

Optimal Thinning of MCMC Output

Marina Riabiz^{1,2}, Wilson Ye Chen³, Jon Cockayne², Pawel Swietach⁴,
Steven A. Niederer¹, Lester Mackey⁵, Chris. J. Oates^{6,2*}

¹King’s College London, UK ²Alan Turing Institute, UK

³University of Sydney, Australia ⁴Oxford University, UK

⁵Microsoft Research, US ⁶Newcastle University, UK

September 21, 2021

Abstract

The use of heuristics to assess the convergence and compress the output of Markov chain Monte Carlo can be sub-optimal in terms of the empirical approximations that are produced. Typically a number of the initial states are attributed to “burn in” and removed, whilst the remainder of the chain is “thinned” if compression is also required. In this paper we consider the problem of retrospectively selecting a subset of states, of fixed cardinality, from the sample path such that the approximation provided by their empirical distribution is close to optimal. A novel method is proposed, based on greedy minimisation of a kernel Stein discrepancy, that is suitable when the gradient of the log-target can be evaluated and approximation using a small number of states is required. Theoretical results guarantee consistency of the method and its effectiveness is demonstrated in the challenging context of parameter inference for ordinary differential equations. Software is available in the **Stein Thinning** package in **Python**, **R** and **MATLAB**.

Keywords: Bayesian computation, greedy optimisation, Markov chain Monte Carlo, reproducing kernel, Stein’s method

1 Introduction

The most popular computational tool for non-conjugate Bayesian inference is Markov chain Monte Carlo (MCMC). Introduced to statistics from the physics literature in Hastings (1970); Geman and Geman (1984); Tanner and Wong (1987); Gelfand and Smith (1990), an enormous amount of research effort has since been expended in the advancement of MCMC

*Address for correspondence: Chris. J. Oates, School of Mathematics, Statistics and Physics, Herschel Building, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK. E-mail: chris.oates@ncl.ac.uk

methodology. Such is the breadth of this topic that we do not attempt a survey here, but instead refer the reader to Robert and Casella (2013); Green et al. (2015) and the references therein to more advanced material. This paper is motivated by the fact that the approaches used for convergence assessment and to post-process the output of MCMC can strongly affect the estimates that are produced.

Let P be a distribution on a measurable space \mathcal{X} and let $(X_i)_{i \in \mathbb{N}}$ be a Markov chain that is P -invariant. The Markov chain sample path provides an empirical approximation

$$\frac{1}{n} \sum_{i=1}^n \delta(X_i) \tag{1}$$

to P , where $\delta(x)$ denotes a point mass centred at $x \in \mathcal{X}$. Our discussion supposes that a practitioner is prepared to simulate a Markov chain up to a maximum number of iterations, n , and that simulating further iterations is not practical; a scenario that is often encountered (e.g. see Section 4.3). In this setting it is common (and indeed recommended) to replace (1) with an alternative estimator

$$\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \tag{2}$$

that is based on a subset of the total MCMC output. The m indices $\pi(j) \in \{1, \dots, n\}$ indicate which states are retained and the identification of a suitable index set π is informed by the following considerations:

Removal of Initial Bias: The distribution of the initial states of the Markov chain may be quite different to P . To mitigate this, it is desirable to identify a “burn-in” $(X_i)_{i=1}^b$ which is then discarded. The burn-in period b is typically selected using convergence diagnostics (Cowles and Carlin, 1996). These are primarily based on the empirical distribution of simple moment, quantile or density estimates across independent chains and making a judgement as to whether the ensemble of chains has converged to the distributional target. The main limitation of convergence diagnostics, as far as we are concerned in this work, is that in taking b large enough to make bias negligible, the number $n - b$ of remaining samples may be rather small, such that the statistical efficiency of the estimator in (2) is sub-optimal as an approximation of P . Nonetheless, a considerable portion of Bayesian pedagogy is devoted to the identification of the burn-in period, as facilitated using diagnostic tests that are built into commercial-grade software such as **WinBUGS** (Lunn et al., 2000), **JAGS** (Plummer, 2003), **R** (R Core Team, 2020), and **Stan** (Carpenter et al., 2017).

Increased Statistical Efficiency: It is often stated that discarding part of the MCMC output leads to a reduction in the statistical efficiency of the estimator (2) compared to (1). This argument, made e.g. in Geyer (1992), applies only when the procedure used to discard part of the MCMC output does not itself depend on the MCMC output and when the length n of the MCMC output is fixed. That estimation efficiency can be *improved* by discarding a portion of the samples in a way that depends on the samples themselves is in fact well-established (see e.g. Dwivedi et al., 2019).

Compression of MCMC Output: A third motivation for estimators of the form (2) is to control the cost of subsequent computation involving the MCMC output. Examples include approximating the expectation of a function f , where either evaluation of f or storage of its output is associated with a computational cost, and *Monte Carlo Maximum Likelihood*, where one constructs an approximate likelihood using MCMC, then performs optimisation on this approximate likelihood (Geyer and Thompson, 1992). In such situations one may want to control the cardinality m of the index set π and to use $(X_{\pi(j)})_{j=1}^m$ as an experimental design on which f is evaluated. The most popular solution is to retain only every t^{th} state visited by the Markov chain, a procedure known as “thinning” of the MCMC output. See also the more sophisticated approach in Paige et al. (2016).

Taking these considerations into account, the most common approach used to select an index set π is based on the identification of a suitable burn-in period b and/or a suitable thinning frequency t , leading to an approximation of the form

$$\frac{1}{\lfloor (n-b)/t \rfloor} \sum_{i=1}^{\lfloor (n-b)/t \rfloor} \delta(X_{b+it}). \quad (3)$$

Here $\lfloor r \rfloor$ denotes the integer part of r . This corresponds to a set of indices π in (2) that discards the burn-in states and retains only every t^{th} iteration from the remainder of the MCMC output. It includes the case where no states are removed when $b = 0$ and $t = 1$. Despite their widespread usage, the interplay between the Markov chain sample path and the heuristics used to select b and t is not widely appreciated. In general it is unclear how much bias may be introduced by employing a post-processing heuristic that is itself based on the MCMC output. Indeed, even the basic question of when the post-processed estimator in (3) is consistent when b and t are chosen based on the MCMC output appears not to have been studied.

In this paper we propose a novel method, called **Stein Thinning**, that selects an index set π , of specified cardinality m , such that the associated discrete approximation in (2) is close to optimal among all approximations supported on the MCMC output. The method is designed to ensure that (2) is a consistent approximation of P . This includes situations when the Markov chain on which it is based is not P -invariant, but we do of course require that the regions of high probability under P are explored. To achieve this we adopt a kernel Stein discrepancy as our optimality criterion. The minimisation of kernel Stein discrepancy is performed using a greedy sequential algorithm and the main contribution of our theoretical analysis is to study the interplay of the greedy algorithm with the randomness inherent to the MCMC output. The proposed **Stein Thinning** method is simple (see Algorithm 1), applicable to most problems where gradients of the log-posterior density can be computed, and implemented as convenient **Python** and **MATLAB** packages that require no additional user input other than the number m of states to be selected (see Appendix S1).

1.1 Related Work

Our work contributes to an active area of research that attempts to cast post-processing of MCMC as an optimisation problem. Mak and Joseph (2018) proposed a method, called **Support Points**, which selects a small number of states in order that an empirical measure supported on those states minimises an “energy distance” to P . However, computation of the energy distance requires access to P , and minimisation of energy distance requires a challenging non-convex optimisation problem to be solved, meaning that in practice approximations are required. Stein discrepancy provides a computable alternative, which was used in Liu and Lee (2017) to optimally weight an arbitrary set $(X_i)_{i=1}^n \subset \mathbb{R}^d$ of states in a manner loosely analogous to importance sampling, at a computational cost of $O(n^3)$. The combined effect of applying the approach of Liu and Lee (2017) to MCMC output was analysed in Hodgkinson et al. (2020), who established situations in which the overall procedure will be consistent.

If a compressed representation of the posterior P is required, but one is not wedded to the use of MCMC for generation of candidate states, then several other methods can be used. Joseph et al. (2015, 2019) proposed a criterion to capture how well an empirical measure based on a point set approximates P and applied repeated numerical optimisation over \mathcal{X} to arrive at a suitable point set. A similar approach was taken in Chen et al. (2018), where a Stein discrepancy was numerically minimised. The reliance of both of these algorithms on non-convex numerical optimisation over \mathcal{X} renders their implementation and analysis difficult. Chen et al. (2019) considered using Markov chains to approximately perform numerical optimisation, allowing a tractable analytic treatment at the expense of a sub-optimal compression of P . An elegant alternative approach is to formulate a convex optimisation problem on the set of probability distributions on \mathcal{X} . In this spirit, Liu and Wang (2016); Liu (2017) identified a gradient flow with P as a fixed point that can be approximately simulated using a particle method. At convergence, one obtains a compressed representation of P , however the theoretical analysis of this approach remains an open and active research topic (see e.g. Duncan et al., 2019).

The present paper differs from the contributions cited, in that (1) our algorithm requires only the output from one run of MCMC, which is a realistic requirement in many situations, and (2) we are able to provide a finite sample size error bound (Theorem 2) and a consistency guarantee (Theorem 3) for **Stein Thinning**, that cover precisely the algorithm that we implement.

1.2 Outline of the Paper

The paper proceeds, in Section 2, to recall the construction of a kernel Stein discrepancy and to present **Stein Thinning**. Then in Section 3 we establish a finite sample size error bound, as well as a widely-applicable consistency result that does not require the Markov chain to be P -invariant. In Section 4 we present an empirical assessment of **Stein Thinning** in the context of parameter inference for ordinary differential equation models. Conclusions are contained in Section 5.

2 Methods

In this section we introduce and analyse **Stein Thinning**. First, in Section 2.1, we recall the construction of a kernel Stein discrepancy and its theoretical properties. The **Stein Thinning** method is presented in Section 2.2, whilst Section 2.3 is devoted to implementational detail.

Before we proceed, we introduce a piece of notation that will often be used and recall the mathematical definition of a reproducing kernel:

Notation: Let \mathcal{P} denote the set of probability distributions P that admit a positive density p , with $\nabla \log p$ Lipschitz on \mathbb{R}^d .

Reproducing Kernel: A *reproducing kernel Hilbert space* (RKHS) of functions on a set \mathcal{X} is a Hilbert space, denoted $\mathcal{H}(k)$, equipped with a function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, called a *kernel*, such that $\forall x \in \mathcal{X}$ we have $k(\cdot, x) \in \mathcal{H}(k)$ and $\forall x \in \mathcal{X}, h \in \mathcal{H}(k)$ we have $h(x) = \langle h, k(\cdot, x) \rangle_{\mathcal{H}(k)}$. In this paper $\langle \cdot, \cdot \rangle_{\mathcal{H}(k)}$ denotes the inner product in $\mathcal{H}(k)$ and the induced norm will be denoted $\| \cdot \|_{\mathcal{H}(k)}$. For further details, see Berlinet and Thomas-Agnan (2004).

2.1 Kernel Stein Discrepancy

To construct a criterion for the selection of states from the MCMC output we require a notion of optimal approximation for probability distributions. To this end, recall that an *integral probability metric* (IPM) (Muller, 1997), based on a set \mathcal{F} of measure-determining functions on a measurable space \mathcal{X} , is defined as

$$D_{\mathcal{F}}(P, Q) := \sup_{f \in \mathcal{F}} \left| \int_{\mathcal{X}} f dP - \int_{\mathcal{X}} f dQ \right|. \quad (4)$$

The fact that \mathcal{F} is measure-determining means that $D_{\mathcal{F}}(P, Q) = 0$ if and only if $P = Q$ is satisfied. Standard choices for \mathcal{F} , e.g. that recover Wasserstein distance as the IPM, cannot be used in the Bayesian context due to the need to compute integrals with respect to P in (4).

In the remainder of Section 2.1 we restrict attention to the setting $P \in \mathcal{P}$. To circumvent intractability of (4), the notion of a *Stein discrepancy* was proposed in Gorham and Mackey (2015). This was based on Stein’s method (Stein, 1972), which consists of finding a set \mathcal{G} of sufficiently differentiable d -dimensional vector fields and a differential operator \mathcal{A}_P , depending on P and acting on elements of \mathcal{G} , such that $\int_{\mathbb{R}^d} \mathcal{A}_P g dP = 0$ for all $g \in \mathcal{G}$. The proposal of Gorham and Mackey (2015) was to take $\mathcal{F} = \mathcal{A}_P \mathcal{G}$ to be the image of \mathcal{G} under \mathcal{A}_P in (4), leading to the *Stein discrepancy*

$$D_{\mathcal{A}_P \mathcal{G}}(P, Q) = \sup_{g \in \mathcal{G}} \left| \int_{\mathbb{R}^d} \mathcal{A}_P g dQ \right|. \quad (5)$$

Theoretical analysis had led to sufficient conditions for $\mathcal{A}_P \mathcal{G}$ to be measure-determining (Gorham and Mackey, 2015). In this paper we focus on a particular form of (5) due to

Liu et al. (2016); Chwialkowski et al. (2016); Gorham and Mackey (2017), called a *kernel* Stein discrepancy (KSD). In this case, \mathcal{A}_P is the *Langevin Stein operator* $\mathcal{A}_P g := p^{-1} \nabla \cdot (pg)$ derived in Gorham and Mackey (2015), where $\nabla \cdot$ denotes the divergence operator in \mathbb{R}^d and $\mathcal{G} := \{g : \mathbb{R}^d \rightarrow \mathbb{R}^d \mid \sum_{i=1}^d \|g_i\|_{\mathcal{H}(k)}^2 \leq 1\}$ is the unit ball in a Cartesian product of RKHS. It follows from construction that the set $\mathcal{A}_P \mathcal{G}$ is the unit ball of another RKHS, denoted $\mathcal{H}(k_P)$, whose kernel is

$$k_P(x, y) := \nabla_x \cdot \nabla_y k(x, y) + \langle \nabla_x k(x, y), \nabla_y \log p(y) \rangle + \langle \nabla_y k(x, y), \nabla_x \log p(x) \rangle + k(x, y) \langle \nabla_x \log p(x), \nabla_y \log p(y) \rangle, \quad (6)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard Euclidean inner product, ∇ denotes the gradient operator and subscripts have been used to indicate the variables being acted on by the differential operators (Oates et al., 2017). Thus KSD is recognised as a maximum mean discrepancy in $\mathcal{H}(k_P)$ (Song, 2008) and is fully characterised by the kernel k_P ; we therefore adopt the shorthand notation $D_{k_P}(Q)$ for $D_{\mathcal{A}_P \mathcal{G}}(P, Q)$.

In the remainder of this section we recall the main properties of KSD. The first is a condition on the kernel k that guarantees elements of $\mathcal{H}(k_P)$ have zero mean with respect to P . In what follows $\|x\| = \langle x, x \rangle^{1/2}$ denotes the Euclidean norm on \mathbb{R}^d . It will be convenient to abuse operator notation, writing $\nabla_x \nabla_y^\top k$ for the Hessian matrix of a bivariate function $(x, y) \mapsto k(x, y)$.

Proposition 1 (Proposition 1 of Gorham and Mackey (2017)). *Let $P \in \mathcal{P}$ and assume that $\int_{\mathbb{R}^d} \|\nabla \log p\| dP < \infty$. Let $(x, y) \mapsto \nabla_x \nabla_y^\top k(x, y)$ be continuous and uniformly bounded on \mathbb{R}^d . Then $\int_{\mathbb{R}^d} h dP = 0$ for all $h \in \mathcal{H}(k_P)$, where k_P is defined in (6).*

The second main property of KSD that we will need is that it can be explicitly computed for an empirical measure $Q = \frac{1}{n} \sum_{i=1}^n \delta(x_i)$, supported on states $x_i \in \mathbb{R}^d$:

Proposition 2 (Proposition 2 of Gorham and Mackey (2017)). *Let $P \in \mathcal{P}$ and let $(x, y) \mapsto \nabla_x \nabla_y^\top k(x, y)$ be continuous on \mathbb{R}^d . Then*

$$D_{k_P} \left(\frac{1}{n} \sum_{i=1}^n \delta(x_i) \right) = \sqrt{\frac{1}{n^2} \sum_{i,j=1}^n k_P(x_i, x_j)}, \quad (7)$$

where k_P was defined in (6).

The third main property is that KSD provides convergence control. Let $Q_n \Rightarrow P$ denote weak convergence of a sequence (Q_n) of measures to P . Theoretical analysis in Gorham and Mackey (2017); Chen et al. (2018); Huggins and Mackey (2018); Chen et al. (2019); Hodgkinson et al. (2020); Gorham et al. (2020) established sufficient conditions for when convergence of (7) to zero implies $\frac{1}{n} \sum_{i=1}^n \delta(x_i) \Rightarrow P$. For our purposes we present one such result, from Chen et al. (2019).

Proposition 3 (Theorem 4 in Chen et al. (2019)). *Let $P \in \mathcal{P}$ be distantly dissipative, meaning that $\liminf_{r \rightarrow \infty} \kappa(r) > 0$ where*

$$\kappa(r) := \inf \left\{ -2 \frac{\langle \nabla \log p(x) - \nabla \log p(y), x - y \rangle}{\|x - y\|^2} : \|x - y\| = r \right\}.$$

Consider the kernel $k(x, y) = (c^2 + \|\Gamma^{-1/2}(x - y)\|^2)^\beta$ for some fixed $c > 0$, a fixed positive definite matrix Γ and a fixed exponent $\beta \in (-1, 0)$. Then $D_{k_P} \left(\frac{1}{n} \sum_{i=1}^n \delta(x_i) \right) \rightarrow 0$ implies $\frac{1}{n} \sum_{i=1}^n \delta(x_i) \Rightarrow P$, where k_P is defined in (6).

The properties just described ensure that KSD is a suitable optimality criterion to consider for the post-processing of MCMC output. However, all discrepancies are associated with finite sample size pathologies; see Matsubara et al. (2021, Section 3.4) for a discussion of the pathologies of KSD. Our attention turns next to the development of algorithms for minimisation of KSD.

2.2 Greedy Minimisation of KSD

The convergence control afforded by Proposition 3 motivates the design of methods that select points $(x_i)_{i=1}^n$ such that (7) is approximately minimised. Continuous optimisation algorithms were proposed for this task in Chen et al. (2018) and Chen et al. (2019). In Chen et al. (2018), deterministic optimisation techniques were considered for low-dimensional problems, whereas in Chen et al. (2019) a Markov chain was used to provide more a practical optimisation strategy when the state space is high-dimensional. In each case greedy sequential strategies were considered, wherein at iteration n a new state x_n is appended to the current sequence (x_1, \dots, x_{n-1}) . Chen et al. (2018) also considered the use of conditional gradient algorithms (so-called *Frank-Wolfe*, or *kernel herding* algorithms) but found that greedy algorithms provided better performance across a range of experiments and therefore we focus on greedy algorithms in this manuscript.

The present paper is distinguished from earlier work in that we do not attempt to solve a continuous optimisation problem for selection of the next point $x_n \in \mathcal{X}$. Such optimisation problems are fundamentally difficult and can at best be approximately solved. Instead, we exactly solve the discrete optimisation problem of selecting a suitable element x_n from supplied MCMC output. In this sense we expect our findings will be more widely applicable than previous work, since we are simply performing post-processing of MCMC output and there exists a variety of commercial-grade software for MCMC. The method that we propose, called **Stein Thinning**, is straight-forward to implement, and is stated in Algorithm 1 for a distribution P on a general measurable space \mathcal{X} . (The convention $\sum_{i=1}^0 = 0$ is employed.)

The algorithm is illustrated on a simple bivariate Gaussian mixture in Figure 1. Observe in this figure that the points selected by the **Stein Thinning** do not belong to the burn-in period (which is visually clear), and that although the MCMC spent a disproportionate amount of time in one of the mixture components, the number of points selected by **Stein Thinning** is approximately equal across the two components of the target. The accuracy of the approximation produced by **Stein Thinning** is, nevertheless, gated by the

Data: The output $(x_i)_{i=1}^n$ from an MCMC method, a kernel k_P for which the conclusion of Proposition 3 holds, and a desired cardinality $m \in \mathbb{N}$.

Result: The indices π of a sequence $(x_{\pi(j)})_{j=1}^m \subset \{x_i\}_{i=1}^n$ where the $\pi(j)$ are elements of $\{1, \dots, n\}$.

for $j = 1, \dots, m$ **do**
 $\pi(j) \in \arg \min_{i=1, \dots, n} \frac{k_P(x_i, x_i)}{2} + \sum_{j'=1}^{j-1} k_P(x_{\pi(j')}, x_i);$
end

Algorithm 1: The proposed method; **Stein Thinning**.

quality of the MCMC output to which it is applied. A detailed empirical assessment is presented in Section 4.

Remark 1 (Tie-breaking). *In the event of a tie, a tie-breaking rule should be used to select the next index. For example, if the minimum in Algorithm 1 is realised by multiple candidate values $\Pi(j) \subseteq \{1, \dots, n\}$, one could adopt a tie-breaking rule that selects the smallest element of $\Pi(j)$ as the value that is assigned to $\pi(j)$. The rule that is used has no bearing on our theoretical analysis in Section 3.*

Remark 2 (Complexity). *The computation associated with iteration j of Algorithm 1 is $O(nr_j)$ where $r_j \leq \min(j, n)$ is the number of distinct indices in $\{\pi(1), \dots, \pi(j-1)\}$; the computational complexity of Algorithm 1 is therefore $O(n \sum_{j=1}^m r_j)$. For typical MCMC algorithms the computational complexity is $O(n)$, so the complexity of **Stein Thinning** is equal to that for MCMC when m is fixed and higher when m is increasing with n , being at most $O(nm^2)$.*

Remark 3 (Re-sampling). *In general the indices in π need not be distinct. That is, Algorithm 1 may prefer to include a duplicate state rather than to include a state which is not useful for representing P . Indeed, if $m > n$ then the sequence $(x_{\pi(j)})_{j=1}^m$ must contain duplicates entries. Theorem 1 in Section 3 clarifies this behaviour.*

Remark 4 (Finite sample error bound). *The approximation produced by **Stein Thinning** satisfies a finite sample error bound*

$$\left| \frac{1}{m} \sum_{j=1}^m f(x_{\pi(j)}) - \int_{\mathbb{R}^d} f(x) dP(x) \right| \leq D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(x_{\pi(j)}) \right) \left\| f - \int_{\mathbb{R}^d} f(x) dP(x) \right\|_{\mathcal{H}(k_P)}$$

following Hickernell (1998, Equation (1.14) applied to $f - \int f dP$). This can be contrasted with the typically asymptotic analysis of MCMC. The practical estimation of the final term in this bound was discussed in Section 4 of South et al. (2021).




Figure 1: Illustration of **Stein Thinning**: (a) Contours of the distributional target P . (b) Markov chain Monte Carlo (MCMC) output, limited to 500 iterations to mimic a challenging computational context, exhibiting burn-in and autocorrelation that must be identified and mitigated. (c) A subset of $m = 40$ states from the MCMC output selected using **Stein Thinning**, which correctly ignores the burn-in period and stratifies states approximately equally across the two components of the target.

2.3 Choice of Kernel

The suitability of KSD to quantify how well Q approximates P is determined by the choice of the kernel k in (6). Several choices are possible and for $P \in \mathcal{P}$, based on Proposition 3 together with extensive empirical assessment, Chen et al. (2019) advocated the pre-conditioned inverse multi-quadric kernel $k(x, y) := (1 + \|\Gamma^{-1/2}(x - y)\|^2)^{-1/2}$ where, compared to Proposition 3, we have fixed $c = 1$ (without loss of generality) and $\beta = -1/2$. The suitability of these choices for **Stein Thinning** is verified in Appendix S5.1. The positive definite matrix Γ remains to be specified and it is natural to take a data-driven approach where the MCMC output is used to select Γ . Provided that a fixed number $n_0 \in \mathbb{N}$ of the states $(X_i)_{i=1}^{n_0}$ from the MCMC output are used in the construction of Γ , the consistency results for **Stein Thinning** that we establish in Section 3 are not affected. To explore different strategies for the selection of Γ , we focus on the following candidates:

- **Median (med)**: The scaled identity matrix $\Gamma = \ell^2 I$, where $\ell = \text{med} := \text{median}\{\|X_i - X_j\| : 1 \leq i < j \leq n_0\}$ is the median Euclidean distance between states (Garreau et al., 2018). In the rare case that $\text{med} = 0$, an exception should be used, such as $\ell = 1$, to ensure a positive definite Γ is used.
- **Scaled median (sclmed)**: The scaled identity matrix $\Gamma = \ell^2 I$, where $\ell = \text{med} / \sqrt{\log(m)}$. This was proposed in Liu and Wang (2016) and can be motivated using the approximation $\sum_{j'=1}^m k_P(x_{\pi(j)}, x_{\pi(j')}) \approx m \exp(-\ell^{-2} \text{med}^2) = 1$. Note the dependence on m means that the preceding theoretical analysis does not apply when this heuristic is used.
- **Sample covariance (smpcov)**: The matrix Γ can be taken as a sample covariance matrix

$$\Gamma = \frac{1}{n_0 - 1} \sum_{i=1}^{n_0} (X_i - \bar{X}) (X_i - \bar{X})^\top, \quad \bar{X} := \frac{1}{n_0} \sum_{i=1}^{n_0} X_i,$$

provided that this matrix is non-singular.

The experiments in Section 4 shed light on which of these settings is the most effective, but we acknowledge that many other settings could also be considered. In what follows, we set $n_0 = \min(n, 10^3)$ for the **med** and **sclmed** settings, to avoid an $O(n^2)$ cost of computing ℓ , and otherwise set $n_0 = n$, so that the whole of the MCMC output is used to select Γ . **Python**, **R** and **MATLAB** packages are provided and their usage is described in Appendix S1.

3 Theoretical Assessment

The theoretical analysis in this section clarifies the limiting behaviour of **Stein Thinning** as $m, n \rightarrow \infty$. Our first main result concerns the behaviour of **Stein Thinning** on a fixed sequence $(x_i)_{i=1}^n$:

Theorem 1. *Let \mathcal{X} be a measurable space and let P be a probability distribution on \mathcal{X} . Let $k_P : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a reproducing kernel with $\int_{\mathcal{X}} k_P(x, \cdot) dP = 0$ for all $x \in \mathcal{X}$. Let $(x_i)_{i=1}^n \subset \mathcal{X}$ be fixed and consider an index sequence π of length m produced by Algorithm 1. Then we have the bound*

$$D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(x_{\pi(j)}) \right)^2 \leq D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(x_i) \right)^2 + \left(\frac{1 + \log(m)}{m} \right) \max_{i=1, \dots, n} k_P(x_i, x_i),$$

where the weights $w^* = (w_1^*, \dots, w_n^*)$ in the first term satisfy

$$w^* \in \arg \min_{\substack{1_n^\top w = 1 \\ w \geq 0}} D_{k_P} \left(\sum_{i=1}^n w_i \delta(x_i) \right) \quad (8)$$

where $1_n^\top = (1, \dots, 1)$ and $w \geq 0$ indicates that $w_i \geq 0$ for $i = 1, \dots, n$.

The proof of Theorem 1 is provided in Appendix S2.1. Its implication is that, given a sequence $(x_i)_{i=1}^n$, **Stein Thinning** produces an empirical distribution that converges in KSD to the optimal weighted empirical distribution $\sum_{i=1}^n w_i^* \delta(x_i)$ based on that sequence. Properties of such optimally weighted empirical measures were studied in Liu and Lee (2017); Hodgkinson et al. (2020), and are not the focus of the present paper, where the case $m \ll n$ is of principal interest.

The role of Theorem 1 is to study the interaction between the greedy algorithm and a given sequence $(x_i)_{i=1}^n$, and this bound is central to our proof of Theorem 2 which deals with the case where $(x_i)_{i=1}^n$ is replaced by MCMC output. Figure 2 illustrates the terms involved in Theorem 1. It is clear that a reduction in KSD is achieved by **Stein Thinning** of the MCMC output.

Remark 5 (Optimal weights). *To further improve the empirical approximation, we can consider an optimally-weighted sum $\sum_{j=1}^m w_{m,j}^* \delta(x_{\pi(j)})$ where the $w_{m,j}^*$ solve a convex optimisation problem analogous to (8). Such weights minimise a quadratic function subject*




Figure 2: Illustration of Theorem 1: The gray curve represents the unprocessed output from MCMC in the example of Figure 1. The black curve represents **Stein Thinning** applied to this same output and, in addition, weighted output of **Stein Thinning** is shown for weights w_m^* subject to $\sum w_{m,j}^* = 1$ and $w_{m,j}^* \geq 0$ (solid red) and weights v_m^* subject only to $\sum v_{m,j}^* = 1$ (solid blue). The dashed horizontal lines are the limiting values of their corresponding solid lines as the number m is increased.

to a linear and a non-negativity constraint and can therefore be precisely computed. If the non-negativity constraint is removed and the indices in π are distinct then

$$v_m^* := \arg \min_{1_m^\top v = 1} D_{k_P} \left(\sum_{j=1}^m v_j \delta(x_{\pi(j)}) \right) = \frac{K_P^{-1} 1_m}{1_m^\top K_P^{-1} 1_m}, \quad (K_P)_{i,j} := k_P(x_{\pi(i)}, x_{\pi(j)}),$$

as derived in Oates et al. (2017). Figure 2 indicates that the benefit of applying weights w_m^* (red curve) to the output of **Stein Thinning** (black curve) is limited, likely because the $x_{\pi(j)}$ were selected in a way that avoids redundancy in the point set. A larger improvement is provided by the weights v_m^* (blue curve), but in this case the associated empirical measure may not be a probability distribution.

Remark 6. The use of a conditional gradient algorithm, instead of a greedy algorithm, in this context amounts to simply removing the term $k_P(x_{\pi(j)}, x_{\pi(j)})$ in Algorithm 1. As discussed in Chen et al. (2018), this term can be thought of as a regulariser that lends stability to the algorithm, avoiding selection of x_i that are far from the effective support of P .

Remark 7. Theorem 1 is formulated at a high level of generality and can be applied on non-Euclidean domains \mathcal{X} . In Barp et al. (2021); Liu and Zhu (2018); Xu and Matsuda (2020); Le et al. (2020) the authors proposed and discussed Stein operators \mathcal{A}_P for the non-Euclidean context.

Next we consider the properties of **Stein Thinning** applied to MCMC output. Let V be a function $V : \mathcal{X} \rightarrow [1, \infty)$ and, for a function $f : \mathcal{X} \rightarrow \mathbb{R}$ and a measure μ on \mathcal{X} , let

$\|f\|_V := \sup_{x \in \mathcal{X}} \frac{|f(x)|}{V(x)}$, $\|\mu\|_V := \sup_{\|f\|_V \leq 1} \left| \int_{\mathcal{X}} f d\mu \right|$. Recall that a ψ -irreducible and aperiodic Markov chain $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$ with n^{th} step transition kernel P^n is *V-uniformly ergodic* (see Theorem 16.0.1 of Meyn and Tweedie, 2012) if and only if $\exists R \in [0, \infty)$, $\rho \in (0, 1)$ such that

$$\|P^n(x, \cdot) - P\|_V \leq RV(x)\rho^n \quad (9)$$

for all initial states $x \in \mathcal{X}$ and all $n \in \mathbb{N}$. The notation \mathbb{E} will be used to denote expectation with respect to the law of the Markov chain in the sequel. Theorem 2 establishes a finite sample size error bound for **Stein Thinning** applied to MCMC output:

Theorem 2. *Let \mathcal{X} be a measurable space and let P be a probability distribution on \mathcal{X} . Let $k_P : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a reproducing kernel with $\int_{\mathcal{X}} k_P(x, \cdot) dP = 0$ for all $x \in \mathcal{X}$. Consider a P -invariant, time-homogeneous Markov chain $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$ generated using a V -uniformly ergodic transition kernel, such that (9) is satisfied with $V(x) \geq \sqrt{k_P(x, x)}$ for all $x \in \mathcal{X}$. Suppose that, for some $\gamma > 0$,*

$$b := \sup_{i \in \mathbb{N}} \mathbb{E} \left[e^{\gamma k_P(X_i, X_i)} \right] < \infty, \quad M := \sup_{i \in \mathbb{N}} \mathbb{E} \left[\sqrt{k_P(X_i, X_i)} V(X_i) \right] < \infty.$$

Let π be an index sequence of length m produced by Algorithm 1 applied to the Markov chain output $(X_i)_{i=1}^n$. Then, with $C = \frac{2R\rho}{1-\rho}$, we have that

$$\mathbb{E} \left[D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \right)^2 \right] \leq \frac{\log(b)}{\gamma n} + \frac{CM}{n} + \left(\frac{1 + \log(m)}{m} \right) \frac{\log(nb)}{\gamma}. \quad (10)$$

The proof of Theorem 2 is provided in Appendix S2.2.

Remark 8. *The upper bound in (10) is asymptotically minimised when (up to log factors) m is proportional to n . In practice we are interested in the case $m \ll n$, so we may for example set $m = \lfloor \frac{n}{1000} \rfloor$ if we aim for substantial compression. It is not claimed that the bound in (10) is tight and indeed empirical results in Section 4 endorse the use of **Stein Thinning** in the small m context.*

Remark 9. *For $P \in \mathcal{P}$ and k_P in (6), based on a radial kernel k , meaning that $k(x, y) = \phi(x - y)$ for some function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying $\nabla \phi(0) = 0$, we have that $k_P(x, x) = -\Delta \phi(0) + \phi(0) \|\nabla \log p(x)\|^2$. The function $x \mapsto \sqrt{k_P(x, x)}$ appearing in the preconditions of Theorem 2 can therefore be understood in terms of $\|\nabla \log p(x)\|$. Further discussion of the preconditions of Theorem 2 is provided in Appendix S2.4.*

Since convergence in mean-square does not in general imply almost sure convergence, we next strengthen the conclusions of Theorem 2. Our final result, Theorem 3, therefore establishes an almost sure convergence guarantee for **Stein Thinning**. Furthermore, the result that follows applies also in the “biased sampler” case, where $(X_i)_{i \in \mathbb{N}}$ is a Q -invariant Markov chain and Q need not equal P :

Theorem 3. Let Q be a probability distribution on \mathcal{X} with P absolutely continuous with respect to Q . Consider a Q -invariant, time-homogeneous Markov chain $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$ generated using a V -uniformly ergodic transition kernel, such that $V(x) \geq \frac{dP}{dQ}(x) \sqrt{k_P(x, x)}$. Suppose that, for some $\gamma > 0$,

$$b := \sup_{i \in \mathbb{N}} \mathbb{E} \left[e^{\gamma \max(1, \frac{dP}{dQ}(X_i)^2) k_P(X_i, X_i)} \right] < \infty, \quad M := \sup_{i \in \mathbb{N}} \mathbb{E} \left[\frac{dP}{dQ}(X_i) \sqrt{k_P(X_i, X_i)} V(X_i) \right] < \infty.$$

Let π be an index sequence of length m produced by Algorithm 1 applied to the Markov chain output $(X_i)_{i=1}^n$. If $m \leq n$ and the growth of n is limited to at most $\log(n) = O(m^{\beta/2})$ for some $\beta < 1$, then $D_{k_P}(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)})) \rightarrow 0$ almost surely as $m, n \rightarrow \infty$. Furthermore, if the preconditions of Proposition 3 are satisfied, then $\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \Rightarrow P$ almost surely as $m, n \rightarrow \infty$.

The proof of Theorem 3 is provided in Appendix S2.3. The interpretation of Theorem 3 is that one may sample states from a Markov chain that is not P -invariant and yet, under the stated assumptions (which ensure that regions of high probability under P are explored), one can use **Stein Thinning** to still obtain a consistent approximation of P . This can be contrasted, for example, with the **Support Points** method of Mak and Joseph (2018), which relies on P being well-approximated by the MCMC output. This completes our theoretical analysis of **Stein Thinning**.

4 Empirical Assessment

In this section we compare the performance of **Stein Thinning** with existing methods for post-processing MCMC output. Our motivation derives from a problem in which we must infer a 38-dimensional parameter in a calcium signalling model defined by a stiff system of 6 coupled ordinary differential equations (ODEs). Posterior uncertainty is required to be propagated through a high-fidelity simulation in a multi-scale and multi-physics model f of the human heart. Here, compression of the MCMC output can be used to construct an approximately optimal experimental design on which f can be evaluated. The calcium model is, however, unsuitable for conducting a thorough *in silico* assessment due to its associated computational cost. Therefore in Section 4.1 we first consider a simpler ODE model, where P can be accurately approximated. Then, as an intermediate example, in Section 4.2 we consider an ODE model that induces stronger correlations among the parameters in P , before addressing the calcium model in Section 4.3.

In Appendix S3 we describe the generic structure of a parameter inference problem for ODEs. In all instances the aim is to post-process the output from MCMC, in order to produce an accurate empirical approximation of the posterior supported on a small number $m \ll n$ of the states that were visited. The following methods were compared:

- The standard approach, which estimates a burn-in period using either the *GR diagnostic* $\hat{b}^{\text{GR}, L}$, $L > 1$, of Gelman and Rubin (1992); Brooks and Gelman (1998); Gelman

et al. (2014) or the more sophisticated *VK diagnostic* $\hat{b}^{\text{VK},L}$, $L \geq 1$, of Vats and Knudson (2018), in each case based on L independent chains as described in Appendix S4, followed by thinning as per (3).

- The **Support Points** algorithm proposed in Mak and Joseph (2018), implemented in the R package `support`.
- The **Stein Thinning** algorithm that we have proposed, with each of the kernel choices described in Section 2.3.

To ensure that our empirical findings are not sensitive to the choice of MCMC method, we implemented four Metropolis–Hastings samplers that differ qualitatively according to the sophistication of their proposal. These were: (i) the Gaussian random walk (**RW**); (ii) the adaptive Gaussian random walk (**ADA-RW**), which uses an estimate of the covariance of the target (Haario et al., 1999); (iii) the Metropolis-adjusted Langevin algorithm (**MALA**), which takes a step in the direction of increasing Euclidean gradient, perturbed by Gaussian noise (Roberts and Tweedie, 1996); (iv) the preconditioned version of **MALA** (**P-MALA**), which employs a preconditioner based on the Fisher information matrix (Girolami and Calderhead, 2011). Full details are in Appendix S3. Metropolis–Hastings algorithms were selected on the basis that we were able to successfully implement them on the challenging calcium signalling model in Section 4.3, which required manually interfacing with the numerical integrator to produce reliable output.

4.1 Goodwin Oscillator

The first example that we consider is a negative feedback oscillator due to Goodwin (1965). The ODE model and the associated $d = 4$ dimensional inference problem are described in Appendix S5.2, where one trace plot for each MCMC method, of length $n = 2 \times 10^6$, are presented in Figure S3.

First we consider the standard approach to post-processing MCMC output, as per (3). From the trace plots in Figure S3, it is clear that a burn-in period $b > 0$ is required. For each method we therefore computed the GR and VK diagnostics, to arrive at candidate values b for the burn-in period. Default settings were used for all diagnostics, which were computed both for the multivariate d -dimensional state vector and for the univariate marginals, as reported in Appendix S5.2. The GR diagnostics were computed using $L = 6$ independent chains and the VK diagnostics were computed using both $L = 1$ and $L = 6$ independent chains; note that when $L > 1$, these diagnostics have access to more information in comparison with **Stein Thinning**, in terms of the number of samples that are available to the method. The estimated values for the burn-in period are reported in Appendix S5.2, Table S4. For all MCMC methods, neither the univariate nor the multivariate GR diagnostics were satisfied, so that $\hat{b}^{\text{GR},6} > n$ and estimation using (3) cannot proceed. The VK diagnostic produced values $\hat{b}^{\text{VK},L} < n$, which typically led to about half of the MCMC output being discarded. Although well-suited for their intended task of minimising bias in MCMC output, the smaller number of states left after burn-in removal may lead to inefficient approximation of P and

derived quantities of interest, strikingly so in the case of the GR diagnostic. The use of an optimality criterion enables **Stein Thinning** to directly address this bias-variance trade-off. Of course, one can in principle run more iterations of MCMC to provide more diversity in the remainder of the sample path after burn-in is removed, but in applications such as the calcium model of Section 4.3 the computational cost associated with each iteration presents a practical limitation in running more iterations of an MCMC method. Effective methods to post-process limited output (or, equivalently, a long output from a poorly mixing Markov chain) are therefore important.

Having identified a burn-in period, the standard approach thins the remainder of the sample path according to (3). In the experiments that follow we focus on the VK diagnostic and consider both the smallest and largest estimates obtained for the burn-in period. The resulting index sets π are displayed, for $m = 20$ and RW (the simplest MCMC method) in Figure 3 (top left panel), and in Appendix S5.2, Figures S6 (**ADA-RW**), S7 (**MALA**), S8 (**P-MALA**). In the same figures (top right panel) we show the set of **Support Points** obtained using algorithm proposed by Mak and Joseph (2018). The remaining panels display the output from **Stein Thinning**. Compared to the standard approach, **Support Points** and **Stein Thinning** produce sets that are more structured.

To assess the performance of these competing methods, we first considered the toy problem of approximating the posterior mean of each parameter in the Goodwin oscillator as an average of m points selected from the MCMC output. Figure 4 displays absolute errors for each method, based on RW; our ground truth was provided by an extended run of MCMC. Results for the other MCMC methods are provided in Appendix S5.2, Figures S9 (**ADA-RW**), S10 (**MALA**), S11 (**P-MALA**). Broadly speaking, **Stein Thinning** tends to provide more accurate estimators compared to the alternatives considered. From Figure 4 it is difficult to see any difference in performance between **med**, **sclmed** and **smpcov**. To gain more insight, in Appendix S5.2 we plot marginal density estimates in Figures S12 (RW), S13 (**ADA-RW**), S14 (**MALA**), S15 (**P-MALA**). It is apparent that **Stein Thinning** improves on the standard approach, whilst **med** and **sclmed** performed slightly better than **smpcov**. This may be because in **smpcov** there are more degrees of freedom in Γ that must be estimated. **Support Points** performed on a par with **Stein Thinning** based on **smpcov**.

To facilitate a more principled assessment, we computed two quantitative measures for how well the resulting empirical distributions approximate the posterior. These were (a) the energy distance (ED; Székely and Rizzo, 2004; Baringhaus and Franz, 2004), given up to an additive constant by

$$\text{ED} := \frac{2}{m} \sum_{j=1}^m \int \|x - x_{\pi(j)}\|_{\Sigma} \, dP(x) - \frac{1}{m^2} \sum_{j,j'=1}^m \|x_{\pi(j)} - x_{\pi(j')}\|_{\Sigma}, \quad (11)$$

where in this paper we used the norm $\|x\|_{\Sigma} := \|\Sigma^{-1/2}x\|$ induced by the covariance matrix of P , with both Σ and (11) being estimated from MCMC output, and (b) the KSD based on **med**, the simplest setting for Γ . ED serves as an objective performance measure, being closely related to the quantity that **Support Points** attempts to minimise (Mak and Joseph (2018) used the $\|\cdot\|$ norm in place of $\|\cdot\|_{\Sigma}$), while KSD is the performance measure that is

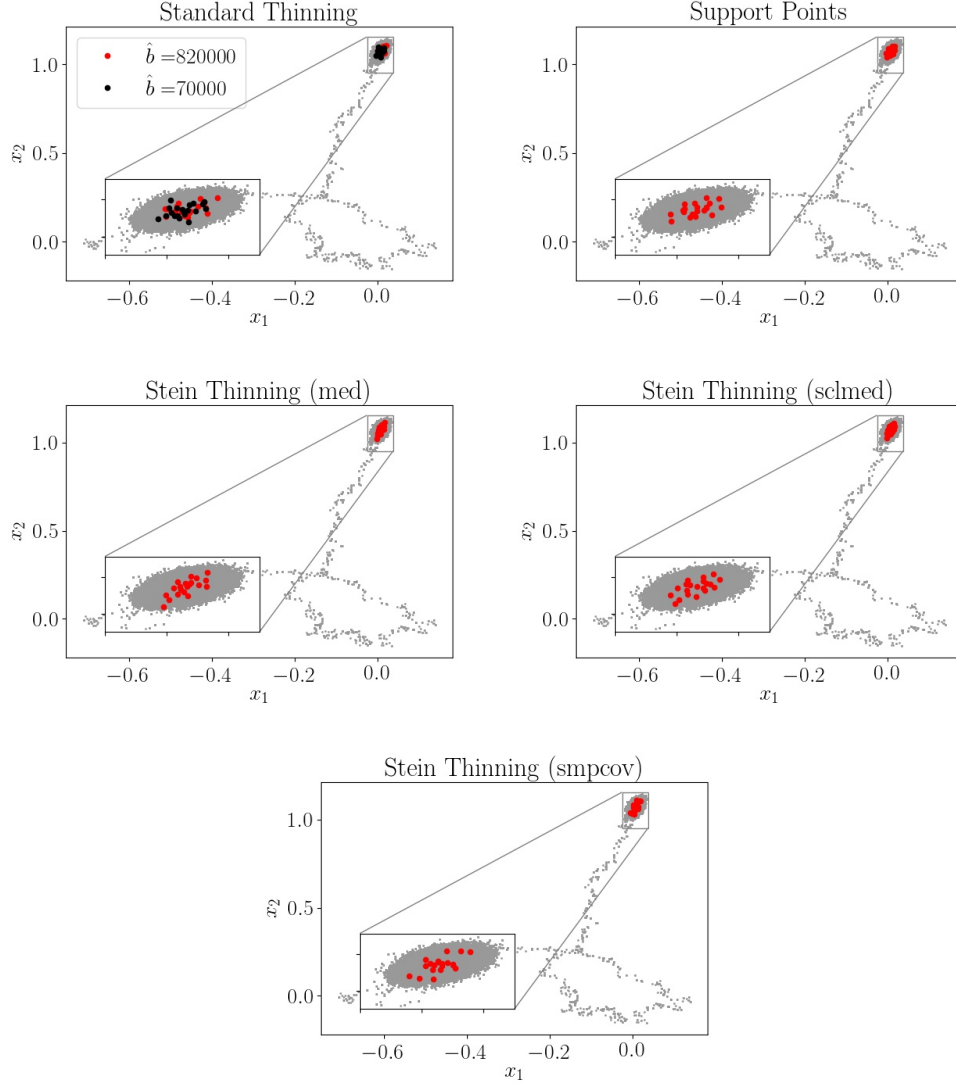


Figure 3: Projections on the first two coordinates of the RW MCMC output from the Goodwin oscillator (gray dots), together with the first $m = 20$ points selected using: the standard approach of discarding burn-in and thinning the remainder (the estimated burn in period is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

being directly optimised in **Stein Thinning**. Our decision to include KSD in the assessment is motivated by three factors; (i) ED is somewhat insensitive to detail, making it difficult to rank competing methods; (ii) the empirical approximation of ED in (11) relies on access to high-quality MCMC output, but this will not be available in Section 4.3; (iii) Stein discrepancies are the only computable performance measures in the Bayesian context, to the best of our knowledge, that have been proven to provide convergence control.

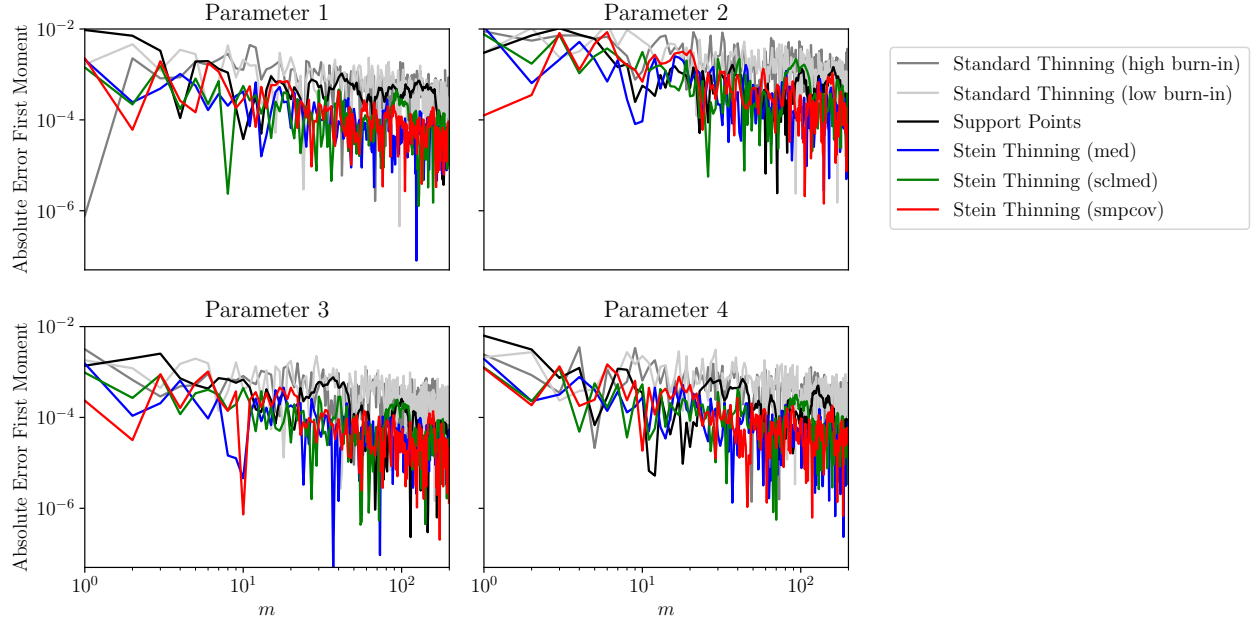


Figure 4: Absolute error of estimates for the posterior mean of each parameter in the Goodwin oscillator, based on output from RW MCMC.




Figure 5: Goodwin oscillator: Energy distance (ED) to the posterior, as per (11), for empirical distributions obtained through traditional burn-in and thinning (grey lines), **Support Points** (black line) and **Stein Thinning** (colored lines), based on output from four different MCMC methods.

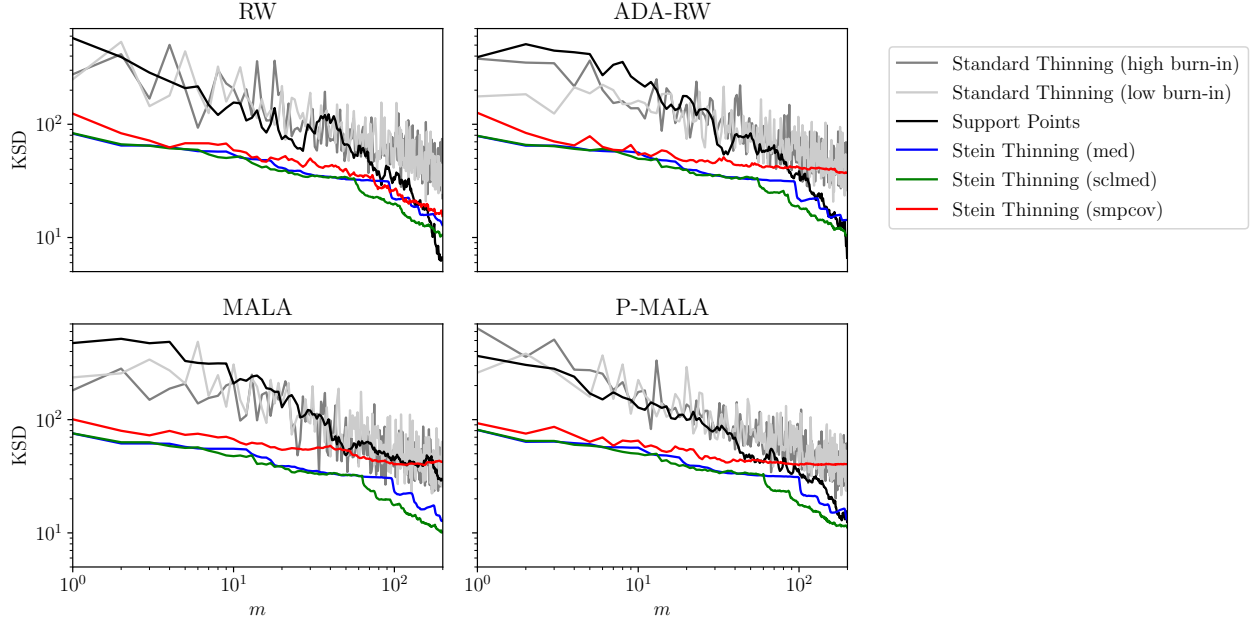


Figure 6: Goodwin oscillator: Kernel Stein discrepancy (KSD) based on `med`, for empirical distributions obtained through traditional burn-in and thinning (grey lines), **Support Points** (black line) and **Stein Thinning** (colored lines), based on output from four different MCMC methods.

The results for ED are shown in Figure 5. Here **Stein Thinning** based on `sclmed` performed at least as well as the other methods considered and, surprisingly, out-performed **Support Points** when applied to MALA and P-MALA output. This may be because MALA and P-MALA provided worse approximations to P compared with RW and ADA-RW (recall that **Support Points** relies on the MCMC output providing an accurate approximation of P). Note that neither ED nor KSD values will tend to 0 as $m \rightarrow \infty$ in this experiment, since the number n of MCMC iterations was fixed. The corresponding results for KSD are presented in Figure 6 and show a clearer performance ordering of the competing methods, with **Stein Thinning** based on `med` and `sclmed` out-performing all other methods for all but the largest values of m considered. The `smpcov` setting performed well for small m but for large m its performance degraded. The performance ordering under KSD was identical across the different MCMC output.

4.2 Lotka–Volterra

The second example that we consider is the predator-prey model of Lotka (1926) and Volterra (1926). A description of the $d = 4$ dimensional inference task, the output from MCMC methods and the implementation of thinning procedures is reserved for Appendix S5.3. Compared to the Goodwin oscillator, the Lotka–Volterra posterior P exhibits stronger correlation among parameters. This has consequences for our assessment, since now all MCMC methods,




Figure 7: Lotka–Volterra model: Energy distance (ED) to the posterior, as per (11), for empirical distributions obtained through traditional burn-in and thinning (grey lines), **Support Points** (black line) and **Stein Thinning** (colored lines), based on output from four different MCMC methods.

and in particular MALA, mixed less well compared to corresponding results for the Goodwin oscillator, as can be seen from the trace plots in Appendix S5.3, Figure S17. Results are reported for ED in Figure 7. It can be seen that **Stein Thinning** based on `med` and `sclmed` performed comparably with **Support Points**, being better for small m in the case of RW and ADA-RW and marginally worse for large m in RW, ADA-RW and P-MALA. Interestingly, the setting `smpcov` was associated with poor performance on output from RW, ADA-RW and especially P-MALA. This may be because, when Γ is poorly conditioned, any error in an estimate for Γ will be amplified when computing Γ^{-1} . However, in the case of MALA, which mixed poorly, the standard approach of burn-in removal and thinning performed poorly and all settings of **Stein Thinning** provided an improvement.

Results for KSD are reported in Figure 8. The performance ordering of competing methods under KSD is similar to that reported in Section 4.1, except for the `smpcov` setting which appears to improve the performance of **Stein Thinning** for larger values of m in the context of MALA. This may be because `smpcov` serves to “whiten” the correlation structure in P , such that the resulting geometry is more favourable for the construction of an empirical approximation. However, this improved performance was not seen on P-MALA. In all cases **Stein Thinning** out-performed **Support Points**.




Figure 8: Lotka-Volterra model: Kernel Stein discrepancy (KSD) based on `med`, for empirical distributions obtained through traditional burn-in and thinning (grey lines), **Support Points** (black line) and **Stein Thinning** (colored lines), based on output from four different MCMC methods.

4.3 Calcium Signalling Model

Our final example a model for calcium signalling in cardiac cells, illustrated in Appendix S5.4, Figure S24. The model describes an electrically-activated intracellular calcium signal that in turn activates the sub-cellular sarcomere, causing the muscle cell to contract and the heart to beat. The intracellular calcium signal is crucial for healthy cardiac function. However, under pathological conditions, dysregulation of this intra-cellular signal can play a central role in the initiation and sustenance of life-threatening arrhythmias. Computational models are increasingly being applied to study this highly-orchestrated multi-scale signalling cascade to determine how changes in cell-scale calcium regulation, encoded in calcium model parameters, impact whole-organ cardiac function (Campos et al., 2015; Niederer et al., 2019; Colman, 2019). The computational cost of simulating from tissue-scale and organ-scale models is high, with single simulations taking thousands of CPU hours (Niederer et al., 2011; Augustin et al., 2016; Strocchi et al., 2020). This limits the capacity to propagate uncertainty in calcium signalling model parameters up to organ-scale simulations, so that at present it remains unclear how uncertainty in calcium signalling parameters impacts the predictions made by a whole-organ model. Our motivation for developing **Stein Thinning** was to obtain a compressed representation of the posterior distribution for the $d = 38$ dimensional parameter of a calcium signalling model, based on a cell-scale experimental dataset, which can subsequently be used as an experimental design to propagate uncertainty through a




Figure 9: Calcium signalling model. Kernel Stein discrepancy (KSD) based on `med`, for empirical distributions obtained using **Support Points** and **Stein Thinning**, based on output from RW MCMC applied to either P or a tempered version of P .

whole-organ model.

This motivating problem entails a second complication in that, compared to the example in Section 4.1 and even the example in Section 4.2, the development of an efficient MCMC method appears to be difficult. Thus, in the experiment that follows, we cannot rely on any of the MCMC methods that we described at the start of Section 4 to provide anything more than a crude approximation of the posterior, at best. This is evidenced by the non-overlapping approximations to the posterior marginals produced when different random seeds are used; see Figures S26 to S29. (Of course, it is possible that a more sophisticated sampling method may be designed for this task, but our aim here is not to develop a new sampling method.) Tempering of the likelihood provides a straightforward route to improve the mixing of MCMC, but the invariant distribution Q will then no longer equal P . Here we explore the potential for **Stein Thinning** to perform bias-correction for such Q -invariant MCMC output, in the spirit of Theorem 3.

Our focus in the remainder is on output from the RW MCMC method. This MCMC method was selected since (a) gradient-free methods can be easier to tune when the posterior is concentrated (Livingstone and Zanella, 2020), and (b) once the sample path has been computed, the associated gradients can be computed in parallel. Both standard and tempered MCMC were performed; in the latter case the likelihood was tempered so that the (biased) target Q was just about tractable for MCMC (see Appendix S5.4). In each case a total of $n = 4 \times 10^6$ iterations of MCMC were performed, representing two weeks' CPU time.

Figure 9 reports the KSD based on `med`, for index sets of cardinality up to $m = 500$; see Appendix S5.4 for results for KSD based on `sclmed` (Figure S30) and `smpcov` (Figure S31). Considering first the tempered MCMC output, the lower values of KSD achieved by **Stein Thinning** are consistent with fact that **Stein Thinning** corrects for bias due to tempering, while **Support Points** does not. Furthermore, **Stein Thinning** of tempered MCMC results in lower values of KSD compared to **Support Points** applied to standard MCMC output, with the latter being negatively affected by the non-convergence of the MCMC. Inspection

of the univariate marginals demonstrates that the combination of tempering and **Stein Thinning** produces approximations that are robust to changes in the random seed, while the approximations produced by standard MCMC with an equivalent computational budget are not; see Figures S26 to S29.

5 Conclusion

In this paper, standard approaches used to post-process and compress output from MCMC were identified as being sub-optimal when one considers the approximation quality of the empirical distribution that is produced. A novel method, **Stein Thinning**, was proposed that seeks a subset of the MCMC output, of fixed cardinality, such that the associated empirical approximation is close to optimal. The theoretical analysis that we have provided for **Stein Thinning** handles the effect of the post-processing procedure jointly with the randomness involved in simulating from the Markov chain, such that consistency of the overall estimator is established.

Although we focused on MCMC, the proposed method can be applied to any computational method that provides a collection of states as output. These include approximate (biased) MCMC methods, where **Stein Thinning** may be able to provide bias correction in the spirit of Theorem 3. On the other hand, the main limitation of **Stein Thinning** is that it requires gradients of the log-target to be computed, which is not always practical.

Our research was motivated by challenging parameter inference problems that arise in ODEs, in particular in cardiac modelling where one is interested in propagating calcium signalling parameter uncertainty through a whole-organ simulation – a task that would naïvely be impractical or impossible using the full MCMC output. Our ongoing research is exploiting **Stein Thinning** in this context and is enabling us to perform scientific investigations that were not feasible beforehand. Furthermore, in a sequel we demonstrate that approximate implementations of **Stein Thinning** can massively reduced its implementation cost (Teymur et al., 2021).

Acknowledgements The authors are grateful for support from the Lloyd’s Register Foundation programme on data-centric engineering and the programme on health and medical sciences at the Alan Turing Institute. MR, SN and CJO were supported by the British Heart Foundation (BHF; SP/18/6/33805). JC was supported by the UKRI Strategic Priorities Fund (EP/T001569/1). PS was supported by the BHF (RG/15/9/31534). SN was supported by the EPSRC (EP/P01268X/1, NS/A000049/1, EP/M012492/1), the BHF (PG/15/91/31812, FS/18/27/33543), the NIHR (II-LB-1116-20001) and the Wellcome Trust (WT 203148/Z/16/Z). The authors thank Matthew Graham, Liam Hodgkinson, Rob Salomone, and the anonymous Editor, Associate Editor, and Reviewers, for helpful comments on the manuscript.

References

- C. M. Augustin, A. Neic, M. Liebmann, A. J. Prassl, S. A. Niederer, G. Haase, and G. Plank. Anatomically accurate high resolution modeling of human whole heart electromechanics: a strongly scalable algebraic multigrid solver method for nonlinear deformation. *Journal of Computational Physics*, 305:622–646, 2016.
- L. Baringhaus and C. Franz. On a new multivariate two-sample test. *Journal of Multivariate Analysis*, 88(1):190–206, 2004.
- A. Barp, C. Oates, E. Porcu, and M. Girolami. A Riemann–Stein kernel method. *Bernoulli*, 2021. To appear.
- A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer Science & Business Media, New York, 2004.
- N. Biswas, P. E. Jacob, and P. Vanetti. Estimating convergence of Markov chains with L-lag couplings. In *Proceedings of the 33rd Conference on Neural Information Processing Systems*, 2019.
- S. P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, 1998.
- B. Calderhead and M. Girolami. Estimating Bayes factors via thermodynamic integration and population MCMC. *Computational Statistics & Data Analysis*, 53(12):4028–4045, 2009.
- F. O. Campos, Y. Shiferaw, A. J. Prassl, P. M. Boyle, E. J. Vigmond, and G. Plank. Stochastic spontaneous calcium release events trigger premature ventricular complexes by overcoming electrotonic load. *Cardiovascular Research*, 107(1):175–183, 2015.
- B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.
- W. Y. Chen, L. Mackey, J. Gorham, F.-X. Briol, and C. J. Oates. Stein points. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- W. Y. Chen, A. Barp, F.-X. Briol, J. Gorham, L. Mackey, M. Girolami, and C. J. Oates. Stein points Markov chain Monte Carlo. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- K. Chwialkowski, H. Strathmann, and A. Gretton. A kernel test of goodness of fit. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

- M. A. Colman. Arrhythmia mechanisms and spontaneous calcium release: Bi-directional coupling between re-entrant and focal excitation. *PLoS Computational Biology*, 15(8), 2019.
- M. K. Cowles and B. P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904, 1996.
- A. Duncan, N. Nüsken, and L. Szpruch. On the geometry of Stein variational gradient descent. *arXiv:1912.00894*, 2019.
- R. Dwivedi, O. N. Feldheim, O. Gurel-Gurevich, and A. Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 174(1-2):103–131, 2019.
- J. M. Flegal, M. Haran, and G. L. Jones. Markov chain Monte Carlo: Can we trust the third significant figure? *Statistical Science*, 23(2):250–260, 2008.
- D. Garreau, W. Jitkrittum, and M. Kanagawa. Large sample analysis of the median heuristic. *arXiv:1707.07269*, 2018.
- A. E. Gelfand and A. F. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.
- A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.
- A. Gelman, W. R. Gilks, and G. O. Roberts. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120, 1997.
- A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*, volume 2. CRC press, 2014.
- S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6):721–741, 1984.
- C. J. Geyer. Practical Markov chain Monte Carlo. *Statistical Science*, 7(4):473–483, 1992.
- C. J. Geyer and E. A. Thompson. Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society, Series B*, 54(3):657–683, 1992.
- M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society, Series B*, 73(2):123–214, 2011.
- B. C. Goodwin. Oscillatory behavior in enzymatic control process. *Advances in Enzyme Regulation*, 3:318–356, 1965.

- J. Gorham and L. Mackey. Measuring sample quality with Stein’s method. In *Proceedings of the 29th Conference on Neural Information Processing Systems*, 2015.
- J. Gorham and L. Mackey. Measuring sample quality with kernels. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- J. Gorham, A. Raj, and L. Mackey. Stochastic Stein discrepancies. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems*, 2020.
- P. J. Green, K. Łatuszyński, M. Pereyra, and C. P. Robert. Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statistics and Computing*, 25(4):835–862, 2015.
- H. Haario, E. Saksman, and J. Tamminen. Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*, 14(3):375–396, 1999.
- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. 1970.
- F. Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of Computation*, 67(221):299–322, 1998.
- R. Hinch, J. Greenstein, A. Tanskanen, L. Xu, and R. Winslow. A simplified local control model of calcium-induced calcium release in cardiac ventricular myocytes. *Biophysical Journal*, 87(6):3723–3736, 2004.
- A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):363–396, 2005.
- L. Hodgkinson, R. Salomone, and F. Roosta. The reproducing Stein kernel approach for post-hoc corrected sampling. *arXiv:2001.09266*, 2020.
- J. Huggins and L. Mackey. Random feature Stein discrepancies. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, 2018.
- G. L. Jones and J. P. Hobert. Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statistical Science*, 16(4):312–334, 2001.
- V. R. Joseph, T. Dasgupta, R. Tuo, and C. Wu. Sequential exploration of complex surfaces using minimum energy designs. *Technometrics*, 57(1):64–74, 2015.
- V. R. Joseph, D. Wang, L. Gu, S. Lyu, and R. Tuo. Deterministic sampling of expensive posteriors using minimum energy designs. *Technometrics*, 61(3):297–308, 2019.
- C. Knudson and D. Vats. *stableGR*, 2020. R package version 1.0.

- H. Le, A. Lewis, K. Bharath, and C. Fallaize. A diffusion approach to Stein’s method on Riemannian manifolds. *arXiv:2003.11497*, 2020.
- C. Liu and J. Zhu. Riemannian Stein variational gradient descent for Bayesian inference. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- Q. Liu. Stein variational gradient descent as gradient flow. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, 2017.
- Q. Liu and J. D. Lee. Black-box importance sampling. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.
- Q. Liu and D. Wang. Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *Proceedings of the 30th Conference on Neural Information Processing Systems*, 2016.
- Q. Liu, J. D. Lee, and M. I. Jordan. A kernelized Stein discrepancy for goodness-of-fit tests and model evaluation. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- S. Livingstone and G. Zanella. The barker proposal: combining robustness and efficiency in gradient-based mcmc. *arXiv:1908.11812*, 2020.
- A. J. Lotka. Elements of physical biology. *Science Progress in the Twentieth Century (1919-1933)*, 21(82):341–343, 1926.
- D. J. Lunn, A. Thomas, N. Best, and D. Spiegelhalter. WinBUGS - a Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing*, 10(4):325–337, 2000.
- S. Mak and V. R. Joseph. Support points. *The Annals of Statistics*, 46(6A):2562–2592, 2018.
- T. Matsubara, J. Knoblauch, F.-X. Briol, and C. J. Oates. Robust generalised Bayesian inference for intractable likelihoods. *arXiv:2104.07359*, 2021.
- S. Meyn and R. Tweedie. Computable bounds for geometric convergence rates of Markov chains. *The Annals of Applied Probability*, 4(4):981–1011, 1994.
- S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer Science & Business Media., 2012.
- A. Muller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.
- S. A. Niederer, L. Mitchell, N. Smith, and G. Plank. Simulating human cardiac electrophysiology on clinical time-scales. *Frontiers in Physiology*, 2:14, 2011.

- S. A. Niederer, J. Lumens, and N. A. Trayanova. Computational models in cardiology. *Nature Reviews Cardiology*, 16(2):100–111, 2019.
- C. J. Oates, T. Papamarkou, and M. Girolami. The controlled thermodynamic integral for Bayesian model evidence evaluation. *Journal of the American Statistical Association*, 111(514):634–645, 2016.
- C. J. Oates, M. Girolami, and N. Chopin. Control functionals for Monte Carlo integration. *Journal of the Royal Statistical Society, Series B*, 79(3):695–718, 2017.
- B. Paige, D. Sejdinovic, and F. D. Wood. Super-sampling with a reservoir. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence*, 2016.
- M. Plummer. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, 2003.
- M. Plummer, N. Best, K. Cowles, and K. Vines. CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11, 2006. URL <https://journal.r-project.org/archive/>.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, 2020. URL <https://www.R-project.org>.
- C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Science & Business Media, 2013.
- G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis–Hastings algorithms. *Statistical Science*, 16(4):351–367, 2001.
- G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71, 2004.
- G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- G. O. Roberts and R. L. Tweedie. Bounds on regeneration times and convergence rates for Markov chains. *Stochastic Processes and Their Applications*, 80(2):211–229, 1999.
- J. S. Rosenthal. Minorization conditions and convergence rates for Markov chain Monte Carlo. *Journal of the American Statistical Association*, 90(430):558–566, 1995.
- L. Song. *Learning via Hilbert space embedding of distributions*. PhD thesis, School of Information Technologies, University of Sydney, 2008.
- L. Song, J. Huang, A. Smola, and K. Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th International Conference on Machine Learning*, 2009.

- L. F. South, T. Karvonen, C. Nemeth, M. Girolami, and C. Oates. Semi-exact control functionals from Sard’s method. *Biometrika*, 2021. To appear.
- C. Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of 6th Berkeley Symposium on Mathematical Statistics and Probability*, pages 583–602. University of California Press, 1972.
- M. Strocchi, M. A. Gsell, C. M. Augustin, O. Razeghi, C. H. Roney, A. J. Prassl, E. J. Vigmond, J. M. Behar, J. S. Gould, C. A. Rinaldi, M. J. Bishop, G. Plank, and S. A. Niederer. Simulating ventricular systolic motion in a four-chamber heart model with spatially varying robin boundary conditions to model the effect of the pericardium. *Journal of Biomechanics*, 101:109645, 2020.
- G. J. Székely and M. L. Rizzo. Testing for equal distributions in high dimension. *InterStat*, 5(16.10):1249–1272, 2004.
- M. A. Tanner and W. H. Wong. The calculation of posterior distributions by data augmentation. *Journal of the American statistical Association*, 82(398):528–540, 1987.
- O. Teymur, J. Gorham, M. Riabiz, and C. J. Oates. Optimal quantisation of probability measures using maximum mean discrepancy. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 2021.
- D. Vats and J. M. Flegal. Lugsail lag windows and their application to MCMC. *arXiv:1809.04541*, 2018.
- D. Vats and C. Knudson. Revisiting the Gelman–Rubin diagnostic. *arXiv:1812.09384*, 2018.
- D. Vats, J. M. Flegal, and G. L. Jones. Multivariate output analysis for Markov chain Monte Carlo. *Biometrika*, 106(2):321–337, 2019.
- V. Volterra. Variazioni e fluttuazioni del numero d’individui in specie animali conviventi. *Memoria della Reale Accademia Nazionale dei Lincei*, 6:31–113, 1926.
- W. Xu and T. Matsuda. A Stein goodness-of-fit test for directional distributions. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2020.

Supplementary Material

This electronic supplement contains details for our **Python**, **R** and **MATLAB** implementations of **Stein Thinning**, proofs for all novel theoretical results reported in Section 3 of the main text, as well as additional material relating to the experimental assessment in Section 4 of the main text. It is structured as follows:

S1 Software	1
S2 Proofs	2
S2.1 Proof of Theorem 1	2
S2.2 Proof of Theorem 2	5
S2.3 Proof of Theorem 3	7
S2.4 Satisfying the Conditions of Theorem 2	13
S3 Experimental Protocol	15
S4 Convergence Diagnostics for MCMC	16
S5 Empirical Assessment: Additional Results	19
S5.1 Choice of Kernel	19
S5.2 Goodwin Oscillator	19
S5.3 Lotka–Volterra	33
S5.4 Calcium Signalling Model	42

S1 Software

To assist with applications of **Stein Thinning** we have provided code in **Python**, **R** and **MATLAB**. The code is available at:

<http://stein-thinning.org/>

In this section we demonstrate **Stein Thinning** in **Python**, but the syntax for **Stein Thinning** in **R** and in **MATLAB** is almost identical. As an illustration of how **Stein Thinning** can be used to post-process output from **Stan**, consider the following simple **Stan** script that produces 1000 correlated samples from a bivariate Gaussian model:

```
from pystan import StanModel
mc = """
parameters {vector[2] x;}
model {x ~ multi_normal([0, 0], [[1, 0.8], [0.8, 1]]);}
"""
sm = StanModel(model_code=mc)
fit = sm.sampling(iter=1000)
```

The bivariate Gaussian model is used for illustration, but regardless of the complexity of the model being sampled the output of `Stan` will be a `fit` object. The sampled points x_i and the gradients $\nabla \log p(x_i)$ can be extracted from the returned `fit` object:

```
import numpy as np
smpl = fit['x']
grad = np.apply_along_axis(fit.grad_log_prob, 1, smpl)
```

One can then perform `Stein Thinning` to obtain a subset of $m = 40$ states by running the following code:

```
from stein_thinning.thinning import thin
idx = thin(smpl, grad, 40)
```

The `thin` function returns a NumPy array containing the row indices in `smpl` (and `grad`) of the selected points. The default usage requires no additional user input and is based on the `sclmed` setting from Section 2.3, informed by the empirical analysis of Section 4. Alternatively, the user can choose to specify which setting to use for computing the preconditioning matrix Γ by setting the option string `pre` to either `'med'`, `'sclmed'`, or `'smpcov'`. For example, the default setting corresponds to

```
idx = thin(smpl, grad, 40, pre='sclmed')
```

The ease with which `Stein Thinning` can be used makes it possible to consider a wide variety of applications, including the ODE models that we considered in Section 4 of the main text.

S2 Proofs

This appendix contains detailed proofs for all novel theoretical results in the main text.

S2.1 Proof of Theorem 1

First we state and prove two elementary results that will be useful:

Lemma 1. *For all $a, b \geq 0$ it holds that $2a\sqrt{a^2 + b} \leq 2a^2 + b$.*

Proof. Since all quantities are non-negative, we may square both sides to get an equivalent inequality $4a^2(a^2 + b) \leq (2a^2 + b)^2$. Expanding the brackets and cancelling terms leads to $0 \leq b^2$, which is guaranteed to hold. \square

Lemma 2. *For all $m \in \mathbb{N}$ it holds that $\sum_{j=1}^m \frac{1}{j} \leq 1 + \log(m)$.*

Proof. Since $x \mapsto \frac{1}{x}$ is convex on $x \in (0, \infty)$, we have that the Riemann sum $\sum_{j=2}^m \frac{1}{j}$ is a lower bound for the Riemann integral $\int_1^m \frac{1}{x} dx = \log(m)$. Thus $\sum_{j=1}^m \frac{1}{j} = 1 + \sum_{j=2}^m \frac{1}{j} \leq 1 + \log(m)$, as required. \square

Now we present the proof of Theorem 1:

Proof of Theorem 1. Let $a_m := m^2 D_{k_P}(\frac{1}{m} \sum_{j=1}^m \delta(x_{\pi(j)}))^2$, $f_m := \sum_{j=1}^m k_P(x_{\pi(j)}, \cdot)$ and also let $S^2 := \max_{i=1, \dots, n} k_P(x_i, x_i)$, so that

$$\begin{aligned} a_m &= \sum_{j=1}^m \sum_{j'=1}^m k_P(x_{\pi(j)}, x_{\pi(j')}) = a_{m-1} + k_P(x_{\pi(m)}, x_{\pi(m)}) + 2 \sum_{j=1}^{m-1} k_P(x_{\pi(j)}, x_{\pi(m)}) \\ &\leq a_{m-1} + S^2 + 2 \min_{y \in \{x_i\}_{i=1}^n} f_{m-1}(y). \end{aligned}$$

Recall that $\mathcal{H}(k_P)$ denotes the reproducing kernel Hilbert space of the kernel k_P and pick an element $h^* \in \mathcal{H}(k_P)$ of the form $h^* := \sum_{i=1}^n w_i^* k_P(x_i, \cdot)$, where the weight vector w^* satisfies (8). From this definition it follows that $\|h^*\|_{\mathcal{H}(k_P)} = D_{k_P}(\sum_{i=1}^n w_i^* \delta(x_i))$, which is the minimal KSD attainable under the constraint (8). Now, let \mathcal{M} denote the convex hull of $\{k_P(x_i, \cdot)\}_{i=1}^n$, so that $h^* \in \mathcal{M} \subset \mathcal{H}(k_P)$ and therefore

$$\min_{y \in \{x_i\}_{i=1}^n} f_{m-1}(y) = \inf_{h \in \mathcal{M}} \langle f_{m-1}, h \rangle_{\mathcal{H}(k_P)} \leq \langle f_{m-1}, h^* \rangle_{\mathcal{H}(k_P)}. \quad (12)$$

Noting that $a_m = \|f_m\|_{\mathcal{H}(k_P)}^2$, we have from (12) and Cauchy-Schwarz that

$$\min_{y \in \{x_i\}_{i=1}^n} f_{m-1}(y) \leq \sqrt{a_{m-1}} \|h^*\|_{\mathcal{H}(k_P)}$$

and therefore

$$a_m \leq a_{m-1} + S^2 + 2\sqrt{a_{m-1}} \|h^*\|_{\mathcal{H}(k_P)}. \quad (13)$$

Letting

$$C_m := \frac{1}{m} (S^2 - \|h^*\|_{\mathcal{H}(k_P)}^2) \sum_{j=1}^m \frac{1}{j}, \quad (14)$$

we will establish by induction that

$$a_m \leq m^2 (\|h^*\|_{\mathcal{H}(k_P)}^2 + C_m). \quad (15)$$

This will in turn prove the result, since

$$\begin{aligned} D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(x_{\pi(j)}) \right)^2 &= \frac{a_m}{m^2} \leq \|h^*\|_{\mathcal{H}(k_P)}^2 + C_m \\ &= D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(x_i) \right)^2 + C_m \\ &\leq D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(x_i) \right)^2 + \left(\frac{1 + \log(m)}{m} \right) S^2, \end{aligned}$$

where the upper bound on C_m follows from the fact that $\|h^*\|_{\mathcal{H}(k_0)} \leq S$, combined with Lemma 2.

The remainder of the proof is dedicated to establishing the induction in (15). The base case $m = 1$ is satisfied since $a_1 = D_{k_P}(\delta(x_{\pi(1)})) = k_P(x_{\pi(1)}, x_{\pi(1)}) \leq S^2$ and $C_1 = S^2 - \|h^*\|_{\mathcal{H}(k_P)}^2$, so that $a_1 \leq \|h^*\|_{\mathcal{H}(k_P)}^2 + C_1$. For the inductive step, we assume that (15) holds when m is replaced by $m - 1$ and aim to derive (15). From (13) and the inductive assumption, we have that

$$\begin{aligned} a_m &\leq a_{m-1} + S^2 + 2\sqrt{a_{m-1}}\|h^*\|_{\mathcal{H}(k_P)} \\ &\leq (m-1)^2(\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1}) + S^2 + 2(m-1)\sqrt{\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1}}\|h^*\|_{\mathcal{H}(k_P)} \\ &= m^2(\|h^*\|_{\mathcal{H}(k_P)}^2 + C_m) + R_m \end{aligned} \quad (16)$$

where

$$\begin{aligned} R_m &:= (m-1)^2C_{m-1} - m^2C_m + (1-2m)\|h^*\|_{\mathcal{H}(k_P)}^2 + S^2 \\ &\quad + 2(m-1)\sqrt{\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1}}\|h^*\|_{\mathcal{H}(k_P)} \end{aligned}$$

The induction (15) will therefore follow from (16) if $R_m \leq 0$. Now, $R_m \leq 0$ if and only if

$$2\sqrt{\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1}}\|h^*\|_{\mathcal{H}(k_P)} \leq \frac{m^2C_m - (m-1)^2C_{m-1}}{m-1} - \frac{S^2 - \|h^*\|_{\mathcal{H}(k_P)}^2}{m-1} + 2\|h^*\|_{\mathcal{H}(k_P)}^2.$$

From Lemma 1 it must hold that

$$2\sqrt{\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1}}\|h^*\|_{\mathcal{H}(k_P)} \leq 2\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1},$$

meaning it is sufficient to show that

$$2\|h^*\|_{\mathcal{H}(k_P)}^2 + C_{m-1} \leq \frac{m^2C_m - (m-1)^2C_{m-1}}{m-1} - \frac{S^2 - \|h^*\|_{\mathcal{H}(k_P)}^2}{m-1} + 2\|h^*\|_{\mathcal{H}(k_P)}^2. \quad (17)$$

Algebraic simplification of (17) reveals that (17) is equivalent to

$$mC_m - (m-1)C_{m-1} \leq \frac{1}{m}(S^2 - \|h^*\|_{\mathcal{H}(k_P)}^2) \quad (18)$$

and, using (14), we verify that (18) is satisfied as an equality. This completes the inductive argument. \square

Remark 10. Our results can also be applied to maximum mean discrepancies, where $D_{\mathcal{F}}$ in (4) is given by $\mathcal{F} = \{f \in \mathcal{H}(k) : \|f\|_{\mathcal{H}(k)} \leq 1\}$ and $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ denotes a reproducing kernel (Song, 2008). Indeed, we can set

$$k_P(x, y) = k(x, y) - \int_{\mathcal{X}} k(x, y) dP(x) - \int_{\mathcal{X}} k(x, y) dP(y) + \int_{\mathcal{X}} \int_{\mathcal{X}} k(x, y) dP(x) dP(y)$$

in order for $\int_{\mathcal{X}} k_P(x, \cdot) dP = 0$ to be satisfied for all $x \in \mathcal{X}$, and observe that this construction ensures $D_{\mathcal{F}}$ and D_{k_P} are identical. Teymur et al. (2021) explores these consequences of our results in detail.

S2.2 Proof of Theorem 2

First we state and prove a technical lemma that will be useful:

Lemma 3. *Let \mathcal{X} be a measurable space and let Q be a probability distribution on \mathcal{X} . Let $k_Q : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a reproducing kernel with $\int_{\mathcal{X}} k_Q(x, \cdot) dQ = 0$ for all $x \in \mathcal{X}$. Consider a Q -invariant, time-homogeneous Markov chain $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$ generated using a V -uniformly ergodic transition kernel, such that $V(x) \geq \sqrt{k_Q(x, x)}$ for all $x \in \mathcal{X}$, with parameters $R \in [0, \infty)$ and $\rho \in (0, 1)$ as in (9). Then with $C = \frac{2R\rho}{1-\rho}$ we have that*

$$\sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_Q(X_i, X_r)] \leq C \sum_{i=1}^{n-1} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right].$$

Proof. First recall that given random variables X, Y taking values in \mathcal{X} , the conditional mean embedding of the distribution $\mathbb{P}[X|Y = y]$ is the function $\mathbb{E}[k_Q(X, \cdot)|Y = y] \in \mathcal{H}(k_Q)$ (Song et al., 2009). By the reproducing property we have $\mathbb{E}[k_Q(X, y)|Y = y] = \langle k_Q(y, \cdot), \mathbb{E}[k_Q(X, \cdot)|Y = y] \rangle_{\mathcal{H}(k_Q)}$, hence $\mathbb{E}[k_Q(X, Y)|Y] = \langle k_Q(Y, \cdot), \mathbb{E}[k_Q(X, \cdot)|Y] \rangle_{\mathcal{H}(k_Q)}$. Thus

$$\begin{aligned} \mathbb{E}[k_Q(X, Y)|Y] &= \langle k_Q(Y, \cdot), \mathbb{E}[k_Q(X, \cdot)|Y] \rangle_{\mathcal{H}(k_Q)} \\ &= \|k_Q(Y, \cdot)\|_{\mathcal{H}(k_Q)} \left\langle \frac{k_Q(Y, \cdot)}{\|k_Q(Y, \cdot)\|_{\mathcal{H}(k_Q)}}, \mathbb{E}[k_Q(X, \cdot)|Y] \right\rangle_{\mathcal{H}(k_Q)} \\ &\leq \|k_Q(Y, \cdot)\|_{\mathcal{H}(k_Q)} \sup_{\|h\|_{\mathcal{H}(k_Q)}=1} \langle h, \mathbb{E}[k_Q(X, \cdot)|Y] \rangle_{\mathcal{H}(k_Q)}. \end{aligned}$$

In what follows it is convenient to introduce a new random variable Z , independent from the Markov chain, such that $Z \sim Q$. Then, since $\mathbb{E}[k_Q(Z, \cdot)] = 0$, we have $\mathbb{E}[h(Z)] = 0$ for any $h \in \mathcal{H}(k_Q)$. Hence we have that

$$\mathbb{E}[k_Q(X, Y)|Y] \leq \sqrt{k_Q(Y, Y)} \sup_{\|h\|_{\mathcal{H}(k_Q)}=1} (\langle h, \mathbb{E}[k_Q(X, \cdot)|Y] \rangle - \mathbb{E}[h(Z)]).$$

Note $|h(x)| \leq \|h\|_{\mathcal{H}(k_Q)} \sqrt{k_Q(x, x)}$, so $\|h\|_{\mathcal{H}(k_Q)} = 1$ implies $|h(x)| \leq \sqrt{k_Q(x, x)}$, thus

$$\mathbb{E}[k_Q(X, Y)|Y] \leq \sqrt{k_Q(Y, Y)} \sup_{|h(x)| \leq \sqrt{k_Q(x, x)}} (\mathbb{E}[h(X)|Y] - \mathbb{E}[h(Z)]).$$

Let Q^n denote the n^{th} step transition kernel of the Markov chain. From V -uniform ergodicity it follows that

$$|\mathbb{E}[h(X_n)|X_0 = y] - \mathbb{E}[h(Z)]| = \left| \int h(x) Q^n(y, dx) - \int h(x) dQ(x) \right| \leq RV(y) \rho^n.$$

Applying this to $Y = X_i$, $X = X_{i+r}$, we find

$$\mathbb{E}[k_Q(X_{i+r}, X_i)|X_i] \leq R \sqrt{k_Q(X_i, X_i)} V(X_i) \rho^r$$

and taking the expectation on both sides yields

$$\mathbb{E}[k_Q(X_{i+r}, X_i)] \leq R \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right] \rho^r. \quad (19)$$

Finally, we can use (19) to obtain that

$$\begin{aligned} \sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_Q(X_r, X_i)] &= 2 \sum_{i=1}^{n-1} \sum_{r=1}^{n-i} \mathbb{E}[k_Q(X_{i+r}, X_i)] \\ &\leq 2R \sum_{i=1}^{n-1} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right] \sum_{r=1}^{n-i} \rho^r. \end{aligned}$$

Thus for $C = 2R \sum_{r=1}^{\infty} \rho^r = \frac{2R\rho}{1-\rho} < \infty$, we have that

$$\sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_Q(X_r, X_i)] \leq C \sum_{i=1}^{n-1} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right]$$

as claimed. \square

We can now prove the main result:

Proof of Theorem 2. Taking expectations of the bound in Theorem 1, we have that

$$\begin{aligned} \mathbb{E} \left[D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \right)^2 \right] &\leq \underbrace{\mathbb{E} \left[D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(X_i) \right)^2 \right]}_{(*)} \\ &\quad + \underbrace{\left(\frac{1 + \log(m)}{m} \right) \mathbb{E} \left[\max_{i=1, \dots, n} k_P(X_i, X_i) \right]}_{(**)}. \end{aligned}$$

In what follows we construct bounds for $(*)$ and $(**)$.

Bounding $(*)$: To bound the term $(*)$, note that

$$D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(X_i) \right) \leq D_{k_P} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right),$$

due to the optimality property of the weights w^* presented in (8). It is therefore sufficient to study the KSD of the un-weighted empirical distribution $\frac{1}{n} \sum_{i=1}^n \delta(X_i)$. To this end, we have that

$$\mathbb{E} \left[D_{k_P} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right)^2 \right] = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[k_P(X_i, X_i)] + \frac{1}{n^2} \sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_P(X_i, X_r)]. \quad (20)$$

To bound the first term in (20) we use Jensen's inequality:

$$\begin{aligned} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[k_P(X_i, X_i)] &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\frac{1}{\gamma} \log e^{\gamma k_P(X_i, X_i)} \right] \\ &\leq \frac{1}{\gamma n^2} \sum_{i=1}^n \log \left(\mathbb{E} [e^{\gamma k_P(X_i, X_i)}] \right) \leq \frac{\log(b)}{\gamma n} \end{aligned}$$

The second term in (20) can be bounded via Lemma 3 with $Q = P$:

$$\frac{1}{n^2} \sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_P(X_i, X_r)] \leq \frac{C}{n^2} \sum_{i=1}^{n-1} \mathbb{E} \left[\sqrt{k_P(X_i, X_i)} V(X_i) \right] \leq \frac{CM(n-1)}{n^2} \leq \frac{CM}{n},$$

where C is defined in Lemma 3.

Bounding ():** We proceed as follows:

$$\begin{aligned} \mathbb{E} \left[\max_{i=1, \dots, n} k_P(X_i, X_i) \right] &= \mathbb{E} \left[\frac{1}{\gamma} \log \max_{i=1, \dots, n} e^{\gamma k_P(X_i, X_i)} \right] \\ &\leq \mathbb{E} \left[\frac{1}{\gamma} \log \sum_{i=1}^n e^{\gamma k_P(X_i, X_i)} \right] \\ &\leq \frac{1}{\gamma} \log \left(\sum_{i=1}^n \mathbb{E} [e^{\gamma k_P(X_i, X_i)}] \right) = \frac{\log(nb)}{\gamma} \end{aligned} \quad (21)$$

Overall Bound: Combining our bounds on (*) and (**) leads to the overall bound

$$\mathbb{E} \left[D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \right)^2 \right] \leq \frac{\log(b)}{\gamma n} + \frac{CM}{n} + \left(\frac{1 + \log(m)}{m} \right) \frac{\log(nb)}{\gamma}$$

as claimed. □

S2.3 Proof of Theorem 3

To facilitate a neat proof of Theorem 3 we first present two useful lemmas, the first of which establishes almost sure convergence in KSD of the empirical distribution based on the full MCMC output:

Lemma 4. *Let Q be a probability distribution on \mathcal{X} . Let $k_Q : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a reproducing kernel with $\int_{\mathcal{X}} k_Q(x, \cdot) dQ = 0$ for all $x \in \mathcal{X}$. Consider a Q -invariant, time-homogeneous Markov chain $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$, generated using a V -uniformly ergodic transition kernel such that $V(x) \geq \sqrt{k_Q(x, x)}$ for all $x \in \mathcal{X}$. Suppose that, for some $\gamma > 0$,*

$$b := \sup_{i \in \mathbb{N}} \mathbb{E} [e^{\gamma k_Q(X_i, X_i)}] < \infty, \quad M := \sup_{i \in \mathbb{N}} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right] < \infty.$$

Then

$$D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right) \rightarrow 0$$

almost surely as $n \rightarrow \infty$.

Proof. Similarly to the proof of Theorem 2, we start by bounding

$$\mathbb{E} \left[D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right)^2 \right] = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[k_Q(X_i, X_i)] + \frac{1}{n^2} \sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_Q(X_i, X_r)]. \quad (22)$$

To bound the first term in (22) we use Jensen's inequality:

$$\begin{aligned} \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}[k_Q(X_i, X_i)] &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\frac{1}{\gamma} \log e^{\gamma k_Q(X_i, X_i)} \right] \\ &\leq \frac{1}{\gamma n^2} \sum_{i=1}^n \log (\mathbb{E}[e^{\gamma k_Q(X_i, X_i)}]) \leq \frac{\log(b)}{\gamma n} \end{aligned}$$

The second term in (22) can be bounded via Lemma 3:

$$\frac{1}{n^2} \sum_{i=1}^n \sum_{r \in \{1, \dots, n\} \setminus \{i\}} \mathbb{E}[k_Q(X_i, X_r)] \leq \frac{C}{n^2} \sum_{i=1}^{n-1} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i) V(X_i)} \right] \leq \frac{CM(n-1)}{n^2} \leq \frac{CM}{n},$$

where C is defined in Lemma 3. This establishes that

$$\mathbb{E} \left[D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right)^2 \right] \leq \frac{\log(b)}{\gamma n} + \frac{CM}{n} =: c_1(n). \quad (23)$$

To simplify notation we adopt the shorthand

$$D_n := D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right)$$

in this proof only. Fix $\epsilon > 0$. If $D_n > \epsilon$ occurs infinitely often (i.o.) then there are infinitely many r such that $\max_{r^2 \leq n < (r+1)^2} D_n^2 > \epsilon$, so that

$$\mathbb{P} [D_n^2 > \epsilon \text{ i.o.}] \leq \mathbb{P} \left[\max_{r^2 \leq n < (r+1)^2} D_n^2 > \epsilon \text{ i.o.} \right]. \quad (24)$$

Now, consider the bound

$$\sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq n < (r+1)^2} D_n^2 > \epsilon \right] \leq \underbrace{\sum_{r=1}^{\infty} \mathbb{P} \left[D_{r^2}^2 > \frac{\epsilon}{2} \right]}_{(*)} + \underbrace{\sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq n < (r+1)^2} |D_n^2 - D_{r^2}^2| > \frac{\epsilon}{2} \right]}_{(**)},$$

where the inequality follows from the fact that, for any $a, b \in \mathbb{R}$, if $a > \epsilon$ then either $b > \frac{\epsilon}{2}$ or $|a - b| > \frac{\epsilon}{2}$. In the remainder we will show that the sums $(*)$ and $(**)$ are finite, so that from the Borel–Cantelli lemma

$$\mathbb{P} \left[\max_{r^2 \leq n < (r+1)^2} D_n^2 > \epsilon \text{ i.o.} \right] = 0. \quad (25)$$

Since (25) holds for all $\epsilon > 0$, it will follow from (24) that $\mathbb{P}[D_n \rightarrow 0] = 1$, as claimed.

Bounding $(*)$: From the Markov inequality and (23) we have that, for any $\epsilon > 0$,

$$\mathbb{P} \left[D_{r^2}^2 > \frac{\epsilon}{2} \right] \leq \frac{2}{\epsilon} \mathbb{E}[D_{r^2}^2] \leq \frac{2}{\epsilon} c_1(r^2).$$

Since $c_1(r^2) = O(1/r^2)$, it follows that

$$(*) = \sum_{r=1}^{\infty} \mathbb{P} \left[D_{r^2}^2 > \frac{\epsilon}{2} \right] \leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} c_1(r^2) < \infty.$$

Bounding $()$:** For $r^2 \leq n < (r+1)^2$,

$$\begin{aligned} |D_n^2 - D_{r^2}^2| &= \frac{1}{n^2} |n^2 D_n^2 - n^2 D_{r^2}^2| \leq \frac{1}{n^2} |n^2 D_n^2 - r^4 D_{r^2}^2| + \frac{1}{n^2} |(n^2 - r^4) D_{r^2}^2| \\ &\leq \frac{1}{r^4} |n^2 D_n^2 - r^4 D_{r^2}^2| + \frac{2(r+1)}{r^2} D_{r^2}^2, \end{aligned}$$

and also that, using the reproducing property and Cauchy-Schwarz,

$$\begin{aligned} |n^2 D_n^2 - r^4 D_{r^2}^2| &= \left| \sum_{i,j=r^2+1}^n k_Q(X_i, X_j) \right| \\ &\leq \sum_{i,j=r^2+1}^n \sqrt{k_Q(X_i, X_i)} \sqrt{k_Q(X_j, X_j)} \\ &\leq (n - r^2)^2 \max_{1 \leq i \leq n} k_Q(X_i, X_i) \leq 4r^2 \max_{1 \leq i \leq n} k_Q(X_i, X_i). \end{aligned}$$

Similarly again to the proof of Theorem 2 we have the bound

$$\begin{aligned} \mathbb{E} \left[\max_{i=1, \dots, n} k_Q(X_i, X_i) \right] &= \mathbb{E} \left[\frac{1}{\gamma} \log \max_{i=1, \dots, n} e^{\gamma k_Q(X_i, X_i)} \right] \\ &\leq \mathbb{E} \left[\frac{1}{\gamma} \log \sum_{i=1}^n e^{\gamma k_Q(X_i, X_i)} \right] \\ &\leq \frac{1}{\gamma} \log \left(\sum_{i=1}^n \mathbb{E} [e^{\gamma k_Q(X_i, X_i)}] \right) = \frac{\log(nb)}{\gamma} \end{aligned} \quad (26)$$

so that, taking expectations, we obtain the bound

$$\begin{aligned}
\mathbb{E} \left[\max_{r^2 \leq n < (r+1)^2} |D_n^2 - D_{r^2}^2| \right] &\leq \frac{4}{r^2} \mathbb{E} \left[\max_{1 \leq i \leq n} k_Q(X_i, X_i) \right] + \frac{2(r+1)}{r^2} \mathbb{E} [D_{r^2}^2] \\
&\leq \frac{4}{r^2} \frac{\log(nb)}{\gamma} + \frac{2(r+1)}{r^2} c_1(r^2) \\
&\leq \frac{8}{r^2} \frac{\log((r+1)b)}{\gamma} + \frac{2(r+1)}{r^2} c_1(r^2) =: c_2(r)
\end{aligned}$$

where $c_2(r) = O(\log(r)/r^2)$. Using the Markov inequality,

$$\begin{aligned}
(**) = \sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq n < (r+1)^2} |D_n^2 - D_{r^2}^2| > \frac{\epsilon}{2} \right] &\leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} \mathbb{E} \left[\max_{r^2 \leq n < (r+1)^2} |D_n^2 - D_{r^2}^2| \right] \\
&\leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} c_2(r) < \infty.
\end{aligned}$$

This completes the proof. \square

Our second lemma is a technical result on almost sure convergence:

Lemma 5. *Let f be a non-negative function on \mathcal{X} . Consider a sequence of random variables $(X_i)_{i \in \mathbb{N}} \subset \mathcal{X}$ such that, for some $\gamma > 0$,*

$$b := \sup_{i \in \mathbb{N}} \mathbb{E} [e^{\gamma f(X_i)}] < \infty$$

If $m \leq n$ and the growth of n is limited to at most $\log(n) = O(m^{\beta/2})$ for some $\beta < 1$, then

$$\left(\frac{\log(m)}{m} \right) \max_{i=1, \dots, n} f(X_i) \rightarrow 0$$

almost surely as $m, n \rightarrow \infty$.

Proof. To simplify notation we adopt the shorthand

$$E_m := \left(\frac{\log(m)}{m} \right) \max_{i=1, \dots, n} f(X_i)$$

in this proof only, where $n = n(m)$. The argument is similar to the proof of Lemma 4. Fix $\epsilon > 0$. If $E_m > \epsilon$ i.o. then there are infinitely many r such that $\max_{r^2 \leq m < (r+1)^2} E_m > \epsilon$, so that

$$\mathbb{P} [E_m > \epsilon \text{ i.o.}] \leq \mathbb{P} \left[\max_{r^2 \leq m < (r+1)^2} E_m > \epsilon \text{ i.o.} \right]. \quad (27)$$

Now, consider the bound

$$\sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq m < (r+1)^2} E_m > \epsilon \right] \leq \underbrace{\sum_{r=1}^{\infty} \mathbb{P} \left[E_{r^2} > \frac{\epsilon}{2} \right]}_{(*)} + \underbrace{\sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq m < (r+1)^2} |E_m - E_{r^2}| > \frac{\epsilon}{2} \right]}_{(**)},$$

In the remainder we will show that the sums $(*)$ and $(**)$ are finite, so that from the Borel–Cantelli lemma

$$\mathbb{P} \left[\max_{r^2 \leq m < (r+1)^2} E_m > \epsilon \text{ i.o.} \right] = 0. \quad (28)$$

Since (28) holds for all $\epsilon > 0$, it will follow from (27) that $\mathbb{P}[E_m \rightarrow 0] = 1$, as claimed.

Bounding $(*)$: Similarly to the proof of Theorem 2, we have the bound

$$\begin{aligned} \mathbb{E} \left[\max_{i=1, \dots, n} f(X_i) \right] &= \mathbb{E} \left[\frac{1}{\gamma} \log \max_{i=1, \dots, n} e^{\gamma f(X_i)} \right] \\ &\leq \mathbb{E} \left[\frac{1}{\gamma} \log \sum_{i=1}^n e^{\gamma f(X_i)} \right] \\ &\leq \frac{1}{\gamma} \log \left(\sum_{i=1}^n \mathbb{E} [e^{\gamma f(X_i)}] \right) = \frac{\log(nb)}{\gamma} \end{aligned}$$

and thus from the Markov inequality we have that, for any $\epsilon > 0$,

$$\mathbb{P} \left[E_m > \frac{\epsilon}{2} \right] \leq \frac{2}{\epsilon} \mathbb{E}[E_m] \leq \frac{2}{\epsilon} c_1(m)$$

where

$$c_1(m) := \frac{\log(m)}{m} \frac{\log(nb)}{\gamma}.$$

The assumption $\log(n) = O(m^{\beta/2})$ for some $\beta < 1$ implies that $\log(nb) \leq \alpha m^{\beta/2} + \log(b)$ for some constant $\alpha \in (0, \infty)$. Thus

$$c_1(r^2) \leq \frac{2 \log(r)}{r^2} \left(\frac{\alpha r^{\beta} + \log(b)}{\gamma} \right).$$

This shows that $c_1(r^2) = O(\log(r)/r^{2-\beta})$, and it follows that

$$(*) = \sum_{r=1}^{\infty} \mathbb{P} \left[E_{r^2} > \frac{\epsilon}{2} \right] \leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} c_1(r^2) < \infty.$$

Bounding $()$:** For the second term we argue that, since $E_m \geq 0$,

$$\begin{aligned} \max_{r^2 \leq m < (r+1)^2} |E_m - E_{r^2}| &\leq 2 \max_{r^2 \leq m < (r+1)^2} E_m \\ &\leq \frac{2 \log(r^2)}{r^2} \max_{i=1, \dots, n((r+1)^2)} f(X_i) \end{aligned}$$

so that, taking expectations,

$$\mathbb{E} \left[\max_{r^2 \leq m < (r+1)^2} |E_m - E_{r^2}| \right] \leq \frac{4 \log(r) \log(n((r+1)^2)b)}{r^2 \gamma} =: c_2(r)$$

Using the bound $\log(nb) \leq \alpha m^{\beta/2} + \log(b)$, the quantity $c_2(r)$ just defined satisfies

$$c_2(r) \leq \frac{4 \log(r)}{r^2} \left(\frac{\alpha(r+1)^\beta + \log(b)}{\gamma} \right),$$

which is $O(\log(r)/r^{2-\beta})$. Using the Markov inequality and the fact that $(a-b)^2 \leq |a^2 - b^2|$,

$$\begin{aligned} (**) &= \sum_{r=1}^{\infty} \mathbb{P} \left[\max_{r^2 \leq m < (r+1)^2} |E_m - E_{r^2}| > \frac{\epsilon}{2} \right] \leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} \mathbb{E} \left[\max_{r^2 \leq m < (r+1)^2} |E_m - E_{r^2}| \right] \\ &\leq \frac{2}{\epsilon} \sum_{r=1}^{\infty} c_2(r) < \infty. \end{aligned}$$

This completes the proof. □

Now we present the proof of Theorem 3:

Proof of Theorem 3. Our starting point is again the bound in Theorem 1:

$$D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \right)^2 \leq \underbrace{D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(X_i) \right)^2}_{(*)} + \underbrace{\left(\frac{1 + \log(m)}{m} \right) \max_{i=1, \dots, n} k_P(X_i, X_i)}_{(**)}.$$

For term (*), note that

$$D_{k_P} \left(\sum_{i=1}^n w_i^* \delta(X_i) \right) \leq D_{k_P} \left(\sum_{i=1}^n w_i \delta(X_i) \right), \quad w_i := \frac{\frac{dP}{dQ}(X_i)}{\sum_{i'=1}^n \frac{dP}{dQ}(X_{i'})}$$

due to the optimality property of the weights w^* presented in (8). Further note that

$$D_{k_P} \left(\sum_{i=1}^n w_i \delta(X_i) \right) = \frac{1}{\frac{1}{n} \sum_{i'=1}^n \frac{dP}{dQ}(X_{i'})} D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right)$$

where $k_Q(x, y) := \frac{dP}{dQ}(x) k_P(x, y) \frac{dP}{dQ}(y)$ is a reproducing kernel such that $\int_{\mathcal{X}} k_Q(x, \cdot) dQ = 0$ for all $x \in \mathcal{X}$. The preconditions of Theorem 3 ensure that $V(x) \geq \sqrt{k_Q(x, x)}$ and

$$\sup_{i \in \mathbb{N}} \mathbb{E} [e^{\gamma k_Q(X_i, X_i)}] \leq b < \infty, \quad M = \sup_{i \in \mathbb{N}} \mathbb{E} \left[\sqrt{k_Q(X_i, X_i)} V(X_i) \right] < \infty.$$

Therefore we may apply Lemma 4 to obtain that

$$D_{k_Q} \left(\frac{1}{n} \sum_{i=1}^n \delta(X_i) \right) \rightarrow 0$$

almost surely as $n \rightarrow \infty$. Moreover, since $\int_{\mathcal{X}} \left| \frac{dP}{dQ} \right| dQ < \infty$, it follows from Meyn and Tweedie (2012, Theorem 17.0.1, part (i)) that

$$\frac{1}{n} \sum_{i=1}^n \frac{dP}{dQ}(X_i) \rightarrow \int_{\mathcal{X}} \frac{dP}{dQ} dQ = 1$$

almost surely as $n \rightarrow \infty$. Standard properties of almost sure convergence thus imply that $(*) \rightarrow 0$ almost surely as $n \rightarrow \infty$.

For term $(**)$, we notice that

$$\sup_{i \in \mathbb{N}} \mathbb{E} [e^{\gamma k_P(X_i, X_i)}] \leq b < \infty$$

and we can therefore use Lemma 5 with $f(x) = k_P(x, x)$ to deduce that $(**) \rightarrow 0$ almost surely as $m, n \rightarrow \infty$.

Thus we have established that

$$D_{k_P} \left(\frac{1}{m} \sum_{j=1}^m \delta(X_{\pi(j)}) \right) \rightarrow 0 \tag{29}$$

almost surely as $m, n \rightarrow \infty$. The final part of the statement of Theorem 3 is immediate from Proposition 3. \square

S2.4 Satisfying the Conditions of Theorem 2

The conditions for Theorem 2 are agnostic to the specific Markov chain used (e.g. Metropolis–Hastings, Gibbs sampling, etc), making it quite general. In this appendix we discuss how explicit sufficient conditions can be obtained if one restricts attention to a specific MCMC method. Here we focus on the *Metropolis-adjusted Langevin algorithm* (MALA; whose definition is recalled in Appendix S3).

Let $q(x, y)$ denote the probability density for the proposal $x \rightarrow y$ in MALA, with step size $\epsilon > 0$ fixed. Let $A(x) \subset \mathbb{R}^d$ be the set of values y which, if a move $x \rightarrow y$ is proposed, then y is always accepted. Let $I(x) := \{y \in \mathbb{R}^d : \|y\| \leq \|x\|\}$ and let $A \Delta B := (A \cup B) \setminus (A \cap B)$. MALA is said to be *inwardly convergent* if

$$\lim_{\|x\| \rightarrow \infty} \int_{A(x) \Delta I(x)} q(x, y) dy = 0, \tag{30}$$

see Section 4 of Roberts and Tweedie (1996). The following result will then be established:

Lemma 6. *Let $P \in \mathcal{P}$ be distantly dissipative on $\mathcal{X} = \mathbb{R}^d$, let $\mathbb{E}_{X \sim P}[\exp(\beta \|X\|^2)] < \infty$ for some $\beta \in (0, \infty)$, and assume that MALA is inwardly convergent. Then, with kernel $k(x, y) = (1 + \|x - y\|^2)^{-1/2}$, the conditions of Theorem 2 are satisfied.*

Proof. This proof exploits Theorem 9 of Chen et al. (2019), which establishes V -uniform ergodicity of MALA for each of $V(x) = \exp(\gamma \|x\|)$ (any $\gamma > 0$), $V(x) = \exp(\gamma \|x\|^2)$ (for $\gamma > 0$ sufficiently small) and $V(x) = 1 + \|x\|^\gamma$ ($\gamma \in \{1, 2\}$). Each choice of V leads to a different set of preconditions for Theorem 2, and the claimed result follows from taking $V(x) = 1 + \|x\|$. Note that w.l.o.g. we can consider $V(x) = C(1 + \|x\|)$ for any fixed $C \in [1, \infty)$, since the constant C cancels in the definition of V -uniform ergodicity. The conditions for Theorem 2 now simplify as follows:

First Condition: The form of k implies that $k_P(x, x) = d + \|\nabla \log p(x)\|^2$ (see Remark 9). Since $\nabla \log p$ is assumed to be Lipschitz we have, for C sufficiently large,

$$\begin{aligned} V(x) \geq \sqrt{k_P(x, x)} &\Leftrightarrow C(1 + \|x\|) \geq \sqrt{d + \|\nabla \log p(x)\|^2} \\ &\Leftrightarrow \|\nabla \log p(x)\|^2 \leq C_1 + C_2 \|x\|^2 \quad (\text{some } C_1, C_2 > 0), \end{aligned}$$

so that the first condition of Theorem 2 is automatically satisfied.

Second Condition: Now, suppose MALA is also \tilde{V} -uniformly ergodic; i.e. $\|P^i(x, \cdot) - P\|_{\tilde{V}} \leq \tilde{R}\tilde{V}(x)\tilde{\rho}^n$ for some $\tilde{R} \in [0, \infty)$, $\tilde{\rho} \in (0, 1)$. Let $X \sim P$ independent of X_i . Then

$$\begin{aligned} \mathbb{E}[\tilde{V}(X_i)] &= \mathbb{E}[\tilde{V}(X)] + \mathbb{E}[\tilde{V}(X_i) - \tilde{V}(X)] \\ &\leq \mathbb{E}[\tilde{V}(X)] + \|P^i(X_0, \cdot) - P\|_{\tilde{V}} \|\tilde{V}\|_{\tilde{V}} \\ &\leq \mathbb{E}[\tilde{V}(X)] + \tilde{R}\tilde{V}(X_0)\tilde{\rho}^i \rightarrow \mathbb{E}[\tilde{V}(X)] \quad \text{as } i \rightarrow \infty \end{aligned}$$

which shows that

$$\mathbb{E}[\tilde{V}(X)] < \infty \Rightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[\tilde{V}(X_i)] < \infty. \quad (31)$$

From Theorem 9 of Chen et al. (2019) we have \tilde{V} -uniform ergodicity for $\tilde{V}(x) = \exp(\gamma C_2 \|x\|^2)$ (for $\gamma > 0$ sufficiently small), which shows that, for all $\gamma > 0$ sufficiently small

$$\begin{aligned} b = \sup_{i \in \mathbb{N}} \mathbb{E}[\exp(\gamma k_P(X_i, X_i))] &< \infty \Leftrightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[\exp(\gamma \|\nabla \log p(X_i)\|^2)] < \infty \\ &\Leftrightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[\exp(\gamma C_2 \|X_i\|^2)] < \infty \\ &\Leftrightarrow \mathbb{E}[\exp(\gamma C_2 \|X\|^2)] < \infty \quad (\text{due to (31)}). \end{aligned} \quad (32)$$

Thus the second condition of Theorem 2 is satisfied if $\mathbb{E}[\exp(\beta \|X\|^2)] < \infty$ for some $\beta \in (0, \infty)$.

Third Condition: A similar argument used for the second condition can again be used,

this time with $\tilde{V}(x) = 1 + \|x\|^s$ for $s \in \{1, 2\}$. Specifically, we have that

$$\begin{aligned}
M = \sup_{i \in \mathbb{N}} \mathbb{E}[\sqrt{k_P(X_i, X_i)} V(X_i)] < \infty &\Leftrightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[\sqrt{d + \|\nabla \log p(X_i)\|^2} (1 + \|X_i\|)] < \infty \\
&\Leftrightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[\sqrt{d + C_1 + C_2 \|X_i\|^2} (1 + \|X_i\|)] < \infty \\
&\Leftrightarrow \sup_{i \in \mathbb{N}} \mathbb{E}[1 + \|X_i\|], \sup_{i \in \mathbb{N}} \mathbb{E}[1 + \|X_i\|^2] < \infty \\
&\Leftrightarrow \mathbb{E}[1 + \|X\|], \mathbb{E}[1 + \|X\|^2] < \infty \quad (\text{due to (31)})
\end{aligned}$$

with the latter being implied by the stronger moment condition in (32). \square

The sufficient conditions presented in Lemma 6 may be explicitly verified, with the possible exception of the inwards convergence condition of Roberts and Rosenthal (2004).

S3 Experimental Protocol

In this appendix we describe the generic structure of a parameter inference problem for a system of ODEs, that forms our empirical test-bed.

Consider the solution u of a system of q coupled ODEs of the form

$$\begin{aligned}
\frac{du_1}{dt} &= F_1(t, u_1, \dots, u_q; x) \\
&\vdots \\
\frac{du_q}{dt} &= F_q(t, u_1, \dots, u_q; x),
\end{aligned} \tag{33}$$

together with the initial condition $u(0) = u^0 \in \mathbb{R}^q$. The functions F_i that define the gradient field are assumed to depend on a number d of parameters, collectively denoted $x \in \mathbb{R}^d$, and the F_i are assumed to be differentiable with respect to u_1, \dots, u_q and x . It is assumed that $u(t)$ exists and is unique on an interval $t \in [0, T]$ for all values $x \in \mathbb{R}^d$. For simplicity in this work we assumed that the initial condition u^0 is not dependent on x and is known. The goal is to make inferences about the parameters x based on noisy observations of the state vector $u(t_i)$ at discrete times t_i ; this information is assumed to be contained in a likelihood of the form

$$\mathcal{L}(x) := \prod_{i=1}^N \phi_i(u(t_i)) \tag{34}$$

where the functions $\phi_i : \mathbb{R}^q \rightarrow [0, \infty)$, describing the nature of the measurement at time t_i , are problem-specific and to be specified. The parameter x is endowed with a prior density $\pi(x)$ and the posterior of interest P admits a density $p(x) \propto \pi(x)\mathcal{L}(x)$. Computation of the gradient $\nabla \log p$ therefore requires computation of $\nabla \log \pi$ and $\nabla \log \mathcal{L}$; the latter can

be performed by augmenting the system in (33) with the sensitivity equations, as described next.

Straight-forward application of the chain rule leads to the following expression for the gradient of the log-likelihood:

$$(\nabla \log \mathcal{L})(x) = - \sum_{i=1}^N \frac{\partial u}{\partial x}(t_i) (\nabla \log \phi_i)(u(t_i)),$$

where $(\partial u / \partial x)_{r,s} := \partial u_r / \partial x_s$ is the matrix of *sensitivities* of the solution u to the parameter x and is time-dependent. Sensitivities can be computed by augmenting the system in (33) and simultaneously solving the *forward sensitivity equations*

$$\frac{d}{dt} \left(\frac{\partial u_r}{\partial x_s} \right) = \frac{\partial F_r}{\partial x_s} + \sum_{l=1}^q \frac{\partial F_r}{\partial u_l} \frac{\partial u_l}{\partial x_s} \quad (35)$$

together with the initial condition $(\partial u_r / \partial x_s)(0) = 0$, which follows from the independence of u^0 and x .

The experiments reported in Section 4 were based on four distinct Metropolis–Hastings MCMC methods, whose details have not yet been described. The generic structure of the proposal mechanism is $x^* = x_{n-1} + H \nabla \log p(x_{n-1}) + G \xi_n$, where the $\xi_n \sim \mathcal{N}(0, I)$ are independent. The matrices H and G are specified in Table S1. Our implementation of these samplers interfaces with the **CVODES** library (Hindmarsh et al., 2005), which presents a practical barrier to reproducibility. Moreover, the CPU time required to obtain MCMC samples was approximately two weeks for the calcium model. Since our research focused on post-processing of MCMC output, rather than MCMC itself, we directly make available the full output from each sampler on each model considered at

<https://doi.org/10.7910/DVN/MDKNWM>.

This Harvard database download link consists of a single ZIP archive (1.5GB) that contains, for each ODE model and each MCMC method, the states $(x_i)_{i=1}^n$ visited by the Markov chain, their corresponding gradients $\nabla \log p(x_i)$ and the values $p(x_i)$ up to an unknown normalisation constant. The **Stein Thinning** software described in S1 can be used to post-process these datasets at minimal effort, enabling our findings to be reproduced.

S4 Convergence Diagnostics for MCMC

Rigorous approaches for selecting a burn-in period b have been proposed by authors including Meyn and Tweedie (1994); Rosenthal (1995); Roberts and Tweedie (1999); see also Jones and Hobert (2001). Unfortunately, these often involve conditions that are difficult to establish (Biswas et al., 2019, discuss how some of the terms appearing in these conditions can be estimated), or, when they hold, they provide loose bounds, implying an unreasonably long burn-in period.

Proposal	H	G	Details
RW	0	ϵI	Step size ϵ selected following Roberts and Rosenthal (2001)
ADA-RW (Haario et al., 1999)	0	$\sqrt{\hat{\Sigma}}$	$\hat{\Sigma}$ is the sample covariance matrix of preliminary MCMC output
MALA (Roberts and Tweedie, 1996)	$\frac{\epsilon^2}{2} I$	ϵI	Step size ϵ selected following Roberts and Rosenthal (2001)
P-MALA (Girolami and Calderhead, 2011)	$\frac{\epsilon^2}{2} M^{-1}(x_{n-1})$	$\epsilon \sqrt{M^{-1}(x_{n-1})}$	$M(x) = F(x) + \Sigma_0^{-1}$ where $F(x)$ is the Fisher information matrix at x and Σ_0 is the prior covariance matrix.

Table S1: Parameters H and G used in the Metropolis–Hastings proposal.

Convergence diagnostics have emerged as a practical solution to the need to test for non-convergence of MCMC. Their use is limited to reducing bias in MCMC output; they are not optimised for the fixed n setting, which requires a bias-variance trade-off. Nevertheless, convergence diagnostics constitute the principal means by which MCMC output is post-processed. In this section we recall standard practice for selection of a burn-in period b in constructing an estimator of the form (3), focussing on the widely-used diagnostics of Gelman and Rubin (1992); Brooks and Gelman (1998); Gelman et al. (2014) (the *GR diagnostic*), as well as the more recent work of Vats and Knudson (2018) (the *VK diagnostic*).

The GR diagnostic is based on running L independent chains, each of length n , with starting points that are over-dispersed with respect to the target. Obtaining initial points with such characterisation is not trivial because the target is not known beforehand; we refer to the original literature for advice on how to select these initial points, but, in practice, it is not uncommon to guess them. When the support of the target distribution is uni-dimensional (or when $d > 1$, but a specific uni-dimensional summary $f(x)$ is used), the GR diagnostics ($\hat{R}^{\text{GR},L}$) is obtained as the square root of the ratio of two estimators of the variance σ^2 of the target. In particular,

$$\hat{R}^{\text{GR},L} := \sqrt{\frac{\hat{\sigma}^2}{s^2}}, \quad (36)$$

where s^2 is the (arithmetic) mean of the sample variances s_l^2 , $l = 1, \dots, L$, of the chains, which typically provides an underestimate of σ^2 , and $\hat{\sigma}^2$ is constructed as an overestimate of the target variance

$$\hat{\sigma}^2 := \frac{n-1}{n} s^2 + \frac{B}{n},$$

where the term B/n is an estimate of the asymptotic variance of the sample mean of the Markov chain. In the original GR diagnostics, this asymptotic variance was estimated as the sample variance of the means \bar{X}_l , $l = 1, \dots, L$, from the L chains, leading to

$$\frac{B}{n} = \frac{1}{L-1} \sum_{l=1}^L \left(\bar{X}_l - \frac{1}{L} \sum_{l'=1}^L \bar{X}_{l'} \right)^2.$$

The improved VK diagnostic, $\hat{R}^{\text{VK},L}$, is formally obtained in the same way as (36), but with more efficient estimators τ^2/n for the asymptotic variance used in place of B/n . A number of options are available here, but the (lugsail) batch mean estimator of Vats and Flegal (2018) is recommended because it is guaranteed to be biased from above, while still being consistent (in our simulations we use batches of size $\sqrt[3]{n}$). This gain in efficiency leads to improved performance of the VK diagnostic over the GR diagnostic, in the sense that it is less sensitive to the randomness in the Markov chains and the number of chains used. In particular, $\hat{R}^{\text{VK},L}$ can be computed using one chain only ($L = 1$), which has clear practical appeal.

For an ergodic Markov chain, $\hat{R}^{\text{GR},L}$ and $\hat{R}^{\text{VK},L}$ converge to 1 as $n \rightarrow \infty$, so that selection of a suitable burn-in period b amounts to observing when these diagnostics are below $1 + \delta$, where δ is a suitable threshold. In the literature on $\hat{R}^{\text{GR},L}$, the somewhat arbitrary choice $\delta = 0.1$ is commonly used, see Gelman et al. (2014, Ch. 11.5) and the survey in Vats and Knudson (2018). In the literature on $\hat{R}^{\text{KV},L}$, Vats and Knudson (2018) showed how δ can be selected by exploiting the relationship between $\hat{R}^{\text{VK},L}$ and the effective sample size (ESS) when estimating the mean of the target. In particular, it is possible to re-write

$$\hat{R}^{\text{VK},L} = \sqrt{\left(\frac{n-1}{n}\right) + \frac{L}{\widehat{\text{ESS}}}} \quad (37)$$

where $\widehat{\text{ESS}}$ is a strongly consistent estimator of the ESS. One can therefore (approximately) select a δ threshold that corresponds to a pre-specified value of the ESS. The literature on error assessment for MCMC provides guidance on how large the ESS ought to be in order that the width of a $(1 - \alpha)\%$ confidence interval for the mean is less than a specified threshold ϵ ; see Jones and Hobert (2001); Flegal et al. (2008); Vats et al. (2019):

$$\widehat{\text{ESS}} \geq M_{\alpha,\epsilon} := \frac{2^2\pi}{(\Gamma(1/2))^2} \frac{\chi_{1-\alpha}^2}{\epsilon^2}, \quad (38)$$

where $\Gamma(\cdot)$ is the Gamma function, $\chi_{1-\alpha}^2$ is the $(1 - \alpha)^{\text{th}}$ quantile of the χ^2 distribution with one degree of freedom. Plugging (38) in (37) leads to the conclusion that, after the first iteration for which $\hat{R}^{\text{VK},L}$ is below $1 + \delta$, where

$$\delta \equiv \delta(L, \alpha, \epsilon) = \sqrt{1 + \frac{L}{M_{\alpha,\epsilon}}} - 1, \quad (39)$$

the chain will provide an estimate of the mean with small Monte Carlo error, when compared to the variability of the target. The default choices $\alpha = 0.05$ and $\epsilon = 0.05$ were suggested in Vats and Knudson (2018), and were used in our work. For experiments reported in this paper we used (39) to select an appropriate threshold for both $\hat{R}^{\text{GR},L}$ and $\hat{R}^{\text{VK},L}$, which leads to estimated burn-in periods that we denote $\hat{b}^{\text{GR},L}$, and $\hat{b}^{\text{VK},L}$, respectively, in the main text.

The above discussion focussed on the univariate case, but generalisations of these convergence diagnostics are available and can be found in Brooks and Gelman (1998) and Vats

and Knudson (2018). All convergence diagnostics in this work were computed using the R packages `coda` (Plummer et al., 2006) and `stableGR` (Knudson and Vats, 2020)¹.

S5 Empirical Assessment: Additional Results

This section first explores the effect of the choice of kernel, then collects together additional empirical results that accompany the assessment in Section 4.

S5.1 Choice of Kernel

This section concerns the selection of the kernel parameter $\beta \in (-1, 0)$, which enters into the kernel $k(x, y) = (c^2 + \|\Gamma^{-1/2}(x - y)\|^2)^\beta$ used within **Stein Thinning**. The experiment of Figure 1 was repeated for different values of $\beta \in \{-0.1, -0.5, -0.9\}$, and for each setting of the Γ preconditioner matrix (`med`, `sclmed`, `smpcov`), with the results displayed in Figure S1. It can be seen that the states selected when $\beta = -0.5$ are qualitatively reasonable. However, poor results were observed for $\beta \in \{-0.1, -0.9\}$. At $\beta = -0.1$, states were selected that were closer together compared to what would intuitively have been expected. At $\beta = -0.9$, stratification of selected states across the two components of the target was not achieved. These results reflect that $\beta \in \{-0.1, -0.9\}$ are edge cases for KSD, which has been shown to enjoy convergence control only for $\beta \in (-1, 0)$ (see Theorem 4 in Chen et al., 2019). These results support the use of $\beta = -0.5$ as a default in **Stein Thinning**.

S5.2 Goodwin Oscillator

The Goodwin oscillator is a phenomenological model for genetic regulatory processes in a cell and is described by g coupled ODEs of the form

$$\begin{aligned}\frac{du_1}{dt} &= \frac{a_1}{1 + a_2 u_g^\rho} - \alpha u_1, \\ \frac{du_2}{dt} &= k_1 u_1 - \alpha u_2, \\ &\vdots \\ \frac{du_g}{dt} &= k_{g-1} u_{g-1} - \alpha u_g,\end{aligned}$$

where the first component u_1 represents the concentration of mRNA, u_2 that of its corresponding protein product, while u_3, \dots, u_g represent concentrations of proteins in a signalling cascade, that can either be present ($g > 2$) or absent ($g = 2$), and with the g^{th} protein having a negative feedback on the production of mRNA, by means the Hill curve in the

¹The GR diagnostic in the software package uses the original definition in Gelman and Rubin (1992), that differs slightly from (36); however, this difference is not expected to strongly affect the simulation results that we present.




Figure S1: Exploring the role of β in **Stein Thinning**. [The experiment in Figure 1 was reproduced, varying the parameter β in the kernel. The results reported in Figure 1 correspond to $\beta = -0.5$.]

first equation. The nontrivial oscillations that result have led to the Goodwin oscillator being used in previous studies to assess the performance of Bayesian computational methods (Calderhead and Girolami, 2009; Oates et al., 2016; Chen et al., 2019). The parameters $a_1, k_1, \dots, k_{g-1} > 0$ represent synthesis rates and $a_2, \alpha > 0$ representing degradation rates. To cast this model in the setting of Section 2 we set $x \in \mathbb{R}^{g+2}$ to be the vector whose entries are $\log(a_1)$, $\log(k_1)$, and so forth, so that we have a $d = g + 2$ dimensional parameter for which inference is performed.

The experiment that we report considers synthetic data $y_i \in \mathbb{R}^g$ generated in the simple case $g = 2$, which are then corrupted by Gaussian noise such that the terms ϕ_i in (34) are equal to

$$\phi_i(u(t_i)) \propto \exp \left(-\frac{1}{2} (y_i - u(t_i))^\top C^{-1} (y_i - u(t_i)) \right) \quad (40)$$

with $C = \text{diag}(0.1^2, 0.05^2)$. The initial condition was $u(0) = (0, 0)$ and the data-generating parameters were $(a_1, a_2, \alpha, k_1) = (1, 3, 0.5, 1)$. The times t_i , $i = 1, \dots, 2400$, at which data were obtained were taken to be uniformly spaced on $[1, 25]$, in order to capture both the oscillatory behaviour of the system and its steady state. This relatively high frequency of observation and corresponding informativeness of the dataset was used to pre-empt a similarly high frequency observation process in the calcium signalling model of Section 4.3. Figure S2 displays the dataset. A standard Gaussian prior $\pi(x)$ was placed on the parameter x and each MCMC method was applied to approximately sample from the posterior P .

Exemplar trace plots for the MCMC methods are presented in Figure S3. The over-dispersed initial states used for the L chains are reported in Table S2, while the univariate and multivariate convergence diagnostics, computed every 1000 iterations, are shown respectively in Figure S4 and Figure S5. The values of the thresholds $\delta(L, \alpha, \epsilon)$ are reported in Table S3. For each MCMC method, the estimated burn-in period is presented in Table S4. The GR diagnostic did not fall below the $1 + \delta$ threshold in the allowed number of iterations, which is consistent with the empirical observations of Vats and Knudson (2018).

The additional results for the Goodwin oscillator that we present in this appendix are as follows:

- Figures 3 (RW), S7 (MALA) and S8 (P-MALA) display point sets of size $m = 20$ selected using traditional burn in and thinning methods, **Support Points** and **Stein Thinning**, based on MCMC output. Note that the gray regions are not *necessarily* regions of high posterior probability; they are the regions explored by the sample path and, moreover, these panels are two-dimensional projections from \mathbb{R}^4 . Therefore we are hesitant to draw strong conclusions from these figures.
- Figures 4 (RW), S10 (MALA) and S11 (P-MALA) display the absolute error in estimating the first moment of each parameter, for each of the competing methods, where an extended run from MCMC provided the ground truth.
- Figures S12 (RW), S13 (ADA-RW), S14 (MALA) and S15 (P-MALA) show marginal density estimates, for each parameter and each of the competing methods, where an extended run from MCMC provided the ground truth.




Figure S2: Data (gray) and ODE solution corresponding to the true data-generating parameters (black) for the Goodwin oscillator.

Chain Number	Initial State for Parameters (a_1, a_2, α, k)
1	(0.5, 1, 3, 2)
2	(0.001, 0.2, 0.1, 10)
3	(10, 0.1, 0.9, 0.1)
4	(0.1, 30, 0.1, 0.3)
5	(2, 2, 2, 2)
6	(5, 5, 1, 1)

Table S2: Initial states, over-dispersed with respect to the posterior, for the $L = 6$ independent Markov chains used in the Goodwin oscillator. The parameters used to generate the data were $(a_1, a_2, \alpha, k) = (1, 3, 0.5, 1)$.

	Univariate Diagnostics	Multivariate Diagnostics
$L = 6$	4.88×10^{-4}	3.56×10^{-4}
$L = 5$	4.07×10^{-4}	2.96×10^{-4}
$L = 1$	8.13×10^{-5}	5.93×10^{-5}

Table S3: The values of the threshold $\delta(L, \alpha, \epsilon)$ used in analysis of the Goodwin and Lotka–Volterra models, with $\alpha = 0.05$, $\epsilon = 0.05$, when changing L , and considering the univariate and multivariate convergence diagnostics $\hat{R}^{\text{VK}, L}$.

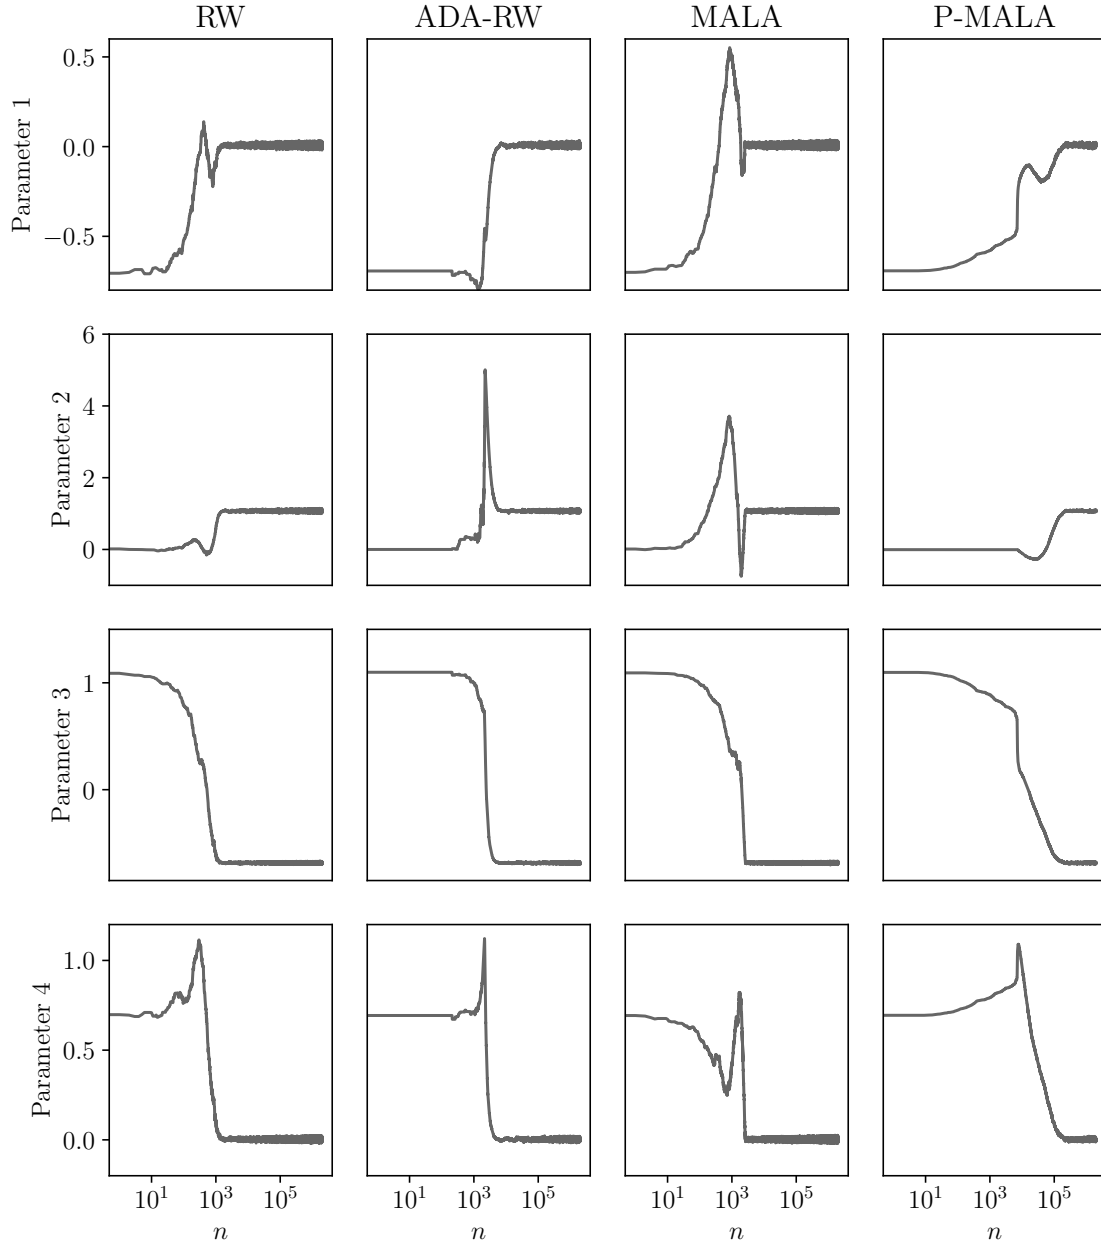


Figure S3: Trace plots for each parameter in the Goodwin oscillator, plotted against the MCMC iteration number. Each row corresponds to one of the four parameters, while each column corresponds to one of the four MCMC methods considered. (Note the logarithmic scale on the horizontal axis, used to better visualise the initial part of the MCMC sample path.)

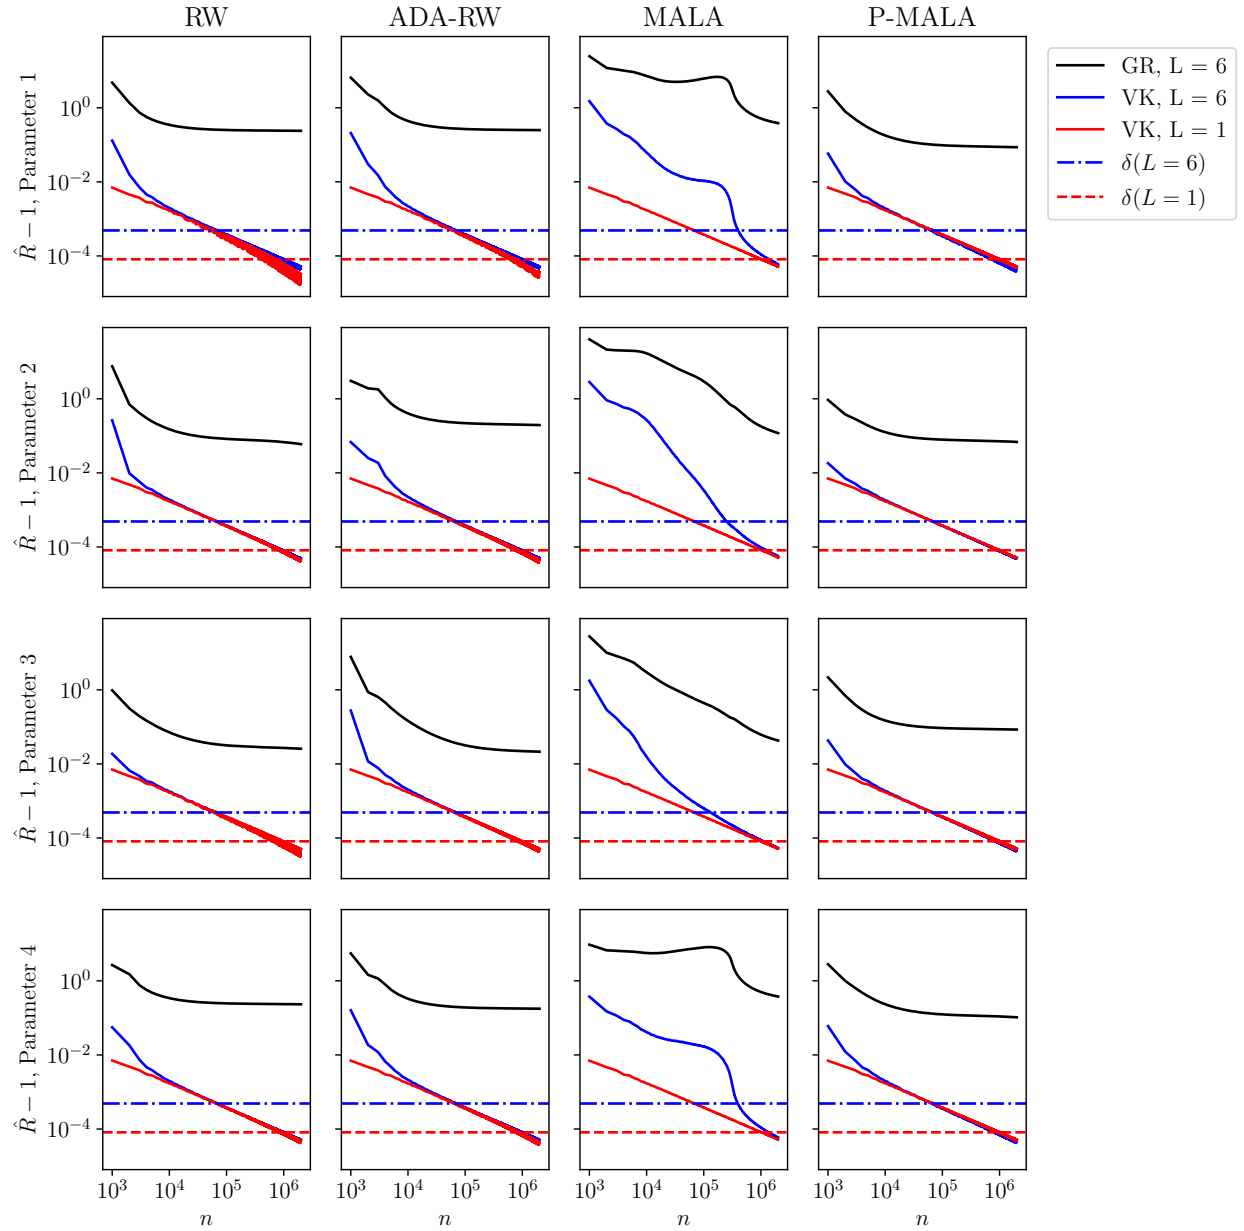


Figure S4: Univariate convergence diagnostics, for the Goodwin oscillator, plotted against the MCMC iteration number. The black line represents the GR diagnostic (based on $L = 6$ chains), while the blue and red lines represent the VK diagnostic (based on $L = 6$ and $L = 1$ chains, respectively). The dash-dotted ($L = 6$) and dashed ($L = 1$) horizontal lines correspond to the critical values $\delta(L, \alpha, \epsilon)$, used to determine the burn-in period; see Table S3.




Figure S5: Multivariate convergence diagnostics for the Goodwin oscillator, plotted against the MCMC iteration number n . The black line is the GR diagnostic (based on $L = 6$ chains), while the blue and red lines are the VK diagnostic (based on $L = 6$ and $L = 1$ chains, respectively). The dotted ($L = 6$) and dashed ($L = 1$) horizontal lines correspond to the threshold δ that is used to determine the burn-in period; see Table S3 in Appendix S5.2.

MCMC Diagnostics	Sampler	$\hat{b}^{\text{GR},6}$	$\hat{b}^{\text{VK},6}$	$\hat{b}^{\text{VK},1}$
Univariate	RW	$> n$	70,000	820,000
	ADA-RW	$> n$	71,000	816,000
	MALA	$> n$	397,000	1,020,000
	P-MALA	$> n$	68,000	987,000
Multivariate	RW	$> n$	93,000	578,000
	ADA-RW	$> n$	107,000	824,000
	MALA	$> n$	316,000	1,615,000
	P-MALA	$> n$	103,000	1,475,000

Table S4: Estimated burn-in period for the Goodwin oscillator, using the GR diagnostic based on L chains, $\hat{b}^{\text{GR},L}$ ($L = 6$), and the VK diagnostic based on L chains, $\hat{b}^{\text{VK},L}$, ($L = 1, 6$). In each case both univariate and multivariate convergence diagnostics are presented; in the univariate case we report the largest value obtained when looking at each of the d parameters individually to estimate the burn-in period. The symbol “ $> n$ ” indicates the case in which a diagnostic did not go below the $1 + \delta$ threshold.

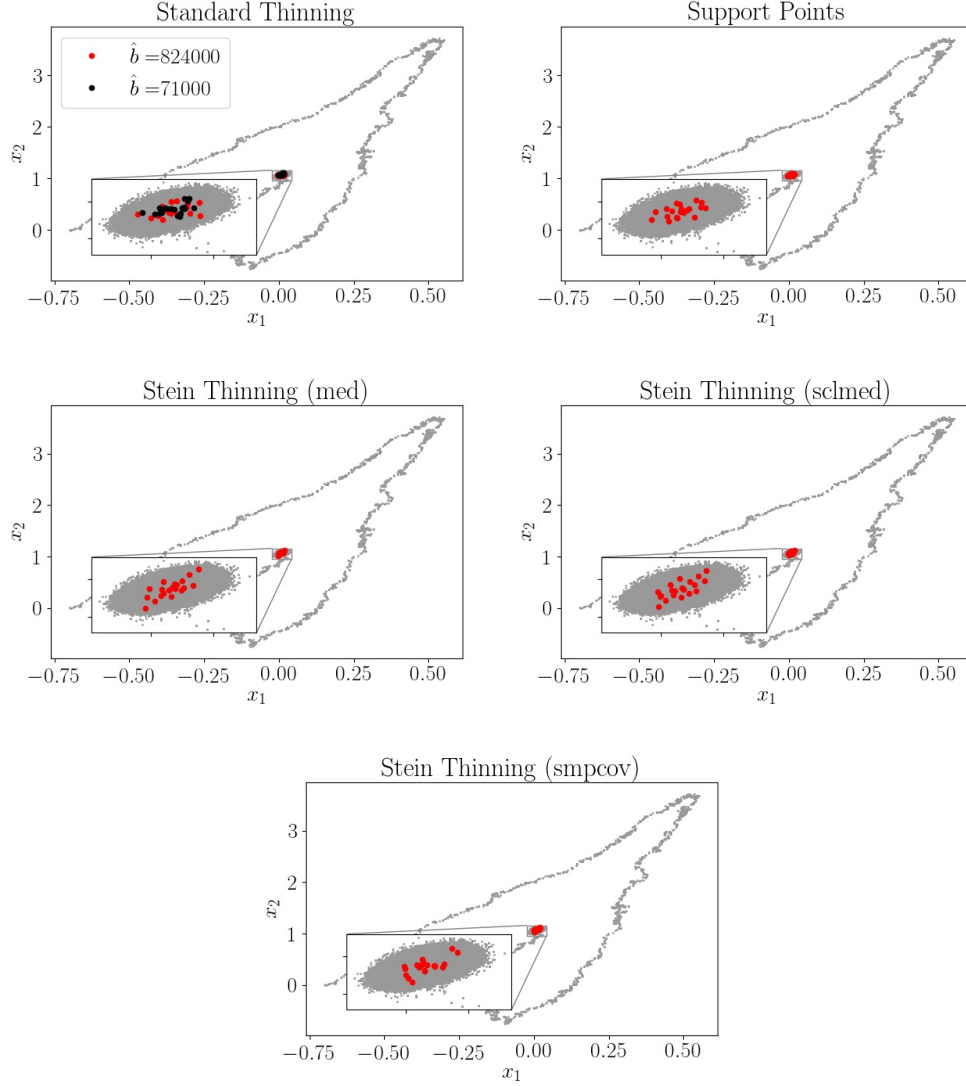


Figure S6: Projections on the first two coordinates of the **ADA-RW** MCMC output for the Goodwin oscillator (grey dots), together with the first $m = 20$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

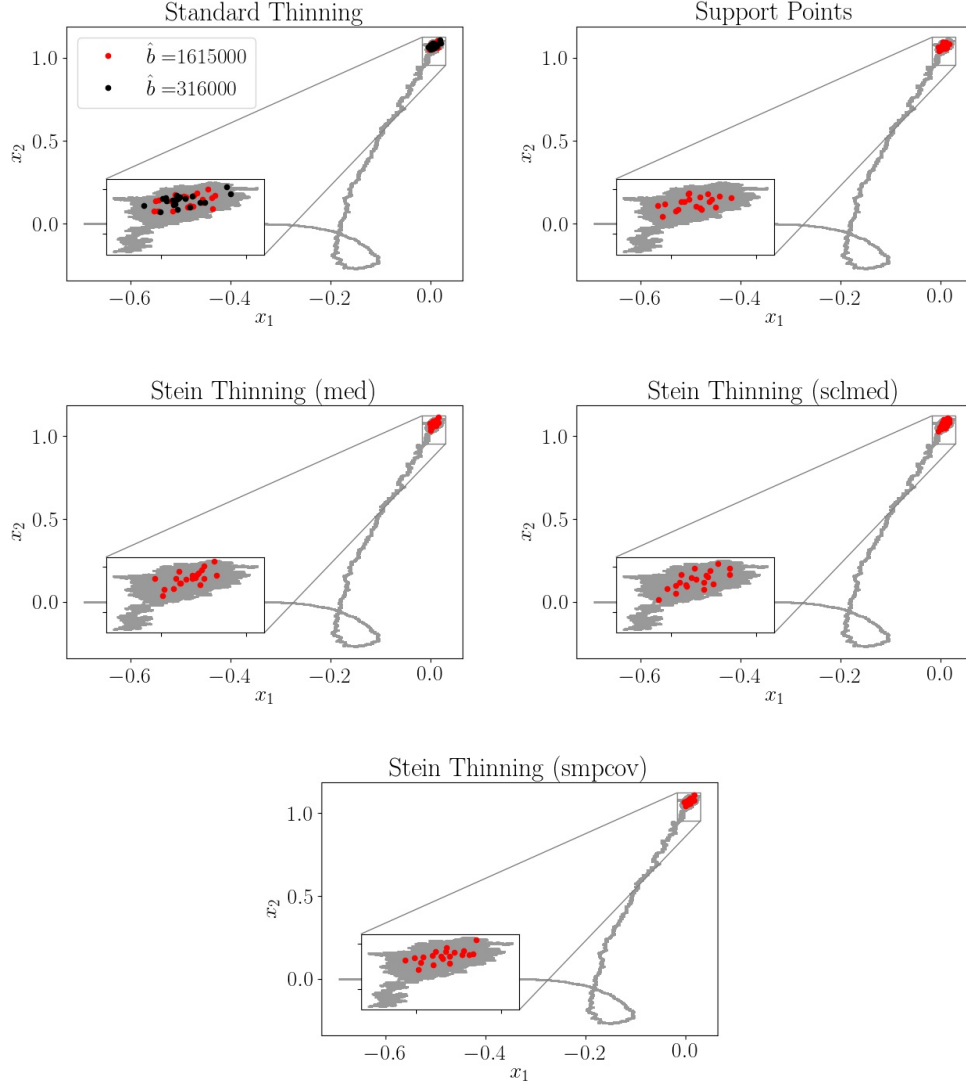


Figure S7: Projections on the first two coordinates of the **MALA** MCMC output for the Goodwin oscillator (grey dots), together with the first $m = 20$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

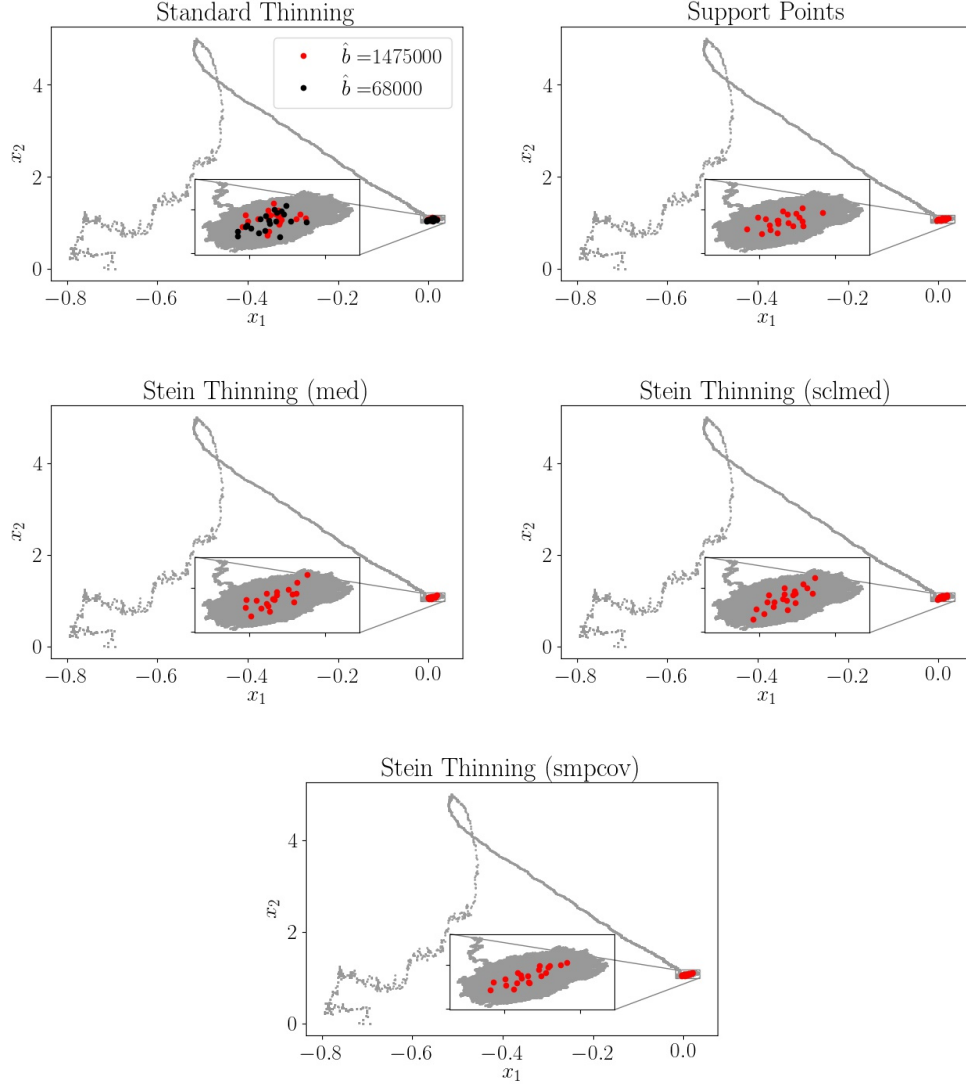


Figure S8: Projections on the first two coordinates of the P-MALA MCMC output for the Goodwin oscillator (grey dots), together with the first $m = 20$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

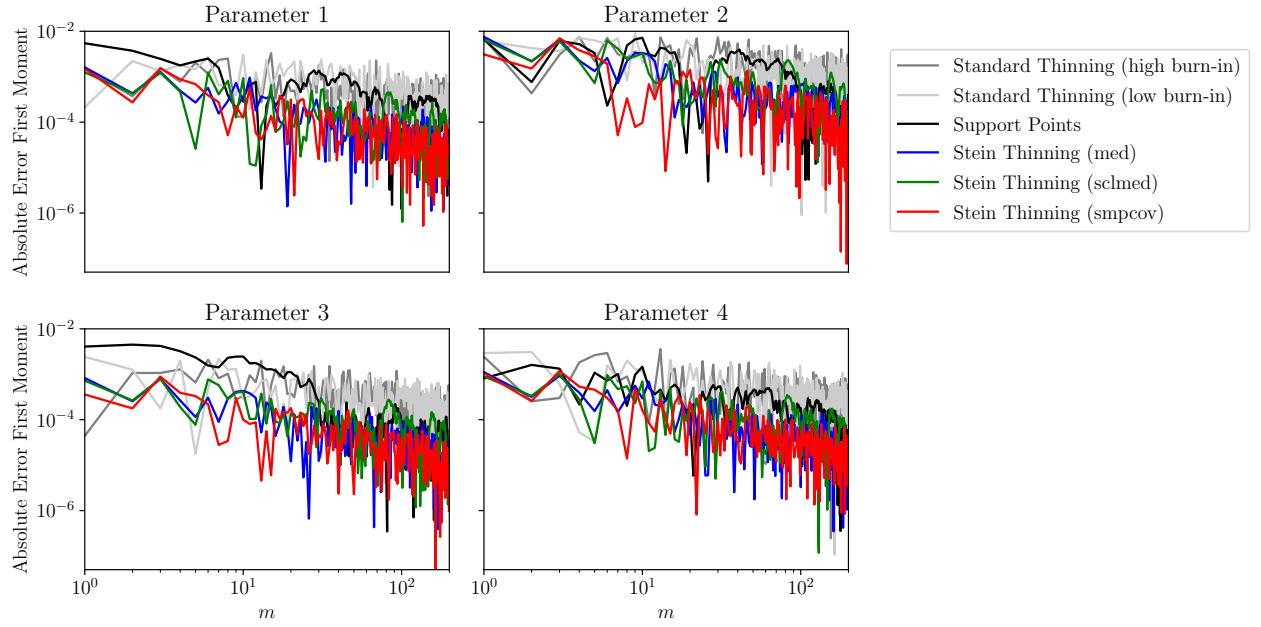


Figure S9: Absolute error of estimates for the posterior mean of each parameter in the Goodwin oscillator, based on output from ADA-RW MCMC.

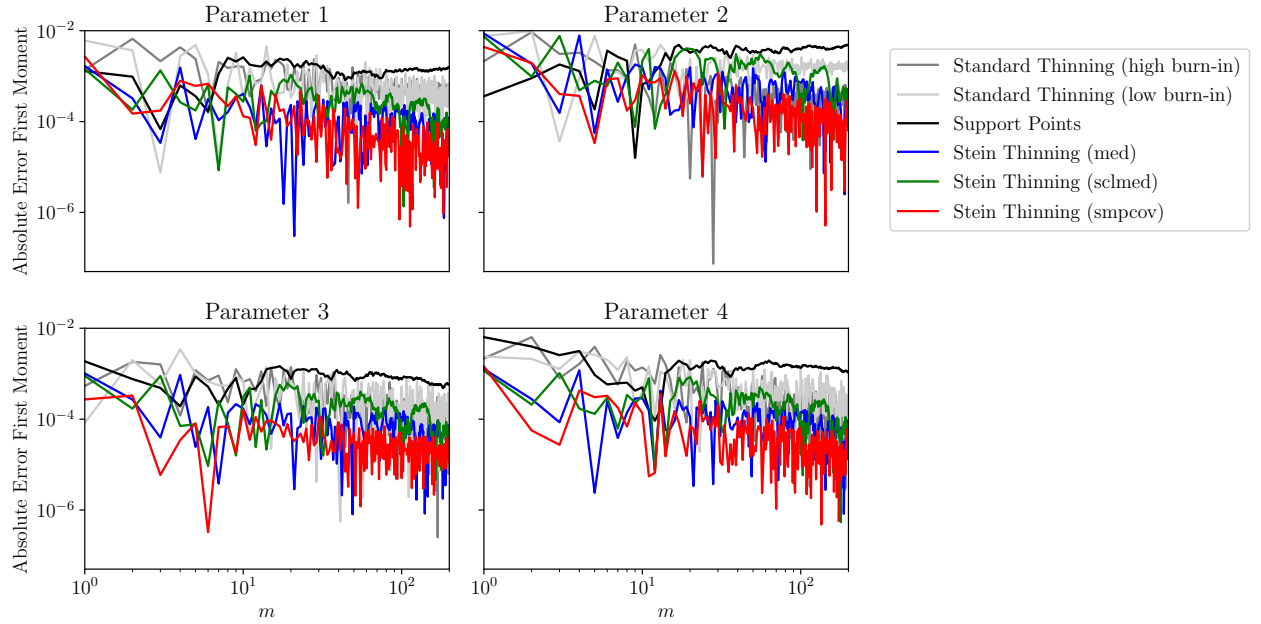


Figure S10: Absolute error of estimates for the posterior mean of each parameter in the Goodwin oscillator, based on output from MALA MCMC.

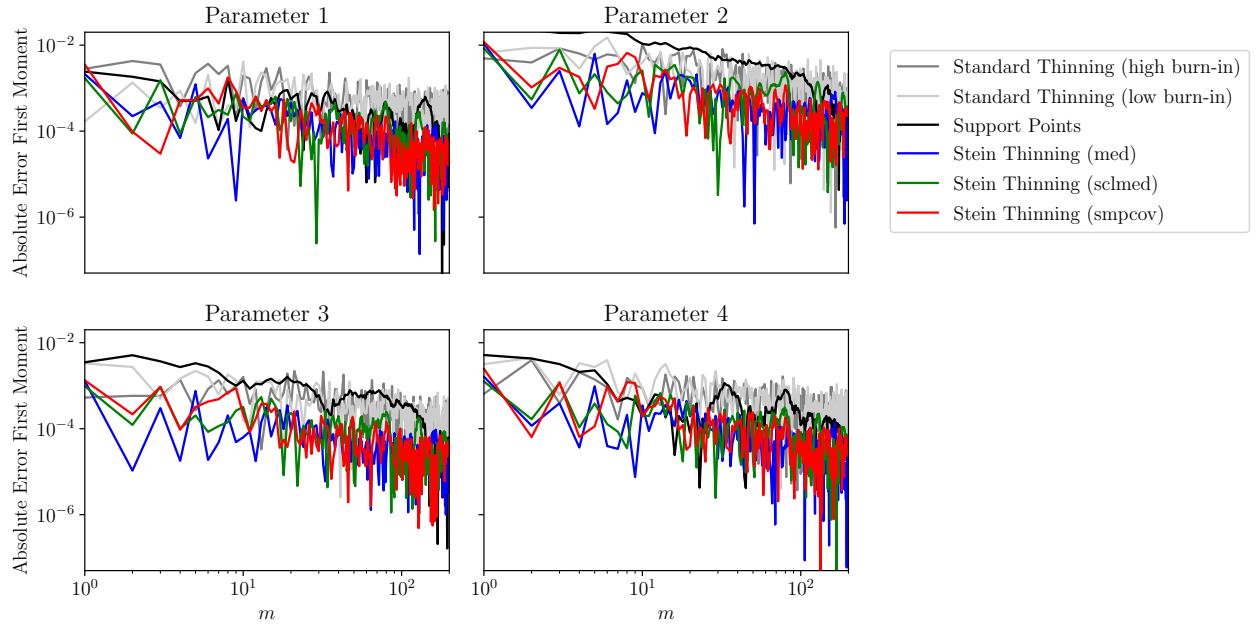


Figure S11: Absolute error of estimates for the posterior mean of each parameter in the Goodwin oscillator, based on output from P-MALA MCMC.

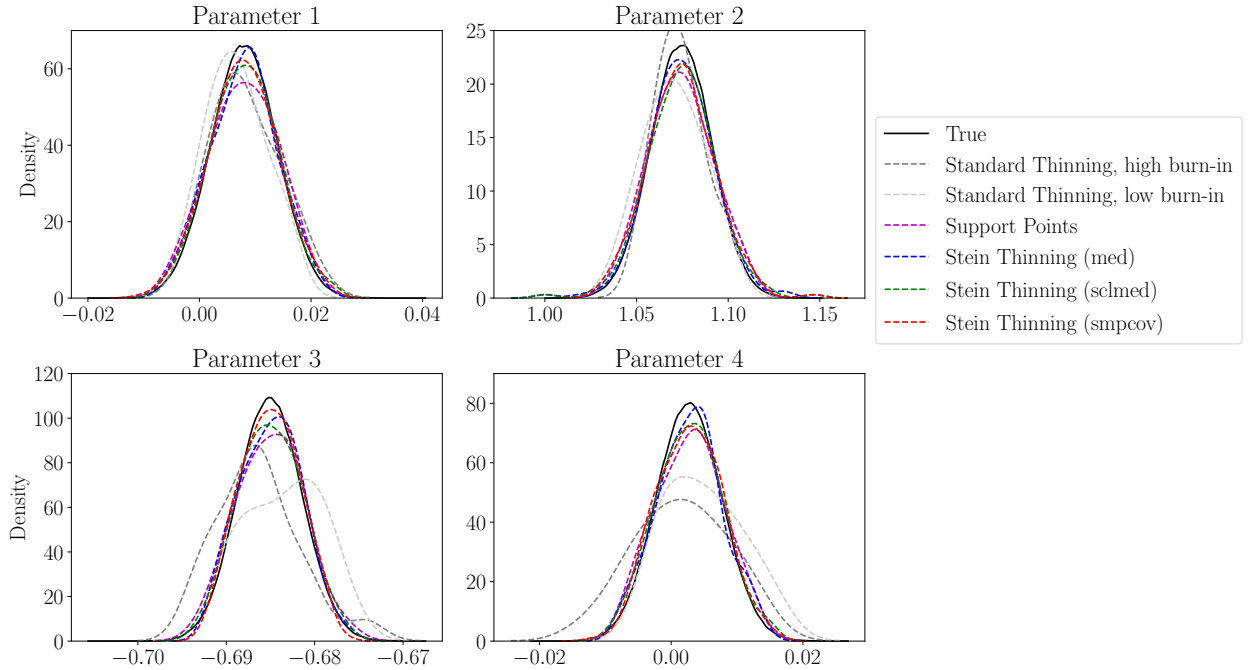


Figure S12: True and estimated marginal densities of the parameters in the Goodwin oscillator, using $m = 20$ points selected from RW MCMC output.

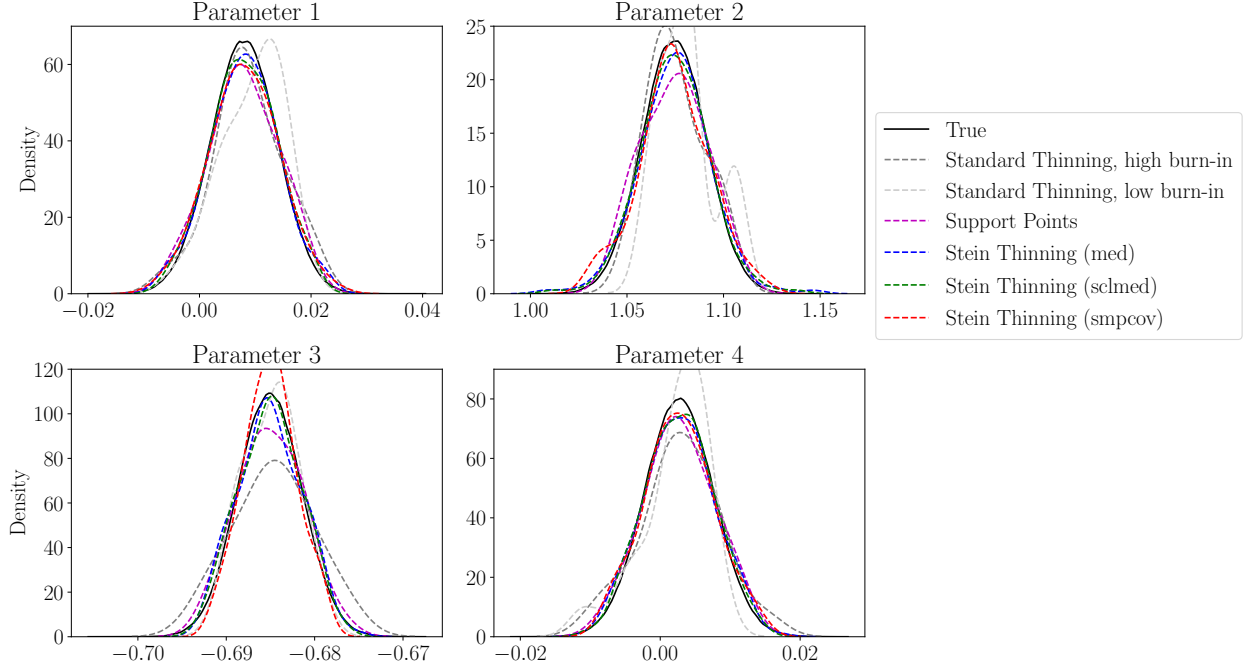


Figure S13: True and estimated marginal densities of the parameters in the Goodwin oscillator, using $m = 20$ points, selected from ADA-RW MCMC output.

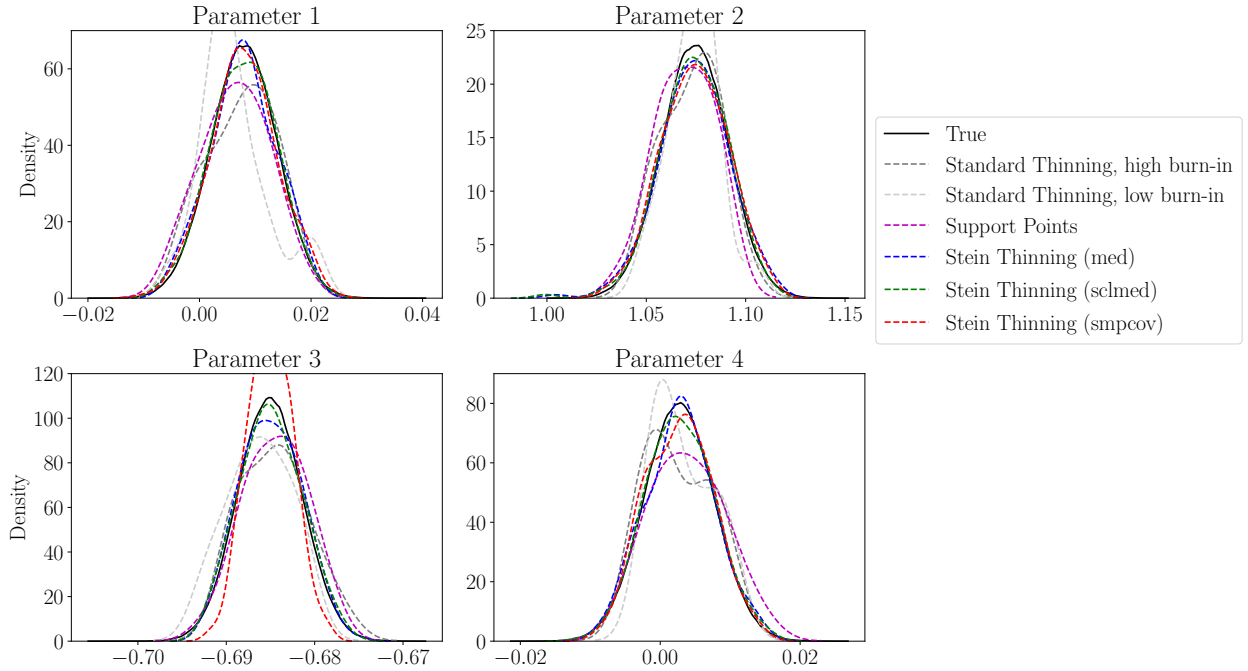


Figure S14: True and estimated marginal densities of the parameters in the Goodwin oscillator, using $m = 20$ points, selected from MALA MCMC output.

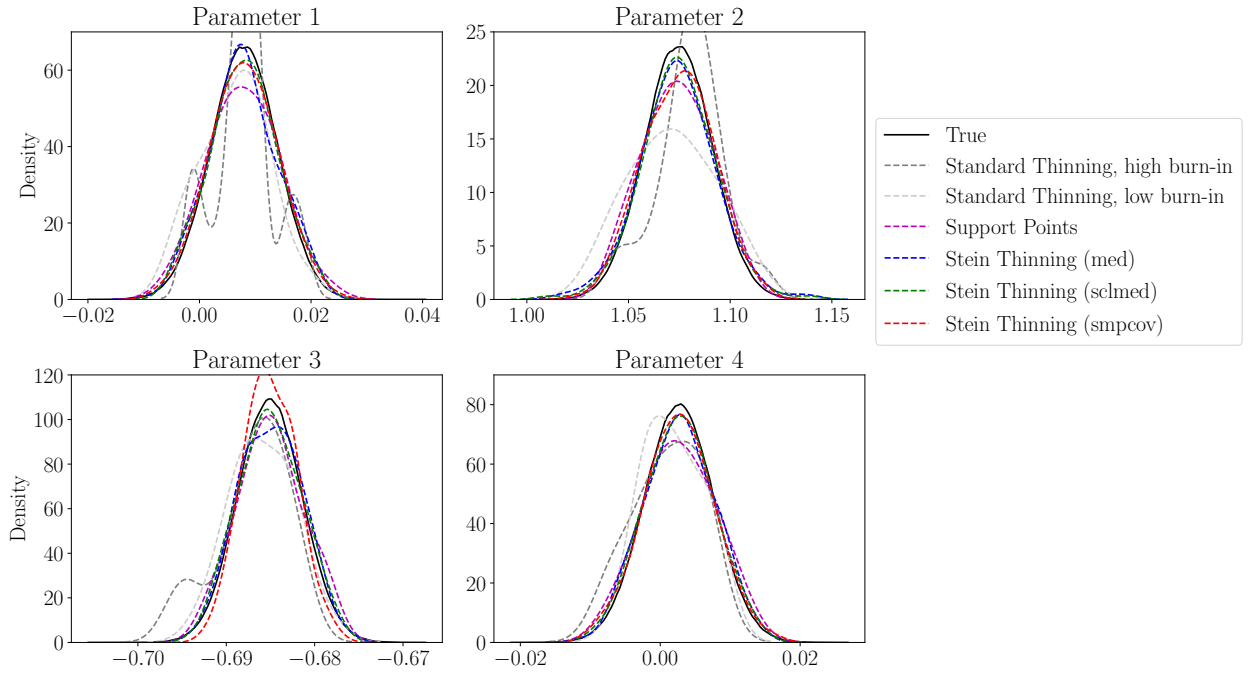


Figure S15: True and estimated marginal densities of the parameters in the Goodwin oscillator, using $m = 20$ points, selected from P-MALA MCMC output.

S5.3 Lotka–Volterra

The Lotka–Volterra model describes the oscillatory evolution of prey (u_1) and predator (u_2) species in a closed environment. The prey has an intrinsic mechanism for growth proportional to its abundance, described by a parameter $\theta_1 > 0$, whilst interaction with the predator leads to a decrease in the prey population at a rate described by a parameter $\theta_2 > 0$. Conversely, the predator has an intrinsic mechanism for decline proportional to its abundance, described by a parameter $\theta_3 > 0$, whilst interaction with the prey leads to an increase in the predator population at a rate described by a parameter $\theta_4 > 0$. The resulting system of ODEs is:

$$\begin{aligned}\frac{du_1}{dt} &= \theta_1 u_1 - \theta_2 u_1 u_2, \\ \frac{du_2}{dt} &= \theta_4 u_1 u_2 - \theta_3 u_2.\end{aligned}$$

To cast this model in the setting of Section 2 we set $x \in \mathbb{R}^4$ to be the vector whose entries are $\log(\theta_1), \dots, \log(\theta_4)$, so that we have a $d = 4$ dimensional parameter for which inference is performed.

The experiment that we report considers synthetic data which are corrupted by Gaussian noise such that the terms ϕ_i in (34) have expression (40), with $C = \text{diag}(0.2^2, 0.2^2)$. The initial condition was $u(0) = (1, 1)$ and the data-generating parameters were $x = \log(\theta)$, with $\theta = (0.67, 1.33, 1, 1)$. The times t_i , $i = 1, \dots, 2400$, at which data were obtained were taken to be uniformly spaced on $[0, 25]$. Figure S16 displays the dataset. A standard Gaussian prior $\pi(x)$ was used.

Exemplar trace plots for the MCMC methods are presented in Figure S17. The over-dispersed initial states used for the L chains are the same as those reported in Table S2, while the univariate and multivariate convergence diagnostics, computed every 1000 iterations, are shown respectively in Figure S18 and Figure S19. The values of the thresholds $\delta(L, \alpha, \epsilon)$ are the same as those reported in Table S3. For each MCMC method, the estimated burn-in period is presented in Table S6.

The additional results for the Lotka–Volterra model that we present in this appendix are as follows:

- Figures S20 (RW), S21 (ADA-RW), S22 (MALA) and S23 (P-MALA) display point sets of size $m = 20$ selected using traditional burn in and thinning methods, **Support Points** and **Stein Thinning**, based on MCMC output.




Figure S16: Data (gray) and ODE solution corresponding to the true data-generating parameters (black) for the Lotka–Volterra model.

Chain Number	Initial State for Parameters (RW, MALA, P-MALA)	Initial State for Parameters (ADA-RW)
1	(0.55, 1, 0.8, 0.8)	(0.55, 1, 0.8, 0.8)
2	(1.5, 1, 0.8, 0.8)	(0.55, 1, 0.8, 1.3)
3	(1.3, 1.33, 0.5, 0.8)	(1.3, 1.33, 0.5, 0.8)
4	(0.55, 3, 3, 0.8)	(0.55, 1, 1.5, 1.5)
5	(0.55, 1, 1.5, 1.5)	(0.55, 1.3, 1, 0.8)

Table S5: Initial states, $\theta = \exp(x)$, over-dispersed with respect to the posterior, for the $L = 5$ independent Markov chains used in the Lotka–Volterra model. The parameters used to generate the data were $\theta = (0.67, 1.33, 1, 1)$.

MCMC Diagnostics	Sampler	$\hat{b}^{\text{GR},5}$	$\hat{b}^{\text{VK},5}$	$\hat{b}^{\text{VK},1}$
Univariate	RW	$> n$	88,000	954,000
	ADA-RW	$> n$	84,000	764,000
	MALA	$> n$	424,000	995,000
	P-MALA	$> n$	90,000	820,000
Multivariate	RW	$> n$	119,00	1,512,000
	ADA-RW	1,797,000	99,000	743,000
	MALA	$> n$	259,000	1,573,000
	P-MALA	$> n$	114,000	1,251,000

Table S6: Estimated burn-in period for the Lotka–Volterra model, using the GR diagnostic based on L chains, $\hat{b}^{\text{GR},L}$ ($L = 5$), and the VK diagnostic based on L chains, $\hat{b}^{\text{VK},L}$, ($L = 1, 5$). In each case both univariate and multivariate convergence diagnostics are presented; in the univariate case we report the largest value obtained when looking at each of the d parameters individually to estimate the burn-in period. The symbol “ $> n$ ” indicates the case in which a diagnostic did not go below the $1 + \delta$ threshold.

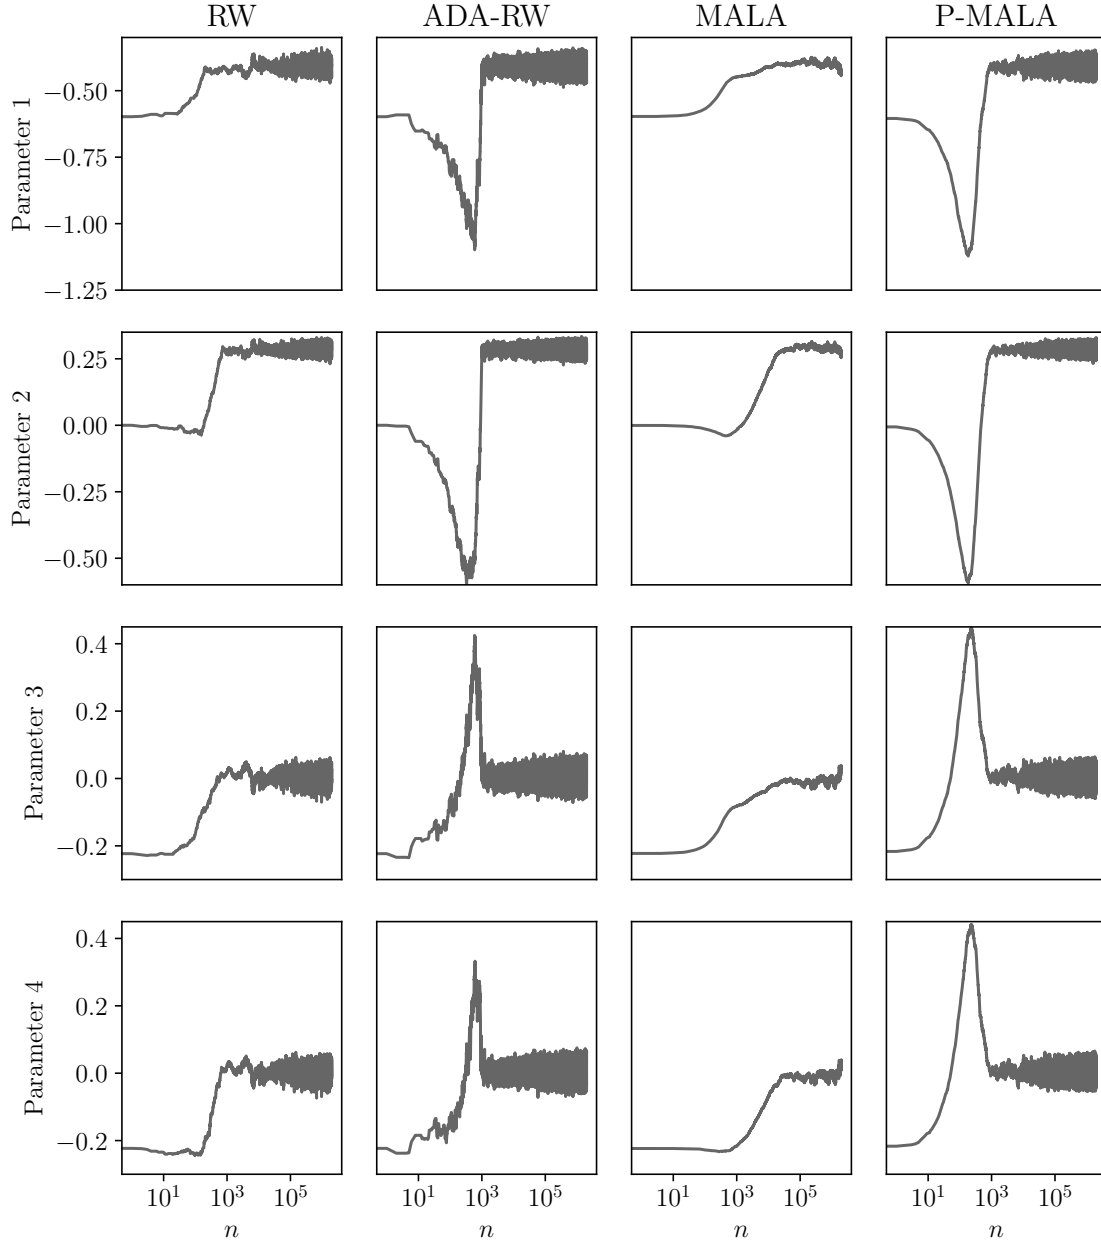


Figure S17: Trace plots for the parameters x_i in the Lotka–Volterra model, plotted against the MCMC iteration number. Each row corresponds to one of the four parameters, while each column corresponds to one of the four MCMC methods considered.




Figure S18: Univariate convergence diagnostics, for the Lotka–Volterra model, plotted against the MCMC iteration number. The black line represents the GR diagnostic (based on $L = 5$ chains), while the blue and red lines represent the VK diagnostic (based on $L = 5$ and $L = 1$ chains, respectively). The dash-dotted ($L = 5$) and dashed ($L = 1$) horizontal lines correspond to the critical values $\delta(L, \alpha, \epsilon)$, used to determine the burn-in period; see Table S3.




Figure S19: Multivariate convergence diagnostics for Lotka–Volterra, plotted against the MCMC iteration number. The black line is the GR diagnostic (based on $L = 5$ chains), while the blue and red lines are the VK diagnostic (based on $L = 5$ and $L = 1$ chains, respectively). The dotted ($L = 5$) and dashed ($L = 1$) horizontal lines correspond to the critical values $\delta(L, \alpha, \epsilon)$, used to determine the burn-in period; see Table S3.

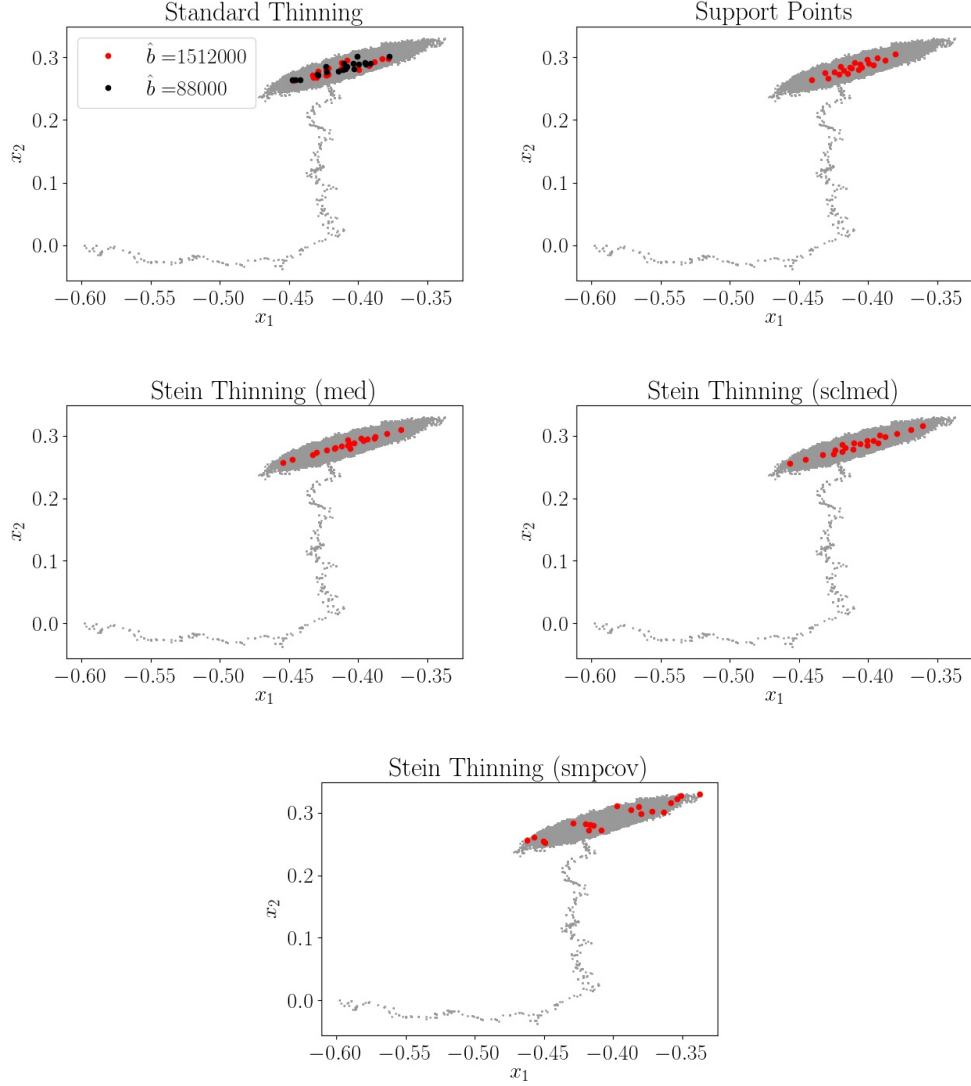


Figure S20: Projections on the first two coordinates of the RW MCMC output for the Lotka–Volterra model (grey dots), together with the first $m = 4$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

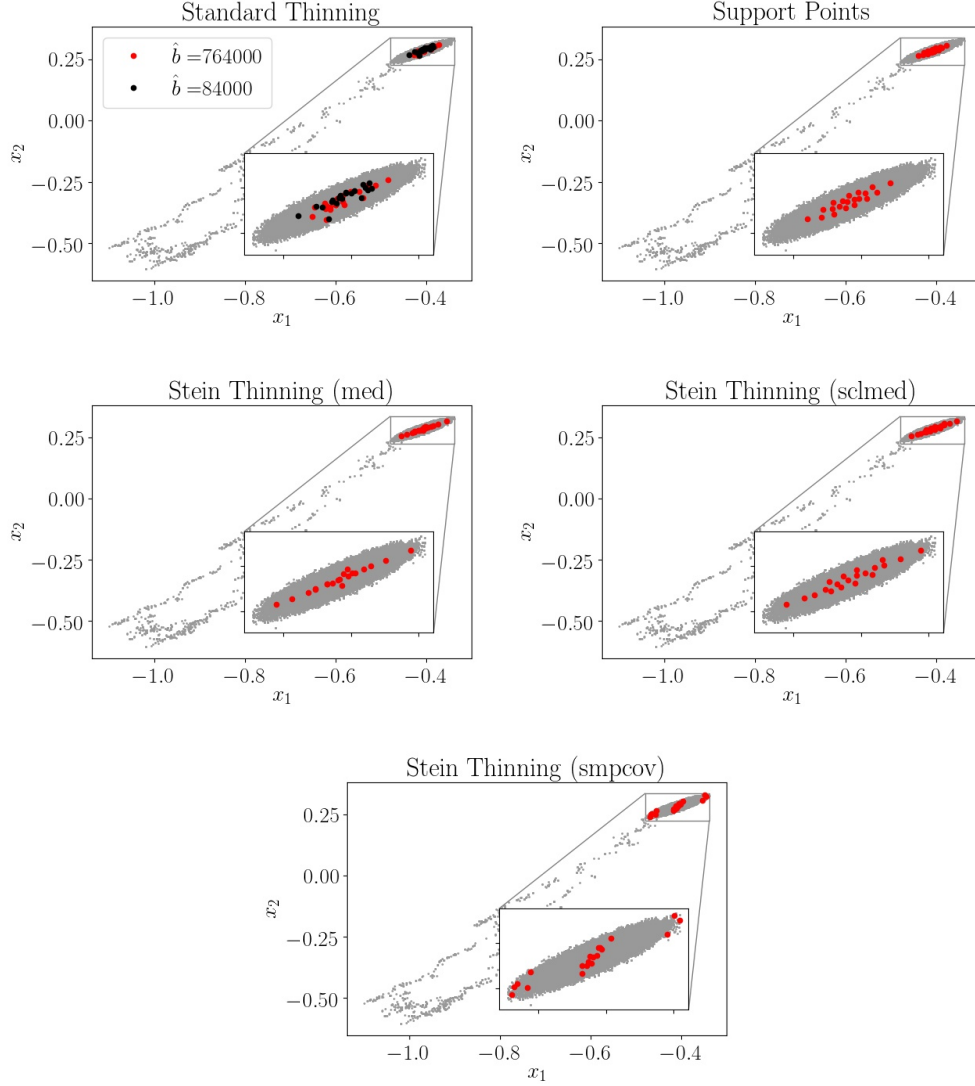


Figure S21: Projections on the first two coordinates of the **ADA-RW** MCMC output for the Lotka–Volterra model (grey dots), together with the first $m = 4$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

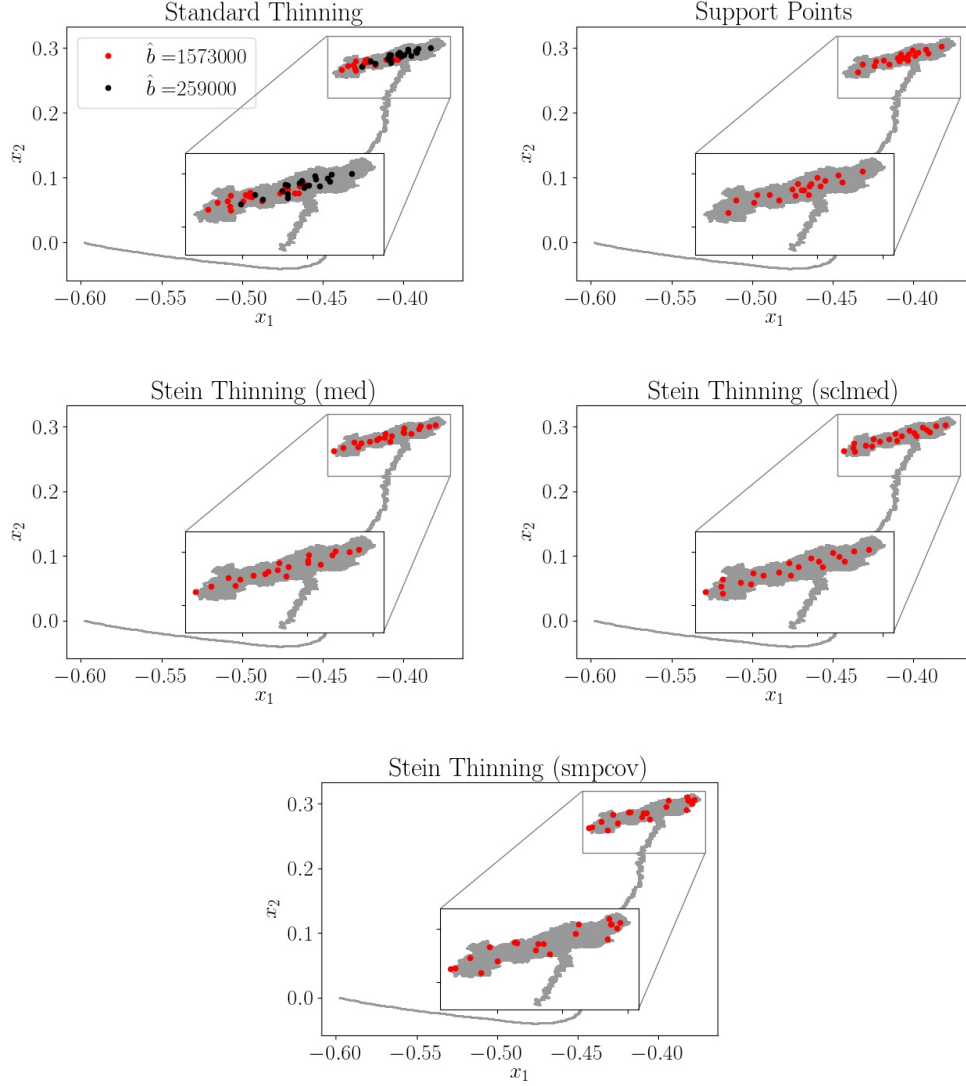


Figure S22: Projections on the first two coordinates of the **MALA** MCMC output for the Lotka–Volterra model (grey dots), together with the first $m = 4$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

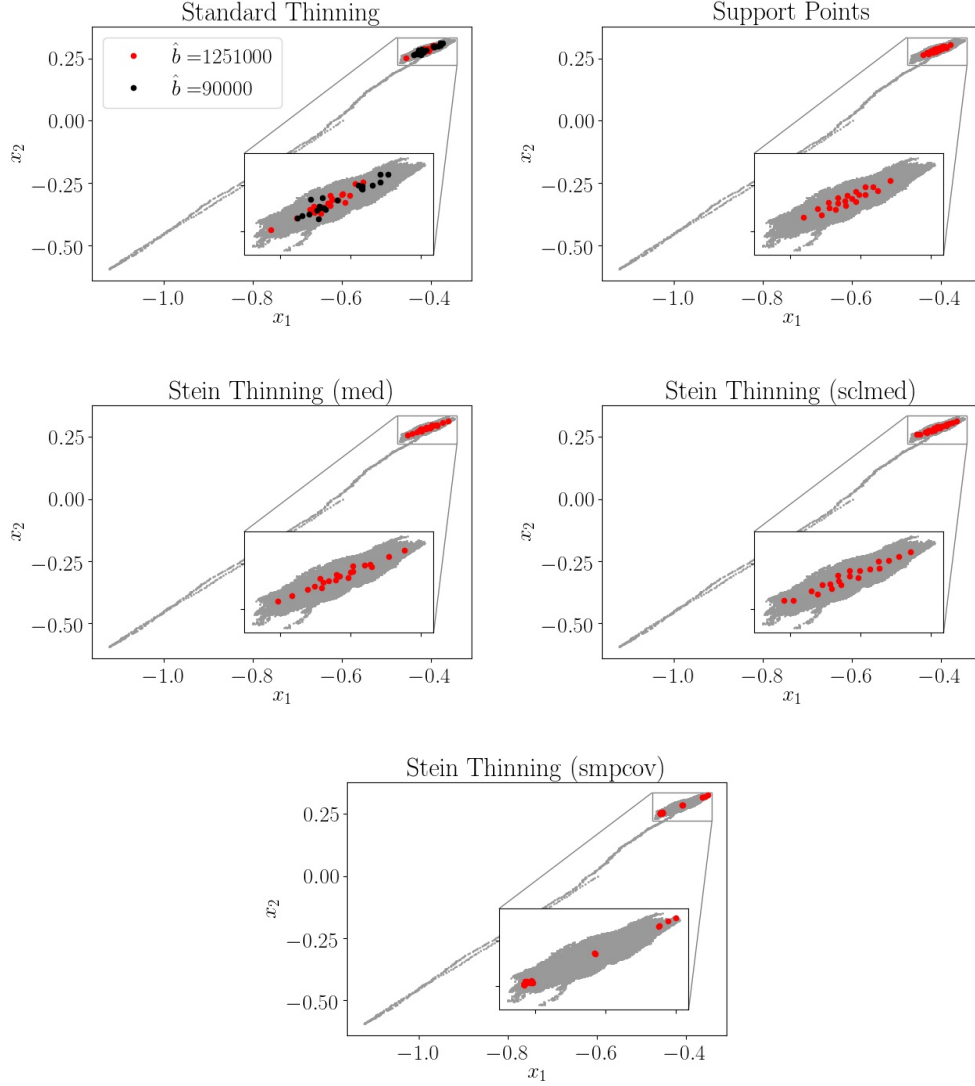


Figure S23: Projections on the first two coordinates of the P-MALA MCMC output for the Lotka–Volterra model (grey dots), together with the first $m = 4$ points selected through: traditional burn-in and thinning (the amount of burn in is indicated in the legend); the **Support Points** method; **Stein Thinning**, for each of the settings **med**, **sclmed**, **smpcov**.

S5.4 Calcium Signalling Model

This appendix contains a detailed biochemical description of the calcium signalling model studied in Section 4.3 of the main text, together with the experimental dataset that we collected.

The Hinch et al. (2004) single cell model simulates the calcium transient evoked by membrane depolarisation in a cardiac cell. The model has a mathematical representation of the extracellular space and the intracellular compartment consisting of the sarcoplasmic reticulum (SR), dyadic space and cytosol. The major sarcolemmal calcium pathways are included: the L-type Ca channel (LCC), the plasmalemmal membrane calcium ATPase (PMCA) and the sodium-calcium exchanger (NCX). Inside the cell, the model has mathematical representations for calcium release from the SR to dyadic space through ryanodine receptors (RyR) and re-sequestration of calcium from the dyadic space into the SR by the SR ATPase (SERCA). Calcium buffering is also featured for the cytosol. A schematic representation of the cell model is given in Figure S24.

Membrane depolarisation is triggered by an electrical event. This causes calcium to enter through LCCs into the dyadic space, producing a local rise in Ca concentration, sufficient to activate RyRs. This process engages a feedback, whereby Ca release from the SR causes more RyR opening events. As the released Ca diffuses into the cytosol, most of it becomes buffered, but some ions remain free and underpin the Ca transient. Recovery following Ca release is driven by SERCA, which re-sequesters Ca into the SR, and NCX and PMCA which extrude calcium across the sarcolemma. This returns the cell to its initial conditions, ready for the next electrical stimulation.

The Hinch model describes the nonlinear, time-dependent interaction of the four Ca handling transporters (LCC, PMCA, RyR and SERCA) and lumped buffering by a system of 7 ODEs whose parameter space is $d = 38$ dimensional. The first three differential equations provide a simplified four-state model describing the interaction between LCC and RyR within the dyadic space; here, only three states are simulated due to a conservation of mass constraint. The remaining four differential equations describe: calcium concentration in the sarcoplasmic reticulum and the cytosol, the calcium bound to cytosolic buffers and calcium current across the cellular membrane. Of these state variables, only the concentration of free calcium in the cytosol and the transmembrane current can be experimentally observed.

To provide a rich dataset for characterising calcium dynamics in a single cardiac myocyte, we applied three experimental protocols in sequence on a single myocyte. During these protocols, we controlled membrane potential and measured membrane currents electrophysiologically and, after appropriate calibration, followed Ca fluorimetrically. The calcium handling proteins were interrogated by relating currents and Ca concentration in response to defined membrane potential manoeuvres, and in the presence of drugs to eliminate various confounding components. The first voltage protocol interrogated LCC currents at different voltages, and measured their response in terms of SR release. In the second protocol, a train of depolarisations then triggered Ca transients which provided information about SR release and their recovery provided a readout of SERCA, NCX and PMCA activities. The third protocol consisted of rapid exposure to caffeine which emptied the SR and short-

circuited SERCA. This provided information about SR load, and the subsequent recovery is a readout of NCX and PMCA. Buffering was calculated from the quotient of measured Ca rise upon caffeine exposure and the amount of Ca released back-calculated from sarcolemmal current generated by NCX. The dataset contains 12998 observations of cytosolic free calcium concentration observed at a 60 Hz sampling frequency, and 22260 transmembrane current observations, both for a duration of 3 minutes. The data are displayed in Figure S25, where the different colours show the three parts of the biological protocol explained above. The calcium signalling model in Figure S24 is represented by a coupled system of 7 ODEs and depends on a $d = 38$ dimensional parameter, which is to be estimated based on the experimental dataset. As just described, the data consist of measurements of calcium concentration in the cytoplasm and transmembrane current whilst the cell was externally stimulated, so that only two of the state variables (in our case, u_5 and u_7) were observed (we denote the observations of these states y^5 and y^7 , respectively). Our likelihood took the simple Gaussian form $\phi_i(u(t_i)) \propto \exp(-\frac{1}{2\sigma_5^2}(y_i^5 - u_5(t_i))^2) + \exp(-\frac{1}{2\sigma_7^2}(y_i^7 - u_7(t_i))^2)$ with $\sigma_5 = 2.07 \times 10^{-8}$ and $\sigma_7 = 1.62 \times 10^{-10}$. The ODE was numerically solved using `CVODES` (Hindmarsh et al., 2005) and sensitivities were computed by solving the forward sensitivity equations; see Appendix S3. Further details of the expert-elicited prior, the data pre-processing procedure and numerical details associated with the ODE solver will be reported in a separate manuscript, in preparation as of 12th July 2021, and are available on request.

In the experiments that follow, `RW` MCMC was used both to target the posterior P and to target a tempered distribution Q . The latter is equivalent to multiplying the measurement error standard deviations σ_5 and σ_7 by 8, and has the effect of rendering Q more diffuse than P , in order that Q is more favourable for MCMC. The specific value of 8 corresponded to the smallest amount of tempering required to achieve convergence within the available computational budget. A total of $n = 4 \times 10^6$ iterations were performed, and in each case the first 10^6 iterations were used to adapt the scale of the Gaussian proposal distribution in the `RW` sampler, so that an acceptance rate close to 0.234 (Gelman et al., 1997) was achieved. The first 10^6 iterations were then discarded.

The additional results for the calcium signalling model that we present in this appendix are as follows:

- Figures S26 to S29 contain kernel density estimates for posterior marginals obtained by **Stein Thinning** applied to tempered `RW` MCMC output, versus standard `RW` MCMC output.
- Figures S30 and S30 present results for KSD based on `sclmed` and `smpcov` settings, to complement Figure 9 in the main text.

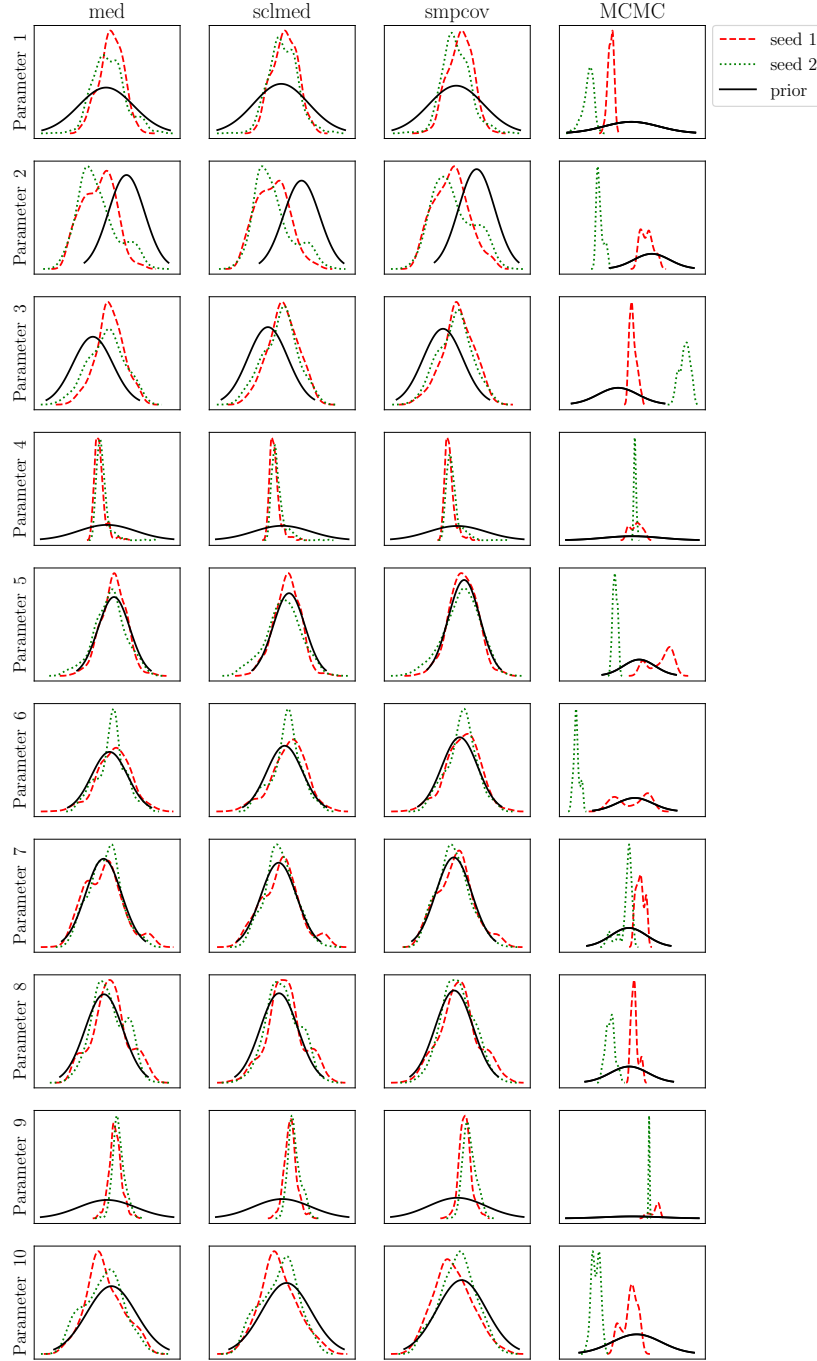


Figure S26: Kernel density estimates for posterior marginals in the calcium signalling model. Stein thinning with **med**, **sclmed** and **smpcov** preconditioners (first three columns) was applied to tempered RW MCMC output (to obtain $m = 500$ points). These can be contrasted with the last column, where kernel density estimates based on standard RW MCMC are displayed. [In each case, two distinct random seeds were used. For reference, the black curve represents the prior marginal.]

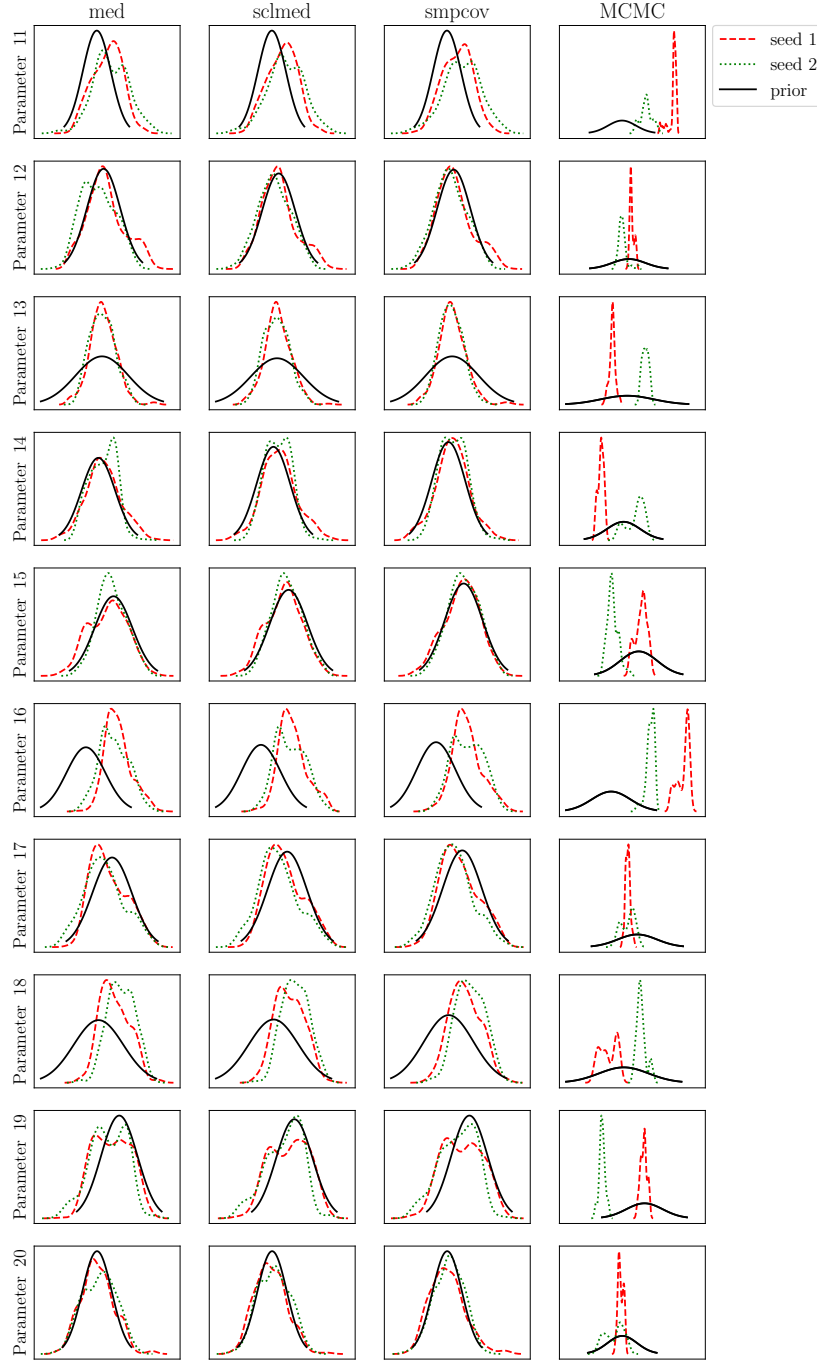


Figure S27: Kernel density estimates for posterior marginals in the calcium signalling model. Stein thinning with **med**, **sclmed** and **smpcov** preconditioners (first three columns) was applied to tempered RW MCMC output (to obtain $m = 500$ points). These can be contrasted with the last column, where kernel density estimates based on standard RW MCMC are displayed. [In each case, two distinct random seeds were used. For reference, the black curve represents the prior marginal.]

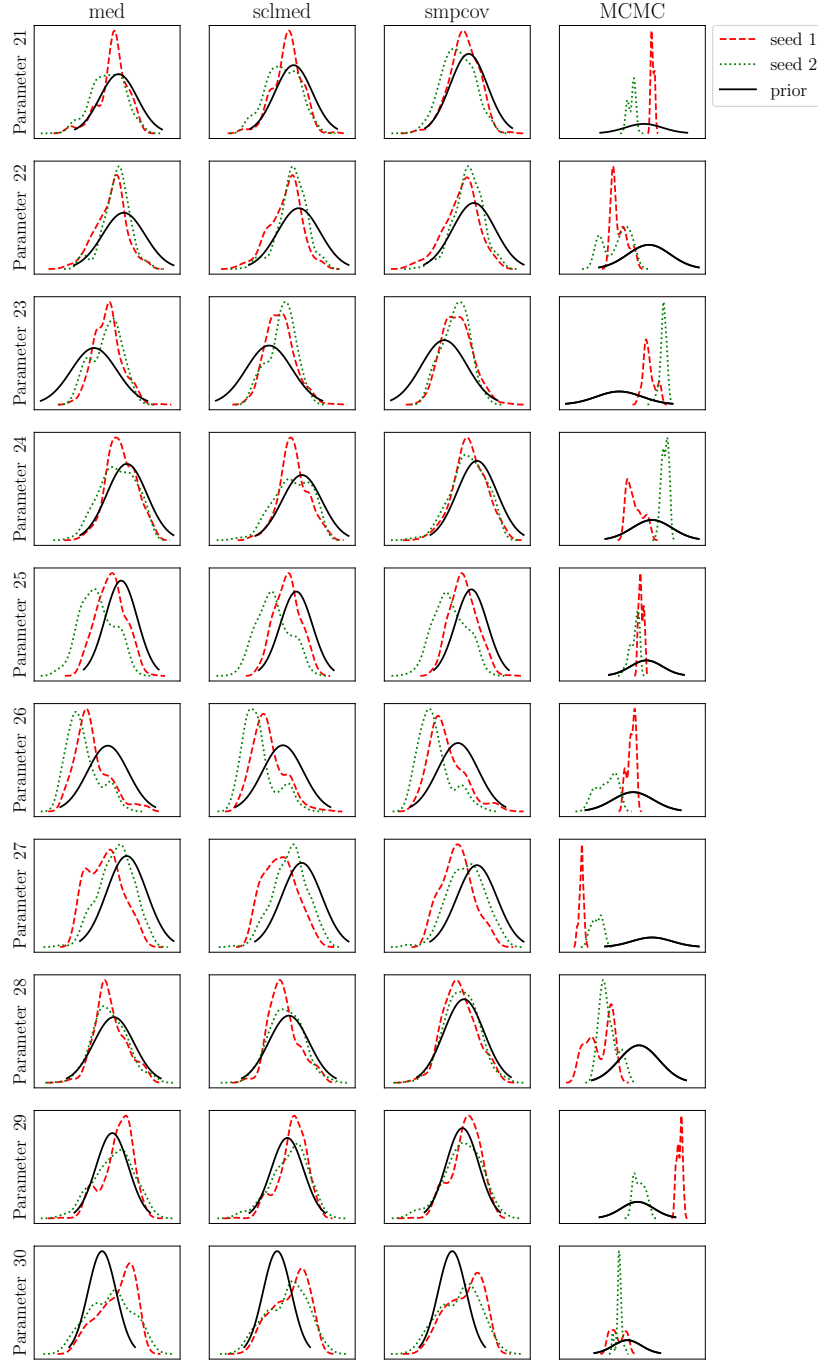


Figure S28: Kernel density estimates for posterior marginals in the calcium signalling model. Stein thinning with **med**, **sclmed** and **smpcov** preconditioners (first three columns) was applied to tempered RW MCMC output (to obtain $m = 500$ points). These can be contrasted with the last column, where kernel density estimates based on standard RW MCMC are displayed. [In each case, two distinct random seeds were used. For reference, the black curve represents the prior marginal.]

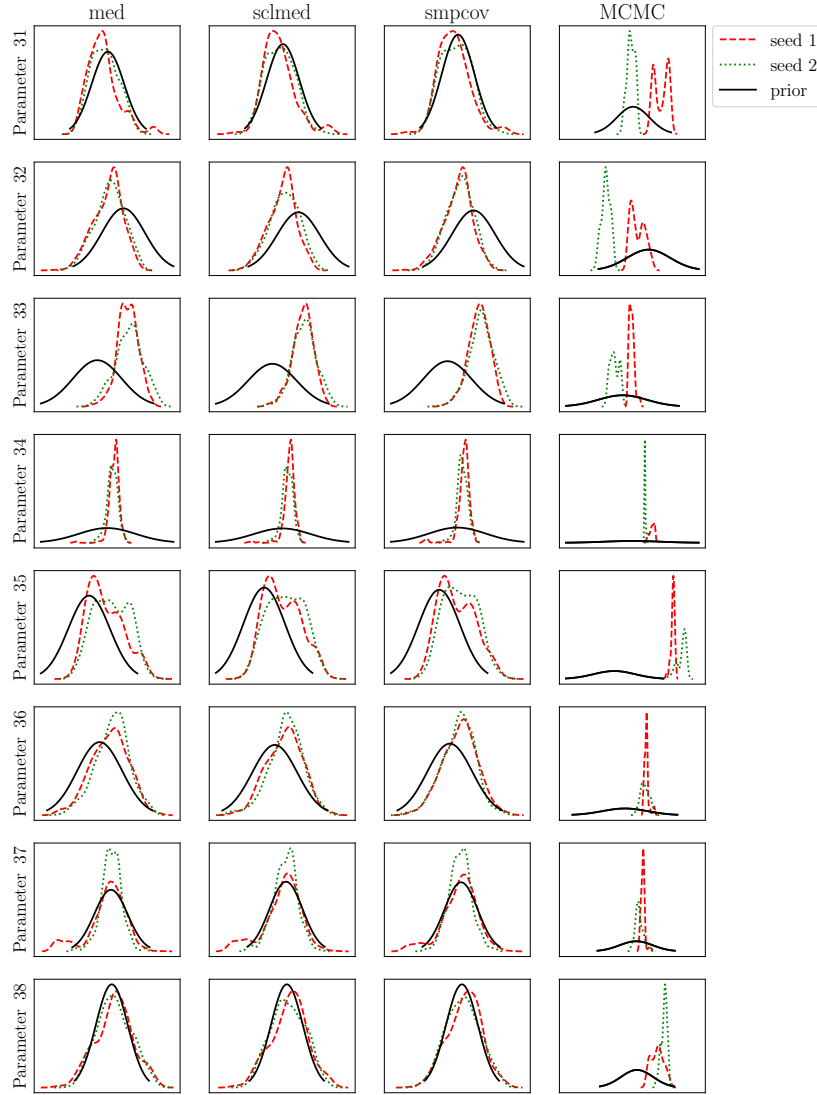


Figure S29: Kernel density estimates for posterior marginals in the calcium signalling model. Stein thinning with **med**, **sclmed** and **smpcov** preconditioners (first three columns) was applied to tempered RW MCMC output (to obtain $m = 500$ points). These can be contrasted with the last column, where kernel density estimates based on standard RW MCMC are displayed. [In each case, two distinct random seeds were used. For reference, the black curve represents the prior marginal.]




Figure S30: Calcium signalling model. Kernel Stein discrepancy (KSD) based on `sclmed`, for empirical distributions obtained using **Support Points** and **Stein Thinning**, based on output from RW MCMC applied to either P or a tempered version of P .




Figure S31: Calcium signalling model. Kernel Stein discrepancy (KSD) based on `smpcov`, for empirical distributions obtained using **Support Points** and **Stein Thinning**, based on output from RW MCMC applied to either P or a tempered version of P .