.1	<i>0.0013</i>
Semi-automatic ABC	<i>0.022</i>	1.0	<i>0.0013</i>
Semi-automatic predictors	0.024	1.2	0.0017
Indirect inference	0.18	<i>0.42</i>	0.0033

<sup>†</sup>Losses within 10% of the smallest values for that parameter are italicized.

these values to  $f(\mathbf{y})$  produced only minor improvements so the results are not reported. The analysis of each data set used  $10^7$  simulated data sets, split in the usual way.

We also applied the indirect inference approach of Heggland and Frigessi (2004). This used auxiliary statistics  $(\bar{y}, \min(y_i), \hat{\theta}_2^{\text{ML}})$  where  $\hat{\theta}_2^{\text{ML}}$  is the maximum likelihood estimate of  $\theta_2$  under an auxiliary model which has a closed form likelihood, namely that corresponding to independent observations from the steady state of the queue. Numerical calculation of  $\hat{\theta}_2^{\text{ML}}$  is expensive so indirect inference used many fewer simulated data sets than ABC but had similar run times.

Table 7 shows the results. Semi-automatic ABC outperforms a comparison analysis using 20 quantiles as the summary statistics, but once a regression correction has been applied to the latter the results become very similar. Here the semi-automatic linear predictors are less accurate when used directly rather than in ABC. Indirect inference is more accurate at estimating  $\theta_2$ , presumably because of the accuracy of  $\hat{\theta}_2^{\text{ML}}$  as an estimate of  $\theta_2$ . However, it is still substantially less accurate for the other two parameters. One advantage of indirect inference is that, as it requires fewer simulations to estimate the parameters accurately, it can more easily accommodate summaries that are expensive to calculate, such as  $\hat{\theta}_2^{\text{ML}}$ .

#### 4.6. Inference of transmission of tuberculosis



Tanaka *et al.* (2006) used ABC to analyse tuberculosis bacteria genotype data sampled in San Francisco over a period from 1991 to 1992. Table 8 shows the data, consisting of 473 bacteria samples split into clusters which share the same genotype on a particular genetic marker. Thus the data consist of 282 bacteria samples that had unique genotypes, 20 pairs of bacteria that had the same genotype, and so on.

The model proposed was based on an underlying continuous time Markov process. Denote the total number of cases at time  $t$  by  $N(t)$ . The process starts at  $t=0$  with  $N(0)=1$ . There are three types of event: birth, death (encompassing recovery of the host) and mutation. The rate of each type of event is the product of  $N(t)$  and the appropriate parameter:  $\alpha$  for birth,  $\delta$  for death and  $\theta$  for mutation. It was assumed that each mutation creates a completely new genotype. Cluster data are a simple random sample of 473 cases taken at the first  $t$  such that  $N(t)=10000$ . The model conditions on such a  $t$  existing in the underlying process.

**Table 8.** Tuberculosis bacteria genotype data

Cluster size	1	2	3	4	5	8	10	15	23	30
Number of clusters	282	20	13	4	2	1	1	1	1	1

1 These data contain no information on time, so, for  $k > 0$ , parameter values  $(\alpha, \delta, \theta)$  and  
 2  $(k\alpha, k\delta, k\theta)$  give the same likelihood. We reparameterize to  $(a, d, \theta)$  where  $a = \alpha/(\alpha + \delta + \theta)$  and  
 3  $d = \delta/(\alpha + \delta + \theta)$ . The likelihood under this parameterization depends only on  $a$  and  $d$ . To reflect  
 4 prior ignorance of  $(a, d)$  we use the prior density  $\pi(a, d, \theta) \propto \pi(\theta) \mathbb{I}(0 \leq d \leq a) \mathbb{I}(a + d < 1)$ , where  
 5  $\pi(\theta)$  is the marginal prior for  $\theta$  that was used in Tanaka *et al.* (2006). The prior restriction  $d \leq a$   
 6 avoids the need for simulations in which  $N(t) = 10000$  is highly unlikely to occur. The other



**Fig. 3.** ABC output for the tuberculosis application (every 1000th state is plotted): (a) comparison; (b) semi-automatic ABC

1 prior restrictions follow from positivity constraints on the original parameters. Under this prior  
 2 and parameterization, the marginal posterior of  $\theta$  is equal to its prior and the problem reduces  
 3 to inference on  $a$  and  $d$ .

4 Tanaka *et al.* (2006) used two summary statistics for their ABC analysis:  $g/473$  and  $H =$   
 5  $1 - \sum_i (n_i/473)^2$ , where  $g$  is the number of distinct clusters in the sample, and  $n_i$  is the number  
 6 of observed samples in the  $i$ th genotype cluster. We retain this choice for our semi-automatic  
 7 ABC pilot and the comparison ABC analysis reported below.

8 As parameters in the pilot output are highly correlated (Fig. 3), we fitted a line to the output  
 9 by linear regression and made a reparameterization,  $(u \ v)^T = M(a \ d)^T$ , where  $M$  is a rotation  
 10 matrix chosen so that on the fitted line  $u$  is constant. The semi-automatic ABC analysis was  
 11 continued using  $(u, v)$  as the parameters of interest. The explanatory variables  $f(\mathbf{y})$  comprised  
 12 the number of clusters of size  $i$  for  $1 \leq i \leq 5$ , the number of clusters of size above 5, the average  
 13 cluster size  $H$ , the size of the three largest clusters and the squares of all these quantities. The  
 14 semi-automatic ABC analysis used  $4 \times 10^6$  simulated data sets in total.

15 Fig. 3 shows the ABC posteriors, indicating that our methodology places less weight on  
 16 the high  $d$  tail of the ABC posterior, reducing the marginal variances from (0.0029, 0.0088)  
 17 (comparison) to (0.0017, 0.0048) (semi-automatic ABC).

18 We further investigated this model through a simulation study. We constructed 50 data sets by  
 19 simulating parameters from the ABC posterior, and simulating data from the model for each of  
 20 these pairs. We then compared ABC with the summary statistics of Tanaka *et al.* (2006) to semi-  
 21 automatic ABC. Averaged across these data sets, semi-automatic ABC reduced the squared  
 22 error loss by around 20–25% for the two parameters. Here we also found that the regression  
 23 correction of Beaumont *et al.* (2002) did not change the accuracy of the method of Tanaka *et al.*  
 24 (2006). Using the linear predictors obtained within semi-automatic ABC to estimate parameters  
 25 (rather than as summary statistics) gave substantially worse estimates, with mean-squared error  
 26 for  $a$  increasing by a factor of over 3.

## 27 28 29 5. Discussion

30 We have argued that ABC can be justified by aiming for calibration of the ABC posterior,  
 31 together with accuracy of estimates of the parameters. We have introduced a new variant of  
 32 ABC, called noisy ABC, which is calibrated. Standard ABC can be justified as an approxima-  
 33 tion to noisy ABC, but one which can be more accurate. Theoretical results suggest that, when  
 34 a single data set is analysed by using a relatively small number of summary statistics, standard  
 35 ABC should be preferred.

36 However, we show that, when we attempt to combine ABC analyses of multiple data sets,  
 37 noisy ABC is to be preferred. Empirical evidence for this comes from our analysis of a stochastic  
 38 kinetic network, where using standard ABC within a sequential ABC algorithm leads to biased  
 39 parameter estimates. This result could be important for ABC analyses of population genetic  
 40 data, if inferences are combined across multiple genomic regions. We believe that the ideal  
 41 approach would be to implement a Rao–Blackwellized version of noisy ABC, where we attempt  
 42 to average out the noise that is added to the summary statistics. If implemented, this would  
 43 lead to a method which is more accurate than noisy ABC for estimating any function of the  
 44 parameters. However, at present we are unsure how to implement such a Rao–Blackwellization  
 45 scheme in a computationally efficient manner.

46 The main focus of this paper was a semi-automatic approach to implementing ABC. The  
 47 main idea is to use simulation to construct appropriate summary statistics, with these sum-  
 48 mary statistics being estimates of the posterior mean of the parameters. This approach is based

on theoretical results which show that choosing summary statistics as the posterior means produces ABC estimators that are optimal in terms of minimizing quadratic loss. We have evaluated our method on several models, comparing both with ABC as it has been implemented in the literature, with indirect inference, and the synthetic likelihood approach of Wood (2010).

The most interesting comparison was between semi-automatic ABC and ABC using the regression correction of Beaumont *et al.* (2002). In several examples, these two approaches gave similarly accurate estimates. However, the semi-automatic ABC seemed to be more robust, with the regression correction actually producing less accurate estimates for the Ricker model, and not improving the accuracy for the tuberculosis model.

There are alternatives to using linear regression to construct the summary statistics, with many dimension reduction techniques being possible (see, for example, Wegmann *et al.* (2009) and Bazin *et al.* (2010)). However, one particular approach motivated by our theory is to use approximate estimates for each parameter. Such an approach has been used in Wilson *et al.* (2009), where estimates of the recombination rate under the incorrect demographic model were used as summary statistics in ABC. It is also the idea behind the approach of Drovandi *et al.* (2011), who summarized the data through parameter estimates under an approximating model. A further alternative is to use sliced inverse regression (Li, 1991) rather than linear regression to produce the summary statistics. The advantage of sliced inverse regression is that it can, where appropriate, generate multiple linear combinations of the explanatory variables, such that the posterior mean is approximately a function of these linear combinations. Each linear combination could then be used as a summary statistic within ABC.

Finally we note that there has been much recent research into the computational algorithms underpinning ABC. As well as the rejection sampling, importance sampling and MCMC algorithms that we have mentioned, there has been work on using sequential Monte Carlo methods (Beaumont *et al.*, 2009; Sisson *et al.*, 2007; Peters *et al.*, 2010) and adaptive methods (Bortot *et al.*, 2007; Del Moral *et al.*, 2009) among others. The ideas in this paper are focusing a separate issue within ABC, and we think that the semi-automatic approach for implementing ABC that we describe can be utilized within whatever computational method is preferred.

## Acknowledgements

We thank Chris Sherlock for many helpful discussions, and the reviewers for their detailed comments.

## Appendix A: Variance of importance sampling–approximate Bayesian computation and Markov chain Monte Carlo–approximate Bayesian computation

Calculating the variance of estimates of posterior means by using importance sampling is complicated by the dependence between the importance sampling weights and values of  $h(\theta)$  (see Liu (1996)). A common approach to quantifying the accuracy of importance sampling is to use an effective sample size. If the weights are normalized to have mean 1, then the effective sample size  $N_{\text{eff}}$  is  $N$  divided by the mean of the square of the weights. Liu (1996) argued that, for most functions  $h(\theta)$ , the variance of the estimator in expression (4) will be approximately expression (5) but with  $N_{\text{acc}}$  replaced by  $N_{\text{eff}}$ .

For a given proposal distribution  $g(\theta)$ , the effective sample size is

$$N_{\text{eff}} = N \frac{\int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) d\theta}{E_{\text{ABC}}\{\pi(\theta)/g(\theta)\}}. \quad (11)$$

It can be shown that the optimal proposal distribution, in terms of maximizing  $N_{\text{eff}}/N$ , is

$$g_{\text{opt}}(\theta|\mathbf{y}_{\text{obs}}) \propto \pi(\theta) p(\theta|\mathbf{s}_{\text{obs}})^{1/2},$$

in which case

$$N_{\text{eff}}^* = N_{\text{acc}} \left[ \frac{1 + \text{var}_{\pi} \{ p(\theta|\mathbf{s}_{\text{obs}})^{1/2} \}}{E_{\pi} \{ p(\theta|\mathbf{s}_{\text{obs}})^{1/2} \}^2} \right],$$

where the variance and expectation on the right-hand side are with respect to  $\pi(\theta)$ . It is immediate that  $N_{\text{eff}}^* \geq N_{\text{acc}}$ , with equality only if  $p(\theta|\mathbf{s}_{\text{obs}})$  does not depend on  $\theta$ . The potential gains of importance sampling occur when  $p(\theta|\mathbf{s}_{\text{obs}})$  varies greatly.

Analysis of the Monte Carlo error within the MCMC ABC algorithm is more difficult. However, consider fixing a proposal kernel  $g(\cdot)$ , which will fix the type of transitions attempted. The Monte Carlo error then will be primarily governed by the average acceptance probability. For simplicity assume that  $K(\cdot)$  is a uniform kernel and that either  $g(\cdot)$  is chosen to have the prior  $\pi(\theta)$  as its stationary distribution or the term  $\pi(\theta) g(\theta_{i-1}|\theta)/\pi(\theta_{i-1}) g(\theta|\theta_{i-1}) \approx 1$  and can be ignored. The average acceptance probability at stationarity is

$$\iint \pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) g(\theta'|\theta) p(\theta'|\mathbf{s}_{\text{obs}}) d\theta d\theta' = \int \pi(\theta) p(\theta|\mathbf{s}_{\text{obs}}) d\theta \iint \pi_{\text{ABC}}(\theta|\mathbf{s}_{\text{obs}}) g(\theta'|\theta) \frac{\pi_{\text{ABC}}(\theta'|\mathbf{s}_{\text{obs}})}{\pi(\theta')} d\theta d\theta'.$$

The integral comes from averaging over the current and proposed values for the MCMC algorithm, with the average acceptance probability for a given proposed value  $\theta'$  being  $p(\theta'|\mathbf{s}_{\text{obs}})$ . The right-hand side comes from using expression (1). The first term on the right-hand side is the average acceptance probability of the rejection algorithm. The second term is 1 if we use an independence sampler  $g(\theta'|\theta)$  and will be much greater than 1 if the ABC posterior is peaked relative to the prior, and if the transition kernel proposes localized moves.

## Appendix B: Proof of lemma 1

Write

$$\begin{aligned} \int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) d\theta &= \iint K\{(\mathbf{s} - \mathbf{s}_{\text{obs}})/h\} \pi(\mathbf{s}|\theta) \pi(\theta) d\theta d\mathbf{s} \\ &= \int h^d K(\mathbf{x}) \pi(\mathbf{s}_{\text{obs}} + h\mathbf{x}) d\mathbf{x}. \end{aligned}$$

The first equality comes from the definition of  $p(\theta|\mathbf{s}_{\text{obs}})$ ; the second by integrating out  $\theta$  and making a change of variable  $\mathbf{x} = (\mathbf{s} - \mathbf{s}_{\text{obs}})/h$ . So

$$\left| h^{-d} \int p(\theta|\mathbf{s}_{\text{obs}}) \pi(\theta) d\theta - \pi(\mathbf{s}_{\text{obs}}) \right| \leq \int K(\mathbf{x}) |\pi(\mathbf{s}_{\text{obs}} + h\mathbf{x}) - \pi(\mathbf{s}_{\text{obs}})| d\mathbf{x}.$$

This bound is shown to be  $o(1)$  under either condition.

Consider first condition (a). Define  $c$  to be the maximum value of  $|\mathbf{x}|$  such that  $K(|\mathbf{x}|) > 0$ . For any  $\varepsilon > 0$ , by continuity of  $\pi(\mathbf{s})$  at  $\mathbf{s}_{\text{obs}}$  we have that there is a  $\delta > 0$  such that  $|\mathbf{x}| < \delta$  implies  $|\pi(\mathbf{s}_{\text{obs}} + \mathbf{x}) - \pi(\mathbf{s}_{\text{obs}})| < \varepsilon$ . Define  $h_{\varepsilon} = \delta/c$ . Then for  $h < h_{\varepsilon}$  we have

$$\int K(\mathbf{x}) |\pi(\mathbf{s}_{\text{obs}} + h\mathbf{x}) - \pi(\mathbf{s}_{\text{obs}})| d\mathbf{x} \leq \varepsilon.$$

This inequality follows as  $h|\mathbf{x}| < \delta$  for all  $\mathbf{x}$  where  $K(\mathbf{x}) > 0$ .

Now consider condition (b). By differentiable continuity, Taylor's theorem gives

$$\pi(\mathbf{s}_{\text{obs}} + h\mathbf{x}) = \pi(\mathbf{s}_{\text{obs}}) + \sum_i h x_i r_i(\mathbf{x}).$$

The remainder factor  $r_i(\mathbf{x})$  is  $|\partial\pi(\mathbf{z})/\partial s_i|$  for some  $\mathbf{z}(\mathbf{x})$ , so, by assumption,  $|r_i(\mathbf{x})| \leq R$ , a finite bound. Thus

$$\int K(\mathbf{x}) |\pi(\mathbf{s}_{\text{obs}} + h\mathbf{x}) - \pi(\mathbf{s}_{\text{obs}})| d\mathbf{x} \leq hR \sum_i \int |x_i| K(\mathbf{x}) d\mathbf{x}.$$

**Appendix C: Proof of theorem 4**

Rearrangement of the loss function gives

$$E\{L(\theta, \hat{\theta}; A) | \mathbf{y}_{\text{obs}}\} - \text{tr}(A\Sigma) = E\{(\tilde{\theta} - \hat{\theta})^T A (\tilde{\theta} - \hat{\theta}) | \mathbf{y}_{\text{obs}}\},$$

where  $\tilde{\theta} = E(\theta | \mathbf{y}_{\text{obs}})$ , the mean under the true posterior, which equals  $S(\mathbf{y}_{\text{obs}})$  for the summary statistics under discussion. Define  $\delta(\mathbf{s}_{\text{obs}}) = \hat{\theta}(\mathbf{s}_{\text{obs}}) - \mathbf{s}_{\text{obs}}$  and make the change of variables  $\mathbf{x} = (\tilde{\theta} - \mathbf{s}_{\text{obs}})/h$ . Now

$$\begin{aligned} E\{L(\theta, \hat{\theta}; A) | \mathbf{y}_{\text{obs}}\} - \text{tr}(A\Sigma) - h^2 \int \mathbf{x}^T A \mathbf{x} K(\mathbf{x}) d\mathbf{x} &= -2h \int \mathbf{x}^T A \delta(\tilde{\theta} + h\mathbf{x}) K(\mathbf{x}) d\mathbf{x} \\ &\quad + \int \delta(\tilde{\theta} + h\mathbf{x})^T A \delta(\tilde{\theta} + h\mathbf{x}) K(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

It is required that the modulus of the right-hand side be  $o(h^2)$ . Let  $R$  be the support of  $K$ . Since this is finite, it suffices to show that

$$\max_{\mathbf{x} \in R} \{\delta(\tilde{\theta} + h\mathbf{x})\} = o(h).$$

To find an expression for  $\delta$ , observe that condition (1) holds for noisy ABC. Taking its expectation and making the change of variable  $\mathbf{s} = S(\mathbf{y})$  give

$$\hat{\theta}(\mathbf{t}) = \frac{\int \theta \pi(\theta) \pi(\mathbf{s} | \theta) K\{(\mathbf{s} - \mathbf{t})/h\} d\theta d\mathbf{s}}{\int \pi(\mathbf{s}) K\{(\mathbf{s} - \mathbf{t})/h\} d\mathbf{s}}.$$

Note that

$$\int \theta \pi(\theta) \pi(\mathbf{s} | \theta) d\theta = \pi(\mathbf{s}) E(\theta | \mathbf{s}) = \mathbf{s} \pi(\mathbf{s}),$$

where the second equality is due to our choice of  $S$ . Thus

$$\delta(\mathbf{t}) = \frac{\int (\mathbf{s} - \mathbf{t}) \pi(\mathbf{s}) K\{(\mathbf{s} - \mathbf{t})/h\} d\mathbf{s}}{\int \pi(\mathbf{s}) K\{(\mathbf{s} - \mathbf{t})/h\} d\mathbf{s}}.$$

Make the change of variables  $\mathbf{y} = (\mathbf{s} - \mathbf{t})/h$  and consider the case  $\mathbf{t} = \tilde{\theta} + h\mathbf{x}$ . Then

$$\delta(\tilde{\theta} + h\mathbf{x}) = \frac{h \int \mathbf{y} \pi\{\tilde{\theta} + h(\mathbf{x} + \mathbf{y})\} K(\mathbf{y}) d\mathbf{y}}{\int \pi\{\tilde{\theta} + h(\mathbf{x} + \mathbf{y})\} K(\mathbf{y}) d\mathbf{y}}.$$

By the argument in Appendix B, continuity of  $\pi(\mathbf{s})$  at  $\mathbf{s} = \tilde{\theta}$  gives denominator  $\pi(\tilde{\theta}) + o(1)$ . Consider the  $i$ th component of the integral in the numerator,

$$\left| \int y_i \pi\{\tilde{\theta} + h(\mathbf{x} + \mathbf{y})\} K(\mathbf{y}) d\mathbf{y} - \int y_i \pi(\tilde{\theta}) K(\mathbf{y}) d\mathbf{y} \right| \leq \max_{\mathbf{y} \in R} |y_i| \left| \int [\pi\{\tilde{\theta} + h(\mathbf{x} + \mathbf{y})\} - \pi(\tilde{\theta})] K(\mathbf{y}) d\mathbf{y} \right|.$$

This integral is  $o(h)$  by the continuity argument just mentioned. Noting that  $\int y_i K(\mathbf{y}) d\mathbf{y} = 0$  by assumption, we have

$$\int \mathbf{y} \pi\{\tilde{\theta} + h(\mathbf{x} + \mathbf{y})\} K(\mathbf{y}) d\mathbf{y} = o(1).$$

Combining these results gives the required bound for  $\delta$ .

## References

- Allingham, D., King, R. A. R. and Mengersen, K. L. (2009) Bayesian estimation of quantile distributions. *Statist. Comput.*, **19**, no. 2.
- Bazin, E., Dawson, K. J. and Beaumont, M. A. (2010) Likelihood-free inference of population structure and local adaptation in a bayesian hierarchical model. *Genetics*, **185**, 587–602.
- Beaumont, M. A., Cornuet, J.-M., Marin, J.-M. and Robert, C. P. (2009) Adaptive approximate Bayesian computation. *Biometrika*, **96**, 983–990.
- Beaumont, M. A., Zhang, W. and Balding, D. J. (2002) Approximate Bayesian computation in population genetics. *Genetics*, **162**, 2025–2035.
- Bernardo, J. M. and Smith, A. F. M. (1994) *Bayesian Theory*. Chichester: Wiley.
- Blum, M. G. B. (2010) Approximate bayesian computation: a nonparametric perspective. *J. Am. Statist. Ass.*, **105**, 1178–1187.
- Blum, M. G. B. and François, O. (2010) Non-linear regression models for Approximate Bayesian Computation. *Statist. Comput.*, **20**, 63–73.
- Bortot, P., Coles, S. and Sisson, S. (2007) Inference for stereological extremes. *J. Am. Statist. Ass.*, **102**, 84–92.
- Boys, R. J., Wilkinson, D. J. and Kirkwood, T. B. L. (2008) Bayesian inference for a discretely observed stochastic kinetic model. *Statist. Comput.*, **18**, 125–135.
- Cornuet, J.-M., Santos, F., Beaumont, M. A., Robert, C. P., Marin, J.-M., Balding, D. J., Guillemaud, T. and Estoup, A. (2008) Inferring population history with DIY ABC: a user-friendly approach to approximate Bayesian computation. *Bioinformatics*, **24**, 2713–2719.
- Del Moral, P., Doucet, A. and Jasra, A. (2009) An adaptive sequential Monte Carlo method for approximate Bayesian computation. To be published.
- Diggle, P. J. and Gratton, R. J. (1984) Monte Carlo methods of inference for implicit statistical models (with discussion). *J. R. Statist. Soc. B*, **46**, 193–227.
- Doucet, A., de Freitas, J. F. G. and Gordon, N. J. (eds) (2001) *Sequential Monte Carlo Methods in Practice*. New York: Springer.
- Doucet, A., Godsill, S. J. and Andrieu, C. (2000) On sequential Monte Carlo sampling methods for Bayesian filtering. *Statist. Comput.*, **10**, 197–208.
- Drovandi, C. C. and Pettitt, A. N. (2009) Likelihood-free Bayesian estimation of quantile distributions. *Technical Report*. Queensland University of Technology, Brisbane.
- Drovandi, C. C., Pettitt, A. N. and Faddy, M. J. (2011) Approximate Bayesian computation using indirect inference. *Appl. Statist.*, **60**, 317–337.
- Gillespie, D. T. (1977) Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, **81**, 2340–2361.
- Gouriéroux, C. and Ronchetti, E. (1993) Indirect inference. *J. Appl. Econometr.*, **8**, s85–s118.
- Grelaud, A., Robert, C., Marin, J. M., Rodolphe, F. and Taly, J. F. (2009) ABC likelihood-free methods for model choice in Gibbs random fields. *Baysn Anal.*, **4**, 317–336.
- Hastie, T., Tibshirani, R. and Friedman, J. (2001) *The Elements of Statistical Learning, Data Mining, Inference, and Prediction*. New York: Springer.
- Haynes, M. (1998) Flexible distributions and statistical models in ranking and selection procedures, with applications. *PhD Thesis*. Queensland University of Technology, Brisbane.
- Heggland, K. and Frigessi, A. (2004) Estimating functions in indirect inference. *J. R. Statist. Soc. B*, **66**, 447–462.
- Joyce, P. and Marjoram, P. (2008) Approximately sufficient statistics and Bayesian computation. *Statist. Applic. Genet. Molec. Biol.*, **7**, article 26.
- Kingman, J. F. C. (1982) The coalescent. *Stoch. Process. Applic.*, **13**, 235–248.
- Li, K.-C. (1991) Sliced inverse regression for dimension reduction. *J. Am. Statist. Ass.*, **86**, 316–327.
- Liepe, J., Barnes, C., Cule, E., Erguler, K., Kirk, P., Toni, T. and Stumpf, M. P. (2010) ABC-SysBio approximate Bayesian computation in Python with GPU support. *Bioinformatics*, **26**, 1797–1799.
- Lindsay, B. G. (1988) Composite likelihood methods. *Contemp. Math.*, **80**, 221–239.
- Liu, J. S. (1996) Metropolis independent sampling with comparisons to rejection sampling and importance sampling. *Statist. Comput.*, **6**, 113–119.
- Liu, J. and West, M. (2001) Combined parameter and state estimation in simulation based filtering. In *Sequential Monte Carlo in Practice* (eds A. Doucet, J. F. G. de Freitas and N. J. Gordon), pp. 197–223, New York: Springer.
- Lopes, J. S., Balding, D. and Beaumont, M. A. (2009) PopABC, a program to infer historical demographic parameters. *Bioinformatics*, **25**, 2747–2749.
- Mardia, K. V., Kent, J. T. and Bibby, J. M. (1979) *Multivariate Analysis*. London: Academic Press.
- Marjoram, P., Molitor, J., Plagnol, V. and Tavaré, S. (2003) Markov chain Monte Carlo without likelihoods. *Proc. Natn. Acad. Sci. USA*, **100**, 15324–15328.
- McKinley, T., Cook, A. R. and Deardon, R. (2009) Inference in epidemic models without likelihoods. *Int. J. Biostatist.*, **5**, 24.
- Padhukasahasram, B., Wall, J. D., Marjoram, P. and Nordborg, M. (2006) Estimating recombination rates from single-nucleotide polymorphisms using summary statistics. *Genetics*, **174**, 1517–1528.

- 1 Peters, G. W., Fan, Y. and Sisson, S. A. (2010) On sequential Monte Carlo, partial rejection control and approxi-  
2 mate Bayesian computation. To be published.
- 3 Prangle, D. (2011) Summary statistics and sequential methods for approximate Bayesian computation. *PhD*  
4 *Thesis*. Lancaster University, Lancaster.
- 5 Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A. and Feldman, M. W. (1999) Population growth of human Y  
6 chromosomes, a study of Y chromosome microsatellites. *Molec. Biol. Evoln*, **16**, 1791–1798.
- 7 Ratmann, O., Andrieu, C., Wiuf, C. and Richardson, S. (2009) Model criticism based on likelihood-free inference,  
8 with an application to protein network evolution. *Proc. Natn. Acad. Sci. USA*, **106**, 10576.
- 9 Ratmann, O., Jorgensen, O., Hinkley, T., Stumpf, M., Richardson, S. and Wiuf, C. (2007) Using likelihood-free  
10 inference to compare evolutionary dynamics of the protein networks of *h. pylori* and *p. falciparum*. *PLoS*  
11 *Comput. Biol.*, **3**, article e230.
- 12 Rayner, G. D. and MacGillivray, H. L. (2002) Numerical maximum likelihood estimation for the g-and-k and  
13 generalized g-and-h distributions. *Statist. Comput.*, **12**, 57–75.
- 14 Ripley, B. (1987) *Stochastic Simulation*. Chichester: Wiley.
- 15 Sisson, S. A., Fan, Y. and Tanaka, M. M. (2007) Sequential Monte Carlo without likelihoods. *Proc. Natn. Acad.*  
16 *Sci. USA*, **104**, 1760–1765; correction; **106** (2009), article 16889.
- 17 Sisson, S. A., Peters, G. W., Briers, M. and Fan, Y. (2010) A note on target distribution ambiguity of likelihood-free  
18 samplers. *Preprint*.
- 19 Tanaka, M. M., Francis, A. R., Luciani, F. and Sisson, S. A. (2006) Using approximate Bayesian computation to  
20 estimate tuberculosis transmission parameters from genotype data. *Genetics*, **173**, 1511–1520.
- 21 Tavaré, S., Balding, D. J., Griffiths, R. C. and Donnelly, P. (1997) Inferring coalescent times from DNA sequence  
22 data. *Genetics*, **145**, 505–518.
- 23 Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Stumpf, M. (2009) Approximate Bayesian computation scheme  
24 for parameter inference and model selection in dynamical systems. *J. R. Soc. Interface*, **6**, 187–202.
- 25 Wegmann, D., Leuenberger, C. and Excoffier, L. (2009) Efficient approximate Bayesian computation coupled  
26 with Markov chain Monte Carlo without likelihood. *Genetics*, **182**, 1207–1218.
- 27 Wilkinson, D. J. (2009) Stochastic modelling for quantitative description of heterogeneous biological systems.  
28 *Nat. Rev. Genet.*, **10**, 122–133.
- 29 Wilkinson, D. J. (2011) Parameter inference for stochastic kinetic models of bacterial gene regulation, a Bayesian  
30 approach to systems biology. In *Bayesian Statistics 9* (eds J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P.  
31 Dawid, D. Heckerman, A. F. M. Smith and M. West), pp. 679–690. Oxford: Oxford University Press.
- 32 Wilkinson, R. D. (2008) Approximate Bayesian computation (ABC) gives exact results under the assumption of  
33 model error. *Arxiv Preprint* arXiv:0811.3355v1.
- 34 Wilson, D. J., Gabriel, E., Leatherbarrow, A. J. H., Cheesebrough, J., Gee, S., Bolton, E., Fox, A., Hart, C. A.,  
35 Diggle, P. J. and Fearnhead, P. (2009) Rapid evolution and the importance of recombination to the gastro-enteric  
36 pathogen *campylobacter jejuni*. *Molec. Biol. Evoln*, **26**, 385–397.
- 37 Wood, S. (2010) Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, **466**, 1102–1104.
- 38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48