

On the geometric ergodicity of Hamiltonian Monte Carlo

SAMUEL LIVINGSTONE^{1,*}, MICHAEL BETANCOURT^{2,**}, SIMON BYRNE^{1,†}
and MARK GIROLAMI^{2,3,4,‡}

¹*Department of Statistical Science, University College, Gower Street, London WC1E 6BT, United Kingdom.
E-mail: *samuel.livingstone@ucl.ac.uk; †simonbyrne@gmail.com*

²*Department of Statistics, University of Warwick, Coventry, CV4 7AL, United Kingdom.
E-mail: **betanalphabet@gmail.com; ‡m.girolami@warwick.ac.uk*

³*Department of Mathematics, South Kensington Campus, Imperial College London, London SW7 2AZ,
United Kingdom. E-mail: m.girolami@imperial.ac.uk*

⁴*The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, United Kingdom*

We establish general conditions under which Markov chains produced by the Hamiltonian Monte Carlo method will and will not be geometrically ergodic. We consider implementations with both position-independent and position-dependent integration times. In the former case, we find that the conditions for geometric ergodicity are essentially a gradient of the log-density which asymptotically points towards the centre of the space and grows no faster than linearly. In an idealised scenario in which the integration time is allowed to change in different regions of the space, we show that geometric ergodicity can be recovered for a much broader class of tail behaviours, leading to some guidelines for the choice of this free parameter in practice.

Keywords: geometric ergodicity; Hamiltonian dynamics; Hamiltonian Monte Carlo; hybrid Monte Carlo; Markov chain Monte Carlo; Markov chains; stochastic simulation

1. Introduction

This paper deals with ergodic properties of Markov chains produced by the *Hamiltonian* (or *Hybrid*) Monte Carlo method (HMC), a technique for approximating high dimensional integrals through stochastic simulation [18]. Iterative algorithms of this type are widely used in (for example) statistics and machine learning [2,22], inverse problems [46], and molecular dynamics [1]. In many of these settings a prior distribution can be constructed for an unknown quantity, and after conditioning on some observed data, Bayes' theorem gives a posterior – to extract relevant information from this typically high-dimensional integrals must be evaluated.

A popular approach to such problems is to simulate a Markov chain whose limiting distribution is the posterior, and compute long-run averages (e.g., [42]). Provided the chain is *ergodic*, then a Law of Large Numbers exists for these. Several *Markov chain Monte Carlo* (MCMC) methods of this nature have been proposed in the literature, and many are well understood theoretically (e.g., [43,44]). HMC has proven an empirical success, with numerous authors noting its superior performance in a variety of settings (e.g., [22]) and high performance software available for its implementation [14]. Comparatively few rigorous results, however, exist to justify this. Indeed,

the absence of such analysis has been noted on more than one occasion [15,17]. The major contribution of this work is to establish general scenarios under which geometric ergodicity can and cannot be established for Markov chains produced by common HMC implementations.

Consider a Borel space $(\mathbf{X}, \mathcal{B})$. In this article we focus on the case $\mathbf{X} = \mathbb{R}^d$. We define a Markov chain $(X_n)_{n \geq 0}$ on $(\mathbf{X}, \mathcal{B})$ through an initial distribution $\delta_x(\cdot)$ and a family of mappings $f_\theta : \mathbf{X} \rightarrow \mathbf{X}$, indexed by θ defined on the Borel space $(\Theta, \mathcal{B}_\theta)$ and with associated law $\gamma(\cdot)$ (e.g., [16]). A transition kernel $P : \mathbf{X} \times \mathcal{B} \rightarrow [0, 1]$ can then be induced through the relation

$$P(x, A) = \int \mathbb{1}_A(f_\theta(x))\gamma(d\theta),$$

for any $A \in \mathcal{B}$. Constructing a Markov chain for which some distribution of interest $\pi(\cdot)$ is *invariant* is not very difficult, owing to the Metropolis–Hastings algorithm [25,36], in which the family $\{f_\theta, \theta \in \Theta\}$ is given by

$$f_\theta(x) := \begin{cases} g_\xi(x) & u < \alpha(x, g_\xi(x)), \\ x & \text{otherwise,} \end{cases}$$

where $\theta = \{\xi, u\}$ in this case, with $u \sim U[0, 1]$, and $\{g_\xi, \xi \in \Xi\}$ is a family of ‘candidate’ maps, with $\xi \sim \mu(\cdot)$. A candidate transition kernel is induced as $Q(x, A) = \int \mathbb{1}_A(g_\xi(x))\mu(d\xi)$ for any $A \in \mathcal{B}$. If $\pi(\cdot)$ and $Q(x, \cdot)$ admit densities $\pi(x)$ and $q(x, y)$, then the ‘acceptance probability’ $\alpha : \mathbf{X} \times \mathbf{X} \rightarrow [0, 1]$ can be defined as follows. Let $S := \{(x, y) \in \mathbf{X}^2 : \pi(x)q(x, y) > 0\}$. Then for $(x, y) \in S$ we set

$$r(x, y) := \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)}, \tag{1}$$

and set $r(x, y) := 0$ otherwise. Then $\alpha(x, y) := 1 \wedge r(x, y)$. A more general definition is given in Proposition 1 of [48]. The resulting chain $(X_n)_{n \geq 0}$ is reversible with respect to $\pi(\cdot)$.

Simple choices for the family $\{g_\xi, \xi \in \Xi\}$ result in Markov chains which are intuitive and convenient to analyse. In the random walk case $g_\xi(x) = x + \xi$, with $\Xi = \mathbf{X}$ and $\mu(\cdot)$ a centred, symmetric distribution [47]. For the Metropolis-adjusted Langevin algorithm (MALA) $g_\xi(x) = x + h\nabla \log \pi(x)/2 + \sqrt{h}\xi$, with $\mu(\cdot)$ a standard Gaussian measure on $\Xi = \mathbf{X}$, $h > 0$ a constant, ∇ the gradient operator and $\pi(x)$ the Lebesgue density of $\pi(\cdot)$. The former is in some sense a naive choice, while the latter is an Euler–Maruyama scheme for the diffusion governed by $dX_t = \nabla \log \pi(X_t) dt + \sqrt{2}dW_t$, for which $\pi(\cdot)$ is invariant under suitable regularity conditions (see, e.g., [44]). In both cases, proposals are *local* (only depending on analytic information at the current point), and x is combined with ξ *linearly*, with added complexity coming only through the (typically nonlinear) α . As a result, simple bounds on α allow stochastic stability properties such as π -irreducibility to be deduced straightforwardly, and rates of convergence for different forms of $\pi(\cdot)$ are also well-understood in both cases [43,44].

The HMC method can also be considered within the above framework, as outlined in [7]. The algorithm is designed to exploit the measure-preserving properties of *Hamiltonian flow* (e.g., [32]), which can be induced provided the state space for the chain is a symplectic manifold (e.g., [31]). The space \mathbf{X} can be made symplectic by doubling the dimension, introducing auxiliary *momentum* variables p which follow some user-specified distribution. A Hamiltonian function

can then be constructed on the resulting *phase space* which preserves a distribution for (x, p) , the x -marginal of which will be $\pi(\cdot)$. At each step of the Markov chain, a fresh value for p is drawn from its conditional distribution given the current x state, and then the relevant Hamiltonian flow is approximated for T units of time to produce the next proposed move. The resulting proposal map is

$$g_\xi(x) = \text{Pr}_x \circ \varphi_T(x, p), \quad (2)$$

where Pr_x denotes the projection operator onto the x coordinate, φ_T the approximate flow for T units of time, and $\xi = \{T, p\}$. Typically the distribution for p is chosen to be a d -dimensional Gaussian. If the law of p does not depend on x , then the Störmer–Verlet (or *leapfrog*) numerical integrator is typically used to approximate the flow, with $\varepsilon > 0$ chosen as the integrator step-size and L the number of ‘leapfrog steps’ (meaning $T = L\varepsilon$). The choice of T is a point of ambiguity; often it is set to be some fixed value, however heuristics have also been suggested for choosing this dynamically (e.g., [26]). For $T = \varepsilon$ (meaning $L = 1$) in fact HMC reduces to MALA. In general, however, for $L > 1$ (2) will be a nonlinear function of p , making analysis of the method challenging, particularly in the case of a dynamic T .

Our main contribution is to establish conditions under which common HMC implementations produce a geometrically ergodic Markov chain. We also establish instances where convergence will *not* be geometric, meaning the sampler may perform poorly in practice. We first consider the case where the choice of *integration time* T is chosen independently of the current position, and show that here the nonlinear terms in $g_\xi(x)$ can be *bounded in probability* as the norm $\|x\| \rightarrow \infty$ under suitable assumptions, meaning that geometric convergence essentially occurs for HMC in the same scenarios as for MALA, when the tails of $\pi(x)$ are uniformly exponential or lighter, but no lighter than that of a Gaussian density. We then consider an idealised scheme in which T is chosen as a function of the current position, and show that in this case geometrically converging chains can be constructed for a much broader class of targets. Although the latter results are in an idealised case, they do offer some practical guidelines for the choice of integration time, which can be used to examine some commonly used heuristics in the literature as well as suggest alternatives.

1.1. Literature review

The HMC method was first introduced in lattice field theory [18], as a *hybrid* of two differing approaches to molecular simulation introduced in [1] and [36], respectively. A statistically-oriented review is given in [38]. Several extensions have been suggested. A *generalized* scheme in which the momentum is only partially refreshed was introduced in [27] (see also [40]). Other extensions have been proposed to allow more directed motion and reduced rejections (e.g., [12]). A dynamic approach to tuning the integration time parameter was introduced through the ‘No-U-Turn Sampler’ of [26], which is now implemented in the Stan software [14]. An extension showing how to implement the sampler on a Riemannian manifold which is *globally diffeomorphic* to \mathbb{R}^d is given in [23] (see also [5]), and to embedded manifolds with closed form geodesics in [11].

Theoretical study of MCMC methods is in the main focused on two themes: *convergence to equilibrium* and *asymptotic variance*. The first is often understood through upper bounding some

suitable discrepancy between the n th iterate of the Markov chain and its limiting distribution, as a function of n . When the discrepancy is taken as either the Total Variation or V -norm distance (for some suitable Lyapunov function $V : \mathbf{X} \rightarrow [1, \infty)$), then the drift and minorisation conditions popularised in [37] can be used to show that the distance to equilibrium decreases geometrically in n (we elaborate in Section 3). If such a bound holds then for reversible chains a central limit theorem exists for long-run averages of $L^2(\pi)$ functionals (e.g., [42]). We take this approach here. Note that such techniques rely crucially on the chain being ψ -irreducible for some σ -finite measure $\psi(\cdot)$.

For HMC, [13] establish that if the *potential energy* $U(x) = -\log \pi(x)$ is bounded above, continuous and has bounded derivative then the algorithm will produce a π -irreducible chain. The result holds for both the exact flow and the leapfrog integrator variants of HMC. Typically the boundedness assumption on $U(x)$ will only be satisfied when \mathbf{X} is compact. The authors also show that π -irreducibility can be established more broadly if the integration time is chosen stochastically. More recently, [9] consider a continuous-time version of HMC in which the integration step-size is randomly sampled from an Exponential distribution. Under the assumption that Hamilton's equations can be exactly integrated, they prove that the algorithm will produce a geometrically ergodic Markov chain whenever the tails of $\pi(x)$ decay at a Gaussian rate or faster. The method of the authors is to relate HMC to *underdamped* Langevin dynamics, the ergodic properties of which are established in [35]. By contrast, we relate HMC to *overdamped* Langevin dynamics, as analysed in [44]. Although at first this may seem less natural, in fact it allows us to paint a broad picture of when the algorithm as used in practice will and will not produce a geometrically ergodic Markov chain. In [45] some practical approximations are given for convergence bounds under a *positive curvature* assumption on the underlying chain. We discuss these further in Section 7. We comment further on connections between HMC and Langevin dynamics in the supplementary material [33].

Asymptotic variances of long-run averages from Markov chains are often considered via analysing the expected squared jump distance $\mathbb{E}[(X_{i+1} - X_i)^2]$; at equilibrium this can then be optimised over the various parameters of the dynamics. Careful study of this quantity can also indicate how algorithm performance depends on d . In the case of HMC such analysis has been performed [3], suggesting that the method scales more favourably than other approaches with dimension, and a larger optimal acceptance rate is attained.

Recently a HMC has been generalised to the context of sampling on spaces of infinite dimension [4]. Due to the frequent singularity of measures in such spaces, it is often necessary to characterise distance to equilibrium here through other metrics than Total Variation. Such analysis is beyond the scope of this paper, though we note that recent work in the context of MALA in [21] and [19] are useful pre-cursors in this direction.

1.2. Notation

Let $(\mathbf{X}, \mathcal{B})$ denote a Borel space. Here we restrict attention to $\mathbf{X} = \mathbb{R}^d$ (and write $\|x\|$ for the Euclidean norm of $x \in \mathbf{X}$). For functions $f, g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ let $f(x) \asymp g(x)$ mean that $\lim_{x \rightarrow \infty} f(x)/g(x) = c$ for some $c < \infty$. Throughout let $\pi(\cdot)$ be a finite 'target' measure, and $\pi(x)$ the corresponding Lebesgue density for some $x \in \mathbf{X}$, and let $\mathcal{L}(\cdot)$ be a distribution defined over \mathbb{Z}_+ . We will denote Lebesgue measure on \mathbb{R}^d with $\mu^L(\cdot)$, the Dirac point mass at

x with $\delta_x(\cdot)$ and the standard Gaussian measure with $\mu^G(\cdot)$. We write $P : \mathbf{X} \times \mathcal{B} \rightarrow [0, 1]$ to denote a Markov transition kernel, meaning $P(x, \cdot)$ is a probability measure for any $x \in \mathbf{X}$ and $P(\cdot, A)$ is measurable for any $A \in \mathcal{B}$. P acts to the left on measures through $\mu P(\cdot) := \int \mu(dx)P(x, dy)$ and to the right on functions through $Pf(x) := \int f(y)P(x, dy)$. We let $P^n(x, \cdot) := \int P^{n-1}(x, dy)P(y, \cdot)$ and say $\pi(\cdot)$ is *invariant* for P if $\pi P(\cdot) = \pi(\cdot)$.

Denote the Total Variation distance between two distributions $\mu(\cdot)$ and $\nu(\cdot)$ on $(\mathbf{X}, \mathcal{B})$ as $\|\mu(\cdot) - \nu(\cdot)\|_{\text{TV}} := \sup_{|f| \leq 1} |\mathbb{E}_\mu f - \mathbb{E}_\nu f|$. We say $\pi(\cdot)$ is a *limiting* distribution for P if $\|P^n(x, \cdot) - \pi(\cdot)\|_{\text{TV}} \rightarrow 0$ as $n \rightarrow \infty$, for π -a.e. $x \in \mathbf{X}$. Recall that an invariant distribution $\pi(\cdot)$ will be the unique limiting measure if P is both π -irreducible and aperiodic (e.g. [47]). We note that the convergence results presented here could equivalently be shown under the V -norm distance [41].

2. Overview of main results

The majority of results in this paper concern the version of HMC which is typically used in practice, in which the ‘integration time’ for a typical proposal is chosen independently of the current position in the chain. In this scenario, we have the following result.

Theorem 2.1. *If Assumptions A1 (on page 3118), A2 (on page 3125) and A3 (on page 3127) hold, then a Markov chain produced by the Hamiltonian Monte Carlo method (outlined in Algorithm 1) will be geometrically ergodic.*

Assumption **A1** introduces a controlled degree of randomness into the integration time parameter, which ensures ergodicity of the HMC transition kernel. Instead of establishing π -irreducibility directly on the multiple step HMC transition, we make a simple stochasticity assumption on the integration time parameter, which allows much of the technical difficulty to be sidestepped. Assumption **A2** imposes conditions on the distribution from which expectations are desired, essentially restricting the tail behaviour to be lighter than a Laplacian but no lighter than a Gaussian distribution. This is to ensure that when the chain is very far from the ‘centre’ of the space then typical proposals will bring it back to regions where probability mass concentrates. Assumption **A3** relates to the Metropolis–Hastings acceptance rate, ensuring that this does not behave undesirably, in the sense that desirable proposals are often rejected. We make these arguments precise in Section 5.

We also present the following conditions under which Markov chains produced using HMC will not be geometrically ergodic.

Theorem 2.2. *If either of the following hold, then HMC will not produce a geometrically ergodic Markov chain:*

- (i) $\lim_{\|x\| \rightarrow \infty} \frac{\|\nabla U(x)\|}{\|x\|} = \infty$ and (18) and (19) are satisfied.
- (ii) There is an $M < \infty$ such that $\|\nabla U(x)\| \leq M$ for all $x \in \mathbf{X}$, and $\mathbb{E}_\pi[e^{s\|x\|}] = \infty$ for every $s > 0$.

The first of these scenarios in essence covers the case where the distribution of interest has lighter tails than those of a Gaussian distribution. In this case explicit numerical solvers for Hamilton's equations typically become unstable in some regions of the state space. The second is concerned with 'heavy tailed' distributions, in which the resulting Hamiltonian flow can be slow, precluding a geometric rate of convergence.

To give some intuition for these results, we consider the Exponential Family class of models first introduced in [44] and denoted $\mathcal{E}(\beta, \alpha)$, in which $\pi(x) \in C^1(\mathbb{R})$ and for all $|x| > M$ for some $M < \infty$ it holds that

$$\pi(x) \propto \exp(-\alpha|x|^\beta)$$

for some $\alpha, \beta > 0$ and any $x \in \mathbb{R}$. Different choices of β correspond to different tail behaviours, with larger values resulting in 'lighter' tails. For $\beta \geq 1$ the density is log-concave, and the specific choices $\beta = 1$ and $\beta = 2$ correspond to Laplace and Gaussian distributions.

Corollary 2.3. *For the exponential family class of models $\mathcal{E}(\beta, \alpha)$, under assumption **A1** the following results hold:*

- (i) *For $1 \leq \beta \leq 2$, the Hamiltonian Monte Carlo method will produce a geometrically ergodic chain (for small enough ε in the $\beta = 2$ case).*
- (ii) *If $\beta < 1$ or $\beta > 2$, then the Hamiltonian Monte Carlo method will not produce a geometrically ergodic chain.*

Proof. See page 3127. □

The results are analogous to those found for the Metropolis-adjusted Langevin algorithm in [44]. A key finding of this work is that when the integration time parameter is chosen in a manner which is independent of the current position, then the two methods essentially coincide in terms of presence or absence of geometric ergodicity. In other words, taking more than a single leapfrog step in the method will not result in a chain 'becoming' geometrically ergodic, even though it may still improve the speed of convergence.

We also consider an idealised version of the method in Section 6, in which the integration time is allowed to depend on the current position in a prescribed way. This scheme was designed to mimic several more recent versions of HMC (e.g., [26]) which are commonly used in modern software packages (e.g., [14]). For a specific one-dimensional class of smooth exponential family models, we find the following theorem.

Theorem 2.4. *For the one-dimensional class of distributions with densities of the form*

$$\pi(x) \propto \exp(-\beta^{-1}(1+x^2)^{\beta/2}),$$

then the idealised Hamiltonian Monte Carlo method introduced in Section 6 will produce a geometrically ergodic Markov chain for any choice of $\beta > 0$.

The positive result in the case where $\beta > 2$ is an artefact of the assumption that Hamilton's equations can be exactly solved in the idealised scheme - this result would disappear if a typical

explicit numerical solver were used instead. However, the findings for the case $\beta < 1$ suggest that there are advantages to using an position-dependent integration time in the presence of heavy tails. We discuss this in more detail in Section 7.

3. Preliminaries

The approach taken here to establishing geometric convergence was popularised in the monograph [37]. A key observation first shown in that work is the following.

Theorem 3.1. *Consider a π -irreducible aperiodic Markov chain with state space $(\mathbf{X}, \mathcal{B})$ and transition kernel P . If there exists a π -a.e. finite Lyapunov function $V : \mathbf{X} \rightarrow [1, \infty]$ with ‘small’ level sets, such that the condition $PV(x) \leq \lambda V(x) + b\mathbb{1}_{C_\omega}(x)$ holds for some $\lambda < 1$, $b < \infty$ and some set $C_\omega := \{x : V(x) \leq \omega\}$ with $\omega < \infty$, then $\exists \rho < 1$ and a π -a.e. finite $M : \mathbf{X} \rightarrow [0, \infty]$ such that*

$$\|P^n(x, \cdot) - \pi(\cdot)\|_{\text{TV}} \leq M(x)\rho^n. \tag{3}$$

Recall that a set $C \in \mathcal{B}$ is called ‘small’ if there is a $t < \infty$, a measure $\nu(\cdot)$ defined on $(\mathbf{X}, \mathcal{B})$ and an $\epsilon > 0$ such that $\forall x \in C$ and $\forall A \in \mathcal{B}$ it holds that $P^t(x, A) \geq \epsilon \nu(A)$ (see, e.g., [42]).

We are concerned here with specific forms of P .

Definition 3.2. We say P is of the *Metropolis–Hastings* type if

$$P(x, dy) = \alpha(x, y)Q(x, dy) + r(x)\delta_x(dy), \tag{4}$$

where Q is a Markov kernel, $\alpha(x, y)$ is defined in (1) and $r(x) = 1 - \int \alpha(x, y)Q(x, dy)$.

The following was shown in [43] when P is of the form (4).

Proposition 3.3. *If $\pi(\cdot)$ and $Q(x, \cdot)$ admit Lebesgue densities $\pi(x)$ and $q(y|x)$, $\pi(x)$ is bounded away from 0 and ∞ on compact sets, and there exists $\delta_q > 0$ and $\epsilon_q > 0$ such that, for every x ,*

$$\|x - y\| \leq \delta_q \implies q(y|x) \geq \epsilon_q,$$

then the Metropolis–Hastings chain with candidate density $q(y|x)$ is π -irreducible and aperiodic, and all compact sets are small.

Corollary 3.4. *If P is of Metropolis–Hastings type and the conditions of Proposition 3.3 are satisfied, then (3) holds if and only if*

$$\limsup_{\|x\| \rightarrow \infty} \frac{PV(x)}{V(x)} < 1, \tag{5}$$

for some Lyapunov function V .

Showing a lack of geometric ergodicity typically requires careful study of the distribution of return times to small sets. The following result of [43], however, provides a straightforward method for doing this for Metropolis–Hastings kernels.

Proposition 3.5. *If P is of Metropolis–Hastings type, then (3) fails to hold if $\text{ess sup } r(x) = 1$.*

Lack of geometric ergodicity can also be established in some cases using the following result of [29].

Proposition 3.6. *If for any $\eta > 0$ there is a $\delta > 0$ such that*

$$P(x, B_\delta(x)) > 1 - \eta, \quad (6)$$

where $B_\delta(x) := \{y \in \mathbf{X} : \|x - y\| < \delta\}$, then P can be geometrically ergodic only if $\mathbb{E}_\pi[e^{\beta\|x\|}] < \infty$ for some $\beta > 0$.

If P is of Metropolis–Hastings type, it is straightforward to verify that $Q(x, B_\delta(x)) > 1 - \eta$ ensures (6), meaning we only need consider the candidate kernel in these cases.

4. Hamiltonian Monte Carlo

We give a brief introduction here. For a more detailed account see [38] or [7]. We consider probability densities of the form $\pi(x) \propto e^{-U(x)}$ for some $U : \mathbf{X} \rightarrow [0, \infty)$. If we view $U(x) = -\log \pi(x)$ as a ‘potential’ energy in a physical system, it is natural to consider the larger phase space and construct the *Hamiltonian*

$$H(x, p) = U(x) + \frac{1}{2} p^t M^{-1} p, \quad (7)$$

where p denotes a d -dimensional ‘momentum’ variable, M a $d \times d$ ‘mass’ matrix and $p^t M^{-1} p/2$ the ‘kinetic’ energy (other forms of kinetic energy are also possible, see, for example, [23]). Provided $U(x)$ is differentiable, we can evolve the coordinates (x_t, p_t) through time in such a way that $H(x_t, p_t) = H(x_{t+s}, p_{t+s})$ for any $t, s \in \mathbb{R}$ using Hamilton’s equations

$$\frac{dp_t}{dt} = -\frac{\partial H}{\partial x}, \quad \frac{dx_t}{dt} = \frac{\partial H}{\partial p}. \quad (8)$$

Solving (8) results in *Hamiltonian flow*. To put this presentation into the framework introduced in Section 1, we can consider constructing a measure-preserving map $f_\theta : \mathbf{X} \rightarrow \mathbf{X}$ by setting the input to be x_0 , choosing a momentum variable p_0 , solving (8) for T units of time and then projecting back down onto \mathbf{X} to produce x_T . The parameters $\theta = \{p_0, T\}$ define the behaviour of a single map f_θ , and how they are chosen define the behaviour of the Markov chain produced by iterating the process of randomly selecting a θ and then applying the resulting map f_θ to the current point to produce the next.

Algorithm 1 Hamiltonian Monte Carlo, single iteration

Require: $x_{i-1}, \epsilon \geq 0, \mathcal{L}(\cdot)$
 Set $x_0 \leftarrow x_{i-1}$, draw $p_0 \sim N(0, M)$, $L \sim \mathcal{L}(\cdot)$, set $T \leftarrow L\epsilon$
 Draw $u \sim U[0, 1]$
 Set $\delta \leftarrow H(x_0, p_0) - H \circ \varphi_T(x_0, p_0)$,
if $\log(u) < \delta$ **then**
 Set $x_i \leftarrow \text{Pr}_x \circ \varphi_T(x_0, p_0)$
else
 Set $x_i \leftarrow x_{i-1}$
end if

Of course, it is often not possible to solve (8) exactly, so numerical methods are needed. Fortunately, the rich geometric structure of Hamiltonian systems allows the construction of *symplectic* integrators, which possess attractive long term numerical stability properties (e.g., [32]), meaning that for appropriate Hamiltonians the approximate solution of (8) is such that $H(x_t, p_t) \approx H(x_0, p_0)$ for all $t < \eta$, where $\eta \gg 0$. The standard choice when the Hamiltonian is of the form (7) is the Störmer–Verlet or *leapfrog* scheme, in which $(x_{L\epsilon}, p_{L\epsilon})$ is generated from (x_0, p_0) using L steps of the recursion

$$\begin{aligned}
 p_{t+\frac{\epsilon}{2}} &= p_t - \frac{\epsilon}{2} \nabla U(x_t), \\
 x_{t+\epsilon} &= x_t + \epsilon M^{-1} p_{t+\frac{\epsilon}{2}}, \\
 p_{t+\epsilon} &= p_{t+\frac{\epsilon}{2}} - \frac{\epsilon}{2} \nabla U(x_{t+\epsilon}),
 \end{aligned}$$

for some step-size $\epsilon > 0$. Although the resulting *approximate* flow map $\varphi_{L\epsilon}(x_0, p_0) := (x_{L\epsilon}, p_{L\epsilon})$ no longer preserves $\pi(\cdot)$, it can be used as a proposal mechanism within the Metropolis–Hastings framework (e.g., [38]). The full method is shown in Algorithm 1.

Remark 4.1. From this point forward, we assume $M = I$ for ease of exposition but without loss of generality.

4.1. The marginal chain

To use the techniques of [37], it is helpful to express the HMC transition in such a way that when $\|x\|$ is large it is clear how the chain will behave. Although it is typically presented as a map on the larger phase space, HMC can simply be thought of as a Markov chain on \mathbf{X} , and we will find this representation useful in relation to the above. In this case the candidate map g_ξ is given by the following proposition, which can be straightforwardly be derived using classical results (see, e.g., [10]).

Proposition 4.2. *The HMC candidate map can be written*

$$x_{L\varepsilon} = x_0 - \frac{L\varepsilon^2}{2} \nabla U(x_0) - \varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon}) + L\varepsilon p_0. \tag{9}$$

where $p_0 \sim N(0, I)$, L is the number of leapfrog steps and ε the integrator step-size. With this choice, the acceptance probability will be

$$\alpha(x_0, x_{L\varepsilon}) = 1 \wedge \frac{\pi(x_{L\varepsilon})}{\pi(x_0)} \exp\left(\frac{1}{2} \|p_0\|^2 - \frac{1}{2} \|p_{L\varepsilon}\|^2\right), \tag{10}$$

where

$$p_{L\varepsilon} = p_0 - \frac{\varepsilon}{2} \nabla U(x_0) - \varepsilon \sum_{i=1}^{L-1} \nabla U(x_{i\varepsilon}) - \frac{\varepsilon}{2} \nabla U(x_{L\varepsilon}). \tag{11}$$

Proposition 4.2 highlights the previously noted relationship between HMC and MALA quite explicitly, as setting $L = 1$ means the third term on the right-hand side of (9) disappears, leaving the MALA proposal $x_0 - \varepsilon^2 \nabla U(x_0)/2 + \varepsilon p_0$. It also highlights why taking $L > 1$ proposes a greater challenge, as for each $x_{i\varepsilon}$ with $i \geq 1$ this term will typically be a nonlinear transformation of x_0 and p_0 . As p_0 is stochastic, then $\varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon})$ will be also, but its distribution will often be intractable.

5. Results for an position-independent integration time

In this section, we make the assumption that the distribution $\mathcal{L}(\cdot)$ for the number of leapfrog steps L does not depend on the current position. This is relaxed in Section 6.

5.1. π -Irreducibility

It is known (e.g., [13]) that establishing π -irreducibility is not so straightforward in the case of HMC as for Metropolis–Hastings methods based on random walks or Langevin diffusions. The canonical example where the system becomes reducible is integrating the harmonic oscillator over precisely one period (e.g., [32]). We show this in the supplementary material [33].

The observation noted here and elsewhere that HMC in the case $L = 1$ corresponds to MALA, for which irreducibility is established in [44], can be exploited to alleviate these issues and establish π -irreducibility of HMC under the following assumption.

A1 *The distribution $\mathcal{L}(\cdot)$ is such that $\mathbb{P}_{\mathcal{L}}[L = 1] > 0$, and for any fixed $(x_0, p_0) \in \mathbb{R}^{2d}$ and $\varepsilon > 0$, and that there is an $s < \infty$ such that $\mathbb{E}_{\mathcal{L}}[e^{s\|x_{L\varepsilon}\|}] < \infty$.*

When assumption **A1** holds then the fact that HMC produces an ergodic Markov chain can be straightforwardly invoked from existing MALA results [44]. The idea of randomising the integration time is commonly recommended for practical applications of the method (e.g., [23,38]),

and more theoretical motivation for doing so is given in [6]. The finite exponential expectation condition is needed to ensure that the Lyapunov function used to prove geometric ergodicity is valid. One simple way to ensure this in practice (under the additional assumptions imposed on the potential U in the next subsection) is that $\mathbb{P}_{\mathfrak{L}}[L > l] = 0$ for some fixed $l < \infty$, though weaker conditions than this are also possible.

Remark 5.1. Assumption **A1** can be viewed as the discrete time analogue to the the exponential integration time assumption made in [9], and in many respects serves a similar purpose. Similar conditions are also exploited to prove π -irreducibility results in [13].

Remark 5.2. In fact, before the final publication of the present work, it was shown in [20] that π -irreducibility can indeed be established without using assumption **A1**, but instead considering the HMC chain using a fixed number of leapfrog steps $L \geq 1$, under suitable assumptions and using appropriate techniques. We refer the interested reader to that work for details.

5.2. Geometric ergodicity

We first present here some seemingly abstract conditions under which the HMC method produces a geometrically ergodic Markov chain. We then give some natural assumptions on the potential $U(x)$ under which these hold.

We present the results of this section conditioned on a fixed choice of the number of leapfrog steps L , for ease of exposition. Note that the required drift conditions shown hold for a fixed L , then under **A1** they will hold when possible values for L are averaged over according to $\mathfrak{L}(\cdot)$, so this does not affect the generality of the results.

Notation. We introduce some further notation for this section. Let $I_{\delta}(x) := \{y \in \mathbf{X} : \|y\| \leq \|x\|^{\delta}\}$ for some $1/2 < \delta < 1$. In the case $\delta = 1$, we will simply write $I(x)$. Let

$$m_{L,\varepsilon}(x_0, p_0) := x_0 - L\varepsilon^2 \nabla U(x_0)/2 - \varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon})$$

denote the ‘mean’ next candidate position $(x_{L\varepsilon} - L\varepsilon p_0)$, and

$$\psi_{L,\varepsilon}(x_0, p_0) := L\varepsilon^2 \nabla U(x_0)/2 + \varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon})$$

denote the proposal ‘drift’ (implying $m_{L,\varepsilon}(x_0, p_0) = x_0 - \psi_{L,\varepsilon}(x_0, p_0)$). We will also sometimes write $h := \varepsilon^2/2$ in a most likely futile attempt to keep things readable.

Theorem 5.3. *The HMC method produces a geometrically ergodic Markov chain if assumption **A1** holds, and in addition both*

$$\limsup_{\|x_0\| \rightarrow \infty, \|p_0\| \leq \|x_0\|^{\delta}} (\|m_{L,\varepsilon}(x_0, p_0)\| - \|x_0\|) < -\sqrt{2}L\varepsilon\eta(d), \tag{12}$$

where $\eta(d) := \Gamma((d + 1)/2) / \Gamma(d/2)$, and

$$\lim_{\|x_0\| \rightarrow \infty} \int_{R(x_0) \cap I(x_0)} Q(x_0, dy) = 0, \tag{13}$$

where $R(x_0) := \{y \in \mathbf{X} : \alpha(x_0, y) < 1\}$ denotes the ‘potential rejection region’ and $I(x_0) := \{y \in \mathbf{X} : \|y\| \leq \|x_0\|\}$ the ‘interior’ of x_0 .

Proof. Take $V(x) = e^{s\|x\|}$ for some $s > 0$ and write $A(x_0) := R(x_0)^c$. Then we can write

$$\begin{aligned} \frac{PV(x_0)}{V(x_0)} &= \int_{A(x_0)} e^{s(\|y\| - \|x_0\|)} Q(x_0, dy) + \int_{R(x_0)} e^{s(\|y\| - \|x_0\|)} \alpha(x_0, y) Q(x_0, dy) \\ &\quad + \int_{R(x_0)} (1 - \alpha(x_0, y)) Q(x_0, dy) \\ &= \int_{\mathbb{R}^d} e^{s(\|y\| - \|x_0\|)} Q(x_0, dy) + \int_{R(x_0)} (1 - e^{s(\|y\| - \|x_0\|)})(1 - \alpha(x_0, y)) Q(x_0, dy) \\ &\leq \int e^{s(\|y\| - \|x_0\|)} Q(x_0, dy) + \int_{R(x_0) \cap I(x_0)} Q(x_0, dy). \end{aligned}$$

The last integral asymptotes to zero as $\|x_0\| \rightarrow \infty$ by (13). Writing $x_{L\varepsilon}(p_0)$ to indicate that $x_{L\varepsilon}$ depends on p_0 , the first integral can be written

$$\int_{I_\delta(x_0)} e^{s(\|x_{L\varepsilon}(p_0)\| - \|x_0\|)} \mu^G(dp_0) + \int_{I_\delta(x_0)^c} e^{s(\|x_{L\varepsilon}(p_0)\| - \|x_0\|)} \mu^G(dp_0). \tag{14}$$

Noting that $\|x_{L\varepsilon}(p_0)\| \leq \|m_{L,\varepsilon}(x_0, p_0)\| + L\varepsilon\|p_0\|$ for large enough $\|x_0\|$ and using (12) above then setting $\xi(x_0) := \sup_{\|p_0\| \leq \|x_0\|^\delta} (\|m_{L,\varepsilon}(x_0, p_0)\| - \|x_0\|)$ we can write

$$\int_{I_\delta(x_0)} e^{s(\|x_{L\varepsilon}\| - \|x_0\|)} \mu^G(dp_0) \leq e^{s\xi(x_0)} \int_{I_\delta(x_0)} e^{sL\varepsilon\|p_0\|} \mu^G(dp_0).$$

The last integral can be bounded above by the moment generating function of a Chi-distributed random variable with d degrees of freedom, and so equals $e^{\log(1+s\sqrt{2}L\varepsilon\eta(d)+o(s))} \leq e^{s\sqrt{2}L\varepsilon\eta(d)+o(s)}$. Therefore by (12), the integral asymptotes to a quantity which is strictly less than one if $s > 0$ is chosen to be suitably small.

It remains to show that the right-hand integral in (14) becomes negligibly small as $\|x_0\| \rightarrow \infty$. It follows from (9) and (A4.4.1) that $\|x_{L\varepsilon}\| \in O(\max(\|x_0\|, \|p_0\|))$. This means that for some constants $C \in \mathbb{R}$ and for $p_0 \in I_\delta(x_0)^c$ and $\|x_0\|$ large enough we can write

$$\begin{aligned} \exp\left(s\|x_{L\varepsilon}\| - s\|x_0\| - \frac{1}{2}\|p_0\|^2\right) &\leq \exp\left(c \max(\|x_0\|, \|p_0\|) - \frac{1}{2}\|p_0\|^2\right) \\ &= \exp\left(\|p_0\| \left(C \frac{\max(\|x_0\|, \|p_0\|)}{\|p_0\|} - \|p_0\|\right)\right) \\ &\leq \exp(\|p_0\|(C\|x_0\|^{1-\delta} - \|p_0\|)). \end{aligned}$$

Provided $\delta > 1/2$, then for $\|x_0\|$ large enough $C\|x_0\|^{1-\delta} - \|x_0\|^\delta < -1$, meaning

$$\exp(\|p_0\|(C\|x_0\|^{1-\delta} - \|p_0\|)) \leq \exp(-\|p_0\|),$$

meaning

$$\int_{I_\delta(x_0)^c} e^{s(\|x_{L\varepsilon}\| - \|x_0\|)} \mu^G(dp_0) \leq \int_{I_\delta(x_0)^c} e^{-\|p_0\|} \mu^G(dp_0) \leq 2e^{-\|x_0\|^\delta},$$

which becomes negligibly small as $\|x_0\| \rightarrow \infty$, as required. \square

Theorem 5.3 is a generalisation of Theorem 4.1 in [44] to the HMC case. The nontriviality involved in this extension is accounting for the randomness induced into $m_{L,\varepsilon}(x_0, p_0)$ from p_0 .

5.2.1. Requirements for (12) to be satisfied

In the case $L = 1$ (12) corresponds to

$$\|x_0 - h\nabla U(x_0)\| - \|x_0\| < -\sqrt{2\varepsilon}\eta(d) \tag{15}$$

whenever $\|x_0\| > M$ for some $M < \infty$. The statements in this section give three simple conditions which establish this are *also* sufficient to establish (12) when $L \geq 2$. The main result is stated below. The crucial consequence of this is that controlling the behaviour of ‘global move’ updates produced by HMC when $L > 1$ can be done through only ‘local’ knowledge, meaning analytic information at the current point x_0 .

Theorem 5.4. *For any $L \geq 1$ (12) holds if the following conditions are met*

- (SC1.1) $\lim_{\|x_0\| \rightarrow \infty} \|\nabla U(x_0)\| = \infty$,
- (SC1.2) $\liminf_{\|x_0\| \rightarrow \infty} \frac{\langle \nabla U(x_0), x_0 \rangle}{\|\nabla U(x_0)\| \|x_0\|} > 0$,
- (SC1.3) $\lim_{\|x_0\| \rightarrow \infty} \frac{\|\nabla U(x_0)\|}{\|x_0\|} = 0$.

In addition, if (SC1.3) is replaced by

$$(SC1.3b) \limsup_{\|x_0\| \rightarrow \infty} \frac{\|\nabla U(x_0)\|}{\|x_0\|} = S_l,$$

for some $S_l < \infty$, then there is an $\varepsilon_0 \in (0, \infty)$ such that the same result holds provided $\varepsilon \in (0, \varepsilon_0)$.

Proof. This is restated as Proposition 5.10 and Proposition 5.11 below, which follow from the preceding lemmas. \square

The conditions of the result are intuitive. Condition (SC1.2) ensures that the gradient asymptotically ‘points inwards’, while (SC1.1) and (SC1.3) ensure that $\|\nabla U(x_0)\|$ grows but at an asymptotically sublinear rate. We begin with a straightforward observation.

Proposition 5.5. *Sufficient conditions such that*

$$\limsup_{\|x_0\| \rightarrow \infty} (\|x_0 - h\nabla U(x_0)\| - \|x_0\|) < 0$$

are:

$$(SC1.2a) \quad \limsup_{\|x_0\| \rightarrow \infty} \left(\frac{h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle}{2\|x_0\|} \right) < -\sqrt{2}\varepsilon\eta(d),$$

$$(SC1.3) \quad \lim_{\|x_0\| \rightarrow \infty} \frac{\|\nabla U(x_0)\|}{\|x_0\|} = 0.$$

Proof. First note that $\|x_0 - h\nabla U(x_0)\| = \sqrt{\|x_0\|^2 + h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle}$. Recall the generalised Bernoulli inequality: if $y > -1$ and $r \in [0, 1]$ then $(1 + y)^r \leq 1 + ry$. Setting $r := 1/2$, $a(x_0) := \|x_0\|^2$ and $b(x_0) := h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle$ then we have

$$\begin{aligned} a(x_0)^r \left(1 + \frac{b(x_0)}{a(x_0)} \right)^r &\leq a(x_0)^r \left(1 + \frac{b(x_0)}{2a(x_0)} \right) \\ &= \|x_0\| + \frac{h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle}{2\|x_0\|}. \end{aligned}$$

This will be strictly less than $\|x_0\|$ in the limit as $\|x_0\| \rightarrow \infty$ provided $b(x_0)/a(x_0) > -1$. Noting that

$$\frac{b(x_0)}{a(x_0)} = \frac{h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle}{\|x_0\|^2} > \frac{h^2 \|\nabla U(x_0)\|^2 - 2h \|\nabla U(x_0)\| \|x_0\|}{\|x_0\|^2},$$

then it suffices to see that under (SC1.3) the right-hand side can be made arbitrarily close to zero by taking $\|x_0\|$ large enough. \square

We can also recover some more intuitive sufficient conditions.

Corollary 5.6. *A more intuitive condition which implies (SC1.2a) conditional on (SC1.3) and (SC1.1) is*

$$(SC1.2) \quad \lim_{\|x_0\| \rightarrow \infty} \frac{\langle \nabla U(x_0), x_0 \rangle}{\|\nabla U(x_0)\| \|x_0\|} > 0.$$

Proof. Using (SC1.1) and (SC1.2) gives

$$\lim_{\|x_0\| \rightarrow \infty} \frac{\langle \nabla U(x_0), x_0 \rangle}{\|x_0\|} = \infty.$$

For large enough $\|x_0\|$ this implies

$$\begin{aligned} &\left(\frac{h^2 \|\nabla U(x_0)\|}{2\|x_0\|} - \frac{h \langle \nabla U(x_0), x_0 \rangle}{2\|\nabla U(x_0)\| \|x_0\|} \right) < -\frac{\sqrt{2}\varepsilon\eta(d)}{\|\nabla U(x_0)\|} \\ \implies \quad &\limsup_{\|x_0\| \rightarrow \infty} \left(\frac{h^2 \|\nabla U(x_0)\|^2 - 2h \langle \nabla U(x_0), x_0 \rangle}{2\|x_0\|} \right) < -\sqrt{2}\varepsilon\eta(d). \end{aligned} \quad \square$$

From now on we refer to (SC1.1), (SC1.2) and (SC1.3) combined as (SC1.1)–(SC1.3). Next, we show that these same conditions are sufficient for (12) to hold.

Lemma 5.7. *Under the following conditions (12) holds:*

- (i) $\liminf_{\|x_0\| \rightarrow \infty, \|p_0\| \leq \|x_0\|^\delta} \left(\frac{\|\psi_{L,\varepsilon}\|^2 - 2\langle \psi_{L,\varepsilon}, x_0 \rangle}{\|x_0\|^2} \right) > -1,$
- (ii) $\limsup_{\|x_0\| \rightarrow \infty, \|p_0\| \leq \|x_0\|^\delta} \left(\frac{\|\psi_{L,\varepsilon}\|^2 - 2\langle \psi_{L,\varepsilon}, x_0 \rangle}{2\|x_0\|} \right) < 0.$

Proof. Using the generalised Bernoulli inequality as above gives the result. \square

Next we relate the conditions of Lemma 5.7 to criteria that only depend on the current point x_0 . The following lemmas give a starting point.

Lemma 5.8. *Provided $\|p_0\| \leq \|x_0\|^\delta$ and (SC1.3) holds then we have the following*

- (i) *For any $\eta > 0$ there is an $M_\eta < \infty$ such that whenever $\|x_0\| > M_\eta$ it holds that $(1 - \eta)\|x_0\| \leq \|x_\varepsilon\| \leq (1 + \eta)\|x_0\|,$*
- (ii) $\|\nabla U(x_\varepsilon)\| = o(\|x_0\|),$
- (iii) $\|p_\varepsilon\| \in o(\|x_0\|).$

Proof. (i) Noting that $\|x_\varepsilon\| = \|x_0 - h\nabla U(x_0) + \varepsilon p_0\|$ gives

$$\begin{aligned} (1 - \delta)\|x_0\| &\leq \|x_0\| - h\|\nabla U(x_0)\| - \varepsilon\|x_0\|^\delta \\ &\leq \|x_0 - h\nabla U(x_0) + \varepsilon p_0\| \\ &\leq \|x_0\| + h\|\nabla U(x_0)\| + \varepsilon\|x_0\|^\delta \\ &\leq (1 + \delta)\|x_0\|. \end{aligned}$$

(ii) We have from (i) and (SC1.3) that for any $\gamma > 0$ there is an $M_\gamma < \infty$ such that whenever $\|x_0\| > M_\gamma/(1 - \delta)$ then $\|\nabla U(x_\varepsilon)\|/\|x_\varepsilon\| < \gamma$. This implies using (i) that $\|\nabla U(x_\varepsilon)\|/\|x_0\| < \gamma(1 - \delta)$, and since $\gamma(1 - \delta)$ can be made arbitrarily small then the result follows.

(iii) $\|p_\varepsilon\| = \|p_0 - \varepsilon\nabla U(x_0)/2 - \varepsilon\nabla U(x_\varepsilon)/2\| \leq \|p_0\| + \varepsilon\|\nabla U(x_0)\|/2 + \varepsilon\|\nabla U(x_\varepsilon)\|/2,$ which is $\in o(\|x_0\|)$ using (i) and (ii) and the fact that $\|p_0\| \leq \|x_0\|^\delta$. \square

Lemma 5.9. *Provided $\|p_0\| \leq \|x_0\|^\delta$ and (SC1.3) holds then for any $L < \infty$ and each $i \in \{0, \dots, L - 1\}$ the following hold*

- (i) *For any $\eta > 0$ there is an $M_\eta < \infty$ such that whenever $\|x_0\| > M_\eta$ it holds that $(1 - \eta)\|x_0\| \leq \|x_{i\varepsilon}\| \leq (1 + \eta)\|x_0\|,$*
- (ii) $\|\nabla U(x_{i\varepsilon})\| \in o(\|x_0\|),$
- (iii) $\|p_{i\varepsilon}\| \in o(\|x_0\|),$
- (iv) $\|\psi_{L,\varepsilon}\| \in o(\|x_0\|).$

Proof. The results follow iteratively for each i using the same approach as in the previous Lemma. For the case $i = 2$ then noting that $\|x_{2\varepsilon}\| = \|x_\varepsilon - h\nabla U(x_\varepsilon) + \varepsilon p_\varepsilon\|$, then (i) in this

case follows from Lemma 5.8. It follows that $\|\nabla U(x_{2\varepsilon})\| \in o(\|x_0\|)$ and $\|p_{2\varepsilon}\| \in o(\|x_0\|)$ by an analogous argument to this lemma. Given this then it can be shown that (i) holds for $i = 3$, and then (ii) and (iii) by the same logic, and the argument can be iterated as many times as is needed. The last claim follows trivially from the second. \square

Proposition 5.10. *Under (SC1.1)–(SC1.3) then for any $L < \infty$ the conditions of Lemma 5.7 are satisfied.*

Proof. First we show (i). Writing $x^* := \arg \max_{i \in 0, \dots, L-1} \{\|\nabla U(x_{i\varepsilon})\|\}$, then we have $\|\psi_{L,\varepsilon}\| \leq L\varepsilon^2 \sum_{i=0}^{L-1} \|\nabla U(x_{i\varepsilon})\| \leq L^2\varepsilon^2 \|\nabla U(x^*)\|$, which implies

$$\frac{\|\psi_{L,\varepsilon}\|^2}{\|x_0\|^2} \leq \frac{L^4\varepsilon^4 \|\nabla U(x^*)\|^2}{(1-\eta)^2 \|x^*\|^2}$$

which can be made arbitrarily small by taking $\|x_0\|$ large enough using (SC1.3). Noting that $\|\psi_{L,\varepsilon}\|/\|x_0\| \geq \langle \psi_{L,\varepsilon}, x_0 \rangle / \|x_0\|^2 \geq -\|\psi_{L,\varepsilon}\|/\|x_0\|$, then an analogous argument can be used to show that $-2\langle \psi_{L,\varepsilon}, x_0 \rangle / \|x_0\|^2$ will also tend to zero as $\|x_0\| \rightarrow \infty$.

(ii) First note from above that $\lim_{\|x_0\| \rightarrow \infty} \|\psi_{L,\varepsilon}\|^2 / (L^2\varepsilon^2 \|\nabla U(x^*)\| \|x_0\|) = 0$. By an analogous argument to that used in the proof of Corollary 5.6, it is clear therefore that (ii) holds if

$$\limsup_{\|x_0\| \rightarrow \infty} \left(\frac{\|\psi_{L,\varepsilon}\|^2}{L^2\varepsilon^2 \|\nabla U(x^*)\| \|x_0\|} - \frac{\langle \psi_{L,\varepsilon}, x_0 \rangle}{L^2\varepsilon^2 \|\nabla U(x^*)\| \|x_0\|} \right) < -\sqrt{2}L\varepsilon\eta(d)$$

which in turn holds if the statement

$$\lim_{\|x_0\| \rightarrow \infty} \frac{\langle \psi_{L,\varepsilon}, x_0 \rangle}{\|x_0\|} = \infty$$

does. The numerator can be decomposed as

$$\begin{aligned} \langle \psi_{L,\varepsilon}, x_0 \rangle &\geq \sum_{i=0}^{L-1} c_i \langle \nabla U(x_{i\varepsilon}), x_0 \rangle \\ &= \sum_{i=0}^{L-1} c_i \langle \nabla U(x_{i\varepsilon}), x_{i\varepsilon} \rangle + \sum_{i=0}^{L-1} c_i \langle \nabla U(x_{i\varepsilon}), x_0 - x_{i\varepsilon} \rangle, \end{aligned}$$

where each $c_i = (L-i)\varepsilon^2$ for $i \geq 1$ and $c_0 = L\varepsilon^2/2$. The second of these terms is $o(\|\nabla U(x^*)\| \times \|x_0\|)$ using the Cauchy–Schwarz inequality and Lemma 5.9 (which shows that $\|x_0 - x_{i\varepsilon}\| \in o(\|x_0\|)$), and so this term vanishes if divided by $\|\nabla U(x^*)\| \|x_0\|$. The first term divided by the same quantity will be strictly positive as each term in the sum is ≥ 0 using (SC1.3) and Lemma 5.9, and at least one of them is > 0 since it will correspond to x^* . Using (SC1.1) establishes that $\langle \psi_{L,\varepsilon}, x_0 \rangle / \|x_0\| \rightarrow \infty$ as $\|x_0\| \rightarrow \infty$, proving the result. \square

The condition (SC1.3) allows clarity in the proofs, but precludes the natural boundary case of distributions with Gaussian tails. The following proposition addresses this.

Proposition 5.11. *If (SC1.1)–(SC1.2) hold and in addition*

$$(SC1.3b) \quad \limsup_{\|x_0\| \rightarrow \infty} \frac{\|\nabla U(x_0)\|}{\|x_0\|} = S_l$$

for some constant $S_l < \infty$, then there is an $\varepsilon_0 \in (0, \infty)$ such that for any choice of $\varepsilon \in (0, \varepsilon_0)$ the conditions of Lemma 5.7 are satisfied.

Proof. We simply note that the term $\langle \psi_{L,\varepsilon}, x_0 \rangle \in O(\varepsilon^2)$, while $\|\psi_{L,\varepsilon}\|^2 \in O(\varepsilon^4)$, so that the proofs of the preceding lemmas can be straightforwardly modified when (SC1.3) is replaced by (SC1.3b) by choosing a small enough value of ε that the inner product dominates the square norm. We omit the details of this. \square

The sensitivity to the choice of ε in this case is well known in this scenario as a potential source of numerical instabilities, and choosing $\varepsilon < 1/S_l$ is recommended to alleviate such issues (e.g., [32]). We conclude this subsection with the following assumption that we require for a geometrically ergodic Markov chain produced by the HMC method, which is a natural conclusion of the preceding results.

A2 *The potential $U(x)$ satisfies either (SC1.1)–(SC1.3), or it satisfies (SC1.1)–(SC1.3b) and ε is chosen to be suitably small that the conditions of Lemma 5.7 are satisfied.*

Remark 5.12. Condition (SC1.1) precludes densities for which $\|\nabla U(x)\| \rightarrow c$ for some $0 < c < \infty$. Often geometric ergodicity will still hold in this case, as we demonstrate in Corollary 2.3, however a different argument is required to that presented above.

It remains to consider (13), which reflects the role of the acceptance rate in the HMC method. We turn to this next.

5.2.2. Discussion of (13)

In [44], the authors note that (13) applied to the MALA transition $x_\varepsilon = x - \varepsilon^2 \nabla U(x)/2 + \varepsilon p_0$ can be viewed as the restriction that for $x_\varepsilon \in I(x_0)$

$$U(x_0) - U(x_\varepsilon) - \hat{U}_1 \geq \frac{\varepsilon^2}{8} (\|\nabla U(x_\varepsilon)\|^2 - \|\nabla U(x_0)\|^2),$$

where $\hat{U}_1 := \langle x_0 - x_\varepsilon, \nabla U(x_\varepsilon) + \nabla U(x_0) \rangle / 2$ denotes the ‘trapezium’ estimate for the line integral $U(x_0) - U(x_\varepsilon) = \int_{x_\varepsilon}^{x_0} \nabla U(z) dz$. We can extend this intuition to HMC and arrive at the following natural generalisation of the same condition.

Proposition 5.13. *The acceptance rate for HMC will satisfy the ‘inwards acceptance’ property (13) if whenever $x_{L\varepsilon} \in I(x_0)$ then in the limit as $\|x_0\| \rightarrow \infty$ it holds that*

$$U(x_0) - U(x_{L\varepsilon}) - \hat{U}_L \geq \frac{1}{2L^2\varepsilon^2} (\|\psi_{L,\varepsilon}^R\|^2 - \|\psi_{L,\varepsilon}\|^2), \tag{16}$$

where $\hat{U}_L := \langle x_0 - x_{L\varepsilon}, \nabla U(x_0) + \nabla U(x_{L\varepsilon}) + 2 \sum_{i=1}^{L-1} \nabla U(x_{i\varepsilon}) \rangle / (2L)$ denotes the quadrature rule estimate for the line integral $\int_{x_0}^{x_{L\varepsilon}} \nabla U(z) dz$ based on L trapezia and the forward and reverse drift components are given by

$$\psi_{L,\varepsilon} := \frac{L\varepsilon^2}{2} \nabla U(x_0) + \varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon}), \quad \psi_{L,\varepsilon}^R := \frac{L\varepsilon^2}{2} \nabla U(x_{L\varepsilon}) + \varepsilon^2 \sum_{i=1}^{L-1} i \nabla U(x_{i\varepsilon}).$$

Proof. We first note that we can write $p_0 = \frac{1}{L\varepsilon}(x_{L\varepsilon} - x_0 + \psi_{L,\varepsilon})$, and that using reversibility of the leapfrog integrator, we can also write $p_{L\varepsilon} = \frac{1}{L\varepsilon}(x_{L\varepsilon} - x_0 - \psi_{L,\varepsilon}^R)$. The log acceptance ratio can therefore be written

$$U(x_0) - U(x_{L\varepsilon}) + \frac{1}{2L^2\varepsilon^2} (\|x_{L\varepsilon} - x_0 + \psi_{L,\varepsilon}\|^2 - \|x_{L\varepsilon} - x_0 - \psi_{L,\varepsilon}^R\|^2).$$

We require this quantity to be ≥ 0 . This is equivalent to the requirement

$$U(x_{L\varepsilon}) - U(x_0) \leq \frac{1}{2L^2\varepsilon^2} (\|x_{L\varepsilon} - x_0 + \psi_{L,\varepsilon}\|^2 - \|x_{L\varepsilon} - x_0 - \psi_{L,\varepsilon}^R\|^2).$$

We can rewrite the right-hand side of the above expression as

$$\frac{1}{2L^2\varepsilon^2} (2\langle x_{L\varepsilon} - x_0, \psi_{L,\varepsilon} + \psi_{L,\varepsilon}^R \rangle + \|\psi_{L,\varepsilon}\|^2 - \|\psi_{L,\varepsilon}^R\|^2),$$

and then note that

$$\psi_{L,\varepsilon} + \psi_{L,\varepsilon}^R = \frac{L\varepsilon^2}{2} \left(\nabla U(x_0) + \nabla U(x_{L\varepsilon}) + 2 \sum_{i=1}^{L-1} \nabla U(x_{i\varepsilon}) \right).$$

Substituting this into the inequality and simplifying gives the result. □

The requirement (16) can sometimes be established using convexity arguments. In the exponential family class of Corollary 2.3, for example, setting $x^i := x_{L\varepsilon} + i(x_0 - x_{L\varepsilon})/L$, when $1 \leq \beta < 4/3$ one can show as $|x_0| \rightarrow \infty$ that $\hat{U}_L \xrightarrow{a.s.} (x_0 - x_{L\varepsilon})(\nabla U(x_0) + \nabla U(x_{L\varepsilon}) + 2 \sum_{i=1}^{L-1} \nabla U(x^i)) / (2L)$, the regular trapezium rule estimate for $\int_{x_0}^{x_{L\varepsilon}} \nabla U(z) dz = U(x_0) - U(x_{L\varepsilon})$. Since $\nabla U(x)$ is concave/convex for x positive/negative then the trapezium rule gives an underestimate for the integral as $|x_0| \rightarrow \infty$, and hence the left-hand side of (16) will be positive, while it is also possible to show that the right-hand side is negative in this case (using arguments given in the proof of Corollary 2.3). We omit the details of this.

There is some discussion in [44] of relaxations of (13) to the requirement that $\alpha(x_0, x_\varepsilon) \geq \delta$ for some $\delta > 0$ if $\|x_\varepsilon\| \leq \|x_0\|$, which are also applicable to the HMC case and would relax the inequality (16) to some degree. In essence, the key role of the ‘inwards acceptance’ property (13) (among the class of potentials which satisfy **A2**) is to limit the degree of oscillation in the tails of the density $e^{-U(x)}$, which can potentially mean that too many proposals $x_{L\varepsilon}$ for which the chosen Lyapunov function $V(x_{L\varepsilon})/V(x_0) < 1$ are rejected to establish a geometric bound

of the form (3). Similar requirements to (13) are needed for many Markov chain Monte Carlo methods which rely on the Metropolis–Hastings construction (e.g. [28,43,44]). The issues are discussed in some detail in the case of the Random Walk Metropolis in [28]. It is possible that choosing the more natural (but less pliable) Lyapunov function $V(x) = e^{sU(x)}$ for some $s > 0$ would remove the need for (13) here, owing to the ergodic nature of the proposal kernel. We leave such explorations for future work.

The preceding discussion leads to the following assumption that we require for geometric ergodicity here.

A3 *The chain satisfies the ‘inwards acceptance’ property (13) which can equivalently be formulated as (16).*

Assumptions **A1–A3** together are sufficient to establish a geometric bound.

Proof of Corollary 2.3. Part (ii) is a direct consequence of Theorem 2.2 For part (i), we consider three cases separately.

First consider $\beta \in (1, 2)$, meaning $1 > \beta - 1 > 0$. Since $\nabla U(x) = \alpha\beta \text{sgn}(x)|x|^{\beta-1}$ then **A2** holds. It remains to establish **A3**. We let $x_0 \rightarrow \infty$ but an analogous argument holds as $x_0 \rightarrow -\infty$ by symmetry. Note that $x_\varepsilon = x_0 - \varepsilon^2\alpha\beta \text{sgn}(x_0)|x_0|^{\beta-1}/2 + \varepsilon p_0$ will clearly satisfy $(1 - \delta)x_0 < x_\varepsilon < x_0$ with probability reaching one in the limit for any $\delta > 0$. Similarly $(1 - \delta)x_0 < x_{2\varepsilon} = x_\varepsilon - \varepsilon^2\alpha\beta \text{sgn}(x_0)|x_0|^{\beta-1} - \varepsilon^2\alpha\beta \text{sgn}(x_\varepsilon)|x_\varepsilon|^{\beta-1}/2 + \varepsilon p_0 < x_\varepsilon$ in the same asymptotic regime. Iterating the argument reveals that $(1 - \delta)x_0 < x_{L\varepsilon} < \dots < x_\varepsilon < x_0$ with probability tending to one as $x_0 \rightarrow \infty$. Hence, a.s. the proposal will be ‘inwards’, as will each intermediate point in the trajectory. To establish geometric convergence we must show that these inwards proposals are accepted with probability tending to one as $x_0 \rightarrow \infty$. A Taylor series expansion of the difference in Hamiltonians for large enough x_0 gives

$$H(x_0, p_0) - H(x_{L\varepsilon}, p_{L\varepsilon}) = \frac{L^2\varepsilon^4}{8}(\alpha\beta)^3(\beta - 1)x_0^{3\beta-4} + o(x^{3\beta-4}).$$

Since the leading order term is strictly positive then the result is proved. A detailed derivation is provided in the supplementary material [33].

In the case $\beta = 1$, then as $x_0 \rightarrow \infty$ the proposal in fact a.s. becomes $x_{L\varepsilon} = x_0 - L\varepsilon^2/2 + L\varepsilon p_0$, which resembles that of a random walk with inwards drift. Here the acceptance rate a.s. becomes one as the leapfrog integrator becomes exact provided the zero boundary is not crossed, and hence the scheme is geometrically ergodic following Theorem 16.0.1 and the argument of Section 16.1.3 in Chapter 16 of [37]. Again a similar argument holds as $x_0 \rightarrow -\infty$.

In the case $\beta = 2$ following Example 3.5 in [3], setting $\theta := \arccos(1 - \alpha\varepsilon^2)$ the proposal becomes $x_{L\varepsilon} = \cos(\theta L)x_0 + \sin(\theta L)p_0/\sqrt{2\alpha(1 - \alpha\varepsilon^2/2)}$, which will be inwards provided $|\cos(\theta L)| < 1$, which will be true for suitably small ε . Similarly provided $p_0 = o(x_0)$ then the difference in Hamiltonian values will be

$$H(x_0, p_0) - H(x_{L\varepsilon}, p_{L\varepsilon}) = \left(1 - \cos^2(\theta L) - 2\alpha\left(1 - \alpha\frac{\varepsilon^2}{2}\right)\sin^2(\theta L)\right)x_0^2 + o(x_0^2).$$

The x_0^2 coefficient will be positive provided $(1 + 2\alpha - \alpha^2\varepsilon^2) \sin^2(\theta L) > 0$ which will also be true for small enough ε , hence as $x_0 \rightarrow \pm\infty$ **A3** holds and since **A2** does also then the result is proven. \square

5.3. Necessary conditions for geometric ergodicity

Next, we highlight the importance of the growth assumptions we have made on the potential, by showing two general scenarios in which HMC will *not* produce geometrically ergodic Markov chains.

5.3.1. Light tails

We begin with the case where the gradient term may grow at a faster than linear rate, meaning that the resulting system of equations (8) is ‘stiff’, in the sense that the derivatives can change very rapidly over small time scales, which can pose a challenge to explicit numerical integrators. We show in Theorem 5.14 that in this scenario a Markov chain produced by the HMC method can exhibit undesirable behaviour.

Theorem 5.14. *If it holds that*

$$\lim_{\|x\| \rightarrow \infty} \frac{\|\nabla U(x)\|}{\|x\|} = \infty, \tag{17}$$

and that there is a fixed $C < \infty$ such that whenever $\|y\| \geq 2\|x\| \geq C$ then

$$\|\nabla U(y)\| \geq 3\|\nabla U(x)\|, \tag{18}$$

and it also holds that

$$\lim_{\|x\| \rightarrow \infty, \|y\| \geq 2^L \|x\|} \left(U(x) - U(y) - \frac{1}{2}\|x\|^2 \right) = -\infty, \tag{19}$$

then the Hamiltonian Monte Carlo method with fixed integration time $T = L\varepsilon$ does not produce a geometrically ergodic Markov chain for any choice $T > 0$.

Proof. Lemmas 5.15 and 5.16 below establish that in this case $\|x_{L\varepsilon}\| \geq 2^L \|x_0\|$, and Lemma 5.17 shows that this will result in $\alpha(x_0, x_{L\varepsilon})$ tending to zero as $\|x_0\| \rightarrow \infty$ provided $\|p_0\| \leq \|x_0\|^\delta$ for some $\delta < 1$, allowing Proposition 3.5 to be evoked. To conclude we simply note that $\mathbb{P}(\|p_0\| \leq \|x_0\|^\delta) \rightarrow 1$ as $\|x_0\| \rightarrow \infty$, establishing the result. \square

The conditions (18) and (19) limit the amount that the potential can oscillate as it approaches ∞ , and are introduced to prevent tail oscillations in gradient from making the behaviour of the method too unpredictable to analyse sensibly. They are very lenient and should be satisfied for the vast majority of statistical models of interest for which (17) holds. Below we establish several intermediate results, the first two of which relate to the values of $\|x_{L\varepsilon}\|$ when $\|x_0\|$ is large in this scenario.

Lemma 5.15. *If (17) holds, then there exists an $\eta < \infty$ such that for all $\|x_0\| > \eta$ and any $\|p_0\| \leq \|x_0\|^\delta$ for some $\delta < 1$, it holds that $\|x_\varepsilon\| > 2\|x_0\|$.*

Proof. Taking norms after a single leapfrog step gives

$$\|x_\varepsilon\| = \left\| x_0 - \frac{\varepsilon^2}{2} \nabla U(x_0) + \varepsilon p_0 \right\| \geq \frac{\varepsilon^2}{2} \|\nabla U(x_0)\| - \|x_0\| - \varepsilon \|p_0\|.$$

Dividing by $\|x_0\|$ gives

$$\frac{\|x_\varepsilon\|}{\|x_0\|} \geq \frac{\varepsilon^2}{2} \frac{\|\nabla U(x_0)\|}{\|x_0\|} - 1 - \varepsilon \frac{\|p_0\|}{\|x_0\|}. \tag{20}$$

Using (17), we can choose an x_0 such that the first term on the right-hand side is larger than $6/\varepsilon^2$, and the last term can be made negligibly small as $\|p_0\| \leq \|x_0\|^\delta$ for some $\delta < 1$, which establishes the result. \square

Lemma 5.16. *If (17) holds, then there exists an $\eta < \infty$ such that for all $\|x_0\| > \eta$ and any $\|p_0\| \leq \|x_0\|^\delta$ for some $\delta < 1$, it holds that $\|x_{2\varepsilon}\| \geq 2^L \|x_0\|$.*

Proof. We proceed iteratively. First, note that

$$x_{2\varepsilon} = x_\varepsilon - \varepsilon^2 \nabla U(x_\varepsilon) - \frac{\varepsilon^2}{2} \nabla U(x_0) + \varepsilon p_0.$$

Using this, we have

$$\frac{\|x_{2\varepsilon}\|}{\|x_\varepsilon\|} \geq \frac{\varepsilon^2 \|\nabla U(x_\varepsilon)\|}{2 \|x_\varepsilon\|} - \frac{\varepsilon^2 \|\nabla U(x_0)\|}{2 \|x_\varepsilon\|} - 1 - \varepsilon \frac{\|p_0\|}{\|x_\varepsilon\|}.$$

Showing the right-hand side is ≥ 2 amounts to upper bounding the middle term, or equivalently lower bounding its reciprocal. We have

$$\frac{2\|x_\varepsilon\|}{\varepsilon^2 \|\nabla U(x_0)\|} \geq \frac{\varepsilon^2 \|\nabla U(x_0)\| - 2\|x_0\| - 2\varepsilon \|p_0\|}{\varepsilon^2 \|\nabla U(x_0)\|} \geq 1 - \delta \tag{21}$$

for some $\delta > 0$ which can be made arbitrarily small by choosing $\|x_0\|$ large enough.

Next, we have

$$\frac{\|x_{3\varepsilon}\|}{\|x_{2\varepsilon}\|} \geq \frac{\varepsilon^2 \|\nabla U(x_{2\varepsilon})\|}{2 \|x_{2\varepsilon}\|} - \frac{\varepsilon^2 \|\nabla U(x_0)\|}{2 \|x_{2\varepsilon}\|} - \varepsilon^2 \frac{\|\nabla U(x_\varepsilon)\|}{\|x_{2\varepsilon}\|} - 1 - \varepsilon \frac{\|p_0\|}{\|x_{2\varepsilon}\|}.$$

Here the right-hand side will be ≥ 2 provided the middle two terms can be bounded above. For the first we lower bound the reciprocal, using (21) gives

$$\frac{2}{\varepsilon^2} \frac{\|x_{2\varepsilon}\|}{\|\nabla U(x_0)\|} \geq \frac{2}{\varepsilon^2} \frac{\|x_\varepsilon\|}{\|\nabla U(x_0)\|} \geq 1 - \delta.$$

For the second we have

$$\frac{\|x_{2\varepsilon}\|}{\varepsilon^2 \|\nabla U(x_\varepsilon)\|} \geq 1 - \frac{\|x_\varepsilon\|}{\varepsilon^2 \|\nabla U(x_\varepsilon)\|} - \frac{\|\nabla U(x_0)\|}{2\|\nabla U(x_\varepsilon)\|} - \frac{\|p_0\|}{\varepsilon \|\nabla U(x_\varepsilon)\|}.$$

The second and last terms on the right hand side can be made arbitrarily small by choosing $\|x_0\|$ large enough. Evoking (18) gives

$$\frac{\|\nabla U(x_0)\|}{2\|\nabla U(x_\varepsilon)\|} \leq \frac{1}{6},$$

which therefore shows that $\|x_{3\varepsilon}\| \geq 2\|x_{2\varepsilon}\|$. An entirely analogous argument can be used to show that $\|x_{i\varepsilon}\| \geq 2\|x_{(i-1)\varepsilon}\|$ for any fixed i , establishing the result. \square

The next result shows that as a result of the fact that $\|x_{L\varepsilon}\| \geq 2^L \|x_0\|$ when $\|x_0\|$ is large enough, then the acceptance rate will approach 0 in the limit as $\|x_0\| \rightarrow \infty$.

Lemma 5.17. *If (17), (18) and (19) hold, then for any $\delta < 1$ it holds that*

$$\lim_{\|x_0\| \rightarrow \infty, \|p_0\| \leq \|x_0\|^\delta} \alpha(x_0, x_{L\varepsilon}) = 0.$$

Proof. Recall that

$$\alpha(x_0, x_{L\varepsilon}) = 1 \wedge \exp\left(U(x_0) - U(x_{L\varepsilon}) + \frac{1}{2}\|p_0\|^2 - \frac{1}{2}\|p_{L\varepsilon}\|^2\right).$$

Note that

$$\begin{aligned} \|p_{L\varepsilon}\| &\geq \frac{\varepsilon}{2} \|\nabla U(x_{L\varepsilon})\| - \varepsilon \sum_{i=1}^{L-1} \|\nabla U(x_{i\varepsilon})\| - \frac{\varepsilon}{2} \|\nabla U(x_0)\| - \|p_0\| \\ &\geq \frac{\varepsilon}{2} \|\nabla U(x_{L\varepsilon})\| - \varepsilon \sum_{i=1}^L \left(\frac{1}{3}\right)^i \|\nabla U(x_{L\varepsilon})\| - \|p_0\| \\ &= \frac{\varepsilon}{2} \left(1 - 2 \sum_{i=1}^L \left(\frac{1}{3}\right)^i\right) \|\nabla U(x_{L\varepsilon})\| - \|p_0\|, \end{aligned}$$

where (18) is used for the second line. The term inside the bracket can be bounded below by some fixed constant $\gamma_L > 0$, for any fixed $L < \infty$. Squaring the result gives

$$\begin{aligned} \|p_{L\varepsilon}\|^2 &\geq \left(\frac{\varepsilon\gamma_L}{2} \|\nabla U(x_{L\varepsilon})\| - \|p_0\|\right)^2 \\ &= \frac{\varepsilon^2\gamma_L^2}{4} \|\nabla U(x_{L\varepsilon})\|^2 + \|p_0\|^2 - \varepsilon\delta_L \|\nabla U(x_{L\varepsilon})\| \|p_0\|, \end{aligned}$$

which implies that

$$\|p_0\|^2 - \|p_{L\varepsilon}\|^2 \leq \varepsilon\gamma_L \|\nabla U(x_{L\varepsilon})\| \left(\|p_0\| - \frac{\varepsilon\gamma_L}{4} \|\nabla U(x_{L\varepsilon})\| \right).$$

Noting that $\|p_0\| \leq \|x_0\|^\delta$ and that for any $M < \infty$ we can choose an $\|x_0\|$ large enough that

$$\frac{\varepsilon\gamma_L}{4} \|\nabla U(x_{L\varepsilon})\| \geq M\|x_{L\varepsilon}\| \geq 2^L\|x_0\|,$$

then it follows that $\|p_0\|^2 - \|p_{L\varepsilon}\|^2 \leq -\|x_0\|^2$. Using this, then simply evoking (19) gives the result. \square

5.3.2. Heavy tails

In the case where $\pi(x)$ has ‘heavier than exponential’ tails in some direction the HMC method can also exhibit slow convergence, as $\liminf_{\|x\| \rightarrow \infty} \|\nabla U(x)\| = 0$. Intuitively the problem here is that when $\|x\|$ is large then the gradient provides insufficient drift back into the ‘centre’ of the space, meaning the chain can exhibit random walk behaviour and hence convergence can be very slow. Theorem 5.18 makes this intuition rigorous.

Theorem 5.18. *If $\|\nabla U(x)\| < M$ for all $x \in \mathbf{X}$, then a necessary condition for the Hamiltonian Monte Carlo method to produce a geometrically ergodic Markov chain is*

$$\int e^{s\|x\|} \pi(dx) < \infty$$

for some $s > 0$.

Proof. From Proposition 3.6, it is sufficient to show that for any $\varepsilon > 0$ there is a $\delta > 0$ such that $Q(x, B_\delta(x)) > 1 - \varepsilon$ for all $x \in \mathbf{X}$. Using equation (9) if x_0 is the current point in the chain then

$$\|x_{L\varepsilon} - x_0\| = \left\| L\varepsilon p_0 - \frac{L\varepsilon^2}{2} \nabla U(x_0) - \varepsilon^2 \sum_{i=1}^{L-1} (L-i) \nabla U(x_{i\varepsilon}) \right\|.$$

Applying the triangle inequality and then the global bound on $\|\nabla U(x)\|$ gives

$$\|x_{L\varepsilon} - x_0\| \leq \frac{L\varepsilon^2}{2} M + \frac{M\varepsilon^2 L(L-1)}{2} + L\varepsilon\|p_0\|.$$

As p_0 follows a centred Gaussian distribution with fixed covariance then Chebyshev’s inequality gives the result. \square

In fact, in this case the lack of geometric ergodicity is a property of the flow itself, rather than being a consequence of numerical instabilities as in Theorem 5.14, as shown by the following result.

Proposition 5.19. *Theorem 5.18 still holds even if an exact integrator is available for Hamilton's equations.*

Proof. Using Hamilton's equations, we have

$$x_T - x_0 = \int_0^T p_s ds = \int_0^T \left[p_0 - \int_0^s \nabla U(x_u) du \right] ds. \quad (22)$$

Taking the norm and using the upper bound gives

$$\|x_T - x_0\| \leq T \|p_0\| - \int_0^T \int_0^s \|\nabla U(x_u)\| du ds \leq T \|p_0\| + CT^2/2,$$

and again Chebyshev's inequality gives the result. \square

The class of models for which $\|\nabla U(x)\|$ is bounded and $\int e^s \|x\| \pi(dx) < \infty$ for some $s > 0$ is comparatively narrow, essentially comprising $U(x) = C\|x\| + b(x)$, where $C < \infty$ and $b: \mathbf{X} \rightarrow \mathbb{R}$ is some appropriately regular function which is bounded both above and below.

6. Results for an position-dependent integration time

An important free parameter in HMC is the integration time T , which we have previously assumed to be independent of the current position. The representation (9) does however suggest that allowing this to change can have some benefits. If the candidate map is viewed as

$$x_{L\varepsilon} = x_0 + \text{DRIFT}(x_0, p_0, T) + L\varepsilon p_0,$$

then if the 'DRIFT' function becomes negligible for large $\|x_0\|$ and fixed T , then it can be increased in magnitude by making T larger. We make this simple intuition rigorous for an idealised algorithm on the particular one-dimensional Exponential Family class of models with densities of the form

$$\pi(x) \propto \exp(-\beta^{-1}(1+x^2)^{\beta/2}), \quad (23)$$

for some fixed $\beta > 0$. Here any contour $C_{x_0, p_0} := \{(x, p) : H(x, p) = H(x_0, p_0)\}$ consists of a single closed path, and the flow is periodic from any fixed starting point. We additionally assume that the period length $\zeta_{x_0, p_0} > 0$ is known, and that we have an exact integrator for Hamilton's equations. This means that we need not concern ourselves with the acceptance probability (we discuss this issue in Section 7).

At iteration i (with $x_0 = x_{i-1}$), the *dynamic* HMC implementation we consider consists of resampling $p_0 \sim N(0, 1)$, and then setting $x_i = \text{Pr}_x \circ \varphi_\tau(x_0, p_0)$, where $\tau \sim U[0, \zeta_{x_0, p_0}]$. In words, we flow along the Hamiltonian for τ units of time, where τ is a uniform random variable with maximum value ζ_{x_0, p_0} (note that $\varphi_{\zeta_{x_0, p_0}}(x_0, p_0) = (x_0, p_0)$).

Firstly, note that π -irreducibility is more straightforward to see here. To reach any set $A \in \mathcal{B}$ with $\pi(A) > 0$, we first consider the single contour C_{x_0, p_0} , and specifically the component of this

contour that is connected to (x_0, p_0) . Let C_{x_0} be the projection of this component onto \mathbf{X} . Then any nonempty set $A' \subset C_{x_0}$ has positive probability of occurring, as the next point is chosen from a density with support all of C_{x_0} . As the contours are composed of single components, and cover the entire space, then for any A , the probability of choosing a contour for which this argument can be applied is greater than zero. We provide a figure in the supplementary material to offer more intuition [33].

We introduce some additional notation in this section. We define the *microcanonical* expectation of a real-valued function $f(x_t, p_t)$, where $(x_t, p_t) = \varphi_t(x_0, p_0)$, i.e. the solution to (8) for t units of time initialised at (x_0, p_0) , as

$$\langle f(x_0, p_0) \rangle := \zeta_{x_0, p_0}^{-1} \int_0^{\zeta_{x_0, p_0}} f(x_s, p_s) ds. \tag{24}$$

This is simply the time expectation of f from uniformly sampling across C_{x_0, p_0} .

We first introduce a result from the Physics literature (e.g., [24]) which relates the kinetic and potential energies.

Theorem 6.1 (Virial Theorem). *Under Hamiltonian flow $(x_s, p_s) = \varphi_s(x_0, p_0)$ we have*

$$\langle x_0 \nabla U(x_0) \rangle = \langle p_0^2 \rangle. \tag{25}$$

Proof. Define the *virial* function $G_t = x_t p_t$. From the fundamental theorem of Calculus, we have

$$\langle \dot{G}_0 \rangle = \frac{G_{\zeta_{x_0, p_0}} - G_0}{\zeta_{x_0, p_0}} = 0,$$

where $\dot{G}_t := dG_t/dt$. In this case

$$\dot{G}_t = x_t \dot{p}_t + p_t \dot{x}_t = -x_t \nabla U(x_t) + p_t^2,$$

meaning

$$\langle x_0 \nabla U(x_0) \rangle = \langle p_0^2 \rangle,$$

as required. □

We can now state and prove the main result of this section.

Theorem 6.2. *For the one-dimensional Exponential Family class of distributions with density given by (23), the dynamic Hamiltonian Monte Carlo method produces a geometrically ergodic Markov chain for any value of $\beta > 0$.*

Proof. Note that by conservation of the Hamiltonian, we have

$$\int \langle U(x_0) + p_0^2/2 \rangle \mu^G(dp_0) = \int H(x_0, p_0) \mu^G(dp_0) = U(x_0) + 1/2. \tag{26}$$

Choose the Lyapunov function $V(x) = U(x) + xU'(x) + 1$. Using Theorem 6.1, we can rewrite the above expression

$$PV(x_0) = U(x_0) + 3/2.$$

Note also that for any $\eta > 0$ there is an $M_\eta < \infty$ such that whenever $|x_0| > M_\eta$

$$(1 + \eta)U(x_0) \geq U(x_0) + 3/2.$$

The proof will be complete if we can find a $\lambda < 1$ such that $(1 + \eta)U(x_0) \leq \lambda V(x_0)$ for suitably large $|x_0|$. Now $x_0 U'(x_0) \rightarrow \beta U(x_0)$ here as $|x_0| \rightarrow \infty$, meaning that there is an $M < \infty$ such that whenever $|x_0| > M$

$$x_0 U'(x_0)/2 \geq \beta U(x_0)/2 - 1.$$

Taking $|x_0| \geq \max(M_\eta, M)$ we can therefore re-write the inequality of interest $(1 + \eta)U(x_0) \leq \lambda V(x_0)$ as

$$(1 + \eta)U(x_0) \leq \lambda(1 + \beta/2)U(x),$$

which will be true if

$$\lambda \geq \frac{1 + \eta}{1 + \beta/2}.$$

Choosing $\eta < \beta/2$ ensures $\lambda < 1$ and also gives the desired inequality $PV(x_0) \leq \lambda V(x_0)$ whenever $|x_0| > \max(M_\eta, M)$, showing that the resulting Markov chain will be geometrically ergodic. \square

7. Discussion

We have established conditions under which geometric ergodicity will and will not hold for Markov chains produced by the Hamiltonian Monte Carlo method. Here we discuss how our results can be extended in various ways, as well as how they translate to standard implementations in widely used software [14].

7.1. Dynamic implementations

Allowing the integration time in HMC to depend on the current point in the chain without an exact integrator will typically mean that some adjustments to $\alpha(x_0, x_T)$ must be made to ensure that $\pi(\cdot)$ is still preserved. The reason is that the approximate flow map φ_T may no longer be reversible, as if $T_1 := T(x_0, p_0)$ and $T_2 := T(x_T, p_T)$ then $\varphi_{T_2}^{-1} \circ \varphi_{T_1}$ will typically not be the identity map if $T_1 \neq T_2$. The two possible ways of changing the integration time $T = L\varepsilon$ are to adjust either L or ε . Increasing L requires more computations per transition, while this is not necessarily true for ε . In the No-U-Turn sampler a binary tree approach is introduced to ensure preservation of detailed balance when L is altered in different parts of the space [26]. We are not aware of any implementations involving adjustment of ε , however it is likely that similar

modifications to $\alpha(x_0, x_T)$ are possible here also. Adjusting ε may be a sensible option in some cases, as the leapfrog method is known to ‘almost’ preserve the modified Hamiltonian

$$\tilde{H}(x, p) = H(x, p) + \left(\frac{1}{12} p^t \nabla^t \nabla U(x) p - \frac{1}{24} \nabla U(x)^t \nabla U(x) \right) \varepsilon^2 + O(\varepsilon^4),$$

as shown for example in [32]. When $\pi(x)$ is not log-concave in the tails and hence the elements of ∇U and $\nabla^t \nabla U$ become negligible as $\|x\| \rightarrow \infty$, this implies that ε can be increased for larger $\|x\|$ without compromising on numerical accuracy.

7.2. Extension to other integrators

The fixed integration time results in Section 5 refer specifically to the leapfrog integrator implementation of HMC (aside from Proposition 5.19). It should be possible to use the same approach when analysing other explicit symplectic integrators, however for schemes which rely on implicit methods (e.g., [23]) then composing multiple steps of the integrator as in Proposition 4.2 cannot be done so cleanly. Implicit methods are needed when the Hamiltonian is nonseparable, and can often resolve stiffness issues such as those characterised in Theorem 5.14.

To construct ergodicity results for the most general version of Hamiltonian Monte Carlo (i.e. without restricting attention to \mathbb{R}^n) we note that there are many ways to construct drift conditions in line with the purely geometric framework introduced in [7]. We also point out that for the one-dimensional Exponential Family, choosing the Riemannian metric $G(x) = \|\nabla^2 U(x)\|$ and employing the approach of [23] is mathematically equivalent to applying the transformation $x' = \text{sgn}(x)\|x\|^{\beta/2}$ with corresponding density

$$\pi(x') \propto \|x'\|^{2/(\beta-1)} \exp(-x'^2/\beta^2).$$

This new density will have Gaussian tails for any $\beta > 0$, suggesting a well-behaved sampler can be constructed. Further discussion on the relationship between geometric Markov chain Monte Carlo methods and parameter transformations is given in [34].

7.3. Honest bounds

Geometric ergodicity is often called a *qualitative* bound, as an explicit upper bound on the geometric rate ρ is not established when using the techniques of [43]. With some modifications, however, quantitative bounds can be constructed (e.g., [30]). We have refrained from doing this here, as these bounds are also often too conservative to be of use in practice [30].

Monte Carlo estimates for nonasymptotic quantitative bounds using the Ricci curvature approach of [39] are applied to Hamiltonian Monte Carlo in [45]. We note that the applicability of these bounds relies on the assumption of *positive curvature* in some Wasserstein distance for the underlying Markov chain. When this distance is chosen to be Total Variation, then this is a strictly stronger condition than geometric ergodicity [see Corollary 22 in [39]], so we feel that our results are a useful pre-cursor to understanding when these estimated bounds are informative in practice.

In the case of MALA, when $\|\nabla U(x)\|$ grows at a faster than linear rate for large $\|x\|$ then it is shown in [8] that useful inferences for functionals concentrated in the centre of the space can be made by setting a small enough value for ε . It is likely that the same analysis can be done with HMC, and that the result would be similar, but we leave such explorations for future work.

7.4. Practitioner guidelines

The main conclusion of our work for practitioners implementing the method in a bespoke manner is to consider the form of $\|\nabla U(x)\|$. If this term either grows very fast or becomes negligibly small when $\|x\|$ is large then it is likely that the Markov chains produced will struggle to explore the tails of $\pi(\cdot)$ effectively. When the gradient grows at a faster than linear rate then a suitably small value for ε must be chosen to counteract this, while when it shrinks then the integration time T must be made sufficiently large. Of course in either scenario if there is a re-parametrisation of the model that may not suffer these difficulties then this should be applied. Users implementing the method in the Stan software [14] should note that both of these instances are captured by standard output diagnostics. Numerical trajectories that become unstable due to large gradients are classed as ‘divergences’, while a failure to move far enough because of negligible gradients is recorded through the ‘maximum tree depth reached’ warning. If this happens and $\pi(\cdot)$ is known to be proper, then the user should set as large a maximum tree depth as is computationally feasible when tail exploration is of keen interest.

Acknowledgements

We thank Alexandros Beskos, Gareth Roberts, Krzysztof Łatuszyński, Gabriel Stoltz and Mark Rowland for useful discussions. SL thanks Nawaf Bou-Rabee for pointing him to [13].

SL was supported by a PhD scholarship from Xerox Research Centre Europe and EPSRC grant EP/K014463/1 for this project. SB was supported by EPSRC fellowship EP/K005723/1, MB is funded by EPSRC grant EP/J016934/1, and SB and MB also received a 2014 EPSRC NCSML Award for PDRA Collaboration for this project. MG is funded by an EPSRC Established Career Research Fellowship, EP/J016934/1, a Royal Society Wolfson Research Merit Award, and EPSRC grants EP/P020720/1, EP/J016934/3, EP/K034154/1.

Supplementary Material

Supplement to “On the geometric ergodicity of Hamiltonian Monte Carlo” (DOI: [10.3150/18-BEJ1083SUPP](https://doi.org/10.3150/18-BEJ1083SUPP); .pdf). We provide additional examples of π -irreducibility, with supporting plots, as well as elaborating on the connections between HMC and Langevin dynamics and giving a more detailed proof of inwards convergence for the exponential family model class.

References

- [1] Alder, B.J. and Wainwright, T.E. (1959). Studies in molecular dynamics. I. General method. *J. Chem. Phys.* **31** 459–466. [MR0115313](https://doi.org/10.1063/1.1740354)

- [2] Andrieu, C., De Freitas, N., Doucet, A. and Jordan, M.I. (2003). An introduction to MCMC for machine learning. *Mach. Learn.* **50** 5–43.
- [3] Beskos, A., Pillai, N., Roberts, G., Sanz-Serna, J.-M. and Stuart, A. (2013). Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli* **19** 1501–1534. [MR3129023](#)
- [4] Beskos, A., Pinski, F.J., Sanz-Serna, J.M. and Stuart, A.M. (2011). Hybrid Monte Carlo on Hilbert spaces. *Stochastic Process. Appl.* **121** 2201–2230. [MR2822774](#)
- [5] Betancourt, M. (2013). A general metric for Riemannian manifold Hamiltonian Monte Carlo. In *Geometric Science of Information. Lecture Notes in Computer Science* **8085** 327–334. Heidelberg: Springer. [MR3126061](#)
- [6] Betancourt, M. (2016). Identifying the Optimal Integration Time in Hamiltonian Monte Carlo. Preprint. Available at [arXiv:1601.00225](#).
- [7] Betancourt, M., Byrne, S., Livingstone, S. and Girolami, M. (2017). The geometric foundations of Hamiltonian Monte Carlo. *Bernoulli* **23** 2257–2298. [MR3648031](#)
- [8] Bou-Rabee, N. and Hairer, M. (2013). Nonasymptotic mixing of the MALA algorithm. *IMA J. Numer. Anal.* **33** 80–110. [MR3020951](#)
- [9] Bou-Rabee, N. and Sanz-Serna, J.M. (2017). Randomized Hamiltonian Monte Carlo. *Ann. Appl. Probab.* **27** 2159–2194. Available at [arXiv:1511.09382](#). [MR3693523](#)
- [10] Bou-Rabee, N. and Sanz-Serna, J.M. (2018). Geometric integrators and the Hamiltonian Monte Carlo method. *Acta Numer.* **27** 113–206. [MR3826507](#)
- [11] Byrne, S. and Girolami, M. (2013). Geodesic Monte Carlo on embedded manifolds. *Scand. J. Stat.* **40** 825–845. [MR3145120](#)
- [12] Campos, C.M. and Sanz-Serna, J.M. (2015). Extra chance generalized hybrid Monte Carlo. *J. Comput. Phys.* **281** 365–374. [MR3281978](#)
- [13] Cancès, E., Legoll, F. and Stoltz, G. (2007). Theoretical and numerical comparison of some sampling methods for molecular dynamics. *ESAIM Math. Model. Numer. Anal.* **41** 351–389. [MR2339633](#)
- [14] Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M.A., Guo, J., Li, P. and Riddell, A. (2016). Stan: A probabilistic programming language. *J. Stat. Softw.*
- [15] Diaconis, P. (2013). Some things we’ve learned (about Markov chain Monte Carlo). *Bernoulli* **19** 1294–1305. [MR3102552](#)
- [16] Diaconis, P. and Freedman, D. (1999). Iterated random functions. *SIAM Rev.* **41** 45–76. [MR1669737](#)
- [17] Diaconis, P., Seiler, C. and Holmes, S. (2014). Connections and extensions: A discussion of the paper by Girolami and Byrne [[MR3145120](#)]. *Scand. J. Stat.* **41** 3–7. [MR3181123](#)
- [18] Duane, S., Kennedy, A.D., Pendleton, B.J. and Roweth, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B* **195** 216–222.
- [19] Durmus, A. and Moulines, É. (2015). Quantitative bounds of convergence for geometrically ergodic Markov chain in the Wasserstein distance with application to the Metropolis adjusted Langevin algorithm. *Stat. Comput.* **25** 5–19. [MR3304898](#)
- [20] Durmus, A., Moulines, E. and Saksman, E. (2017). On the convergence of Hamiltonian Monte Carlo. Preprint. Available at [arXiv:1705.00166](#).
- [21] Eberle, A. (2014). Error bounds for Metropolis–Hastings algorithms applied to perturbations of Gaussian measures in high dimensions. *Ann. Appl. Probab.* **24** 337–377. [MR3161650](#)
- [22] Gelman, A., Carlin, J.B., Stern, H.S. and Rubin, D.B. (2014). *Bayesian Data Analysis* **2**. Taylor & Francis.
- [23] Girolami, M. and Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **73** 123–214. With discussion and a reply by the authors. [MR2814492](#)
- [24] Goldstein, H. (1965). *Classical Mechanics*. Pearson Education India.

- [25] Hastings, W.K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109. [MR3363437](#)
- [26] Hoffman, M.D. and Gelman, A. (2014). The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15** 1593–1623. [MR3214779](#)
- [27] Horowitz, A.M. (1991). A generalized guided Monte Carlo algorithm. *Phys. Lett. B* **268** 247–252.
- [28] Jarner, S.F. and Hansen, E. (2000). Geometric ergodicity of Metropolis algorithms. *Stochastic Process. Appl.* **85** 341–361. [MR1731030](#)
- [29] Jarner, S.F. and Tweedie, R.L. (2003). Necessary conditions for geometric and polynomial ergodicity of random-walk-type Markov chains. *Bernoulli* **9** 559–578. [MR1996270](#)
- [30] Jones, G.L. and Hobert, J.P. (2001). Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statist. Sci.* **16** 312–334. [MR1888447](#)
- [31] Lee, J.M. (2012). *Introduction to Smooth Manifolds*. Springer.
- [32] Leimkuhler, B. and Reich, S. (2004). *Simulating Hamiltonian Dynamics*. Cambridge Monographs on Applied and Computational Mathematics **14**. Cambridge: Cambridge Univ. Press. [MR2132573](#)
- [33] Livingstone, S., Betancourt, M., Byrne, S. and Girolami, M. (2018). Supplement to “On the geometric ergodicity of Hamiltonian Monte Carlo.” DOI:[10.3150/18-BEJ1083SUPP](#).
- [34] Livingstone, S. and Girolami, M. (2014). Information-geometric Markov chain Monte Carlo methods using diffusions. *Entropy* **16** 3074–3102. [MR3234224](#)
- [35] Mattingly, J.C., Stuart, A.M. and Higham, D.J. (2002). Ergodicity for SDEs and approximations: Locally Lipschitz vector fields and degenerate noise. *Stochastic Process. Appl.* **101** 185–232. [MR1931266](#)
- [36] Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H. and Teller, E. (1953). Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21** 1087–1092.
- [37] Meyn, S.P. and Tweedie, R.L. (2012). *Markov Chains and Stochastic Stability*. Springer Science & Business Media.
- [38] Neal, R.M. (2011). MCMC using Hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handb. Mod. Stat. Methods 113–162. Boca Raton, FL: CRC Press. [MR2858447](#)
- [39] Ollivier, Y. (2009). Ricci curvature of Markov chains on metric spaces. *J. Funct. Anal.* **256** 810–864. [MR2484937](#)
- [40] Ottobre, M., Pillai, N.S., Pinski, F.J. and Stuart, A.M. (2016). A function space HMC algorithm with second order Langevin diffusion limit. *Bernoulli* **22** 60–106. [MR3449777](#)
- [41] Roberts, G.O. and Rosenthal, J.S. (1997). Geometric ergodicity and hybrid Markov chains. *Electron. Commun. Probab.* **2** 13–25. [MR1448322](#)
- [42] Roberts, G.O. and Rosenthal, J.S. (2004). General state space Markov chains and MCMC algorithms. *Probab. Surv.* **1** 20–71. [MR2095565](#)
- [43] Roberts, G.O. and Tweedie, R.L. (1996). Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika* **83** 95–110. [MR1399158](#)
- [44] Roberts, G.O. and Tweedie, R.L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363. [MR1440273](#)
- [45] Seiler, C., Rubinstein-Salzedo, S. and Holmes, S. (2014). Positive curvature and Hamiltonian Monte Carlo. In *Advances in Neural Information Processing Systems* 586–594.
- [46] Stuart, A.M. (2010). Inverse problems: A Bayesian perspective. *Acta Numer.* **19** 451–559. [MR2652785](#)
- [47] Tierney, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.* **22** 1701–1762. With discussion and a rejoinder by the author. [MR1329166](#)
- [48] Tierney, L. (1998). A note on Metropolis–Hastings kernels for general state spaces. *Ann. Appl. Probab.* **8** 1–9. [MR1620401](#)

Received November 2017 and revised September 2018