



Institute for
New Economic Thinking
AT THE OXFORD MARTIN SCHOOL

Black-box Bayesian inference for economic agent-based models

Joel Dyer, Patrick Cannon, J. Doyne Farmer & Sebastian Schmon

1st February 2022

INET Oxford Working Paper No. 2022-05



Preprint

Black-box Bayesian inference for economic agent-based models

JOEL DYER^{1,2}, PATRICK CANNON³, J. DOYNE FARMER^{1,2,4}, and SEBASTIAN SCHMON^{3,5}

¹Institute for New Economic Thinking, Oxford

²Mathematical Institute, University of Oxford

³Improbable

⁴Santa Fe Institute

⁵Durham University

February 1, 2022

Abstract

Simulation models, in particular agent-based models, are gaining popularity in economics. The considerable flexibility they offer, as well as their capacity to reproduce a variety of empirically observed behaviors of complex systems, give them broad appeal, and the increasing availability of cheap computing power has made their use feasible. Yet a widespread adoption in real-world modelling and decision-making scenarios has been hindered by the difficulty of performing parameter estimation for such models. In general, simulation models lack a tractable likelihood function, which precludes a straightforward application of standard statistical inference techniques. A number of recent works (Grazzini et al., 2017; Platt, 2020, 2021) have sought to address this problem through the application of *likelihood-free* inference techniques, in which parameter estimates are determined by performing some form of comparison between the observed data and simulation output. However, these approaches are (a) founded on restrictive assumptions, and/or (b) typically require many hundreds of thousands of simulations. These qualities make them unsuitable for large-scale simulations in economics and can cast doubt on the validity of these inference methods in such scenarios. In this paper, we investigate the efficacy of two classes of simulation-efficient black-box approximate Bayesian inference methods that have recently drawn significant attention within the probabilistic machine learning community: neural posterior estimation and neural density ratio estimation. We present a number of benchmarking experiments in which we demonstrate that neural network based black-box methods provide state of the art parameter inference for economic simulation models, and crucially are compatible with generic multivariate time-series data. In addition, we suggest appropriate assessment criteria for use in future benchmarking of approximate Bayesian inference procedures for economic simulation models.

1. Introduction

Simulation models are increasingly employed across the sciences. Their primary advantage is the flexibility they afford modellers: models may be specified mechanistically and at the microscopic level without concern for analytic tractability of the full system behaviour. This flexibility is embraced in the *agent-based modelling* paradigm wherein individual agents can be described by realistic and heterogeneous decision and interaction rules. Examining the macroscopic properties of the system is then done simply through simulation, freeing the modeller from the burden of

reasoning about the global behaviour of the system themselves and facilitating the study of radically more complex systems and emergent phenomena than is otherwise possible.

In general, agent-based models (ABMs) take as input some parameter vector θ , and return a (possibly multivariate) time-series \mathbf{x} . Typically, ABMs are *stochastic* simulators, with the property that running the simulator repeatedly with a fixed parameter θ will produce different outputs with each run. The probability (density) with which the simulator produces a particular output \mathbf{x} , conditional on an input parameter vector θ , is given by its likelihood function, written $p(\mathbf{x} \mid \theta)$, which in some sense encapsulates the statistical properties of the ABM. We denote simulation of a dataset \mathbf{x} from the ABM as a draw from its probability density function, $\mathbf{x} \sim p(\mathbf{x} \mid \theta)$.

An essential prerequisite for the successful application of ABMs in real-world decision making scenarios is the ability to perform parameter estimation, that is, to tune θ such that the simulated data are good matches to observed data, which we denote with \mathbf{y} . Traditional statistical workhorses like maximum likelihood estimation or Bayesian inference rely on pointwise evaluations of the likelihood function. In practice, $p(\mathbf{x} \mid \theta)$ cannot be obtained or evaluated in a reasonable time for arbitrary simulation models, and is thus only implicitly defined by the software specifying the behaviour of the simulator. This presents a barrier to the use of ABMs in real-world use.

To overcome this, a number of approaches to performing statistical inference have been developed in which exact density evaluations are replaced by evaluations of approximate densities or cost functions that are constructed using simulations from the model. Some of these simulation-based inference (SBI) methods have been explored within the ABM community. Perhaps the most prevalent of them is simulated minimum distance (SMD), in which an estimate $\hat{\theta}$ is obtained by minimising some loss function $f(\mathbf{y}, \theta)$ between the observed data \mathbf{y} and simulated data $\mathbf{x} \sim p(\cdot \mid \theta)$ over some search space Θ :

$$\hat{\theta} = \arg \min_{\theta \in \Theta} f(\mathbf{y}, \theta). \quad (1)$$

This general class of estimators includes the maximum likelihood estimator – corresponding to choosing, for example, $f(\mathbf{y}, \theta) = -\log p(\mathbf{y} \mid \theta)$ or, where the likelihood is intractable, some estimate thereof (see e.g. [Diggle and Gratton, 1984](#); [Kukacka and Barunik, 2017](#)). It also includes the Method of Simulated Moments (MSM) ([Franke, 2009](#)) – in which f takes the form

$$f(\mathbf{y}, \theta) = (g(\mathbf{y}) - \hat{g}_\theta)' W (g(\mathbf{y}) - \hat{g}_\theta), \quad (2)$$

where $g(\mathbf{y})$ denotes a set of moments derived from \mathbf{y} , \hat{g}_θ denotes the same set of moments derived from $R \geq 1$ simulations at θ , W is a suitably chosen weight matrix, and $'$ denotes the transpose. As a final example, it includes Indirect Inference (II) ([Gourieroux et al., 1993](#)), which follows a similar approach to MSM but replaces the moments with estimated parameters of a tractable auxiliary model. Some further loss functions have been proposed more recently in the context of ABM estimation (see [Platt, 2020](#), for a recent review).

A major drawback of SMD and optimisation-based approaches more generally is that only parameter point estimates are produced, while in many contexts it is desirable to obtain meaningful uncertainty quantification regarding appropriate values for θ . Bayesian inference is an alternative inferential paradigm which naturally provides this meaningful notion of uncertainty, and has seen some recent but otherwise limited attention within the ABM community (see e.g. [Grazzini et al., 2017](#); [Lux, 2018](#); [Platt, 2020, 2021](#); [Lux, 2021](#)). Here, the primary object of interest is the parameter posterior distribution $p(\theta \mid \mathbf{y})$, obtained via Bayes theorem:

$$p(\theta \mid \mathbf{y}) = \frac{p(\mathbf{y} \mid \theta)}{p(\mathbf{y})} p(\theta). \quad (3)$$

In Equation (3), $p(\theta)$ is referred to as the *prior distribution* over parameters $\theta \in \Theta$, which encodes the experimenter’s initial beliefs regarding appropriate values for θ before data is observed, while the simulation model itself appears via its associated data likelihood function

$p(\mathbf{y} | \boldsymbol{\theta})$. Importantly, Bayesian inference procedures seek to accurately approximate the entire distribution $p(\boldsymbol{\theta} | \mathbf{y})$, rather than any single point estimate derived from this distribution. In this way, accurate and meaningful uncertainty estimates may also be obtained which correctly account for the experimenter’s initial beliefs, as well as the evidence provided by the data \mathbf{y} .

A well-known investigation into Bayesian estimation methods for ABMs is [Grazzini et al. \(2017\)](#), which explores various means of constructing a surrogate likelihood $\tilde{p}(\mathbf{y} | \boldsymbol{\theta})$ to account for the intractability of the true likelihood $p(\mathbf{y} | \boldsymbol{\theta})$. Two immediate drawbacks of the methods [Grazzini et al. \(2017\)](#) discuss are that they (a) are difficult to reconcile with the fact that ABMs are frequently dynamical models generating time-series output, and (b) entail a huge computational burden, since each likelihood evaluation in the employed posterior sampling algorithm requires at least one simulation from the model, and it is typically necessary to make many hundreds of thousands of such evaluations. Recent work by [Platt \(2021\)](#) suggests estimating the model’s transition density via a mixture density network. While the approach offers some improvements through the use of a more flexible density estimator and the incorporation of some temporal dependencies, it suffers from similar drawbacks, assuming time-homogeneity and requiring a computationally expensive estimation step for every single sample from the approximate posterior distribution. In summary, there is a need for parameter inference methods that fulfil two main criteria:

1. they must be simulation-efficient in order to remain applicable to large-scale simulation models such as ABMs;
2. and they must be able to deal with non-homogenous/non-stationary temporal data, both simulated and observed.

To this end, we seek with this paper to analyse the utility of two classes of parameter inference methods that have seen significant activity within the computational statistics and probabilistic machine learning literature in recent years: neural posterior estimation ([Papamakarios and Murray, 2016](#); [Lueckmann et al., 2017](#); [Greenberg et al., 2019](#)), and neural density ratio estimation ([Thomas et al., 2021](#); [Hermans et al., 2020](#); [Durkan et al., 2020](#)). Such methods have been employed successfully in a variety of applied domains, including high-energy physics ([Brehmer et al., 2018](#)), cosmology ([Alsing et al., 2019](#)), and neuroscience ([Gonçalves et al., 2020](#)). In addition, as we will demonstrate below, they work flexibly with potentially multivariate time-series data, requiring minimal model assumptions, and typically generate more accurate parameter inferences with a significantly reduced simulation budget in comparison to common alternatives, such as those reported in [Grazzini et al. \(2017\)](#).

We further seek to establish a higher standard of assessment criteria to be used in future benchmarking experiments for approximate Bayesian parameter inference for economic simulation models through the use of common integral probability metrics and inference validation procedures used widely in the statistics and machine learning communities. In Bayesian inference, the object of interest is the posterior density, whose entire shape is important since this captures degrees of beliefs regarding the unknown $\boldsymbol{\theta}$. However, the literature on Bayesian estimation of ABMs has largely forgone attempts to examine the degree of overlap between the estimated and ground-truth posterior densities or the extent to which the inference procedure has learnt the dynamics of the data, often in favour of a simple comparison of certain point estimates of $\boldsymbol{\theta}$ ([Grazzini et al., 2017](#); [Platt, 2021](#); [Shiono, 2021](#)). Since the objective of Bayesian inference is to derive a posterior probability density function which naturally quantifies uncertainty, it is important to adopt assessment criteria that quantify the degree of closeness between the estimated and target posterior densities in benchmarking experiments.

In summary, our contributions with this article are to provide:

1. an overview of recent advances in black-box simulation-based Bayesian inference for simulation models, and a motivation for their use in the specific case of agent-based models;

2. a systematic benchmarking of these state-of-the-art methods against popular alternatives within the literature on parameter inference for agent-based economic simulation models;
3. a discussion and demonstration of approaches to validating approximate Bayesian inference algorithms, and their feasibility for agent-based models under different simulation-based inference schemes.

2. Review of Bayesian inference methods for agent-based models

Bayesian inference has seen relatively little attention from the economic ABM community as a means to parameter estimation. In this section, we summarise the key literature in this space. Importantly, we emphasise throughout that a common feature of these approaches is that the task of inference – that is, of constructing the posterior distribution – is inherently linked to the act of simulating from the ABM and building an additional approximate *generative* model (sometimes also called an *emulator*), in the sense that an approximation to the likelihood function is constructed. We will see in later chapters that this contrasts with a new generation of estimation methods in which inference is decoupled from the act of simulating and can be performed in a *discriminative* manner, typically resulting in more efficient inferences.

2.1. Approximate Bayesian computation

Bayesian estimation of economic ABMs has gained popularity with the work of [Grazzini et al. \(2017\)](#), in which the authors discuss three approaches to approximate Bayesian parameter inference. In this section, we briefly outline the three approaches, and show that each method can be seen as explicitly or implicitly approximating the unknown likelihood function associated with the ABM by

$$\hat{p}(\mathbf{y} \mid \boldsymbol{\theta}) \propto \int K_\epsilon(\mathbf{y}, \mathbf{x}) p(\mathbf{x} \mid \boldsymbol{\theta}) d\mathbf{x}, \quad (4)$$

where K_ϵ is a (possibly unnormalised) kernel function, with parameter(s) ϵ , providing a measure of “distance” between observed data \mathbf{y} and simulated data \mathbf{x} . This permits an approximation $\hat{p}(\boldsymbol{\theta} \mid \mathbf{y})$ of the posterior density as

$$\hat{p}(\boldsymbol{\theta} \mid \mathbf{y}) \propto \int K_\epsilon(\mathbf{y}, \mathbf{x}) p(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}) d\mathbf{x}. \quad (5)$$

Such approaches – targeting (5) – are comprehensively subsumed under “approximate Bayesian computation (ABC)” which, as noted by [Grazzini et al. \(2017\)](#), is a technique that has received significant attention within the computational statistics community over the last two decades as a means to constructing approximate posterior densities for models with intractable likelihood functions (see e.g. [Tavaré et al., 1997](#); [Pritchard et al., 1999](#); [Beaumont et al., 2002](#); [Sunnåker et al., 2013](#); [Schmon et al., 2020](#)). The approach has been successfully applied within a variety of fields, most prominently in ecology, epidemiology, and systems biology ([Toni et al., 2009](#); [Liepe et al., 2014](#); [Beaumont, 2010](#); [Ju et al., 2021](#)). While the approaches discussed by [Grazzini et al. \(2017\)](#) can be cast in this general form, a judicious choice of a *particular* kernel K_ϵ based on reasonable assumptions is crucial. In particular, different kernel choices are called for in different scenarios. Some possibilities are outlined below.

2.1.1. Parametric density estimation

The simplest approach discussed by [Grazzini et al. \(2017\)](#) involves assuming that the ABM has entered a statistical equilibrium, such that the observations $y_t \in \mathbb{R}^d$ in the time-series $\mathbf{y} := (y_1, y_2, \dots, y_T) \in \mathbb{R}^{d \times T}$ are fluctuations around some stationary value m^* :

$$y_t = m^* + \varepsilon_t, \quad (6)$$

where (ε_t) are *iid* noise terms with density g_ϵ and parameters ϵ . Although more complex distributions are available, g_ϵ may for example be a zero-mean Gaussian distribution, and ϵ the elements of the covariance matrix.

Under these assumptions, recalling that the elements of the time-series are assumed to be independent, the parameters m^* and ϵ can be estimated by the means and covariances of the elements of the time-series to give \hat{m}^* and $\hat{\epsilon}$. More precisely, first fixing θ and drawing a sample $\mathbf{x} \sim p(\mathbf{x} | \theta)$, one can calculate the estimates $\hat{m}^*(\mathbf{x})$ and $\hat{\epsilon}(\mathbf{x})$ using, for example, maximum likelihood estimation. Adopting as the measure of distance the probability of the true data being observed under this approximated process yields the choice of kernel

$$K_\epsilon(\mathbf{y}, \mathbf{x}) = \prod_{t=1}^T g_{\hat{\epsilon}(\mathbf{x})}(y_t - \hat{m}^*(\mathbf{x})). \quad (7)$$

Finally by taking $R \geq 1$ *iid* simulated datasets¹ $\mathbf{x}^{(r)} \sim p(\mathbf{x} | \theta)$, $r = 1, \dots, R$, and calculating their associated estimators $\hat{m}^*(\mathbf{x}^{(r)})$ and $\hat{\epsilon}(\mathbf{x}^{(r)})$, the likelihood at θ , i.e. Equation (4), is approximated with the Monte Carlo average

$$\hat{p}_\theta(\mathbf{y}) \approx \frac{1}{R} \sum_{r=1}^R \prod_{t=1}^T g_{\hat{\epsilon}(\mathbf{x}^{(r)})}(y_t - \hat{m}^*(\mathbf{x}^{(r)})), \quad (8)$$

where we write \hat{p}_θ instead of $\hat{p}(\cdot | \theta)$. Alternatively, the R simulations may be pooled to generate single estimates $\hat{m}^*(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(R)})$ and $\hat{\epsilon}(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(R)})$ at θ . This is closely related to *synthetic likelihood* approaches (Wood, 2010; Price et al., 2018) which use similar Gaussian distributions, but usually rely on summary statistics.

In either case, the resulting approximate likelihood \hat{p}_θ from Equation (8) can then be used downstream for parameter estimation, either directly via e.g. maximum likelihood or through the corresponding approximate posterior, simulated e.g. through Markov chain Monte Carlo (MCMC). This approach suffers from clear limitations, however. Firstly, it treats the data points in the observed time-series as independent – that is, as lacking a natural ordering – which destroys important information when the observed \mathbf{y} and simulated \mathbf{x} are time-series. Secondly, choosing an appropriate and sufficiently flexible family of parametric densities g_ϵ to construct the likelihood approximation is non-trivial, with poor choices leading to erroneous Bayesian inference.

2.1.2. Non-parametric density estimation

An alternative approach, which partially addresses the second limitation described in Section 2.1.1, is to forgo the assumption of a parametric family of densities and instead use a non-parametric method for density estimation. Grazzini et al. (2017) describe the use of kernel density estimation (KDE) for this purpose. Here, the data points in the time-series are once again assumed to be independent and fluctuating about some stationary value m^* as in the parametric approach described above. Then, an estimate of the likelihood function is obtained by applying KDE to $R \geq 1$ *iid* simulations of length S , $\mathbf{x}^{(r)} := (x_1^{(r)}, \dots, x_S^{(r)}) \sim p(\mathbf{x} | \theta)$, $r = 1, \dots, R$, providing an unbiased estimate of the approximated likelihood function in Equation (4) as

$$\hat{p}_\theta(\mathbf{y}) \approx \frac{1}{R} \sum_{r=1}^R K_\epsilon(\mathbf{y}, \mathbf{x}^{(r)}) := \frac{1}{R} \sum_{r=1}^R \prod_{t=1}^T \hat{p}_\epsilon(y_t | \theta, \mathbf{x}^{(r)}), \quad (9)$$

where $\hat{p}_\epsilon(y_t | \theta, \mathbf{x}^{(r)})$ is the estimate of the conditional density $p(y_t | \theta, \mathbf{x}^{(r)})$ obtained via KDE:

$$\hat{p}_\epsilon(y_t | \theta, \mathbf{x}) = \frac{1}{S} \sum_{s=1}^S \kappa_\epsilon(y_t - x_s), \quad (10)$$

¹While R can be chosen to be any natural number, it is usually efficient to take $R = 1$ when averages of estimators are concerned and run the Markov chain for an accordingly higher number of iterations (Bornn et al., 2017; Sherlock et al., 2017).

for some probability kernel κ_ϵ with bandwidth parameter(s) ϵ . In particular, the authors employ a Gaussian kernel with bandwidth chosen using Silverman’s method (Silverman, 1986), such that

$$\kappa_\epsilon(y_t - x_s) = \frac{1}{\epsilon} \kappa\left(\frac{\|y_t - x_s\|_2}{\epsilon}\right). \quad (11)$$

Alternatively, as in Section 2.1.1, the R simulations may be pooled and a single KDE model fit to the combined dataset to obtain $\hat{p}(\mathbf{y} \mid \boldsymbol{\theta})$. While these non-parametric approaches to density estimation are arguably more flexible than the assumption of a parametric family, they are well known to suffer from the *curse of dimensionality*, limiting their applicability to low dimensional data only, even under the unrealistic assumption that the y_t are assumed independent.

2.1.3. Classical approximate Bayesian computation

The authors discuss a third common approach to simulation-based Bayesian inference for intractable simulation models, namely classical ABC. A typical ABC algorithm proceeds by proposing a parameter value $\boldsymbol{\theta}$ from some proposal distribution, for example the prior density, and determining whether to accept or reject $\boldsymbol{\theta}$ on the basis of some meaningful notion of distance $d(\mathbf{x}, \mathbf{y})$ between a simulation $\mathbf{x} \sim p(\mathbf{x} \mid \boldsymbol{\theta})$ and the observed data \mathbf{y} . A common choice for the distance d is the Euclidean distance between suitable summary statistics of the data; that is, $d(\mathbf{x}, \mathbf{y}) := \|\mathbf{s}(\mathbf{x}) - \mathbf{s}(\mathbf{y})\|$ where \mathbf{s} is a vector of, for example, the mean, standard deviation, and lag-1 autocorrelation of the data. While many versions of ABC exist, the simplest – Rejection ABC – involves repeatedly proposing candidate parameter values $\boldsymbol{\theta} \sim p(\boldsymbol{\theta})$ and rejecting in the event that $d(\mathbf{x}, \mathbf{y}) > \epsilon$, where $\mathbf{x} \sim p(\mathbf{x} \mid \boldsymbol{\theta})$. Rejection ABC can therefore be cast in the form of Equation (5) by taking

$$K_\epsilon(\mathbf{y}, \mathbf{x}) = \mathbb{1}[d(\mathbf{x}, \mathbf{y}) \leq \epsilon], \quad (12)$$

providing an estimate of the (unnormalised) approximate likelihood function in Equation (4) as

$$\hat{p}_\boldsymbol{\theta}(\mathbf{y}) \propto \frac{1}{R} \sum_{r=1}^R \mathbb{1}[d(\mathbf{x}^{(r)}, \mathbf{y}) \leq \epsilon]. \quad (13)$$

As discussed by Grazzini et al. (2017), ABC enjoys a number of desirable properties, in particular (under mild regularity conditions) consistency with the true posterior distribution as $\epsilon \rightarrow 0$, provided no information loss is incurred by the choice of distance d . However, a major challenge to its use in practice is designing such a distance function: the success or failure of the approach often hinges on the engineering of summary statistics that are close to sufficient². A number of works have investigated strategies for summary statistic selection in ABC (e.g. Fearnhead and Prangle, 2012; Blum et al., 2013; Wiqvist et al., 2019; Chen et al., 2020) and there has been some success in alternative, summary-free distances (see e.g. Park et al., 2016; Bernton et al., 2019; Dyer et al., 2021a). It is however in general not possible to derive a low-dimensional sufficient statistic for arbitrary simulation models, and a poor choice of summary statistics can lead to an unacceptable loss of information from the original data, resulting in a posterior approximation of little value.

2.1.4. Unifying the approaches

In Algorithm 1, we show how the approximate posterior (5), with any of the three possibilities for K_ϵ discussed by Grazzini et al. (2017), may be targeted using a variant of the popular Metropolis-Hastings (MH) algorithm (e.g. Hastings, 1970). If the likelihood estimate $\hat{p}_\boldsymbol{\theta}(\mathbf{y})$ is an unbiased, non-negative estimate of some desired target $p_\boldsymbol{\theta}(\mathbf{y})$, then Algorithm 1 is an example

²A Bayesian sufficient statistic $\mathbf{s}(\mathbf{x})$ for a model $p(\mathbf{x} \mid \boldsymbol{\theta})$ is one for which $p(\boldsymbol{\theta} \mid \mathbf{s}(\mathbf{x})) = p(\boldsymbol{\theta} \mid \mathbf{x})$. Low-dimensional sufficient statistics for arbitrary probabilistic models are generally unobtainable (Pitman, 1936; Koopman, 1936; Darmois, 1935).

Algorithm 1: Metropolis-Hastings sampling scheme for ABC

Input: Prior distribution $p(\cdot)$, observation \mathbf{y} , proposal distribution $q(\cdot | \boldsymbol{\theta})$, initial value $\boldsymbol{\theta}_0$, number of iterations n ;
Result: Empirical posterior $\sum_{i=1}^n \delta_{\boldsymbol{\theta}_i}$
for $r = 1, \dots, R$ **do**
 | Simulate $\mathbf{x}^{(r)} \sim p(\cdot | \boldsymbol{\theta}_0)$;
end
for $i = 1, \dots, n$ **do**
 Sample $\boldsymbol{\theta} \sim q(\cdot | \boldsymbol{\theta}_{i-1})$;
 for $r = 1, \dots, R$ **do**
 | Simulate $\tilde{\mathbf{x}}^{(r)} \sim p(\cdot | \boldsymbol{\theta})$;
 end
 Evaluate $\hat{p}_{\boldsymbol{\theta}}(\mathbf{y})$ using $\{\tilde{\mathbf{x}}^{(r)}\}_{r=1}^R$ according to either Equation (8), (9), or (13) as desired;
 Set $\hat{p}_{\boldsymbol{\theta}_i}(\mathbf{y}) = \hat{p}_{\boldsymbol{\theta}}(\mathbf{y})$, $\boldsymbol{\theta}_i = \boldsymbol{\theta}$ and $\{\mathbf{x}^{(r)}\}_{r=1}^R = \{\tilde{\mathbf{x}}^{(r)}\}_{r=1}^R$ with probability

$$\alpha = \min \left\{ 1, \frac{\hat{p}_{\boldsymbol{\theta}}(\mathbf{y})p(\boldsymbol{\theta})q(\boldsymbol{\theta}_{i-1} | \boldsymbol{\theta})}{\hat{p}_{\boldsymbol{\theta}_{i-1}}(\mathbf{y})p(\boldsymbol{\theta}_{i-1})q(\boldsymbol{\theta} | \boldsymbol{\theta}_{i-1})} \right\},$$

 otherwise set $\hat{p}_{\boldsymbol{\theta}_i}(\mathbf{y}) = \hat{p}_{\boldsymbol{\theta}_{i-1}}(\mathbf{y})$ and $\boldsymbol{\theta}_i = \boldsymbol{\theta}_{i-1}$.
end

of a pseudo-marginal MH algorithm (Andrieu and Roberts, 2009). Note that while the sampling procedure in Algorithm 1 exactly targets the posterior distributions described in Grazzini et al. (2017), many alternative posterior sampling algorithms exist (see e.g. Beaumont et al., 2009).

We once again emphasise that a major disadvantage to applying the methods described by Grazzini et al. (2017) to ABMs is that estimating the posterior involves simulating $R \geq 1$ times from the ABM at each proposed parameter value $(\boldsymbol{\theta}_i)_{i=1, \dots, n}$. Since ABMs are typically expensive to simulate, and n is required to be many hundreds of thousands in order to accurately estimate the targeted posterior, such approaches can rapidly lead to a prohibitively large number of required simulations.

2.2. Latent variable models

A class of models closely related to approximate Bayesian computation is the class of so-called *latent variable models*. Here, the real data, \mathbf{y} , is assumed to be a noisy observation of an unobserved (latent) process \mathbf{x} . Thus, in contrast to previous approaches, some discrepancy between \mathbf{y} and \mathbf{x} is to be expected.

Under the most basic scenario, the data obtained from the simulator are identically $x_t \sim p(x | \boldsymbol{\theta})$ and $p(\mathbf{x} | \boldsymbol{\theta}) = \prod_{t=1}^T p(x_t | \boldsymbol{\theta})$. The observed data is then assumed to have an error distribution, g_ε , say, such that the contribution made by one observation y_t to the likelihood function is thus

$$p(y_t | \boldsymbol{\theta}) = \int g_\varepsilon(y_t | x_t)p(x_t | \boldsymbol{\theta})dx_t \quad (14)$$

and $p(\mathbf{y} | \boldsymbol{\theta}) = \prod_{t=1}^T p(y_t | \boldsymbol{\theta})$. We draw the attention of the reader to the conceptual similarity to (4); however, (14) is no longer an approximation but is the exact likelihood under the assumption of measurement error. (Similarly, ABC can be seen as exact if K_ε describes the observational error, see e.g. Wilkinson 2013.) A simple unbiased estimator of (14) is

$$\hat{p}(\mathbf{y} | \boldsymbol{\theta}) = \prod_{t=1}^T \frac{1}{R} \sum_{r=1}^R g_\varepsilon(y_t | x_t^{(r)}), \quad x_t^{(r)} \sim p(x | \boldsymbol{\theta})$$

for some $R \geq 1$ ³, which can then also be used in Algorithm 1, for example, to target the associated approximate posterior. The independence assumption underlying the simulator data is, however, not realistic for ABMs and other time-series models, and incorporating the temporal aspect of data requires further structural assumptions, such as those of hidden Markov models.

2.2.1. Hidden Markov Models and particle filters

Particle filters, an instance of a broader class of sequential Monte Carlo (SMC) methods (see Ju et al., 2021, for a recent overview of SMC and their application to epidemiological ABMs) relax the independence assumption imposed on the output of the simulation model and have previously been employed for Bayesian parameter inference for economic ABMs (Lux, 2018, 2021). However, such methods also involve making assumptions about the structure of the model – specifically, the assumption of a state space structure with a particular error distribution. Furthermore, previous works have noted the significant computational burden required for SMC methods (Malleson et al., 2020; Lux, 2021), which can arise due to the fact that a number of particles that grows exponentially with the model dimensions is often required to avoid failure modes such as particle collapse.

2.3. Neural methods

As an alternative to kernel density estimation, neural density estimators have previously been employed to perform Bayesian estimation of ABMs. Platt (2021) presents a method which diverges from the approaches of Grazzini et al. (2017) in two main regards. Firstly, in contrast to the previous independence assumption, it assumes an autoregressive structure for the time-series model to better capture temporal correlations in the data. In particular, the approach assumes a time-series model that is Markov of order L , that is, $p(x_t | x_{1:t-1}, \theta) = p(x_t | x_{t-L:t-1}, \theta)$, where $x_{1:t} = (x_1, \dots, x_t)$. In words, the distribution of the state at time t depends only on the previous L states (and θ). In light of this assumption, the full likelihood can be factorised as

$$p_{\theta}(\mathbf{y}) = p(y_{1:L} | \theta) \prod_{t=L+1}^T p(y_t | y_{t-1}, \dots, y_{t-L}, \theta). \quad (15)$$

Secondly, as a consequence of the autoregressive model structure, the method employs a *conditional* density estimator q_{ϕ} to approximate the transition density function $p(x_t | x_{t-L:t-1}, \theta)$. The resulting likelihood approximation can be written

$$\hat{p}_{\theta}(\mathbf{y}) = \prod_{t=L+1}^T q_{\phi}(\theta)(y_t | y_{t-1}, \dots, y_{t-L}), \quad (16)$$

where the contribution of the first L terms is ignored. A mixture density network (Bishop, 1994) assumes the role of the conditional density estimator $q_{\phi}(\theta)$ with parameters $\phi(\theta)$ trained for each θ . Denoting the L states preceding y_t as $\mathbf{y}_{<t}$ for brevity, the particular form of the mixture density network is

$$q_{\phi}(y_t | \mathbf{y}_{<t}) = \sum_{k=1}^K \alpha_k(\mathbf{y}_{<t}) \mathcal{N}(y_t | \mu_k(\mathbf{y}_{<t}), \Sigma_k(\mathbf{y}_{<t})).$$

where α_k are the mixing coefficients, and μ_k, Σ_k are the mean and covariance matrix respectively, all of which are parameterised by neural networks and trained via maximum likelihood on R iid model samples $\mathbf{x}^{(r)} \sim p(\mathbf{x} | \theta)$, $r = 1, \dots, R$. The resulting likelihood estimate can then

³Note that opposed to earlier cases the choice $R = 1$ is generally not optimal for products of unbiased estimators, see Doucet et al. (2015); Sherlock et al. (2015); Schmon et al. (2021).

once again be used also with Algorithm 1, for example, to target the associated approximate posterior distribution.

While the model structure described above accounts for the sequential nature of simulations generated by ABMs to some extent, the computational burden associated with the simulation of sufficient training data and the training of a new conditional density estimator at each θ renders it largely infeasible (Shiono, 2021). In other work, Shiono (2021) examines a particular neural approach, with an invertible architecture and the ability to accommodate time-series data, for the specific purpose of estimating the posterior density for a New-Keyensian ABM. In contrast to both Platt (2021) and Shiono (2021), we take in this work a broad perspective on the subject of neural simulation-based inference for ABMs, by thoroughly contextualising the problem of Bayesian parameter estimation for economic ABMs and by incorporating both benchmarking tasks and an exploration of suitable assessment criteria.

3. A new generation of simulation-based inference methods

In this section, we provide an overview of the two main simulation-based inference (SBI) methods we make use of in this paper: neural posterior estimation (NPE) and neural ratio estimation (NRE), both of which – as their names suggest – employ neural networks⁴ to obtain an estimate of the posterior density $p(\theta | \mathbf{y})$. We motivate the use of these methods for agent-based models (ABMs) by framing them as *discriminative* approaches to simulation-based parameter inference, and contrasting them with the *generative* approaches described in Section 2 which seek to approximate the model’s distribution (likelihood) using some form of probabilistic model. We then present the core elements of these methods, and further provide a discussion on how they may be leveraged flexibly and automatically to accommodate inference tasks involving high-dimensional and potentially multivariate time-series data, as is often the case in economic applications.

3.1. Motivating black-box discriminative approaches to parameter inference

The methods we endorse in this article – NPE and NRE – can be motivated by the fact that they are:

1. simulation-efficient alternatives to more traditional approaches to SBI, such as the approximate Bayesian computation (ABC) approaches described in Section 2;
2. generic SBI methods that treat the simulator as a black-box, thus making minimal assumptions about the structure and output of the simulation model;
3. discriminative approaches to SBI, in the sense that inference does not require a probabilistic model for the data \mathbf{x} : instead of *explaining* the mechanisms that create the data, we only need to *distinguish* between realisations that arise from different parameter values.

Simulation efficiency The algorithms described in Section 2 share a common pattern. For a fixed parameter θ , *iid* simulations $\mathbf{x}^{(r)} \sim p(\mathbf{x} | \theta)$, $r = 1, \dots, R$, are sampled to produce a proxy likelihood $\hat{p}(\mathbf{x} | \theta)$. Then, to generate n approximate posterior samples via Markov chain Monte Carlo (MCMC), this procedure is performed at n different values for θ , where n must typically be a large number – often a few hundred thousand, if not orders or magnitude larger – to ensure a low Monte Carlo error. Consequently, at least nR simulations from the ABM are required in total for inference under these algorithms. Since ABMs can be very expensive to simulate, and the act

⁴Note that neural networks are not essential here – the procedures we describe require only flexible function approximators. We frame the discussion around neural networks primarily because they are a convenient choice of flexible function approximators in many settings.

of simulating from the ABM remains the primary bottleneck in Bayesian estimation procedures for ABMs, this simulation demand can quickly become infeasible.

In contrast, the methods for which we advocate here differ by eliminating the need to simulate when sampling from the posterior. Instead, they decouple the act of simulating from the task of constructing the posterior by employing powerful function approximators to learn – in essence – *global* posterior density estimators on the basis of a limited number of simulations from across the parameter space of the ABM; that is, they learn functions $h : \mathcal{Y} \times \Theta \rightarrow \mathbb{R}$ which approximate the posterior density $p(\boldsymbol{\theta} \mid \mathbf{y})$ across the space of all possible values for $\boldsymbol{\theta}$ and \mathbf{y} . This has the potential to significantly reduce the simulation burden associated with approximate Bayesian inference procedures, because the pointwise estimates of $p(\boldsymbol{\theta} \mid \mathbf{y})$ can borrow strength from, and share information between, one another; in contrast, the algorithms described in Section 2 consider each pointwise evaluation of $p(\boldsymbol{\theta} \mid \mathbf{y})$ as standalone density estimation tasks, such that no information can be shared between them. In this way, a large number of approximate posterior samples can be generated from an algorithm trained on what can in practice be a far smaller number of model samples than is required for the algorithms described in Section 2.

Finally, the approaches we endorse here are equipped with a further efficiency benefit: *amortisation*. This phrase captures the fact that the global density estimators can be used to generate samples from an approximate posterior $\hat{p}(\boldsymbol{\theta} \mid \mathbf{y})$ for any data \mathbf{y} without the need to further simulate from the ABM, which results from the fact that they are trained on the full space of possible values for $\boldsymbol{\theta}$ and \mathbf{y} . In contrast, applying the methods described in Section 2 to a new dataset \mathbf{y}' would entail another nR simulations from the ABM, multiplying the already high simulation costs.

Black-box inference methods We saw in Sections 2.1.1 and 2.1.2 that many inference approaches come with restrictive assumptions, for instance, stationarity of the simulated and observed time-series. In general, it is useful to dispense with these assumptions, since they can be difficult to verify and limit the applicability of these methods for arbitrary simulators. Instead, it can be preferable to employ black-box methods that make minimal assumptions about the model and are therefore generically applicable to arbitrary simulators. Doing so enables the modeller to concentrate resources on model design and implementation, rather than on developing bespoke inference algorithms for each new simulator. Furthermore, assumptions such as stationarity are known to be particularly poorly suited to certain economic simulation models such as ABMs, since these models are known and are even designed to produce non-equilibrium dynamics. Employing inference procedures that are able to handle such dynamics is therefore essential to the task of estimating generic ABMs.

Discriminative approaches to parameter inference Discriminative tasks in machine learning are typically simpler than generative tasks, since generative tasks address larger problems than pure discrimination. This is intuitive: for example, it is typically easier for humans and computers alike to distinguish between images of cats and dogs than to generate them. Analogously, it is generally a simpler task to discriminate between complex time-series data than it is to generate such time-series. The approaches described in Section 2 adopt a generative approach: they each – either explicitly or implicitly – seek to derive an approximation to the simulator’s likelihood function using a probabilistic model, and thus seek to model the (probability density function of the) simulation output itself. Formally, this may be understood as learning the (stochastic) map $\boldsymbol{\theta} \mapsto \mathbf{x}$. NPE and NRE, in contrast, do not seek to model the simulation output: instead, they map from instances of the simulation output to certain target values which we will describe in more detail below. This can be formally thought of as learning the map $\mathbf{x} \mapsto \boldsymbol{\theta}$. Such an approach to parameter estimation therefore embodies a fundamental departure from the approaches described in Section 2. It has the potential to be particularly beneficial for ABMs, which are known to be able to produce complex, non-equilibrium dynamics that are especially difficult to model and generate.

3.2. Neural posterior estimation

In this section, we provide an introduction to neural conditional density estimators and their use in posterior estimation tasks. The core idea underlying this class of simulation-based Bayesian inference techniques is to use such conditional density estimators to model the parameter posterior density directly. While various conditional density estimators can be used for this purpose, e.g. mixture density networks (Bishop, 1994; Papamakarios and Murray, 2016), we focus here on the case of **normalising flows** (Tabak and Vanden-Eijnden, 2010; Tabak and Turner, 2013; Rezende and Mohamed, 2015) due to their widespread use in various density estimation tasks, including in the areas of image generation (e.g. Kingma and Dhariwal, 2018) and physics (e.g. Noé et al., 2019).

3.2.1. Normalising flows

Normalising flows involve transforming a simple base distribution into a more complicated one, often with the use of neural networks. Consider a random variable $\mathbf{U} \sim p_{\mathbf{U}}$, where $p_{\mathbf{U}}$ is a probability distribution chosen to be a “simple” base distribution, in the sense that it is easy to both generate samples from $p_{\mathbf{U}}$ and to evaluate $p_{\mathbf{U}}(\mathbf{u})$ for any \mathbf{u} . Now consider the transformed random variable $\mathbf{X} = g(\mathbf{U})$, where g is a differentiable, invertible function with differentiable inverse $f := g^{-1}$. Then, by the change of variables formula, we have

$$p_{\mathbf{X}}(\mathbf{x}) = p_{\mathbf{U}}(f(\mathbf{x})) \left| \det J_f(\mathbf{x}) \right| \quad (17)$$

where J_f is the Jacobian of f . Sampling from $p_{\mathbf{X}}$ then simply involves sampling a value \mathbf{u} from the base distribution $p_{\mathbf{U}}$ and immediately obtaining $\mathbf{x} = g(\mathbf{u})$. Evaluating $p_{\mathbf{X}}(\mathbf{x})$ also then simply involves evaluation of the right-hand side of Equation (17). The above may be extended easily to a composition, or *flow*, $g = g_n \circ g_{n-1} \circ \dots \circ g_1$ of differentiable and invertible functions g_i with inverses f_i , for which we now have that

$$p_{\mathbf{X}}(\mathbf{x}) = p_{\mathbf{U}}(f_1(\dots f_{n-1}(f_n(\mathbf{x}))\dots)) \left| \prod_{i=1}^n \det J_{f_i}(\mathbf{x}) \right|. \quad (18)$$

Given samples from $p_{\mathbf{X}}$, the question of how to choose a base distribution $p_{\mathbf{U}}$ and a series of functions g_1, \dots, g_n such that the distribution of the samples is modelled accurately arises. The idea behind normalising flows is to *learn* this sequence of transformations, that is, to have each f_i be a flexible transformation f_{ϕ_i} with trainable parameters ϕ_i . The resulting density estimator q_{ϕ} can be written

$$q_{\phi}(\mathbf{x}) = p_{\mathbf{U}}(f_{\phi_1}(\dots f_{\phi_{n-1}}(f_{\phi_n}(\mathbf{x}))\dots)) \left| \prod_{i=1}^n \det J_{f_{\phi_i}}(\mathbf{x}) \right|, \quad \text{where } \phi = (\phi_n, \dots, \phi_1). \quad (19)$$

The simple base distribution is then often taken to be a standard normal distribution⁵ of the same dimension as the target distribution, and the flexible transformations f_{ϕ_i} are typically neural networks with a structure designed to guarantee the required invertibility and differentiability. The trainable parameters $\phi = \{\phi_i : 1 \leq i \leq n\}$ are optimised by maximising the log-likelihood of the data $\mathbf{x}^{(r)} \stackrel{iid}{\sim} p_{\mathbf{X}}, r = 1, \dots, R$:

$$\hat{\phi} = \arg \max_{\phi} \sum_{r=1}^R \log q_{\phi}(\mathbf{x}^{(r)}). \quad (20)$$

Normalizing flows are typically implemented in machine learning libraries that support automatic differentiation such as `pytorch` (Paszke et al., 2019) or `TensorFlow` Abadi et al. (2016)

⁵For this reason, the composition $f = f_1 \circ \dots \circ f_{n-1} \circ f_n$ is often referred to as the *normalising flow*, since it is a flow of transformations resulting (typically) in a Gaussian base distribution.




Figure 1: Schematic of a normalising flow. Sampling and density evaluation are performed via the processes illustrated at the top and bottom of the figure, respectively.

allowing the effective use of gradient-based optimisation techniques, such as Adam (Kingma and Ba, 2014). We provide a schematic of the sampling (flow, right-pointing arrows) and density evaluation (normalising flow, left-pointing arrows) processes in Figure 1, in which the simple, left-most distribution is morphed over successive steps into the more complex, right-most distribution.

3.2.2. Normalising flows for neural posterior estimation

Normalising flows may also be extended to the case of *conditional* density estimation (see e.g. Papamakarios and Murray, 2016; Lueckmann et al., 2017; Papamakarios et al., 2019; Greenberg et al., 2019). In this way, we can use a normalising flow to estimate the posterior density $p(\boldsymbol{\theta} \mid \mathbf{x})$ associated with a simulation model by following the same procedure as above and finding the neural network parameters as

$$\hat{\phi} = \arg \max_{\phi \in \Phi} \sum_{r=1}^R \log q_{\phi} \left(\boldsymbol{\theta}^{(r)} \mid \mathbf{x}^{(r)} \right), \quad (21)$$

where $(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)}) \stackrel{iid}{\sim} p(\mathbf{x}, \boldsymbol{\theta})$, $r = 1, \dots, R$. Importantly, this framework allows us to learn a *single* conditional density estimator for the posterior density function across data \mathbf{x} , rather than having to undergo the expensive procedure of training a separate density estimator for each point evaluation of the posterior density as in e.g. Platt (2021).

3.3. Neural density ratio estimation

In this section, we describe an alternative common approach to SBI termed neural ratio estimation (NRE). The core idea underlying this method contrasts with NPE in that while NPE seeks to model the posterior density directly, NRE instead seeks to estimate the following ratio:

$$\frac{p(\mathbf{x}, \boldsymbol{\theta})}{p(\mathbf{x})p(\boldsymbol{\theta})} = \frac{p(\mathbf{x} \mid \boldsymbol{\theta})}{p(\mathbf{x})} = \frac{p(\boldsymbol{\theta} \mid \mathbf{x})}{p(\boldsymbol{\theta})} =: r(\mathbf{x}, \boldsymbol{\theta}). \quad (22)$$

This ratio is frequently referred to as the **likelihood-to-evidence ratio**. In the event that Bayesian inference is the goal, this then permits evaluation of the posterior density as

$$p(\boldsymbol{\theta} \mid \mathbf{x}) = r(\mathbf{x}, \boldsymbol{\theta})p(\boldsymbol{\theta}). \quad (23)$$

In the following, we discuss how this ratio may be efficiently estimated with well-calibrated probabilistic classifiers trained on simulated data-parameter pairs.

3.3.1. Likelihood-to-evidence ratio estimation via probabilistic binary classification

We outline the approach to approximate Bayesian inference via density ratio estimation as described by Hermans et al. (2020). Here, a probabilistic binary classifier is trained to distinguish between two sets of training examples:

1. a set of “genuine” examples drawn from the joint distribution, $(\mathbf{x}, \boldsymbol{\theta}) \stackrel{iid}{\sim} p(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta})$, which are assigned a class label $c = 1$;
2. a set of “false” examples drawn from the product of the marginals, $(\mathbf{x}, \boldsymbol{\theta}) \stackrel{iid}{\sim} p(\mathbf{x})p(\boldsymbol{\theta})$, with class label $c = 0$.

The difference between the two cases is that the \mathbf{x} are generated by the $\boldsymbol{\theta}$ to which they are paired in the former set of examples, while in the latter set of examples the \mathbf{x} bear no relation to the $\boldsymbol{\theta}$ they are paired with.

The function of a probabilistic binary classifier trained on such data is to model the class probability $d(\mathbf{x}, \boldsymbol{\theta}) := p(c = 1 | \mathbf{x}, \boldsymbol{\theta}) \in [0, 1]$; hard classification labels (i.e. decisions regarding the predicted value of c) are obtained when the continuous-valued $d(\mathbf{x}, \boldsymbol{\theta})$ is combined with a decision rule, for example the rule that $c = 1$ should be predicted whenever $d(\mathbf{x}, \boldsymbol{\theta}) > 0.5$. One can show (see Appendix B of Hermans et al., 2020) that the *optimal* estimate of $d(\mathbf{x}, \boldsymbol{\theta})$ is the value

$$d^*(\mathbf{x}, \boldsymbol{\theta}) := \frac{p(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta}) + p(\mathbf{x})p(\boldsymbol{\theta})}. \quad (24)$$

Thus, a good probabilistic classifier trained to distinguish between the two possible types of pairs $(\mathbf{x}, \boldsymbol{\theta})$ will learn a good estimate $\hat{d}(\mathbf{x}, \boldsymbol{\theta})$ of this ratio. Such a probabilistic classifier will allow us to evaluate the posterior density in this way: by noticing that one can rearrange Equation (24) as

$$p(\boldsymbol{\theta} | \mathbf{x}) = \frac{d^*(\mathbf{x}, \boldsymbol{\theta})}{1 - d^*(\mathbf{x}, \boldsymbol{\theta})}p(\boldsymbol{\theta}) =: r^*(\mathbf{x}, \boldsymbol{\theta})p(\boldsymbol{\theta}), \quad (25)$$

where $r^*(\mathbf{x}, \boldsymbol{\theta})$ is the corresponding estimate of the likelihood-to-evidence ratio $p(\mathbf{x} | \boldsymbol{\theta})/p(\mathbf{x})$. In practice, of course, only an approximation $\hat{d}(\mathbf{x}, \boldsymbol{\theta})$ will be obtained to the ratio in Equation (24), yielding a correspondingly imperfect estimate $\hat{r}(\mathbf{x}, \boldsymbol{\theta})$ of the likelihood-to-evidence ratio. Nonetheless, it is known that training expressive, high-capacity neural network classifiers via the cross-entropy loss

$$\ell \left(\left\{ c^{(i)}, \mathbf{x}^{(i)}, \boldsymbol{\theta}^{(i)} \right\}_{i=1}^N \right) = -\frac{1}{N} \sum_{i=1}^N \left[c^{(i)} \log \hat{d}(\mathbf{x}^{(i)}, \boldsymbol{\theta}^{(i)}) + (1 - c^{(i)}) \log (1 - \hat{d}(\mathbf{x}^{(i)}, \boldsymbol{\theta}^{(i)})) \right]$$

can yield classifiers with good probability estimates, and thus good estimates of $r^*(\mathbf{x}, \boldsymbol{\theta})$.

3.3.2. A generalisation to multi-class classification

In more recent work, Durkan et al. (2020) demonstrate that the above approach to density ratio estimation can be generalised to the problem of training a probabilistic classifier to identify the *correct* $(\mathbf{x}, \boldsymbol{\theta}^{(i)})$ pair from a batch $\{(\mathbf{x}, \boldsymbol{\theta}^{(b)})\}_{b=1}^B$ of B pairs that otherwise contain only “incorrect” pairs, where “correct” and “incorrect”, respectively, correspond to having been drawn from the joint distribution $p(\mathbf{x}, \boldsymbol{\theta})$ and from the product of the marginal distributions $p(\mathbf{x})p(\boldsymbol{\theta})$. In this case, the optimal estimate of the correct class probability is

$$p(c = i | \mathbf{x}, \{\boldsymbol{\theta}^{(b)}\}_{b=1}^B) = \frac{p(\boldsymbol{\theta}^{(i)} | \mathbf{x}) \prod_{b \neq i} p(\boldsymbol{\theta}^{(b)})}{\sum_{b=1}^B p(\boldsymbol{\theta}^{(b)} | \mathbf{x}) \prod_{b' \neq b} p(\boldsymbol{\theta}^{(b')})} \quad (26)$$

$$= \frac{p(\boldsymbol{\theta}^{(i)} | \mathbf{x})/p(\boldsymbol{\theta}^{(i)})}{\sum_{b=1}^B p(\boldsymbol{\theta}^{(b)} | \mathbf{x})/p(\boldsymbol{\theta}^{(b)})}. \quad (27)$$

Thus, by comparison with Equation (27), training a neural network f_ϕ on the loss function

$$\ell(\phi) = -\log \frac{\exp f_\phi(\mathbf{x}, \boldsymbol{\theta}^{(i)})}{\sum_{b=1}^B \exp f_\phi(\mathbf{x}, \boldsymbol{\theta}^{(b)})} \quad (28)$$

for each $(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x} | \boldsymbol{\theta})p(\boldsymbol{\theta})$ will induce $f_\phi(\mathbf{x}, \boldsymbol{\theta})$ to learn the value $f_\phi(\mathbf{x}, \boldsymbol{\theta}) = \log(p(\boldsymbol{\theta} | \mathbf{x})/p(\boldsymbol{\theta}))$, thus recovering an estimate of the desired density ratio. This can be extended to a batch of R “correct” data-parameter pairs as

$$\mathcal{L}(\phi) = -\frac{1}{R} \sum_{r=1}^R \log \frac{\exp f_\phi(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)})}{\sum_{b_r=1}^{B_r} \exp f_\phi(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(b_r)})}, \quad (29)$$

where the terms in the denominator are labelled with b_r to account for the possibility that the contrasting (“incorrect”) set of parameters may be different for different “correct” pairs $(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)})$. The authors further demonstrate that this can yield more accurate density ratio estimators. For this reason, we adopt this approach throughout this article.

3.4. Sampling from the neural posterior

After successful training, the posterior density estimator $q_\phi(\boldsymbol{\theta} | \mathbf{x})$ obtained from NPE is a parametric approximation of the true posterior distribution, which can then be used to generate *iid* samples $\boldsymbol{\theta} \sim q_\phi(\boldsymbol{\theta} | \mathbf{x})$ or to evaluate the posterior density on a pointwise basis in the ways described in Section 3.2.1. While ratio estimation similarly permits a pointwise evaluation of the posterior distribution as $p(\boldsymbol{\theta}) \exp(f_\phi(\mathbf{x}, \boldsymbol{\theta}))$, it requires a further inference step using, for example, Metropolis–Hastings to generate samples from the same posterior. However, the upfront training of the ratio estimator eliminates the need for expensive sampling from the ABM, reducing the run-time significantly in comparison to alternative approaches such as those described in Section 2.

3.5. Round-based training

The procedures described in Sections 3.2 and 3.3 are framed as single density (ratio) estimation tasks. This means that a single dataset of R data-parameter pairs $(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)})$, $r = 1, \dots, R$ is created in the following way:

$$\begin{aligned} \boldsymbol{\theta}^{(r)} &\sim p(\boldsymbol{\theta}), && \text{sample from the prior;} \\ \mathbf{x}^{(r)} &\sim p(\mathbf{x} | \boldsymbol{\theta}^{(r)}), && \text{sample from the model using the draws from the prior.} \end{aligned}$$

Subsequently, NRE and NPE algorithms are trained on the *whole dataset* to find either the likelihood-to-evidence ratio or the posterior directly following, respectively, (21) or (29). Density estimators trained in this way are referred to as *amortised* density estimators, which reflects the fact that they may be used to construct posterior distributions for any data \mathbf{x} without further simulation from the ABM or the need to train multiple density estimators.

The disadvantage of this one-stage approach is that all training simulations from the ABM are generated by parameters drawn from the prior distribution. In many cases, however, interest lies only in the posterior for the *particular observation* \mathbf{y} , which can be concentrated on specific subregions of the entire space covered by the prior density. It can then be preferable to narrow down the search space of good candidate values for the parameter $\boldsymbol{\theta}$ to subregions of the parameter space that could most plausibly have generated \mathbf{y} . By doing so, the density (ratio) estimator will be presented with a less varied range of dynamics between which it must learn to distinguish, allowing it to develop a more refined approximation of the density (ratio) in regions of high posterior density. This can facilitate more rapid learning and potentially reduce the number of training examples that must be simulated by the ABM.

Significant effort has thus been extended towards the constructions of “round”-based approaches to training density (ratio) estimators for SBI (see e.g. Papamakarios et al., 2019; Greenberg et al., 2019; Hermans et al., 2020; Durkan et al., 2020). The idea is to split the total budget of R simulations into subsets of size N_1, N_2, \dots such that $\sum_i N_i = R$. In the first step, N_1 data points are created as before and a posterior approximation $q_{\hat{\phi}}(\boldsymbol{\theta} \mid \mathbf{x})$ is constructed. Subsequently, in round $i \geq 2$, we create new data $(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)})$, $r = 1, \dots, N_i$ using

$$\begin{aligned} \boldsymbol{\theta}^{(r)} &\sim q_{\hat{\phi}}^{(i-1)}(\boldsymbol{\theta} \mid \mathbf{x}), && \text{sample from round-}(i-1) \text{ posterior;} \\ \mathbf{x}^{(r)} &\sim p(\mathbf{x} \mid \boldsymbol{\theta}^{(r)}), && \text{sample from the model using the draws from the round-}(i-1) \text{ posterior.} \end{aligned}$$

The posterior $q_{\hat{\phi}}(\boldsymbol{\theta} \mid \mathbf{x})$ may now be retrained using (a combination of the old and) the new samples, and the process can be repeated for as many rounds as necessary. An identical process can be followed for density ratio estimation, with the exception that parameters in round i are drawn using the likelihood-to-evidence ratio obtained from round $i-1$ and, for example, Metropolis–Hastings.

While the round-based approach is appealing, it does not come without additional challenges. In particular, the joint distribution of the example pairs $(\mathbf{x}^{(r)}, \boldsymbol{\theta}^{(r)})$ is no longer $p(\mathbf{x}, \boldsymbol{\theta})$ for data sampled after the first round. In practice, this needs to be accounted for during training by the use of additional weighing factors, such as those appearing in Equation (30) (cf. Equation (21)). We refer the interested reader to the following papers for further details on this matter: Papamakarios and Murray (2016); Lueckmann et al. (2017); Greenberg et al. (2019).

Training schemes for NPE and NRE which make use of this round-based training approach are referred to as *sequential* neural posterior estimation (SNPE) and *sequential* neural ratio estimation (SNRE), respectively, in which the prefix “sequential” reflects the round-based training design of the now non-amortised density (ratio) estimator.

3.6. Summarising the simulation output

As described above, NPE and NRE take as input either \mathbf{x} alone or $(\mathbf{x}, \boldsymbol{\theta})$ pairs. Given that \mathbf{x} is a (possibly multivariate) time-series for many ABMs and is thus a high-dimensional object, it can be beneficial to incorporate useful inductive biases⁶ into the network architecture that account for the sequential nature of \mathbf{x} and reduce its dimensionality. To this end, the density (ratio) estimator can be prefixed with a so-called embedding network with trainable parameters φ , whose function is to consume the original high-dimensional dataset \mathbf{x} and express this as a lower-dimensional summary statistic vector $\mathbf{s}_{\varphi}(\mathbf{x})$. The parameters of the embedding network and of the density (ratio) estimator may then be learned concurrently using the same loss function, offering a means to automatically learn descriptive low-dimensional features of the raw input data during the *same* posterior or density ratio estimation task in an end-to-end fashion.

This quality is a further attractive feature of NPE and NRE. Finding low-dimensional summary statistics of the data is necessary in many approaches to SBI⁷, and the question of how to summarise data naturally arises in response to this. Popular solutions that have been explored to-date include: the experimenter themselves devising a collection of hand-crafted summary statistics that they believe adequately describes the data, which is an arduous task that can lead to a difficult-to-quantify loss of information; or learning summary statistics by constructing a separate learning task to the existing problem of density (ratio) estimation (see e.g. Fearnhead

⁶Inductive biases are usually realised as constraints on the form of the network architecture in the case of neural networks. For example, when certain symmetries are known to be present in the data, this inductive bias may be incorporated with the use of (partially) exchangeable neural networks (Zaheer et al., 2017; Wqvist et al., 2019). Imposing such constraints helps to restrict the space over which appropriate functions are searched for and facilitates effective learning and generalisation from training data.

⁷This is the prototypical approach to ABC, for example, although methods that obviate the need to summarise high-dimensional data have been proposed and explored in recent years (see e.g. Park et al., 2016; Bernton et al., 2019; Dyer et al., 2021a).

Algorithm 2: Training SNPE (see [Greenberg et al., 2019](#), Algorithm 1).

Input: prior distribution $p(\boldsymbol{\theta})$, simulator $p(\mathbf{x} \mid \boldsymbol{\theta})$, observation \mathbf{y} , conditional density estimator q_ϕ , number of rounds M , number of simulations per round N ;

Result: Trained conditional density estimator

Set $\tilde{p}_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta})$, dataset $\mathcal{D} = \{\}$;

for $m = 0, \dots, M - 1$ **do**

 Sample $\boldsymbol{\theta}^{(n)} \stackrel{iid}{\sim} \tilde{p}_m(\boldsymbol{\theta})$, $n = 1, \dots, N$;

 Simulate $\mathbf{x}^{(n)} \sim p(\mathbf{x} \mid \boldsymbol{\theta}^{(n)})$, $n = 1, \dots, N$;

 Append the dataset:

$$\mathcal{D} := \mathcal{D} \cup \bigcup_{n=1}^N \left\{ \left(\mathbf{x}^{(n)}, \boldsymbol{\theta}^{(n)} \right) \right\};$$

until *convergence* **do**

 Evaluate the loss function

$$\mathcal{L}(\phi) = - \sum_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}} \log \left(q_\phi(\boldsymbol{\theta} \mid \mathbf{x}) \frac{\tilde{p}_m(\boldsymbol{\theta})}{p(\boldsymbol{\theta})} \frac{1}{Z(\mathbf{x}, \phi)} \right); \quad (30)$$

 Update trainable parameters ϕ on the basis of $\mathcal{L}(\phi)$

end

 Set $\tilde{p}_{m+1}(\boldsymbol{\theta}) := q_\phi(\mathbf{y}, \boldsymbol{\theta})$.

end

and [Prangle, 2012](#); [Chen et al., 2020](#)), imposing an additional burden on the experimenter. NPE and NRE, in contrast, offer a convenient approach to incorporating inductive biases and to learning summary statistics in an end-to-end fashion without the additional complications associated with alternative SBI techniques.

3.7. Training procedures for density (ratio) estimators

To summarise and conclude this section, we follow [Greenberg et al. \(2019\)](#) and [Durkan et al. \(2020\)](#) and provide in Algorithms 2 and 3 procedures for training neural posterior and density ratio estimators, respectively, for SBI over multiple rounds. In both cases, we assume that the density (ratio) estimator includes any embedding network used to learn summary statistics concurrently with the density (ratio) estimate, and thus consider ϕ to contain the parameters of both the estimator and the embedding network.

4. Experiments for tractable examples

In this section, we present experiments in which we compare the ability of NPE and NRE to estimate parameter posterior distributions for economic simulation models with tractable likelihood functions against the non-parametric density estimation method described in [Grazzini et al. \(2017\)](#), which we term KDE and outline in Section 2.1.2. We provide details of the neural network architecture and training hyperparameters in Appendix B.

4.1. Performance metrics

Each example presented in this section will be equipped with a tractable transition density, and therefore a tractable likelihood function. Such models are thus amenable to Bayesian inference

Algorithm 3: Training SNRE (see Durkan et al., 2020, Algorithm 1).

Input: prior distribution $p(\boldsymbol{\theta})$, simulator $p(\mathbf{x} \mid \boldsymbol{\theta})$, observation \mathbf{y} , density ratio estimator f_ϕ , number of rounds M , number of simulations per round N , minibatch size B , contrasting set size K ;

Result: Trained density ratio estimator

Set $\tilde{p}_0(\boldsymbol{\theta}) = p(\boldsymbol{\theta})$, dataset $\mathcal{D} = \{\}$;

for $m = 0, \dots, M - 1$ **do**

Sample $\boldsymbol{\theta}^{(n)} \stackrel{iid}{\sim} \tilde{p}_m(\boldsymbol{\theta})$, $n = 1, \dots, N$;

Simulate $\mathbf{x}^{(n)} \sim p(\mathbf{x} \mid \boldsymbol{\theta}^{(n)})$, $n = 1, \dots, N$;

Append the dataset:

$$\mathcal{D} := \mathcal{D} \cup \bigcup_{n=1}^N \left\{ \left(\mathbf{x}^{(n)}, \boldsymbol{\theta}^{(n)} \right) \right\};$$

until *convergence* **do**

Sample minibatch $\left\{ \left(\mathbf{x}^{(b)}, \boldsymbol{\theta}^{(b)} \right) \right\}_{b=1}^B$ uniformly from \mathcal{D} ;

For each minibatch pair, draw $0 < K < B$ parameters $\tilde{\boldsymbol{\theta}}^{(k)}$, $k = 1, \dots, K$, from elsewhere in the training data \mathcal{D} ;

Evaluate the finite-sample multinomial logistic loss

$$\mathcal{L}(\phi) = -\frac{1}{B} \sum_{b=1}^B \log \frac{\exp \left(f_\phi \left(\mathbf{x}^{(b)}, \boldsymbol{\theta}^{(b)} \right) \right)}{\exp \left(f_\phi \left(\mathbf{x}^{(b)}, \boldsymbol{\theta}^{(b)} \right) \right) + \sum_{k=1}^K \exp \left(f_\phi \left(\mathbf{x}^{(b)}, \tilde{\boldsymbol{\theta}}^{(k)} \right) \right)}; \quad (31)$$

Update trainable parameters ϕ on the basis of $\mathcal{L}(\phi)$

end

Set $\tilde{p}_{m+1}(\boldsymbol{\theta}) \propto \exp(f_\phi(\mathbf{y}, \boldsymbol{\theta}))$.

end

via standard means, such as MCMC, to obtain an approximate ground-truth posterior density⁸.

To assess the accuracy of the estimated posteriors in this case, we compare the full ground-truth posterior density with the full posterior estimated with each of the implemented simulation-based approaches. This contrasts with previous studies of Bayesian parameter estimation for ABMs, in which point estimates alone are often used to assess the quality of the tested inference procedures. For example, Platt (2021) uses a prior-weighted Euclidean distance between the estimated posterior mean and generating parameters (that is, the $\boldsymbol{\theta}$ that generated the pseudo-observation \mathbf{y}) as a performance metric. However, such metrics provide a very limited and at times misleading view on the outcome of the inference process, because (a) ‘‘closeness’’ is not measured in the geometry of the target distribution, as visualised in Figure 2; (b) it is possible that an approximate posterior can produce a posterior mean close to the generating parameter but simultaneously under- or over-estimate the width of the distribution or otherwise yield poor uncertainty quantification, and (c) finite datasets are not guaranteed to yield posterior densities that concentrate on the generating parameter.

To compare the ground-truth and simulation-based posteriors, we compute two integral probability metrics which each correspond to a notion of *dissimilarity* between the ground-truth and estimated posteriors:

⁸While the true posterior density is targeted with such a sampling scheme, the ground-truth obtained in this fashion remains an approximation due to the Monte Carlo error associated with the finite sample size.




Figure 2: Visualisation: why standard Euclidean distances are misleading when gauging the performance of Bayesian inference algorithms. The parameter θ_2 is ostensibly closer to the “true parameter” (grey) in this toy posterior than θ_1 . However, θ_1 has higher posterior density, i.e. it is a more credible parameter given the observed data.

Wasserstein distance This is a distance measure derived from optimal transport theory (Kantorovich, 1960) and used widely in various machine learning contexts, e.g. generative adversarial networks (Arjovsky et al., 2017). For some distance ρ_0 on Θ and two probability measures μ and μ' , the p -Wasserstein metric between two sets of samples $\{\theta_i\}_{i=1}^n \stackrel{iid}{\sim} \mu$ and $\{\theta'_i\}_{i=1}^m \stackrel{iid}{\sim} \mu'$ is computed as

$$\mathcal{W}_p(\{\theta_i\}_{i=1}^n, \{\theta'_j\}_{j=1}^m)^p = \inf_{\gamma \in \Gamma_{n,m}} \sum_{i=1}^n \sum_{j=1}^m \rho_0(\theta_i, \theta'_j)^p \gamma_{ij} \quad (32)$$

where $\Gamma_{n,m}$ is the set of $n \times m$ matrices with non-negative entries, columns summing to m^{-1} , and rows summing to n^{-1} . Throughout, we use the Euclidean distance for ρ_0 .

Maximum mean discrepancy (MMD) This metric is once again a metric on probability distributions that draws from the theory of reproducing kernel Hilbert spaces (Gretton et al., 2006, 2012) and is used widely within the machine learning and simulation-based inference community to assess the dissimilarity between probability distributions (Papamakarios et al., 2019; Lueckmann et al., 2021). Here, under a suitable choice of kernel⁹ k chosen by the experimenter, the discrepancy between two probability distributions P and Q is taken to be

$$\text{MMD}(P, Q) = \left\| \mathbb{E}_{\theta \sim P} [k(\theta, \cdot)] - \mathbb{E}_{\theta' \sim Q} [k(\theta', \cdot)] \right\|_{\mathcal{H}}^2, \quad (33)$$

where \mathcal{H} is the reproducing kernel Hilbert space (RKHS) associated with k and $\mathbb{E}_{\theta \sim P} [k(\theta, \cdot)]$ is the so-called *mean embedding* of P via kernel k . When only samples $\{\theta_i\}_{i=1}^n \stackrel{iid}{\sim} P$ and $\{\theta'_i\}_{i=1}^m \stackrel{iid}{\sim} Q$ are available, an unbiased estimator of this metric can be computed as

$$\widehat{\text{MMD}}(P, Q) = \frac{1}{m(m-1)} \sum_{\substack{i,j \\ j \neq i}} k(\theta'_i, \theta'_j) + \frac{1}{n(n-1)} \sum_{\substack{i,j \\ i \neq j}} k(\theta_i, \theta_j) - \frac{2}{nm} \sum_{i,j} k(\theta_i, \theta'_j). \quad (34)$$

Throughout, we use a Gaussian kernel as k ,

$$k(\theta, \theta') = \exp\left(-\frac{\|\theta - \theta'\|_2^2}{2\sigma^2}\right), \quad (35)$$

where the scale parameter $\sigma^2 = \text{median}\{\|\tilde{\theta}_i - \tilde{\theta}_j\|_2^2\}$, following Briol et al. (2019), and $\{\tilde{\theta}_i\}_{i=1}^n$ are the samples drawn from the ground-truth posteriors.

⁹That is: a symmetric, positive semi-definite function on Θ .

4.2. Example 1: Brock and Hommes (1998)

We consider a variant of the model proposed by (Brock and Hommes, 1998), which has previously been used in ABM calibration experiments (Platt, 2020). The model dynamics can be expressed as the following system of coupled equations:

$$x_{t+1} = \frac{1}{R} \left[\sum_{h=1}^H n_{h,t+1} (g_h x_t + b_h) + \epsilon_{t+1} \right], \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2), \quad (36)$$

$$n_{h,t+1} = \frac{\exp(\beta U_{h,t})}{\sum_{h'=1}^H \exp(\beta U_{h',t})}, \quad (37)$$

$$U_{h,t} = (x_t - R x_{t-1}) (g_h x_{t-2} + b_h - R x_{t-1}), \quad (38)$$

where R, β, σ are parameters. We follow Platt (2020) and assume that $H = 4, R = 1.0, \sigma = 0.04, g_1 = b_1 = b_4 = 0$ and $g_4 = 1.01$. β will be 120 or 10, depending on the experiment, and we note below which value is used. We consider the task of estimating the posterior $p(\boldsymbol{\theta} \mid \mathbf{y})$, where $\boldsymbol{\theta} = (g_2, b_2, g_3, b_3)$, $\mathbf{y} := (y_1, y_2, \dots, y_T) \sim p(\mathbf{x} \mid \boldsymbol{\theta}^*)$ is the pseudo-observation, $T = 100$, and $\boldsymbol{\theta}^*$ is the parameter setting used to generate \mathbf{y} . We consider two experimental setups which we describe further in the following two subsections. For each experiment, we note that by rewriting the above system of equations, we are able to find the transition density for observation y_{t+1} as

$$p(y_{t+1} \mid y_{1:t}, \boldsymbol{\theta}) = \mathcal{N} \left(f(y_{t-2:t}, \boldsymbol{\theta}), \frac{\sigma^2}{R^2} \right) \quad (39)$$

where

$$f(y_{t-2:t}, \boldsymbol{\theta}) = \frac{1}{R} \sum_{h=1}^H \frac{\exp[\beta (y_t - R y_{t-1}) (g_h y_{t-2} + b_h - R y_{t-1})]}{\sum_{h'=1}^H \exp[\beta (y_t - R y_{t-1}) (g_{h'} y_{t-2} + b_{h'} - R y_{t-1})]} (g_h y_t + b_h). \quad (40)$$

In this way, we are able to obtain approximate ground truth posteriors with standard MCMC techniques such as Metropolis-Hastings (MH).

For both NPE and NRE, we use a round-based training approach i.e. SNPE and SNRE: we train over 10 rounds and generate 1000 simulations in each round. For KDE, we take $R = 1$ and sample from the posterior with MH (see Appendix A.1 for further details).

Summarising data As discussed in Section 3.6, there exists a number of approaches to representing high-dimensional time-series data \mathbf{x} as a low-dimensional vector $\mathbf{s}(\mathbf{x})$ of summary statistics to facilitate SBI. In the experiments for the Brock & Hommes model, we demonstrate NPE and NRE using summary statistics obtained in two different ways:

1. by summarising them manually as the mean value, variance, maximum, minimum, median, 25th quantile, 75th quantile, and the autocorrelations of the x_t to lags 1, 2, and 3¹⁰. We denote these summary statistics with $\tilde{\mathbf{s}}$ and refer to them as *naive* or *hand-crafted* summary statistics;
2. by learning them as part of the training process with an embedding network \mathbf{s}_φ with trainable parameters φ (see Section 3.6). Through the experiments we present below, we will demonstrate that NPE and NRE can be used flexibly with various embedding networks, and that the experimenter is free to choose from the plethora of candidate networks that incorporate useful inductive biases for time-series data (e.g. Wong et al., 2018; Dyer et al., 2021b; Kidger et al., 2020). Below, we refer to summary statistics obtained in this fashion as *learned* summary statistics.

¹⁰We note however that these are not the only choices available to the experimenter.

Parameter set	Metric	Estimation method				
		KDE	SNPE	SNPE*	SNRE	SNRE*
1	WASSERSTEIN	0.690	0.477	0.336	0.241	<i>0.299</i>
	MMD	1.015	0.789	<i>0.552</i>	0.451	0.781
2	WASSERSTEIN	0.304	0.154	0.306	<i>0.164</i>	0.291
	MMD	0.127	0.036	0.133	<i>0.041</i>	0.118
Simulation budget		1.5×10^5	10^4	10^4	10^4	10^4

Table 1: (Brock & Hommes) Discrepancies between the approximate ground-truth posterior and the posteriors estimated with KDE, SNPE, and SNRE. **Bold** and *italics* indicate best and second-best, respectively. For the neural methods, * indicates that the naive hand-crafted summary statistics described in the main text were used, otherwise summary statistics are learned from the simulated data.

4.2.1. Parameter set 1

We take $\beta = 120$ and consider the task of estimating the posterior density $p(\boldsymbol{\theta} \mid \mathbf{y})$, where $\boldsymbol{\theta}^* := (g_2^*, b_2^*, g_3^*, b_3^*) = (0.9, 0.2, 0.9, -0.2)$. We use the following uniform priors: $g_2, b_2, g_3 \sim \mathcal{U}(0, 1)$, while $b_3 \sim \mathcal{U}(-1, 0)$.

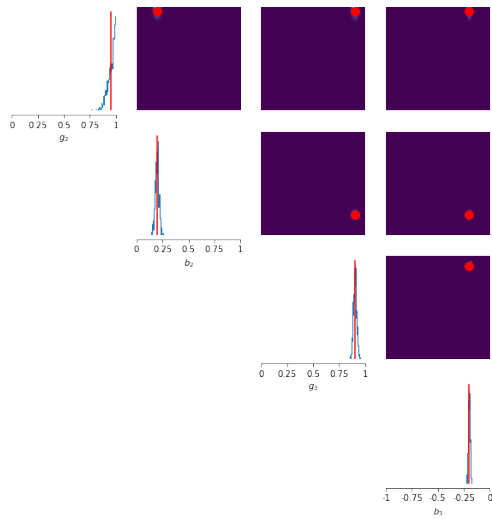
In Figure 3, we show the posteriors obtained using different posterior estimation methods: Figure 3a shows the approximate ground-truth posterior density obtained with the true likelihood function and MH; Figure 3b shows the posterior obtained with KDE and MH; Figures 3c and 3d show the posterior estimated via SNPE with naive and learned summary statistics, respectively; and Figures 3e and 3f show the posterior obtained via SNRE and naive and learned summary statistics, respectively, which were also sampled using MH. Here, we use an embedding network consisting of two stacked Elman recurrent units with hidden state of size 32, followed by a single linear layer of size 16. In each figure, the marginals and joint bivariate densities are located on the diagonal and upper diagonal, respectively, while the red lines/dots locate the mean of the true posterior.

We see from the approximate ground-truth in Figure 3a that the marginal posteriors are sharply peaked on the generating parameters for b_2, g_3 , and b_3 , while the ground truth marginal for g_2 is shifted towards higher values. While these features are recovered reasonably well for Figures 3c–3f (the most notable exception being the posteriors for g_2), we see that they are recovered poorly with KDE. This performance gap is also manifested in the WASSERSTEIN and MMD metrics reported in Table 1, where lower values indicate better estimates of the posterior. In summary, KDE both requires a far larger simulation budget to estimate a single posterior, while simultaneously generating worse estimates of that posterior. In contrast, SNPE and SNRE is able to achieve superior estimates of the ground-truth posterior distribution, despite the 10-fold reduction in the simulation budget they have been afforded.


4.2.2. Parameter set 2

We now take $\beta = 10$ and consider the task of estimating the posterior density $p(\boldsymbol{\theta} \mid \mathbf{y})$, where $\mathbf{y} := (y_1, y_2, \dots, y_T) \sim p(\mathbf{x} \mid \boldsymbol{\theta}^*)$, $T = 100$, and $\boldsymbol{\theta}^* := (g_2^*, b_2^*, g_3^*, b_3^*) = (-0.7, -0.4, 0.5, 0.3)$. In this case, we use the following priors: $g_2, b_2 \sim \mathcal{U}(-1, 0)$ and $g_3, b_3 \sim \mathcal{U}(0, 1)$.

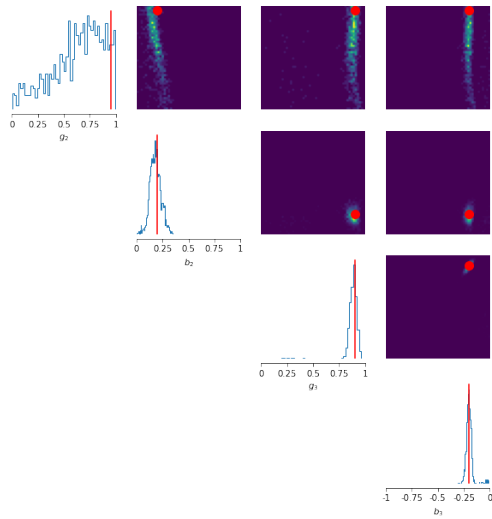
In Figure 4, we show the posteriors obtained using different posterior estimation methods: Figure 4a shows the approximate ground-truth posterior density obtained with the true likelihood function and MH; Figure 4b shows the posterior obtained with KDE and MH; Figures 4c and 4d show the posterior estimated via SNPE with naive and learned summary statistics, respectively; and Figures 4e and 4f show the posterior obtained via SNRE and naive and learned




(a) True posterior




(b) KDE posterior, 1.5×10^5 simulations




(c) Sequential neural posterior estimation with naive summary statistics, 10^4 simulations



(d) Sequential neural posterior estimation with learned summary statistics, 10^4 simulations



(e) Sequential neural ratio estimation with naive summary statistics, 10^4 simulations



(f) Sequential neural ratio estimation with learned summary statistics, 10^4 simulations

Figure 3: (Brock & Hommes, parameter set 1) Posteriors obtained with each posterior estimation method. The marginal posterior distributions are located on the diagonals, while the bivariate joint distributions for each parameter pair are located on the upper diagonal. Red lines/dots indicate the mean of the true posterior.

summary statistics, respectively, which were also sampled with MH. In this experiment, we now use an embedding network consisting of two stacked gated recurrent units (GRUs) with hidden state of size 32, followed by a single linear layer of size 16. This results in an embedding network with approximately 10,000 trainable parameters.

We see from the approximate ground-truth in Figure 4a that the marginal posteriors for b_2 and b_3 remain sharply peaked on the generating parameters, but that the ground truth marginals for g_2 and g_3 are diffuse and sloped. The diffuseness of the posterior with respect to these two dimensions is also visually apparent in the joint density plots in the upper diagonal. Upon inspection of the estimated posteriors, we see that the overall shape is recovered well by SNPE and SNRE with learned summary statistics, but less accurately by KDE and SNPE and SNRE when using the naive hand-crafted summary statistics described in Section 4.2. This is once again reflected by the WASSERSTEIN and MMD scores, reported in Table 1, in which we see significantly decreased values for SNPE and SNRE with learned summary statistics with respect to the alternatives, despite a 10-fold decrease in the simulation budget they have been afforded. It is nonetheless noteworthy that even with the naive hand-crafted summary statistics, SNPE and SNRE achieve comparable and slightly favourable performance than KDE, again with a 10-fold decrease in the allotted simulation budget.

4.3. Example 2: Multivariate Geometric Brownian Motion

In Section 4.2, we demonstrated that NPE and NRE can be used to generate Bayesian parameter posteriors that better match the ground truth posterior than KDE for the univariate Brock & Hommes model, despite the fact that the latter method entailed 10 times as many simulations as the former two methods based on neural networks. In this section, we seek to demonstrate that this remains the case even for models that generate multivariate time-series as output.

To this end, we consider a multivariate geometric Brownian motion (MVGBM), which is used in a variety of applications in financial time-series modelling. The model is a stochastic differential equation in which each of the d components, labelled $i \in \{1, 2, \dots, d\}$, of the path $\mathbf{X}_t = (X_t^1, X_t^2, \dots, X_t^d) \in \mathbb{R}_+^d$ evolve according to

$$dX_t^i = X_t^i \left(\left[b_i - \frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 \right] dt + \sum_{j=1}^d \sigma_{ij} dW_t^j \right), \quad (41)$$

where the b_i are drift coefficients, the σ_{ij} are volatility coefficients, and W_t^i is a Brownian motion.

We consider the case of $d = 3$ and the task of estimating the posterior for the parameters $\boldsymbol{\theta} = (b_1, b_2, b_3)$ given an observation $\mathbf{y} \sim p(\mathbf{x} | \boldsymbol{\theta}^*)$ of $T = 100$ points spaced equally with spacing $\Delta t = 1/(T - 1)$, where $\boldsymbol{\theta}^* = (0.2, -0.5, 0.0)$. We take priors $b_i \sim \mathcal{U}(-1, 1)$ for each $i = 1, 2, 3$. We note that this model once again permits both exact simulations and samples from the exact posterior since the transition density is once again tractable¹¹ and can be written as

$$\mathbf{X}_{t+\Delta t} \sim \mathcal{N}(\mathbf{X}_t + (\boldsymbol{\theta} - \boldsymbol{\gamma}) \Delta t, \boldsymbol{\sigma} \boldsymbol{\sigma}^T \Delta t), \quad (42)$$


where

$$\boldsymbol{\gamma} = \frac{1}{2} \left[\sum_{j=1}^d \sigma_{1j}^2, \sum_{j=1}^d \sigma_{2j}^2, \sum_{j=1}^d \sigma_{3j}^2 \right]'. \quad (43)$$

In our experiments, we take


$$\boldsymbol{\sigma} = \begin{pmatrix} 0.5 & 0.1 & 0.0 \\ 0.0 & 0.1 & 0.3 \\ 0.0 & 0.0 & 0.2 \end{pmatrix}. \quad (44)$$

¹¹Assuming that $\boldsymbol{\sigma}$ is of full rank.



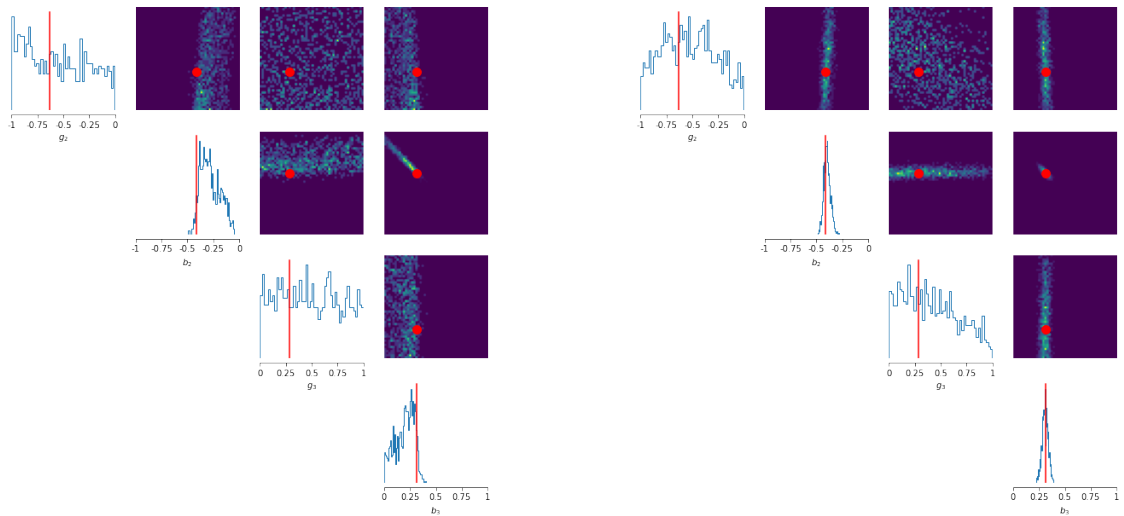
(a) True posterior

(b) KDE posterior, 10^5 simulations



(c) Sequential neural posterior estimation with naive summary statistics, 10^4 simulations

(d) Sequential neural posterior estimation with learned summary statistics, 10^4 simulations



(e) Sequential neural ratio estimation with naive summary statistics, 10^4 simulations

(f) Sequential neural ratio estimation with learned summary statistics, 10^4 simulations

Figure 4: (Brock & Hommes, parameter set 2) Posteriors obtained with each posterior estimation method. The marginal posterior distributions are located on the diagonals, while the bivariate joint distributions for each parameter pair are located on the upper diagonal. Red lines/dots indicate the mean of the true posterior.

Metric	Estimation method		
	KDE	NPE	NRE
WASSERSTEIN	0.364	0.099	<i>0.107</i>
MMD	0.137	<i>0.005</i>	0.004
Simulation budget	10^6	10^3	10^3

Table 2: (Multivariate geometric Brownian motion) Discrepancies between the approximate ground-truth posterior and the posteriors estimated with NPE, NRE, and KDE. Smaller values indicate more accurate posteriors; **bold** and *italics* indicate best and second-best, respectively.

In Figure 5a we show the approximate ground truth posterior obtained with MH and the transition density described above, while in Figures 5b, 5c, and 5d we show the posteriors obtained with KDE and MH, NPE with learned summary statistics, and NRE with MH also with learned summary statistics, respectively. The corresponding simulation budgets are 10^6 , 10^3 , and 10^3 , respectively. To learn the summary statistics, we once again use the embedding network described in Section 4.2.2 and learn summary statistics and the density (ratio) estimator concurrently.

We see that the approximate ground-truth posteriors are relatively diffuse, with peaks approximately coinciding with the true posterior mean, shown with red lines/dots. The shape and degree of diffuseness is captured accurately by NPE and NRE. In contrast, the posterior obtained with KDE is insufficiently diffuse and biased, and thus a significantly worse estimate of the ground-truth posterior. These observations are corroborated by the corresponding WASSERSTEIN and MMD metrics, reported in Table 2 along with the corresponding simulation budgets. In summary, NPE and NRE achieve significantly more accurate posterior estimates here with a 1000-fold decrease in the simulation budget, and are able to flexibly accommodate multivariate time-series as input for the inference problem.

5. Validating approximate Bayesian inference

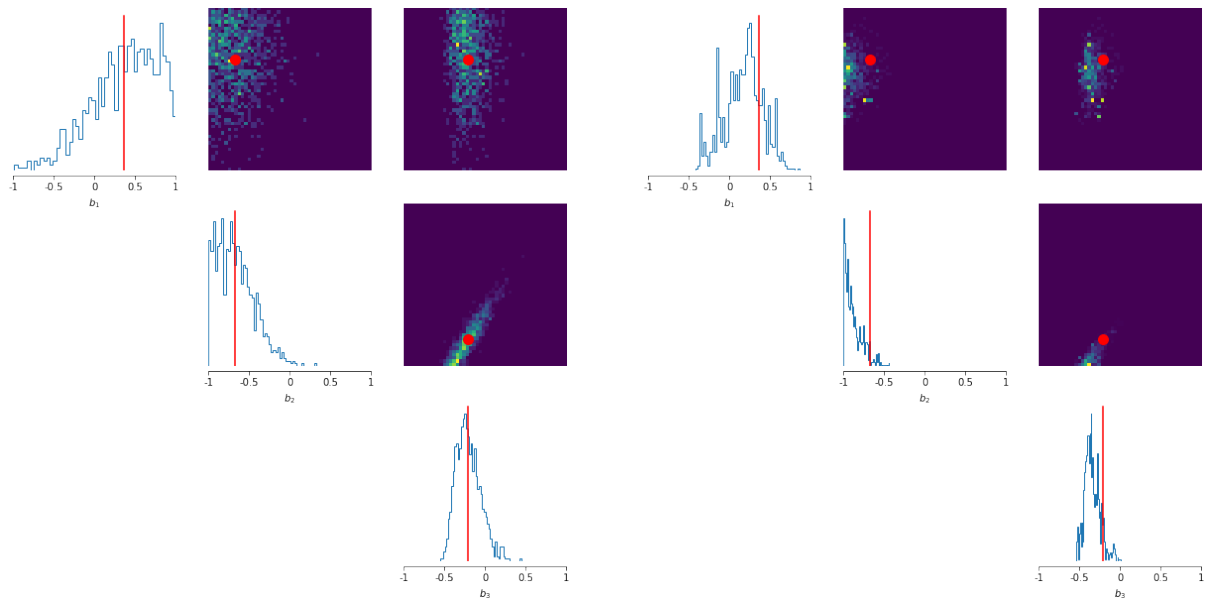
As previously mentioned, a major benefit of the Bayesian inferential paradigm is the fact that uncertainty quantification is a built-in feature captured via the posterior distribution. Its utility relies, however, on the ability to capture the correct degree of diffuseness in the posterior, such that the uncertainty quantification in the recovered posterior is meaningful. This presents a challenge in approximate Bayesian inference, since by definition the experimenter does not know the correct shape of the posterior when exact inference is intractable. In this section, we address this issue by outlining one widely used approach to verifying the accuracy of approximate Bayesian inference pipelines in the absence of any ground-truth posterior densities.

5.1. Simulation-based calibration

Simulation-based calibration is a general purpose method for validating approximate Bayesian inference pipelines. The core idea is to use the fact that the *data-averaged posterior* should be identical to the prior distribution $p(\theta)$. That is, a perfect Bayesian inference pipeline should uphold the following equality:

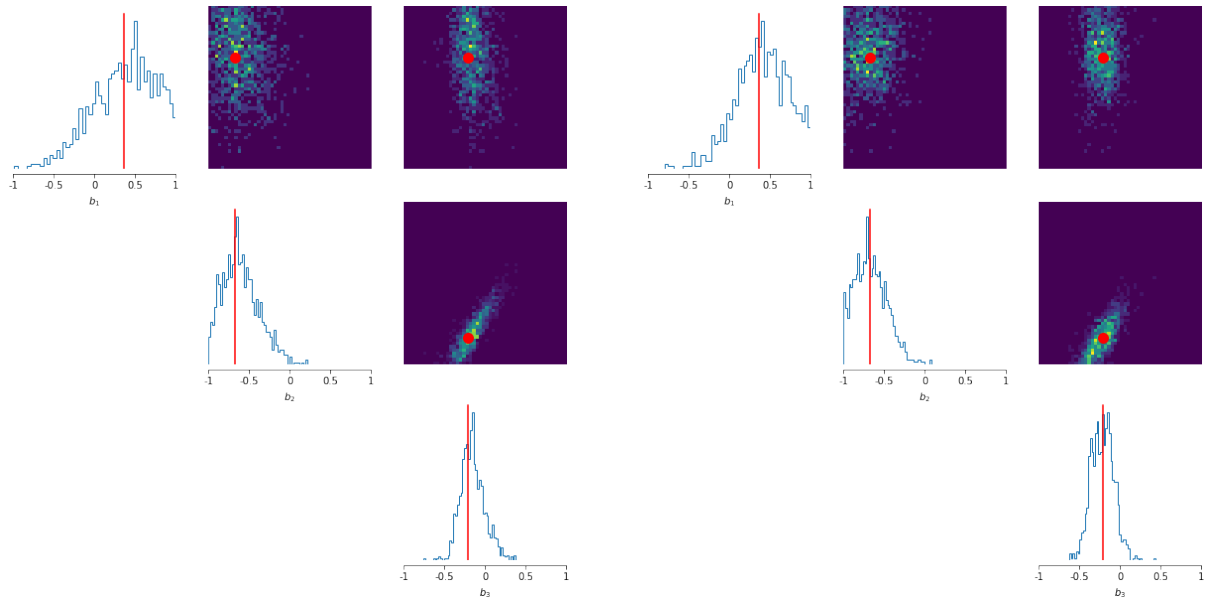
$$\int_{\mathcal{Y} \times \Theta} p(\theta | \mathbf{y}) p(\mathbf{y} | \tilde{\theta}) p(\tilde{\theta}) d\tilde{\theta} d\mathbf{y} = p(\theta). \quad (45)$$

A deviation from this equality signifies some error in the posterior sampling procedure; thus, it has been proposed by Talts et al. (2020) that testing how close our Bayesian workflow is to



(a) True posterior

(b) KDE posterior, 10^6 simulations



(c) Neural posterior estimation with automatically learned summary statistics, 10^3 simulations

(d) Neural ratio estimation with automatically learned summary statistics, 10^3 simulations

Figure 5: (Multivariate geometric Brownian motion) Posteriors obtained with each estimation method. The marginal posterior distributions are located on the diagonals, while the bivariate joint distributions for each parameter pair are located on the upper diagonal. Red lines/dots indicate the mean of the true posterior.

satisfying this equality is a test of the accuracy of the Bayesian pipeline. This may be performed by repeating the following steps P times:

1. Generate $\tilde{\boldsymbol{\theta}} \sim p(\boldsymbol{\theta})$ from the prior distribution
2. Generate $\tilde{\mathbf{y}} \sim p(\mathbf{y} | \tilde{\boldsymbol{\theta}})$ from the likelihood function (i.e. the simulator)
3. Generate L uncorrelated posterior samples $\{\boldsymbol{\theta}_i\}_{i=1}^L \sim p(\boldsymbol{\theta} | \tilde{\mathbf{y}})$
4. Compute and store the rank statistic $r(\{\boldsymbol{\theta}_i\}_{i=1}^L, \tilde{\boldsymbol{\theta}}) \in \{0, \dots, L\}$ for $\tilde{\boldsymbol{\theta}}$ within $\{\boldsymbol{\theta}_i\}_{i=1}^L \cup \{\tilde{\boldsymbol{\theta}}\}$

It can be shown (Theorem 1, [Talts et al., 2020](#)) that the rank statistics obtained as above should follow a discrete Uniform distribution on $\{0, \dots, L\}$ for any joint distribution $p(\mathbf{y} | \boldsymbol{\theta})p(\boldsymbol{\theta})$ for one-dimensional $\boldsymbol{\theta}$, while for d -dimensional $\boldsymbol{\theta}$ each of the d components' rank histograms should be uniform. Inspecting the discrepancy between the true distribution of rank statistics and the desired Uniform distribution thus gives an indication of the accuracy of the Bayesian pipeline; furthermore and conversely, particular patterns of deviation from the desired uniformity are interpretable and provide insight into the specific way in which the obtained posteriors are inaccurate ([Talts et al., 2020](#)). Performing this procedure, and inspecting the resultant rank histograms, is referred to as performing simulation-based calibration (SBC).

While this procedure is in principle possible for KDE, the computational burden is immense even for cheap simulation models, and becomes infeasible for large-scale ABMs. This is because the number of simulations required to generate n samples from the posterior would, in general, be at least¹² nR , such that the total number of simulations required to perform SBC for KDE increases to PnR . Even for a moderately sized SBC task with $P \simeq 10^3$ and the most optimistic $R = 1$, the total number of simulations required can easily reach the order of 10^8 , since it is typically necessary to take n to be many hundreds of thousands to obtain a reasonably accurate estimate of the posterior. In fact, larger values of n are likely to be necessary, since ABMs are usually highly parameterised and thus can have large parameter spaces.

In contrast, such an approach to verifying the accuracy of posteriors derived from NPE and NRE is feasible due to the fact that such approaches involve the prior training of a global density (ratio) estimator, obviating any simulations during the inference (posterior sampling) phase. Thus the strength of NPE and NRE derives from not only their enhanced performance and improved simulation-efficiency, but also from the fact that such methods permit accuracy checks such as SBC that are otherwise infeasible for alternative parameter estimation techniques when the simulator is expensive, as is often the case for large-scale ABMs.

5.1.1. Example: [Franke and Westerhoff \(2012\)](#)


To further demonstrate the ability of NPE and NRE to recover accurate posterior estimates and informative summary statistics in a highly automated manner, we perform SBC using a posterior and density ratio estimator trained on the Franke & Westerhoff model ([Franke and Westerhoff, 2012](#)), which has previously been used in benchmarking experiments for ABM estimation methods ([Platt, 2021](#)). In particular, we consider the model version referred to as the ‘‘Wealth & Predisposition’’ model which, similarly to the Brock & Hommes model (see [Section 4.2](#)), may be written as a system of coupled equations which models the asset price dynamics that result from a heterogeneous system of traders:

$$p_t = p_{t-1} + \mu \left(n_{t-1}^f d_{t-1}^f + n_{t-1}^c d_{t-1}^c \right), \quad (46)$$


$$d_t^f = \phi (p^* - p_t) + \sigma_f \epsilon_t^f, \quad (47)$$

$$d_t^c = \chi (p_t - p_{t-1}) + \sigma_c \epsilon_t^c, \quad (48)$$

¹²In practice, this number may be larger due to the necessity of performing trial runs to estimate the parameters of the proposal distribution.



(a) Simulation-based calibration histogram for neural posterior estimation.



(b) Simulation-based calibration histogram for neural density ratio estimation.

Figure 6: (Franke & Westerhoff) Rank histograms generated according to the simulation-based calibration procedure in Section 5.1 using the trained posterior density estimator (top) and density ratio estimator (bottom).

$$n_t^f = 1 - n_t^c = \frac{1}{1 + \exp(-\beta a_{t-1})}, \quad (49)$$

$$a_t = \alpha_w (w_t^f - w_t^c) + \alpha_0, \quad (50)$$

$$w_t^j = \eta w_{t-1}^j + (1 - \eta) g_t^j, \quad j \in \{f, c\} \quad (51)$$

$$g_t^j = [e^{p_t} - e^{p_{t-1}}] d_{t-2}^j, \quad j \in \{f, c\} \quad (52)$$

where the $\epsilon_t^j \sim \mathcal{N}(0, 1)$ are standard normal random variables and $j = \{f, c\}$ labels fundamentalist and chartist traders, respectively. We take as output from the model the time series of log returns $r_t = p_t - p_{t-1}$, and assume the task of training a global, amortised density and density ratio estimator via NPE and NRE, respectively, using 10^4 simulations of length $T = 100$ from the simulator defined by Equations (46)–(52) with the following parameter settings: $\mu = 0.01$, $\beta = 1$, $\phi = 1$, $\chi = 0.9$, $\alpha_0 = 2.1$, $\sigma_f = 0.752$. The parameter vector for which we seek the posterior and density ratio estimators is taken to be $\theta = (\alpha_w, \eta, \sigma_c)$, and we obtain *iid* samples from the posterior associated with the density ratio estimator via a sampling-importance-resampling scheme described in Appendix A.2. (Samples may be drawn *iid* from the NPE posterior by construction.)

We show in Figure 6 the rank histograms obtained via the simulation-based calibration procedure in Section 5.1 with NPE (Figure 6a) and NRE (Figure 6b). To construct these histograms, we take $P = 5,000$ and the following uniform priors: $\alpha_w \sim \mathcal{U}(0, 15000)$, $\eta \sim \mathcal{U}(0, 1)$, and $\sigma_c \sim \mathcal{U}(0, 5)$. The black line and gray band denotes the expected height and the expected variation in the heights, respectively, of the bars at this P and with this number of bins. We see that the bars tend to lie within the expected range, although some notable deviations exist: for example, the rank histograms for α_w for both NPE and NRE exhibit a minor bias towards larger rank values, which suggests that the marginal posteriors for α_w in both cases is slightly biased towards lower values on average than the true posteriors (Talts et al., 2020). However, we emphasise that a major benefit of NPE and NRE is that they allow the modeller to make statements of this sort in the first place, whereas the number of simulations required to make

such statements in the case of more traditional techniques such as KDE would quickly become prohibitively large.

6. Conclusion

In this paper, we investigated the use of two recently developed approaches to Bayesian parameter estimation for intractable simulation models: neural posterior estimation (NPE) for approximating the posterior density directly, and neural ratio estimation (NRE) for approximating the likelihood-to-evidence ratio. We motivated the use of these parameter inference methods as simulation-efficient black-box discriminative approaches to Bayesian estimation for complex time-series simulators such as agent-based models in economics, which contrast with existing methods that (a) entail an often prohibitively large simulation burden, (b) can make restrictive assumptions about the form of the simulator, for example by ignoring important structural features such as temporal dependencies, and (c) are inherently generative, and thus assume a task that is typically more difficult than discrimination alone. We argue that this latter point is likely to be particularly pertinent for the case of agent-based models in the social sciences, since they are known to be able to generate complex and stochastic non-equilibrium dynamics that can be difficult to model and generate.

We further reviewed existing alternatives, and benchmarked NPE and NRE against the most popular of these. In these examples, we see that NPE and NRE generally achieve superior recovery of the approximate ground-truth posterior distributions, despite requiring simulation budgets that are orders of magnitude lower than traditional alternatives. In addition, we demonstrated that a verification of the accuracy of the density (ratio) estimators is possible with the introduced methods via simulation-based calibration (SBC), due to the amortisation of the estimators. In contrast, such checks would necessitate an unreasonably large number of simulations for existing methods in the literature for Bayesian estimation of economic simulation models, for example the methods discussed by [Grazzini et al. \(2017\)](#). For these reasons, we argue that such simulation-efficient black-box Bayesian inference as NPE and NRE may enable economists and social scientists alike to more readily exploit the potential benefits of the agent-based modelling paradigm.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- Justin Alsing, Tom Charnock, Stephen Feeney, and Benjamin Wandelt. Fast likelihood-free cosmology with neural density estimators and active learning. *Monthly Notices of the Royal Astronomical Society*, 488(3):4440–4458, 07 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz1960. URL <https://doi.org/10.1093/mnras/stz1960>.
- Christophe Andrieu and Gareth O Roberts. The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/arjovsky17a.html>.
- Mark A Beaumont. Approximate bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406, 2010.
- Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- Mark A Beaumont, Jean-Marie Cornuet, Jean-Michel Marin, and Christian P Robert. Adaptive approximate bayesian computation. *Biometrika*, 96(4):983–990, 2009.
- Espen Bernton, Pierre E. Jacob, Mathieu Gerber, and Christian P. Robert. Approximate Bayesian computation with the Wasserstein distance. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 81(2):235–269, 2019. ISSN 14679868. doi: 10.1111/rssb.12312.
- Christopher M Bishop. Mixture density networks. 1994.
- Michael GB Blum, Maria Antonieta Nunes, Dennis Prangle, and Scott A Sisson. A comparative review of dimension reduction methods in approximate bayesian computation. *Statistical Science*, 28(2):189–208, 2013.
- Luke Bornn, Natesh S Pillai, Aaron Smith, and Dawn Woodard. The use of a single pseudo-sample in approximate bayesian computation. *Statistics and Computing*, 27(3):583–590, 2017.
- Johann Brehmer, Kyle Cranmer, Gilles Louppe, and Juan Pavez. Constraining effective field theories with machine learning. *Phys. Rev. Lett.*, 121:111801, Sep 2018. doi: 10.1103/PhysRevLett.121.111801. URL <https://link.aps.org/doi/10.1103/PhysRevLett.121.111801>.
- François Xavier Briol, Alessandro Barp, Andrew B. Duncan, and Mark Girolami. Statistical inference for generative models with maximum mean discrepancy. *arXiv*, pages 1–57, 2019.
- William A. Brock and Cars H. Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8):1235–1274, 1998. ISSN 0165-1889. doi: [https://doi.org/10.1016/S0165-1889\(98\)00011-6](https://doi.org/10.1016/S0165-1889(98)00011-6). URL <https://www.sciencedirect.com/science/article/pii/S0165188998000116>.

- Yanzhi Chen, Dinghuai Zhang, Michael Gutmann, Aaron Courville, and Zhanxing Zhu. Neural Approximate Sufficient Statistics for Implicit Models. pages 1–14, 2020. URL <http://arxiv.org/abs/2010.10079>.
- G Darmais. Sur les lois de probabilités à estimation exhaustive. In *C. R. Acad. Sci. Paris (in French)*, volume 200, page 1265–1266, 1935.
- Peter J Diggle and Richard J Gratton. Monte carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 46(2):193–212, 1984.
- Arnaud Doucet, Michael K Pitt, George Deligiannidis, and Robert Kohn. Efficient implementation of markov chain monte carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313, 2015.
- Conor Durkan, Iain Murray, and George Papamakarios. On Contrastive Learning for Likelihood-free Inference. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2771–2781. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/durkan20a.html>.
- Joel Dyer, Patrick Cannon, and Sebastian M Schmon. Approximate Bayesian Computation with Path Signatures. *arXiv preprint arXiv:2106.12555*, 2021a.
- Joel Dyer, Patrick W Cannon, and Sebastian M Schmon. Deep signature statistics for likelihood-free time-series models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021b.
- Paul Fearnhead and Dennis Prangle. Constructing summary statistics for approximate Bayesian computation: Semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 74(3):419–474, 2012. ISSN 13697412. doi: 10.1111/j.1467-9868.2011.01010.x.
- Reiner Franke. Applying the method of simulated moments to estimate a small agent-based asset pricing model. *Journal of Empirical Finance*, 16(5):804–815, 2009.
- Reiner Franke and Frank Westerhoff. Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36(8):1193–1211, 2012.
- Andrew Gelman, Gareth O Roberts, Walter R Gilks, et al. Efficient metropolis jumping rules. *Bayesian statistics*, 5(599-608):42, 1996.
- Pedro J Gonçalves, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, Tim P Vogels, et al. Training deep neural density estimators to identify mechanistic models of neural dynamics. *Elife*, 9:e56261, 2020.
- Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of applied econometrics*, 8(S1):S85–S118, 1993.
- Jakob Grazzini, Matteo G Richiardi, and Mike Tsionas. Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 77:26–47, 2017.
- David S. Greenberg, Marcel Nonnenmacher, and Jakob H. Macke. Automatic posterior transformation for likelihood-free inference. *36th International Conference on Machine Learning, ICML 2019*, 2019-June:4288–4304, 2019.

- Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. A kernel method for the two-sample-problem. *Advances in neural information processing systems*, 19:513–520, 2006.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012. URL <http://jmlr.org/papers/v13/gretton12a.html>.
- W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- Joeri Hermans, Volodimir Begy, and Gilles Louppe. Likelihood-free MCMC with amortized approximate ratio estimators. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4239–4248. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/hermans20a.html>.
- Nianqiao Ju, Jeremy Heng, and Pierre E Jacob. Sequential monte carlo algorithms for agent-based models of disease transmission. *arXiv preprint arXiv:2101.12156*, 2021.
- L. V. Kantorovich. Mathematical methods of organizing and planning production. *Management Science*, 6(4):366–422, 1960. ISSN 00251909, 15265501. URL <http://www.jstor.org/stable/2627082>.
- Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *arXiv preprint arXiv:2005.08926*, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.
- B. O. Koopman. On distributions admitting a sufficient statistic. *Transactions of the American Mathematical Society*, 39(3):399–409, 1936. ISSN 00029947. URL <http://www.jstor.org/stable/1989758>.
- Jiri Kukacka and Jozef Barunik. Estimation of financial agent-based models with simulated maximum likelihood. *Journal of Economic Dynamics and Control*, 85:21–45, 2017.
- Juliane Liepe, Paul Kirk, Sarah Filippi, Tina Toni, Chris P Barnes, and Michael PH Stumpf. A framework for parameter estimation and model selection from experimental data in systems biology using approximate bayesian computation. *Nature protocols*, 9(2):439–456, 2014.
- Jan-Matthis Lueckmann, Pedro J Goncalves, Giacomo Bassetto, Kaan Öcal, Marcel Nonnenmacher, and Jakob H Macke. Flexible statistical inference for mechanistic models of neural dynamics. *Advances in Neural Information Processing Systems*, 30, 2017.
- Jan-Matthis Lueckmann, Jan Boelts, David Greenberg, Pedro Goncalves, and Jakob Macke. Benchmarking simulation-based inference. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 343–351. PMLR, 13–15 Apr 2021.
- Thomas Lux. Estimation of agent-based models using sequential monte carlo methods. *Journal of Economic Dynamics and Control*, 91:391–408, 2018.

- Thomas Lux. Bayesian estimation of agent-based models via adaptive particle markov chain monte carlo. *Computational Economics*, pages 1–27, 2021.
- Nick Malleon, Kevin Minors, Le-Minh Kieu, Jonathan A Ward, Andrew West, and Alison Heppenstall. Simulating crowds in real time with agent-based modelling and a particle filter. *Journal of Artificial Societies and Social Simulation*, 23(3), 2020.
- Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147, 2019. doi: 10.1126/science.aaw1147. URL <https://www.science.org/doi/abs/10.1126/science.aaw1147>.
- George Papamakarios and Iain Murray. Fast ε -free inference of simulation models with bayesian conditional density estimation. In *Advances in neural information processing systems*, pages 1028–1036, 2016.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 2335–2344, 2017.
- George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848. PMLR, 2019.
- Mijung Park, Wittawat Jitkrittum, and Dino Sejdinovic. K2-ABC: Approximate bayesian computation with kernel embeddings. *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016*, 41:398–407, 2016.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- Edwin James George Pitman. Sufficient statistics and intrinsic accuracy. In *Mathematical Proceedings of the cambridge Philosophical society*, volume 32, pages 567–579. Cambridge University Press, 1936.
- Donovan Platt. A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, 113:103859, 2020. ISSN 0165-1889. doi: <https://doi.org/10.1016/j.jedc.2020.103859>. URL <https://www.sciencedirect.com/science/article/pii/S0165188920300294>.
- Donovan Platt. Bayesian estimation of economic simulation models using neural networks. *Computational Economics*, pages 1–52, 2021.
- Leah F Price, Christopher C Drovandi, Anthony Lee, and David J Nott. Bayesian synthetic likelihood. *Journal of Computational and Graphical Statistics*, 27(1):1–11, 2018.
- Jonathan K Pritchard, Mark T Seielstad, Anna Perez-Lezaun, and Marcus W Feldman. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular biology and evolution*, 16(12):1791–1798, 1999.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/rezende15.html>.

- Gareth O Roberts, Andrew Gelman, and Walter R Gilks. Weak convergence and optimal scaling of random walk metropolis algorithms. *The annals of applied probability*, 7(1):110–120, 1997.
- Sebastian M Schmon and Philippe Gagnon. Optimal scaling of random walk metropolis algorithms using bayesian large-sample asymptotics. *Statistics and Computing*, forthcoming, 2022.
- Sebastian M Schmon, Patrick W Cannon, and Jeremias Knoblauch. Generalized posteriors in approximate bayesian computation. <https://arxiv.org/abs/2011.08644>, 2020.
- Sebastian M Schmon, George Deligiannidis, Arnaud Doucet, and Michael K Pitt. Large-sample asymptotics of the pseudo-marginal method. *Biometrika*, 108(1):37–51, 2021.
- Chris Sherlock, Alexandre H Thiery, Gareth O Roberts, and Jeffrey S Rosenthal. On the efficiency of pseudo-marginal random walk metropolis algorithms. *The Annals of Statistics*, 43(1):238–275, 2015.
- Chris Sherlock, Alexandre H Thiery, and Anthony Lee. Pseudo-marginal metropolis–hastings sampling using averages of unbiased estimators. *Biometrika*, 104(3):727–734, 2017.
- Takashi Shiono. Estimation of agent-based models using bayesian deep learning approach of bayesflow. *Journal of Economic Dynamics and Control*, 125:104082, 2021.
- Bernard W Silverman. *Density Estimation for Statistics and Data Analysis*, volume 26. CRC Press, 1986.
- Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate bayesian computation. *PLoS computational biology*, 9(1):e1002803, 2013.
- Esteban G Tabak and Cristina V Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- Esteban G Tabak and Eric Vanden-Eijnden. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233, 2010.
- Sean Talts, Michael Betancourt, Daniel Simpson, Aki Vehtari, and Andrew Gelman. Validating bayesian inference algorithms with simulation-based calibration, 2020.
- Simon Tavaré, David J Balding, Robert C Griffiths, and Peter Donnelly. Inferring coalescence times from dna sequence data. *Genetics*, 145(2):505–518, 1997.
- Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonçalves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, 2020. doi: 10.21105/joss.02505. URL <https://doi.org/10.21105/joss.02505>.
- Owen Thomas, Ritabrata Dutta, Jukka Corander, Samuel Kaski, and Michael U. Gutmann. Likelihood-Free Inference by Ratio Estimation. *Bayesian Analysis*, pages 1 – 31, 2021. doi: 10.1214/20-BA1238. URL <https://doi.org/10.1214/20-BA1238>.
- Tina Toni, David Welch, Natalja Strelkowa, Andreas Ipsen, and Michael PH Stumpf. Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6(31):187–202, 2009.
- Richard David Wilkinson. Approximate bayesian computation (abc) gives exact results under the assumption of model error. *Statistical applications in genetics and molecular biology*, 12(2):129–141, 2013.

- Samuel Wıqvıst, Pierre-Alexandre Matteı, Umberto Picchini, and Jes Frelsen. Partially exchangeable networks and architectures for learning summary statistics in approximate Bayesian computation. In *International Conference on Machine Learning*, pages 6798–6807. PMLR, 2019.
- Wing Wong, Bai Jiang, Tung-yu Wu, and Charles Zheng. Learning summary statistic for approximate bayesian computation via deep neural network. *Statistica Sinica*, 2018. ISSN 1017-0405. doi: 10.5705/ss.202015.0340. URL <http://dx.doi.org/10.5705/ss.202015.0340>.
- Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf>.

A. Posterior sampling

A.1. Sampling with Metropolis-Hastings

The Metropolis-Hastings algorithm is a classical algorithm for generating samples from some density $p(\boldsymbol{\theta})$. Starting from $\boldsymbol{\theta}_0$ and given a proposal distribution $q(\cdot | \boldsymbol{\theta})$ which proposes successive values in the chain, it entails repeating the following steps: at each step $t \geq 1$,

1. propose $\boldsymbol{\theta} \sim q(\cdot | \boldsymbol{\theta}_t)$;
2. accept $\boldsymbol{\theta}$ (i.e. set $\boldsymbol{\theta}_{t+1} := \boldsymbol{\theta}$) with probability

$$\alpha = \min \left\{ 1, \frac{p(\boldsymbol{\theta})q(\boldsymbol{\theta}_t | \boldsymbol{\theta})}{p(\boldsymbol{\theta}_t)q(\boldsymbol{\theta} | \boldsymbol{\theta}_t)} \right\},$$

else set $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t$.

Since such a procedure generates correlated samples from the posterior, it is customary to *thin* the samples by retaining every n th sample for some integer $n \geq 1$ chosen as required.

To sample from the posteriors obtained via KDE and NRE in Section 4, we use Metropolis-Hastings (MH) with a normal proposal distribution $q(\cdot | \boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \ell^2 \Sigma)$, and perform a trial run of 50,000 steps using an isotropic Gaussian to estimate the covariance matrix Σ of q before a further 100,000 steps are run. The use of such pilot runs is long standing-practice for random walk MH algorithms and can be motivated theoretically, for example, using Bayesian asymptotics (Schmon and Gagnon, 2022). In addition, we set $\ell = 2/\sqrt{d}$, where d is the parameter dimension following the guidelines of Gelman et al. (1996); Roberts et al. (1997); Schmon and Gagnon (2022). All chains are initialised at the parameter value which generated the observation. We thin the resultant chains by retaining every 100th value, resulting in 1,000 approximately uncorrelated samples from the respective posteriors.

A.2. Sampling with sampling-importance-resampling

Sampling-importance-resampling (SIR) is an approach to obtaining samples from a target distribution $f(\boldsymbol{\theta})$ given samples from a different distribution $g(\boldsymbol{\theta})$ which proceeds as follows:

1. generate samples $\{\boldsymbol{\theta}_i\}_{i=1}^N \stackrel{iid}{\sim} g(\boldsymbol{\theta})$;
2. compute weights $w_i = f(\boldsymbol{\theta}_i)/g(\boldsymbol{\theta}_i)$ for all i ;
3. resample $\boldsymbol{\theta}_i$ with probability proportional to w_i .

In particular, this may be used in density ratio estimation to approximate the posterior by setting $g(\boldsymbol{\theta}) := p(\boldsymbol{\theta})$, $f(\boldsymbol{\theta}) := p(\boldsymbol{\theta} | \mathbf{y})$, and using that the density ratio estimator estimates the ratio

$$\frac{f(\boldsymbol{\theta})}{g(\boldsymbol{\theta})} = \frac{p(\boldsymbol{\theta} | \mathbf{y})}{p(\boldsymbol{\theta})} = \frac{p(\mathbf{y} | \boldsymbol{\theta})}{p(\mathbf{y})}. \quad (53)$$

B. Neural network architectures and training

For each neural network method, the z -scores of all variables are taken prior to passing them into the networks. For all neural posterior estimation tasks, we use a masked autoregressive flow (Papamakarios et al., 2017) with 5 flow transforms, each with 2 blocks and 50 hidden features; for all neural density ratio estimation tasks, we use a residual network with two layers of size 50. To train the network weights, we use Adam (Kingma and Ba, 2014), along with a training batch size of 50 and learning rate of 5×10^{-4} . We furthermore reserve 10% of the data for validation, and stop training when the validation error does not improve over 20 epochs to avoid overfitting. Throughout, we use the `sbi` python package (Tejero-Cantero et al., 2020).