

Article

Geometric Ergodicity of the Random Walk Metropolis with Position-Dependent Proposal Covariance

Samuel Livingstone 

Department of Statistical Science, University College London, London WC1E 6BT, UK;
samuel.livingstone@ucl.ac.uk

Abstract: We consider a Metropolis–Hastings method with proposal $\mathcal{N}(x, hG(x)^{-1})$, where x is the current state, and study its ergodicity properties. We show that suitable choices of $G(x)$ can change these ergodicity properties compared to the Random Walk Metropolis case $\mathcal{N}(x, h\Sigma)$, either for better or worse. We find that if the proposal variance is allowed to grow unboundedly in the tails of the distribution then geometric ergodicity can be established when the target distribution for the algorithm has tails that are heavier than exponential, in contrast to the Random Walk Metropolis case, but that the growth rate must be carefully controlled to prevent the rejection rate approaching unity. We also illustrate that a judicious choice of $G(x)$ can result in a geometrically ergodic chain when probability concentrates on an ever narrower ridge in the tails, something that is again not true for the Random Walk Metropolis.

Keywords: Monte Carlo; MCMC; Markov chains; computational statistics; bayesian inference



Citation: Livingstone, S. Geometric Ergodicity of the Random Walk Metropolis with Position-Dependent Proposal Covariance. *Mathematics* **2021**, *9*, 341. <https://doi.org/10.3390/math9040341>

Academic Editor :
Panagiotis-Christos Vassiliou
Received: 19 January 2021
Accepted: 4 February 2021
Published: 8 February 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Markov chain Monte Carlo (MCMC) methods are techniques for estimating expectations with respect to some probability measure $\pi(\cdot)$, which need not be normalised. This is done by sampling a Markov chain which has limiting distribution $\pi(\cdot)$, and computing empirical averages. A popular form of MCMC is the Metropolis–Hastings algorithm [1,2], where at each time step a ‘proposed’ move is drawn from some candidate distribution, and then accepted with some probability, otherwise the chain stays at the current point. Interest lies in finding choices of candidate distribution that will produce sensible estimators for expectations with respect to $\pi(\cdot)$.

The quality of these estimators can be assessed in many different ways, but a common approach is to understand conditions on $\pi(\cdot)$ that will result in a chain which converges to its limiting distribution at a *geometric* rate. If such a rate can be established, then a Central Limit Theorem will exist for expectations of functionals with finite second absolute moment under $\pi(\cdot)$ if the chain is reversible.

A simple yet often effective choice is a symmetric candidate distribution centred at the current point in the chain (with a fixed variance), resulting in the *Random Walk Metropolis* (RWM) (e.g., [3]). The convergence properties of a chain produced by the RWM are well-studied. In one dimension, essentially convergence is geometric if $\pi(x)$ decays at an exponential or faster rate in the tails [4], while in higher dimensions an additional curvature condition is required [5]. Slower rates of convergence have also been established in the case of heavier tails [6].

Recently, some MCMC methods were proposed which generalise the RWM, whereby proposals are still centred at the current point x and symmetric, but the variance changes with x [7–11]. An extension to infinite-dimensional Hilbert spaces is also suggested in Reference [12]. The motivation is that the chain can become more ‘local’, perhaps making larger jumps when out in the tails, or mimicking the local dependence structure of $\pi(\cdot)$ to propose more intelligent moves. Designing MCMC methods of this nature is particularly relevant for modern Bayesian inference problems, where posterior distributions are

often high dimensional and exhibit nonlinear correlations [13]. We term this approach the *Position-dependent Random Walk Metropolis* (PDRWM), although technically this is a misnomer, since proposals are no longer random walks. Other choices of candidate distribution designed with distributions that exhibit nonlinear correlations were introduced in Reference [13]. Although powerful, these require derivative information for $\log \pi(x)$, something which can be unavailable in modern inference problems (e.g., [14]). We note that no such information is required for the PDRWM, as shown by the particular cases suggested in References [7–11]. However, there are relations between the approaches, to the extent that understanding how the properties of the PDRWM differ from the standard RWM should also aid understanding of the methods introduced in Reference [13].

In this article, we consider the convergence rate of a Markov chain generated by the PDRWM to its limiting distribution. Our main interest lies in whether this generalisation can change these *ergodicity* properties compared to the standard RWM with fixed covariance. We focus on the case in which the candidate distribution is Gaussian, and illustrate that such changes can occur in several different ways, either for better or worse. Our aim is not to give a complete characterisation of the approach, but rather to illustrate the possibilities through carefully chosen examples, which are known to be indicative of more general behaviour.

In Section 2 necessary concepts about Markov chains are briefly reviewed, before the PDRWM is introduced in Section 3. Some results in the one-dimensional case are given in Section 4, before a higher-dimensional model problem is examined in Section 5. Throughout $\pi(\cdot)$ denotes a probability measure (we use the terms probability measure and distribution synonymously), and $\pi(x)$ its density with respect to Lebesgue measure dx .

Since an early version of this work appeared online, some contributions to the literature were made that are worthy of mention. A Markov kernel constructed as a state-dependent mixture is introduced in Reference [15] and its properties are studied in some cases that are similar in spirit to the model problem of Section 5. An algorithm called *Directional Metropolis–Hastings*, which encompasses a specific instance of the PDRWM, is introduced and studied in Reference [16], and a modification of the same idea is used to develop the *Hop* kernel within the *Hug and Hop* algorithm of Reference [17]. Kamatani considers an algorithm designed for the infinite-dimensional setting in Reference [18] of a similar design to that discussed in Reference [12] and studies the ergodicity properties.

2. Markov Chains and Geometric Ergodicity

We will work on the Borel space $(\mathcal{X}, \mathcal{B})$, with $\mathcal{X} \subset \mathbb{R}^d$ for some $d \geq 1$, so that each $X_t \in \mathcal{X}$ for a discrete-time Markov chain $\{X_t\}_{t \geq 0}$ with time-homogeneous transition kernel $P : \mathcal{X} \times \mathcal{B} \rightarrow [0, 1]$, where $P(x, A) = \mathbb{P}[X_{i+1} \in A | X_i = x]$ and $P^n(x, A)$ is defined similarly for X_{i+n} . All chains we consider will have invariant distribution $\pi(\cdot)$, and be both π -irreducible and aperiodic, meaning $\pi(\cdot)$ is the limiting distribution from π -almost any starting point [19]. We use $|\cdot|$ to denote the Euclidean norm.

In Markov chain Monte Carlo the objective is to construct estimators of $\mathbb{E}_\pi[f]$, for some $f : \mathcal{X} \rightarrow \mathbb{R}$, by computing

$$\hat{f}_n = \frac{1}{n} \sum_{i=1}^n f(X_i), \quad X_i \sim P^i(x_0, \cdot).$$

If $\pi(\cdot)$ is the limiting distribution for the chain then P will be *ergodic*, meaning $\hat{f}_n \xrightarrow{a.s.} \mathbb{E}_\pi[f]$ from π -almost any starting point. For finite n the quality of \hat{f}_n intuitively depends on how quickly $P^n(x, \cdot)$ approaches $\pi(\cdot)$. We call the chain *geometrically ergodic* if

$$\|P^n(x, \cdot) - \pi(\cdot)\|_{TV} \leq M(x)\rho^n, \quad (1)$$

from π -almost any $x \in \mathcal{X}$, for some $M > 0$ and $\rho < 1$, where $\|\mu(\cdot) - \nu(\cdot)\|_{TV} := \sup_{A \in \mathcal{B}} |\mu(A) - \nu(A)|$ is the total variation distance between distributions $\mu(\cdot)$ and $\nu(\cdot)$ [19].

For π -reversible Markov chains geometric ergodicity implies that if $\mathbb{E}_\pi[f^2] < \infty$ for some $f : \mathcal{X} \rightarrow \mathbb{R}$, then

$$\sqrt{n}(\hat{f}_n - \mathbb{E}_\pi[f]) \xrightarrow{d} \mathcal{N}(0, v(P, f)), \quad (2)$$

for some asymptotic variance $v(P, f)$ [20]. Equation (2) enables the construction of asymptotic confidence intervals for \hat{f}_n .

In practice, geometric ergodicity does not guarantee that \hat{f}_n will be a sensible estimator, as $M(x)$ can be arbitrarily large if the chain is initialised far from the typical set under $\pi(\cdot)$, and ρ may be very close to 1. However, chains which are not geometrically ergodic can often either get ‘stuck’ for a long time in low-probability regions or fail to explore the entire distribution adequately, sometimes in ways that are difficult to diagnose using standard MCMC diagnostics.

Establishing Geometric Ergodicity

It is shown in Chapter 15 of Reference [21] that Equation (1) is equivalent to the condition that there exists a *Lyapunov* function $V : \mathcal{X} \rightarrow [1, \infty)$ and some $\lambda < 1, b < \infty$ such that

$$PV(x) \leq \lambda V(x) + b\mathbb{I}_C(x), \quad (3)$$

where $PV(x) := \int V(y)P(x, dy)$. The set $C \subset \mathcal{X}$ must be *small*, meaning that for some $m \in \mathbb{N}, \varepsilon > 0$ and probability measure $\nu(\cdot)$

$$P^m(x, A) \geq \varepsilon \nu(A), \quad (4)$$

for any $x \in C$ and $A \in \mathcal{B}$. Equations (3) and (4) are referred to as *drift* and *minorisation* conditions. Intuitively, C can be thought of as the centre of the space, and Equation (3) ensures that some one dimensional projection of $\{X_t\}_{t \geq 0}$ drifts towards C at a geometric rate when outside. In fact, Equation (3) is sufficient for the return time distribution to C to have geometric tails [21]. Once in C , (4) ensures that with some probability the chain forgets its past and hence *regenerates*. This regeneration allows the chain to couple with another initialised from $\pi(\cdot)$, giving a bound on the total variation distance through the *coupling inequality* (e.g., [19]). More intuition is given in Reference [22].

Transition kernels considered here will be of the *Metropolis–Hastings* type, given by

$$P(x, dy) = \alpha(x, y)Q(x, dy) + r(x)\delta_x(dy), \quad (5)$$

where $Q(x, dy) = q(y|x)dy$ is some candidate kernel, α is called the acceptance rate and $r(x) = 1 - \int \alpha(x, y)Q(x, dy)$. Here we choose

$$\alpha(x, y) = 1 \wedge \frac{\pi(y)q(x|y)}{\pi(x)q(y|x)}, \quad (6)$$

where $a \wedge b$ denotes the minimum of a and b . This choice implies that P satisfies detailed balance for $\pi(\cdot)$ [23], and hence the chain is π -reversible (note that other choices for α can result in non-reversible chains, see Reference [24] for details).

Roberts and Tweedie [5], following on from Reference [21], introduced the following regularity conditions.

Theorem 1. (Roberts and Tweedie). Suppose that $\pi(x)$ is bounded away from 0 and ∞ on compact sets, and there exists $\delta_q > 0$ and $\varepsilon_q > 0$ such that for every x

$$|x - y| \leq \delta_q \Rightarrow q(y|x) \geq \varepsilon_q.$$

Then the chain with kernel (5) is μ^{Leb} -irreducible and aperiodic, and every nonempty compact set is small.

For the choices of Q considered in this article these conditions hold, and we will restrict ourselves to forms of $\pi(x)$ for which the same is true (apart from a specific case in Section 5). Under Theorem 1 then (1) only holds if a Lyapunov function $V : \mathcal{X} \rightarrow [1, \infty]$ with $\mathbb{E}_\pi[V] < \infty$ exists such that

$$\limsup_{|x| \rightarrow \infty} \frac{PV(x)}{V(x)} < 1. \quad (7)$$

when P is of the Metropolis–Hastings type, (7) can be written

$$\limsup_{|x| \rightarrow \infty} \int \left[\frac{V(y)}{V(x)} - 1 \right] \alpha(x, y) Q(x, dy) < 0. \quad (8)$$

In this case, a simple criterion for lack of geometric ergodicity is

$$\limsup_{|x| \rightarrow \infty} r(x) = 1. \quad (9)$$

Intuitively this implies that the chain is likely to get ‘stuck’ in the tails of a distribution for large periods.

Jarner and Tweedie [25] introduce a necessary condition for geometric ergodicity through a *tightness* condition.

Theorem 2. (Jarner and Tweedie). *If for any $\varepsilon > 0$ there is a $\delta > 0$ such that for all $x \in \mathcal{X}$*

$$P(x, B_\delta(x)) > 1 - \varepsilon,$$

where $B_\delta(x) := \{y \in \mathcal{X} : d(x, y) < \delta\}$, then a necessary condition for P to produce a geometrically ergodic chain is that for some $s > 0$

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

The result highlights that when $\pi(\cdot)$ is heavy-tailed the chain must be able to make very large moves and still be capable of returning to the centre quickly for (1) to hold.

3. Position-Dependent Random Walk Metropolis

In the RWM, $Q(x, dy) = q(y - x)dy$ with $q(y - x) = q(x - y)$, meaning (6) reduces to $\alpha(x, y) = 1 \wedge \pi(y)/\pi(x)$. A common choice is $Q(x, \cdot) = \mathcal{N}(x, h\Sigma)$, with Σ chosen to mimic the global covariance structure of $\pi(\cdot)$ [3]. Various results exist concerning the optimal choice of h in a given setting (e.g., [26]). It is straightforward to see that Theorem 2 holds here, so that the tails of $\pi(x)$ must be uniformly exponential or lighter for geometric ergodicity. In one dimension this is in fact a sufficient condition [4], while for higher dimensions additional conditions are required [5]. We return to this case in Section 5.

In the PDRWM $Q(x, \cdot) = \mathcal{N}(x, hG(x)^{-1})$, so (6) becomes

$$\alpha(x, y) = 1 \wedge \frac{\pi(y)|G(y)|^{\frac{1}{2}}}{\pi(x)|G(x)|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x - y)^T[G(y) - G(x)](x - y)\right).$$

The motivation for designing such an algorithm is that proposals are more able to reflect the local dependence structure of $\pi(\cdot)$. In some cases this dependence may vary greatly in different parts of the state-space, making a global choice of Σ ineffective [9].

Readers familiar with differential geometry will recognise the volume element $|G(x)|^{1/2}dx$ and the linear approximations to the distance between x and y taken at each point through $G(x)$ and $G(y)$ if \mathcal{X} is viewed as a Riemannian manifold with metric G .

We do not explore these observations further here, but the interested reader is referred to Reference [27] for more discussion.

The choice of $G(x)$ is an obvious question. In fact, specific variants of this method have appeared on many occasions in the literature, some of which we now summarise.

1. *Tempered Langevin diffusions* [8] $G(x) = \pi(x)I$. The authors highlight that the diffusion with dynamics $dX_t = \pi^{-\frac{1}{2}}(X_t)dW_t$ has invariant distribution $\pi(\cdot)$, motivating the choice. The method was shown to perform well for a bi-modal $\pi(x)$, as larger jumps are proposed in the low density region between the two modes.
2. *State-dependent Metropolis* [7] $G(x) = (1 + |x|)^{-b}$. Here the intuition is simply that $b > 0$ means larger jumps will be made in the tails. In one dimension the authors compare the expected squared jumping distance $\mathbb{E}[(X_{i+1} - X_i)^2]$ empirically for chains exploring a $\mathcal{N}(0, 1)$ target distribution, choosing b adaptively, and found $b \approx 1.6$ to be optimal.
3. *Regional adaptive Metropolis–Hastings* [7,11]. $G(x)^{-1} = \sum_{i=1}^m \mathbb{I}(x \in \mathcal{X}_i) \Sigma_i$. In this case the state-space is partitioned into $\mathcal{X}_1 \cup \dots \cup \mathcal{X}_m$, and a different proposal covariance Σ_i is learned adaptively in each region $1 \leq i \leq m$. An extension which allows for some errors in choosing an appropriate partition is discussed in [11].
4. *Localised Random Walk Metropolis* [10]. $G(x)^{-1} = \sum_{k=1}^m \check{q}_\theta(k|x) \Sigma_k$. Here $\check{q}_\theta(k|x)$ are weights based on approximating $\pi(x)$ with some mixture of Normal/Student's t distributions, using the approach suggested in Reference [28]. At each iteration of the algorithm a mixture component k is sampled from $\check{q}_\theta(\cdot|x)$, and the covariance Σ_k is used for the proposal $Q(x, dy)$.
5. *Kernel adaptive Metropolis–Hastings* [9]. $G(x)^{-1} = \gamma^2 I + \nu^2 M_x H M_x^T$, where $M_x = 2[\nabla_x k(z_1, x), \dots, \nabla_x k(z_n, x)]$ for some kernel function k and n past samples $\{z_1, \dots, z_n\}$, $H = I - (1/n) \mathbf{1}_{n \times n}$ is a centering matrix (the $n \times n$ matrix $\mathbf{1}_{n \times n}$ has 1 as each element), and γ, ν are tuning parameters. The approach is based on performing nonlinear principal components analysis on past samples from the chain to learn a local covariance. Illustrative examples for the case of a Gaussian kernel show that $M_x H M_x^T$ acts as a weighted empirical covariance of samples z , with larger weights given to the z_i which are closer to x [9].

The latter cases also motivate any choice of the form

$$G(x)^{-1} = \sum_{i=1}^n w(x, z_i) (z_i - x)^T (z_i - x)$$

for some past samples $\{z_1, \dots, z_n\}$ and weight function $w : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty)$ with $\sum_i w(x, z_i) = 1$ that decays as $|x - z_i|$ grows, which would also mimic the local curvature of $\pi(\cdot)$ (taking care to appropriately regularise and diminish adaptation so as to preserve ergodicity, as outlined in Reference [10]).

Some of the above schemes are examples of adaptive MCMC, in which a candidate from among a family of Markov kernels $\{P_\theta : \theta \in \Theta\}$ is selected by learning the parameter $\theta \in \Theta$ during the simulation [10]. Additional conditions on the adaptation process (i.e., the manner in which θ is learned) are required to establish ergodicity results for the resulting stochastic processes. We consider the decisions on how to learn θ appropriately to be a separate problem and beyond the scope of the present work, and instead focus attention on establishing geometric ergodicity of the base kernels P_θ for any fixed $\theta \in \Theta$. We note that this is typically a pre-requisite for establishing convergence properties of any adaptive MCMC method [10].

4. Results in One Dimension

Here we consider two different general scenarios as $|x| \rightarrow \infty$, i) $G(x)$ is bounded above and below, and ii) $G(x) \rightarrow 0$ at some specified rate. Of course there is also the possibility that $G(x) \rightarrow \infty$, though intuitively this would result in chains that spend a long time in the tails of a distribution, so we do not consider it (if $G(x) \rightarrow \infty$ then chains will in fact

exhibit the *negligible moves* property studied in Reference [29]). Proofs to Propositions in Sections 4 and 5 can be found in Appendix A.

We begin with a result that emphasizes that a growing variance is a necessary requirement for geometric ergodicity in the heavy-tailed case.

Proposition 1. *If $G(x) \geq \sigma^{-2}$ for some $\sigma^{-2} > 0$, then unless $\int e^{\eta|x|} \pi(dx) < \infty$ for some $\eta > 0$ the PDRWM cannot produce a geometrically ergodic Markov chain.*

The above is a simple extension of a result that is well-known in the RWM case. Essentially the tails of the distribution should be exponential or lighter to ensure fast convergence. This motivates consideration of three different types of behaviour for the tails of $\pi(\cdot)$.

Assumption 1. *The density $\pi(x)$ satisfies one of the following tail conditions for all $y, x \in \mathcal{X}$ such that $|y| > |x| > t$, for some finite $t > 0$.*

1. $\pi(y)/\pi(x) \leq \exp\{-a(|y| - |x|)\}$ for some $a > 0$
2. $\pi(y)/\pi(x) \leq \exp\{-a(|y|^\beta - |x|^\beta)\}$ for some $a > 0$ and $\beta \in (0, 1)$
3. $\pi(y)/\pi(x) \leq (|x|/|y|)^p$ for some $p > 1$.

Naturally Assumption 1 implies 2 and Assumption 2 implies 3. If Assumption 1 is not satisfied then $\pi(\cdot)$ is generally called *heavy-tailed*. When $\pi(x)$ satisfies Assumption 2 or 3 but not 1, then the RWM typically fails to produce a geometrically ergodic chain [4]. We show in the sequel, however, that this is not always the case for the PDRWM. We assume the below assumptions for $G(x)$ to hold throughout this section.

Assumption 2. *The function $G : \mathcal{X} \rightarrow (0, \infty)$ is bounded above by some $\sigma_b^{-2} < \infty$ for all $x \in \mathcal{X}$, and bounded below for all $x \in \mathcal{X}$ with $|x| < t$, for some $t > 0$.*

The heavy-tailed case is known to be a challenging scenario, but the RWM will produce a geometrically ergodic Markov chain if $\pi(x)$ is log-concave. Next we extend this result to the case of sub-quadratic variance growth in the tails.

Proposition 2. *If $\exists r < \infty$ such that $G(x) \propto |x|^{-\gamma}$ whenever $|x| > r$, then the PDRWM will produce a geometrically ergodic chain in both of the following cases:*

1. $\pi(x)$ satisfies Assumption 1 and $\gamma \in [0, 2)$
2. $\pi(x)$ satisfies Assumption 2 for some $\beta \in (0, 1)$ and $\gamma \in (2(1 - \beta), 2)$

The second part of Proposition 2 is not true for the RWM, for which Assumption 2 alone is not sufficient for geometric ergodicity [4].

We do not provide a complete proof that the PDRWM will not produce a geometrically ergodic chain when only Assumption 3 holds and $G(x) \propto |x|^{-\gamma}$ for some $\gamma < 2$, but do show informally that this will be the case. Assuming that in the tails $\pi(x) \propto |x|^{-p}$ for some $p > 1$ then for large x

$$\alpha(x, x + cx^{\gamma/2}) = 1 \wedge \left(\frac{x}{x + cx^{\gamma/2}} \right)^{p+\gamma/2} \exp \left(-\frac{c^2 x^\gamma}{2h} \left[\frac{1}{(x + cx^{\gamma/2})^\gamma} - \frac{1}{x^\gamma} \right] \right). \quad (10)$$

The first expression on the right hand side converges to 1 as $x \rightarrow \infty$, which is akin to the case of fixed proposal covariance. The second term will be larger than one for $c > 0$ and less than one for $c < 0$. So the algorithm will exhibit the same ‘random walk in the tails’ behaviour which is often characteristic of the RWM in this scenario, meaning that the acceptance rate fails to enforce a geometric drift back into the centre of the space.

When $\gamma = 2$ the above intuition will not necessarily hold, as the terms in Equation (10) will be roughly constant with x . When only Assumption 3 holds, it is, therefore, tempting

to make the choice $G(x) = x^{-2}$ for $|x| > r$. Informally we can see that such behaviour may lead to a favourable algorithm if a small enough h is chosen. For any fixed $x > r$ a typical proposal will now take the form $y = (1 + \xi\sqrt{h})x$, where $\xi \sim N(0, 1)$. It therefore holds that

$$y = e^{\xi\sqrt{h}}x + r(x, h, \xi), \quad (11)$$

where for any fixed x and ξ the term $r(x, h, \xi)/\sqrt{h} \rightarrow 0$ as $h \rightarrow 0$. The first term on the right-hand side of Equation (11) corresponds to the proposal of the *multiplicative Random Walk Metropolis*, which is known to be geometrically ergodic under Assumption 3 (e.g., [3]), as this equates to taking a logarithmic transformation of x , which ‘lightens’ the tails of the target density to the point where it becomes log-concave. So in practice we can expect good performance from this choice of $G(x)$. The above intuition does not, however, provide enough to establish geometric ergodicity, as the final term on the right-hand side of (11) grows unboundedly with x for any fixed choice of h . The difference between the acceptance rates of the multiplicative Random Walk Metropolis and the PDRWM with $G(x) = x^{-2}$ will be the exponential term in Equation (10). This will instead become polynomial by letting the proposal noise ξ follow a distribution with polynomial tails (e.g., student’s t), which is known to be a favourable strategy for the RWM when only Assumption 3 holds [6]. One can see that if the heaviness of the proposal distribution is carefully chosen then the acceptance rate may well enforce a geometric drift into the centre of the space, though for brevity we restrict attention to Gaussian proposals in this article.

The final result of this section provides a note of warning that lack of care in choosing $G(x)$ can have severe consequences for the method.

Proposition 3. *If $G(x)x^2 \rightarrow 0$ as $|x| \rightarrow \infty$, then the PDRWM will not produce a geometrically ergodic Markov chain.*

The intuition for this result is straightforward when explained. In the tails, typically $|y - x|$ will be the same order of magnitude as $\sqrt{G(x)^{-1}}$, meaning $|y - x|/|x|$ grows arbitrarily large as $|x|$ grows. As such, proposals will ‘overshoot’ the typical set of the distribution, sending the sampler further out into the tails, and will therefore almost always be rejected. The result can be related superficially to a lack of geometric ergodicity for Metropolis–Hastings algorithms in which the proposal mean is comprised of the current state translated by a drift function (often based in $\nabla \log \pi(x)$) when this drift function grows faster than linearly with $|x|$ (e.g., [30,31]).

5. A Higher-Dimensional Case Study

An easy criticism of the above analysis is that the one-dimensional scenario is sometimes not indicative of the more general behaviour of a method. We note, however, that typically the geometric convergence properties of Metropolis–Hastings algorithms do carry over somewhat naturally to more than one dimension when $\pi(\cdot)$ is suitably regular (e.g., [5,32]). Because of this we expect that the growth conditions specified above could be supplanted onto the determinant of $G(x)$ when the dimension is greater than one (leaving the details of this argument for future work).

A key difference in the higher-dimensional setting is that $G(x)$ now dictates both the *size* and *direction* of proposals. In the case $G(x)^{-1} = \Sigma$, some additional regularity conditions on $\pi(x)$ are required for geometric ergodicity in more than one dimension, outlined in References [5,32]. An example is also given in Reference [5] of the simple two-dimensional density $\pi(x, y) \propto \exp(-x^2 - y^2 - x^2y^2)$, which fails to meet these criteria. The difficult models are those for which probability concentrates on a ridge in the tails, which becomes ever narrower as $|x|$ increases. In this instance, proposals from the RWM are less and less likely to be accepted as $|x|$ grows. Another well-known example of this phenomenon is the *funnel* distribution introduced in Reference [33].

To explore the behaviour of the PDRWM in this setting, we design a model problem, the *staircase* distribution, with density

$$\mathfrak{s}(x) \propto 3^{-\lfloor x_2 \rfloor} \mathbb{I}_R(x), \quad R := \{y \in \mathbb{R}^2; y_2 \geq 1, |y_1| \leq 3^{1-\lfloor y_2 \rfloor}\}, \quad (12)$$

where $\lfloor z \rfloor$ denotes the integer part of $z > 0$. Graphically the density is a sequence of cuboids on the upper-half plane of \mathbb{R}^2 (starting at $y_2 = 1$), each centred on the vertical axis, with each successive cuboid one third of the width and height of the previous. The density resembles an ever narrowing staircase, as shown in Figure 1.




Figure 1. The staircase distribution, with density given by Equation (12).

We denote by Q_R the proposal kernel associated with the Random Walk Metropolis algorithm with fixed covariance $h\Sigma$. In fact, the specific choice of h and Σ does not matter provided that the result is positive-definite. For the PDRWM we denote by Q_P the proposal kernel with covariance matrix

$$hG(x)^{-1} = \begin{pmatrix} 3^{-2\lfloor x_2 \rfloor} & 0 \\ 0 & 1 \end{pmatrix},$$

which will naturally adapt the scale of the first coordinate to the width of the ridge.

Proposition 4. *The Metropolis–Hastings algorithm with proposal Q_R does not produce a geometrically ergodic Markov chain when $\pi(x) = \mathfrak{s}(x)$.*

The design of the PDRWM proposal kernel Q_P in this instance is such that the proposal covariance reduces at the same rate as the width of the stairs, therefore naturally adapting the proposal to the width of the ridge on which the density concentrates. This state-dependent adaptation results in a geometrically ergodic chain, as shown in the below result.

Proposition 5. *The Metropolis–Hastings algorithm with proposal Q_P produces a geometrically ergodic Markov chain when $\pi(x) = \mathfrak{s}(x)$.*

6. Discussion

In this paper we have analysed the ergodic behaviour of a Metropolis–Hastings method with proposal kernel $Q(x, \cdot) = \mathcal{N}(x, hG(x)^{-1})$. In one dimension we have characterised the behaviour in terms of growth conditions on $G(x)^{-1}$ and tail conditions on the target distribution, and in higher dimensions a carefully constructed model problem is discussed. The fundamental question of interest was whether generalising an existing Metropolis–Hastings method by allowing the proposal covariance to change with position

can alter the ergodicity properties of the sampler. We can confirm that this is indeed possible, either for the better or worse, depending on the choice of covariance. The take home points for practitioners are (i) lack of sufficient care in the design of $G(x)$ can have severe consequences (as in Proposition 3), and (ii) careful choice of $G(x)$ can have much more beneficial ones, perhaps the most surprising of which are in the higher-dimensional setting, as shown in Section 5.

We feel that such results can also offer insight into similar generalisations of different Metropolis–Hastings algorithms (e.g., [13,34]). For example, it seems intuitive that any method in which the variance grows at a faster than quadratic rate in the tails is unlikely to produce a geometrically ergodic chain. There are connections between the PDRWM and some extensions of the Metropolis-adjusted Langevin algorithm [34], the ergodicity properties of which are discussed in Reference [35]. The key difference between the schemes is the inclusion of the drift term $G(x)^{-1} \nabla \log \pi(x)/2$ in the latter. It is this term which in the main governs the behaviour of the sampler, which is why the behaviour of the PDRWM is different to this scheme. Markov processes are also used in a wide variety of application areas beyond the design of Metropolis–Hastings algorithms (e.g., [36]), and we hope that some of the results established in the present work prove to be beneficial in some of these other settings.

We can apply these results to the specific variants discussed in Section 3. Provided that sensible choices of regions/weights are made and that an adaptation scheme which obeys the diminishing adaptation criterion is employed, the Regional adaptive Metropolis–Hastings, Locally weighted Metropolis and Kernel-adaptive Metropolis–Hastings samplers should all satisfy $G(x) \rightarrow \Sigma$ as $|x| \rightarrow \infty$, meaning they can be expected to inherit the ergodicity properties of the standard RWM (the behaviour in the centre of the space, however, will likely be different). In the State-dependent Metropolis method provided $b < 2$ the sampler should also behave reasonably. Whether or not a large enough value of b would be found by a particular adaptation rule is not entirely clear, and this could be an interesting direction of further study. The Tempered Langevin diffusion scheme, however, will fail to produce a geometrically ergodic Markov chain whenever the tails of $\pi(x)$ are lighter than that of a Cauchy distribution. To allow reasonable tail exploration when this is the case, two pragmatic options would be to upper bound $G(x)^{-1}$ manually or use this scheme in conjunction with another, as there is evidence that the sampler can perform favourably when exploring the centre of a distribution [8]. None of the specific variants discussed here are able to mimic the local curvature of the $\pi(x)$ in the tails, so as to enjoy the favourable behaviour exemplified in Proposition 5. This is possible using Hessian information as in Reference [13], but should also be possible in some cases using appropriate surrogates.

Funding: This research was supported by a UCL IMPACT PhD scholarship co-funded by Xerox Research Centre Europe and EPSRC.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not Applicable

Acknowledgments: The author thanks Alexandros Beskos, Krzysztof Łatuszyński and Gareth Roberts for several useful discussions, Michael Betancourt for proofreading the paper, and Mark Girolami for general supervision and guidance.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Proofs

Proof of Proposition 1. In this case, for any choice of $\varepsilon > 0$ there is a $\delta > 0$ such that $Q(x, B_\delta(x)) > 1 - \varepsilon$. Noting that $P(x, B_\delta(x)) \geq Q(x, B_\delta(x))$ when P is of Metropolis–Hastings type, Theorem 2 can be applied directly. \square

Proof of Proposition 2. For the log-concave case, take $V(x) = e^{s|x|}$ for some $s > 0$, and let B_A denote the integral (8) over the set A . We first break up \mathcal{X} into $(-\infty, 0] \cup (0, x - cx^{\frac{\gamma}{2}}] \cup (x - cx^{\frac{\gamma}{2}}, x + cx^{\frac{\gamma}{2}}] \cup (x + cx^{\frac{\gamma}{2}}, x + cx^\gamma] \cup (x + cx^\gamma, \infty)$ for some $x > 0$ and fixed constant $c \in (0, \infty)$, and show that the integral is strictly negative on at least one of these sets, and can be made arbitrarily small as $x \rightarrow \infty$ on all others. The $-\infty$ case is analogous from the tail conditions on $\pi(x)$. From the conditions we can choose $x > r$ and therefore write $G(x)^{-1} = \eta x^\gamma$ for some fixed $\eta < \infty$.

On $(-\infty, 0]$, we have

$$\begin{aligned} B_{(-\infty, 0]} &= e^{-sx} \int_{-\infty}^0 e^{s|y|} \alpha(x, y) Q(x, dy) - \int_{-\infty}^0 \alpha(x, y) Q(x, dy), \\ &\leq e^{-sx} \int_0^\infty e^{sy} Q(-x, dy). \end{aligned}$$

The integral is now proportional to the moment generating function of a truncated Gaussian distribution (see Appendix B), so is given by

$$e^{-sx + h\eta x^\gamma s^2/2} \left[1 - \Phi \left(x^{1-\gamma/2} / \sqrt{h\eta} - \sqrt{h\eta} s x^{\gamma/2} \right) \right].$$

A simple bound on the error function is $\sqrt{2\pi} x \Phi^c(x) < e^{-x^2/2}$ [37], so setting $\vartheta = x^{1-\gamma/2} / \sqrt{h\eta} - \sqrt{h\eta} s x^{\gamma/2}$ we have

$$\begin{aligned} B_{(-\infty, 0]} &\leq \frac{1}{\sqrt{2\pi}} \exp \left(-2sx + \frac{h\eta s^2}{2} x^\gamma - \frac{1}{2} \left(\frac{1}{h\eta} x^{2-\gamma} - 2sx + h\eta s^2 x^\gamma \right) + \log \vartheta \right), \\ &= \frac{1}{\sqrt{2\pi}} \exp \left(-sx - \frac{1}{2h\eta} x^{2-\gamma} + \log \vartheta \right). \end{aligned}$$

which $\rightarrow 0$ as $x \rightarrow \infty$, so can be made arbitrarily small.

On $(0, x - cx^{\gamma/2}]$, note that $e^{s(|y|-|x|)} - 1$ is clearly negative throughout this region provided that $c < x^{1-\gamma/2}$, which can be enforced by choosing x large enough for any given $c < \infty$. So the integral is straightforwardly bounded as $B_{(0, x - cx^{\gamma/2}]} \leq 0$ for all $x \in \mathcal{X}$.

On $(x - cx^{\gamma/2}, x + cx^{\gamma/2}]$, provided $x - cx^{\gamma/2} > r$ then for any y in this region we can either upper or lower bound $\alpha(x, y)$ with the expression

$$\exp \left(-a(y - x) + \frac{\gamma}{2} \log \left| \frac{x}{y} \right| - \frac{1}{2h\eta} \left[(x - y)^2 y^{-\gamma} - (x - y)^2 x^{-\gamma} \right] \right).$$

A Taylor expansion of $y^{-\gamma}$ about x gives

$$y^{-\gamma} = x^{-\gamma} - \gamma x^{-\gamma-1} (y - x) + \gamma(\gamma + 1) x^{-\gamma-2} (y - x)^2 + \dots$$

and multiplying by $(y - x)^2$ gives

$$(y - x)^2 y^{-\gamma} = \frac{(y - x)^2}{x^\gamma} - \gamma \frac{(y - x)^3}{x^{\gamma+1}} + \gamma(\gamma + 1) \frac{(y - x)^4}{x^{\gamma+2}} + \dots$$

If $|y - x| = cx^{\gamma/2}$ then this is:

$$\frac{c^2 x^\gamma}{x^\gamma} - \gamma \frac{c^3 x^{3\gamma/2}}{x^{\gamma+1}} + \gamma(\gamma + 1) \frac{c^4 x^{2\gamma}}{x^{\gamma+2}} + \dots$$

As $\gamma < 2$ then $3\gamma/2 < \gamma + 1$, and similarly for successive terms, meaning each gets smaller as $|x| \rightarrow \infty$. So we have for large x , $y \in (x - cx^{\gamma/2}, x + cx^{\gamma/2})$ and any $\delta > 0$

$$(y - x)^2 y^{-\gamma} \geq \frac{(y - x)^2}{x^\gamma} - \gamma \frac{(y - x)^3}{x^{\gamma+1}} - 2h\eta \delta. \quad (\text{A1})$$

So we can analyse how the acceptance rate behaves. First note that for fixed $\epsilon > 0$

$$\alpha(x, x + \epsilon) \leq \exp\left(-a\epsilon + \frac{\gamma}{2} \log\left|\frac{x}{x + \epsilon}\right| + \frac{1}{2h} \gamma \frac{\epsilon^3}{x^{\gamma+1}} + \delta\right) \rightarrow \exp(-a\epsilon + \delta),$$

recalling that δ can be made arbitrarily small. In fact, it holds that the $e^{-a\epsilon}$ term will be dominant for any ϵ for which $\epsilon^3/x^{\gamma+1} \rightarrow 0$, i.e., any $\epsilon = o(x^{\gamma+1/3})$. If $\gamma < 2$ then $\epsilon = cx^{\gamma/2}$ satisfies this condition. So for any $y > x$ in this region we can choose an x such that

$$\alpha(x, y) \leq \exp(-a(y - x) + \delta_x),$$

where $\delta_x \rightarrow 0$ as $x \rightarrow \infty$. Similarly we have (for any fixed $\epsilon > 0$)

$$\alpha(x, x - \epsilon) \geq \exp\left(a\epsilon + \frac{\gamma}{2} \log\left|\frac{x}{x - \epsilon}\right| - \frac{1}{2h} \gamma \frac{\epsilon^3}{x^{\gamma+1}} - \delta\right) \rightarrow \exp(a\epsilon - \delta).$$

So by a similar argument we have $\alpha(x, y) > 1$ here when $x \rightarrow \infty$. Combining gives

$$B_{(x-cx^{\gamma/2}, x+cx^{\gamma/2})} \leq \int_0^{cx^{\gamma/2}} \left[e^{(s-a)z+\delta_x} - e^{-az+\delta_x} + e^{-sz} - 1 \right] q_x(dz),$$

where $q_x(\cdot)$ denotes a zero mean Gaussian distribution with the same variance as $Q(x, \cdot)$. Using the change of variables $z' = z/(h\eta x^{\gamma/2})$ we can write the above integral

$$\int_0^{\frac{c}{h\eta}} \left[e^{(s-a)h\eta x^{\gamma/2} z' + \delta_x} - e^{-ah\eta x^{\gamma/2} z' + \delta_x} + e^{-sh\eta x^{\gamma/2} z'} - 1 \right] \mu(dz)$$

where $\mu(\cdot)$ denotes a Gaussian distribution with zero mean and variance one. Provided $s < a$, then by dominated convergence as $x \rightarrow \infty$ this asymptotes to

$$-\int_0^{\frac{c}{h\eta}} \mu(dz) = -\frac{1}{2} \operatorname{erf}\left(\frac{c}{\sqrt{2}h\eta}\right) < 0,$$

where $\operatorname{erf}(z) := (2/\sqrt{\pi}) \int_0^z e^{-t^2} dt$ is the Gaussian error function.

On $(x + cx^{\gamma/2}, x + cx^{\gamma})$ we can upper bound the acceptance rate as

$$\alpha(x, y) \leq \frac{\pi(y)}{\pi(x)} \exp\left(\frac{1}{2} \log \frac{|G(y)|}{|G(x)|} + \frac{G(x)}{2h} (x - y)^2\right)$$

If $y \geq x$ and $x > x_0$ we have

$$\alpha(x, y) \leq \exp\left(-a(|y| - |x|) + \frac{1}{2h\eta} \frac{(x - y)^2}{x^{\gamma}}\right).$$

For $|y - x| = cx^{\ell}$ this becomes

$$\alpha(x, y) \leq \exp\left(-acx^{\ell} + \frac{c^2}{2h\eta} x^{2\ell-\gamma}\right)$$

So provided $\gamma > \ell$ the first term inside the exponential will dominate the second for large enough x . In the equality case we have

$$\alpha(x, y) \leq \exp\left(\left(\frac{c^2}{2h\eta} - a\right) cx^{\gamma}\right),$$

so provided we choose c such that $a > c^2/(2h\eta)$ then the acceptance rate will also decay exponentially. Because of this we have

$$\begin{aligned} B_{(x+cx^{\gamma/2}, x+cx^{\gamma})} &\leq \int_{x+cx^{\gamma/2}}^{x+cx^{\gamma}} e^{s(y-x)} \alpha(x, y) Q(x, dy), \\ &\leq e^{(c^2/(2h\eta)+s-a)cx^{\gamma/2}} Q(x, (x+cx^{\gamma/2}, x+cx^{\gamma})), \end{aligned}$$

so provided $a > c^2/(2h\eta) + s$ then this term can be made arbitrarily small.

On $(x+cx^{\gamma}, \infty)$ using the same properties of truncated Gaussians we have

$$\begin{aligned} B_{(x+cx^{\gamma}, \infty)} &\leq e^{-sx} \int_{x+cx^{\gamma}}^{\infty} e^{sy} Q(x, dy), \\ &= e^{s^2 h \eta x^{\gamma}/2} \Phi^c \left(\left(\frac{c}{\sqrt{h\eta}} - \sqrt{h\eta} s \right) x^{\gamma} \right), \end{aligned}$$

which can be made arbitrarily small provided that s is chosen to be small enough using the same simple bound on Φ^c as for the case of $B_{(-\infty, 0]}$.

Combining gives that the integral (8) is bounded above by $-\text{erf}(c/\sqrt{2h^2\eta^2})/2$, which is strictly less than zero as c, h and η are all positive. This completes the proof under Assumption 1.

Under Assumption 2 the proof is similar. Take $V(x) = e^{s|x|^{\beta}}$, and divide \mathcal{X} up into the same regions. Outside of $(x - cx^{\gamma/2}, x + cx^{\gamma/2})$ the same arguments show that the integral can be made arbitrarily small. On this set, note that in the tails

$$(x + cx^{\ell})^{\beta} - x^{\beta} = \beta c x^{\ell+\beta-1} + \frac{\beta(\beta-1)}{2} c^2 x^{2\ell+\beta-2} + \dots$$

For $y - x = cx^{\ell}$, then for $\ell < 1 - \beta$ this becomes negligible. So in this case we further divide the typical set into $(x, x + cx^{1-\beta}] \cup (x + cx^{1-\beta}, x + cx^{\gamma/2})$. On $(x - cx^{1-\beta}, x + cx^{1-\beta})$ the integral is bounded above by $e^{-c_1} Q(x, (x - cx^{1-\beta}, x + cx^{1-\beta})) \rightarrow 0$, for some suitably chosen $c_1 > 0$. On $(x - cx^{\gamma/2}, x - cx^{1-\beta}) \cup (x + cx^{1-\beta}, x + cx^{\gamma/2})$ then for $y > x$ we have $\alpha(x, y) \leq e^{-c_2(y^{\beta}-x^{\beta})}$, so we can use the same argument as in the log-concave case to show that the integral will be strictly negative in the limit. \square

Proof of Proposition 3. First note that in this case for any $g : \mathbb{R} \rightarrow (0, \infty)$ such that as $|x| \rightarrow \infty$ it holds that $g(x)/|x| \rightarrow \infty$ but $g(x)\sqrt{G(x)} \rightarrow 0$, then

$$Q(x, \{x - g(x), x + g(x)\}) = \Phi\left(g(x)\sqrt{G(x)}\right) - \Phi\left(-g(x)\sqrt{G(x)}\right) \rightarrow 0$$

as $|x| \rightarrow \infty$. The chain therefore has the property that $\mathbb{P}(\{|X_{i+1}| > g(X_i)/2\} \cup \{X_{i+1} = X_i\})$ can be made arbitrarily close to 1 as $|X_i|$ grows, which leads to two possible behaviours. If the form of $\pi(\cdot)$ enforces such large jumps to be rejected then $r(x) \rightarrow 1$ and lack of geometric ergodicity follows from (9). If this is not the case then the chain will be transient (this can be made rigorous using a standard Borel–Cantelli argument, see e.g., the proof of Theorem 12.2.2 on p. 299 of [21]). \square

Proof of Proposition 4. It is sufficient to construct a sequence of points $x_p \in \mathbb{R}^2$ such that $|x_p| \rightarrow \infty$ as $p \rightarrow \infty$, and show that $r(x_p) \rightarrow 1$ in the same limit, then apply (9). Take $x_p = (0, p)$ for $p \in \mathbb{N}$. In this case

$$r(x_p) = 1 - \int \alpha(x_p, y) Q_R(x_p, dy)$$

Note that for every $\epsilon > 0$ there is a $\delta < \infty$ such that $Q(x_p, B_{\delta}^c(x_p)) < \epsilon$ for all x_p , where $B_{\delta}(x) := \{y \in \mathbb{R}^2 : |y - x| \leq \delta\}$. The set $A(x_p, \delta) := B_{\delta}(x_p) \cap R$ denotes

the possible values of $y \in B_\delta(x)$ for which the acceptance rate is non-zero. Note that $A(x_p, \delta) \subset S(x_p, \delta) := \{y \in B_\delta(x_p) : |y_1| \leq 3^{1-\lfloor p-\delta \rfloor}\}$, which is simply a strip that can be made arbitrarily narrow for any fixed δ by taking p large enough. Combining these ideas gives

$$\begin{aligned} \int \alpha(x_p, y) Q_R(x_p, dy) &\leq \int_{A(x_p, \delta)} \alpha(x_p, y) Q_R(x_p, dy) + \epsilon \\ &\leq Q_R(x_p, S(x_p, \delta)) + \epsilon. \end{aligned}$$

Both of the quantities on the last line can be made arbitrarily small by choosing p suitably large. Thus, $r(x_p) \rightarrow 1$ as $|x_p| \rightarrow \infty$, as required. \square

Proof of Proposition 5. First note that $\inf_{x \in R} Q_P(x, R)$ is bounded away from zero, unlike in the case of Q_R , owing to the design of Q_P . The acceptance rate here simplifies, since for any $y \in R$

$$\frac{\mathfrak{s}(y)|G(y)|^{\frac{1}{2}}}{\mathfrak{s}(x)|G(x)|^{\frac{1}{2}}} = 1,$$

meaning only the expression $\exp\left(-\frac{1}{2}(y-x)^T[G(y)-G(x)](y-x)\right)$ needs to be considered. In this case the expression is simply

$$\exp\left(-\frac{1}{2}(3^{2\lfloor y_2 \rfloor} - 3^{2\lfloor x_2 \rfloor})(y_1 - x_1)^2\right).$$

Provided that $x_1 \neq y_1$, then when $1 \leq \lfloor y_2 \rfloor < \lfloor x_2 \rfloor$ this expression is strictly greater than 1, whereas in the reverse case it is strictly less than one. The resulting Metropolis–Hastings kernel P using proposal kernel Q_P will therefore satisfy $\int y_2 P(x, dy) < x_2$ for large enough x_2 , and hence geometric ergodicity follows by taking the Lyapunov function $V(x) = e^{s|x_2|}$ (which can be used here since the domain of x_1 is compact) and following an identical argument to that given on pages 404–405 of Reference [21] for the case of the proof of geometric ergodicity of the random walk on the half-line model for suitably small $s > 0$, taking the small set $C := [0, 1] \times [1, r]$ for suitably large $r < \infty$ and $\nu(\cdot) = \int \mathfrak{s}(x) dx$. \square

Appendix B. Needed Facts about Truncated Gaussian Distributions

Here we collect some elementary facts used in the article. For more detail see e.g., [38]. If X follows a truncated Gaussian distribution $\mathcal{N}_{[a,b]}^T(\mu, \sigma^2)$ then it has density

$$f(x) = \frac{1}{\sigma Z_{a,b}} \phi\left(\frac{x-\mu}{\sigma}\right) \mathbb{I}_{[a,b]}(x),$$

where $\phi(x) = e^{-x^2/2}/\sqrt{2\pi}$, $\Phi(x) = \int_{-\infty}^x \phi(y) dy$ and $Z_{a,b} = \Phi((b-\mu)/\sigma) - \Phi((a-\mu)/\sigma)$. Defining $B = (b-\mu)/\sigma$ and $A = (a-\mu)/\sigma$, we have

$$\mathbb{E}[X] = \mu + \frac{\phi(A) - \phi(B)}{Z_{a,b}} \sigma$$

and

$$\mathbb{E}[e^{tX}] = e^{\mu t + \sigma^2 t^2/2} \left[\frac{\Phi(B - \sigma t) - \Phi(A - \sigma t)}{Z_{a,b}} \right].$$

In the special case $b = \infty, a = 0$ this becomes $e^{\mu t + \sigma^2 t^2/2} \Phi(\sigma t) / Z_{a,b}$.

References

1. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H.; Teller, E. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **1953**, *21*, 1087–1092. [\[CrossRef\]](#)
2. Hastings, W.K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **1970**, *57*, 97–109. [\[CrossRef\]](#)

3. Sherlock, C.; Fearnhead, P.; Roberts, G.O. The random walk Metropolis: Linking theory and practice through a case study. *Stat. Sci.* **2010**, *25*, 172–190. [\[CrossRef\]](#)
4. Mengersen, K.L.; Tweedie, R.L. Rates of convergence of the Hastings and Metropolis algorithms. *Ann. Stat.* **1996**, *24*, 101–121. [\[CrossRef\]](#)
5. Roberts, G.O.; Tweedie, R.L. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika* **1996**, *83*, 95–110. [\[CrossRef\]](#)
6. Jarner, S.F.; Roberts, G.O. Convergence of Heavy-tailed Monte Carlo Markov Chain Algorithms. *Scand. J. Stat.* **2007**, *34*, 781–815. [\[CrossRef\]](#)
7. Roberts, G.O.; Rosenthal, J.S. Examples of adaptive MCMC. *J. Comput. Graph. Stat.* **2009**, *18*, 349–367. [\[CrossRef\]](#)
8. Roberts, G.O.; Stramer, O. Langevin diffusions and Metropolis–Hastings algorithms. *Methodol. Comput. Appl. Probab.* **2002**, *4*, 337–357. [\[CrossRef\]](#)
9. Sejdinovic, D.; Strathmann, H.; Garcia, M.L.; Andrieu, C.; Gretton, A. Kernel Adaptive Metropolis-Hastings. In Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 21–26 June 2014; Xing, E.P., Jebara, T., Eds.; PMLR: Beijing, China, 2014; Volume 32, pp. 1665–1673.
10. Andrieu, C.; Thoms, J. A tutorial on adaptive MCMC. *Stat. Comput.* **2008**, *18*, 343–373. [\[CrossRef\]](#)
11. Craiu, R.V.; Rosenthal, J.; Yang, C. Learn from thy neighbor: Parallel-chain and regional adaptive MCMC. *J. Am. Stat. Assoc.* **2009**, *104*, 1454–1466. [\[CrossRef\]](#)
12. Rudolf, D.; Sprungk, B. On a generalization of the preconditioned Crank–Nicolson Metropolis algorithm. *Found. Comput. Math.* **2018**, *18*, 309–343. [\[CrossRef\]](#)
13. Girolami, M.; Calderhead, B. Riemann manifold langevin and hamiltonian monte carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2011**, *73*, 123–214. [\[CrossRef\]](#)
14. Brooks, S.; Gelman, A.; Jones, G.; Meng, X.L. *Handbook of Markov Chain Monte Carlo*; CRC Press: Boca Raton, FL, USA, 2011.
15. Maire, F.; Vandekerckhove, P. On Markov chain Monte Carlo for sparse and filamentary distributions. *arXiv* **2018**, arXiv:1806.09000.
16. Mallik, A.; Jones, G.L. Directional Metropolis-Hastings. *arXiv* **2017**, arXiv:1710.09759.
17. Ludkin, M.; Sherlock, C. Hug and Hop: A discrete-time, non-reversible Markov chain Monte Carlo algorithm. *arXiv* **2019**, arXiv:1907.13570.
18. Kamatani, K. Ergodicity of Markov chain Monte Carlo with reversible proposal. *J. Appl. Probab.* **2017**, 638–654. [\[CrossRef\]](#)
19. Roberts, G.O.; Rosenthal, J.S. General state space Markov chains and MCMC algorithms. *Probab. Surv.* **2004**, *1*, 20–71. [\[CrossRef\]](#)
20. Roberts, G.O.; Rosenthal, J.S. Geometric ergodicity and hybrid Markov chains. *Electron. Comm. Probab.* **1997**, *2*, 13–25. [\[CrossRef\]](#)
21. Meyn, S.P.; Tweedie, R.L. *Markov Chains and Stochastic Stability*; Cambridge University Press: Cambridge, UK, 2009.
22. Jones, G.L.; Hobert, J.P. Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Stat. Sci.* **2001**, *16*, 312–334. [\[CrossRef\]](#)
23. Tierney, L. Markov chains for exploring posterior distributions. *Annal. Stat.* **1994**, *22*, 1701–1728. [\[CrossRef\]](#)
24. Bierkens, J. Non-reversible Metropolis–Hastings. *Stat. Comput.* **2016**, *26*, 1213–1228. [\[CrossRef\]](#)
25. Jarner, S.F.; Tweedie, R.L. Necessary conditions for geometric and polynomial ergodicity of random-walk-type Markov chains. *Bernoulli* **2003**, *9*, 559–578. [\[CrossRef\]](#)
26. Roberts, G.O.; Rosenthal, J.S. Optimal scaling for various Metropolis-Hastings algorithms. *Stat. Sci.* **2001**, *16*, 351–367. [\[CrossRef\]](#)
27. Livingstone, S.; Girolami, M. Information-geometric Markov chain Monte Carlo methods using diffusions. *Entropy* **2014**, *16*, 3074–3102. [\[CrossRef\]](#)
28. Andrieu, C.; Moulines, É. On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Probab.* **2006**, *16*, 1462–1505. [\[CrossRef\]](#)
29. Livingstone, S.; Faulkner, M.F.; Roberts, G.O. Kinetic energy choice in Hamiltonian/hybrid Monte Carlo. *Biometrika* **2019**, *106*, 303–319. [\[CrossRef\]](#)
30. Roberts, G.O.; Tweedie, R.L. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **1996**, *2*, 341–363. [\[CrossRef\]](#)
31. Livingstone, S.; Betancourt, M.; Byrne, S.; Girolami, M. On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli* **2019**, *25*, 3109–3138. [\[CrossRef\]](#)
32. Jarner, S.F.; Hansen, E. Geometric ergodicity of Metropolis algorithms. *Stoch. Process. Their Appl.* **2000**, *85*, 341–361. [\[CrossRef\]](#)
33. Neal, R.M. Slice sampling. *Annal. Stat.* **2003**, 705–741. [\[CrossRef\]](#)
34. Xifara, T.; Sherlock, C.; Livingstone, S.; Byrne, S.; Girolami, M. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Stat. Probab. Lett.* **2014**, *91*, 14–19. [\[CrossRef\]](#)
35. Łatuszyński, K.; Roberts, G.O.; Thiery, A.; Wolny, K. Discussion on ‘Riemann manifold Langevin and Hamiltonian Monte Carlo methods’ (by Girolami, M. and Calderhead, B.). *J. R. Stat. Soc. Ser. B Statist. Methodol.* **2011**, *73*, 188–189.
36. Chen, S.; Tao, Y.; Yu, D.; Li, F.; Gong, B. Distributed learning dynamics of Multi-Armed Bandits for edge intelligence. *J. Syst. Archit.* **2020**, 101919. Available online: <https://www.sciencedirect.com/science/article/abs/pii/S1383762120301806> (accessed on 29 May 2015). [\[CrossRef\]](#)

-
37. Cook, J.D. Upper and Lower Bounds on the Normal Distribution Function; Technical Report. 2009. Available online: <http://www.johndcook.com/normalbounds.pdf> (accessed on 29 May 2015).
 38. Johnson, N.L.; Kotz, S. *Distributions in Statistics: Continuous Univariate Distributions*; Houghton Mifflin: Boston, MA, USA, 1970; Volume 1.