

APPROXIMATIONS OF GEOMETRICALLY ERGODIC REVERSIBLE MARKOV CHAINS

JEFFREY NEGREA * ** AND
JEFFREY S. ROSENTHAL *, *** *University of Toronto*

Abstract

A common tool in the practice of Markov chain Monte Carlo (MCMC) is to use approximating transition kernels to speed up computation when the desired kernel is slow to evaluate or is intractable. A limited set of quantitative tools exists to assess the relative accuracy and efficiency of such approximations. We derive a set of tools for such analysis based on the Hilbert space generated by the stationary distribution we intend to sample, $L_2(\pi)$. Our results apply to approximations of reversible chains which are geometrically ergodic, as is typically the case for applications to MCMC. The focus of our work is on determining whether the approximating kernel will preserve the geometric ergodicity of the exact chain, and whether the approximating stationary distribution will be close to the original stationary distribution. For reversible chains, our results extend the results of Johndrow *et al.* (2015) from the uniformly ergodic case to the geometrically ergodic case, under some additional regularity conditions. We then apply our results to a number of approximate MCMC algorithms.

Keywords: Markov chain; Monte Carlo; perturbations

2010 Mathematics Subject Classification: Primary 60J05
Secondary 60J22; 65C05

1. Introduction

The use of Markov chain Monte Carlo (MCMC) arises from the need to sample from probabilistic models when simple Monte Carlo is not possible. The procedure is to simulate a positive recurrent Markov process where the stationary distribution is the measure one intends to sample, so that the dynamics of the process converge to the distribution required. Temporally correlated samples may then be used to approximate various expectations; see e.g. [7] and the many references therein. Examples of common applications may be found in hierarchical models, spatio-temporal models, random networks, finance, bioinformatics, etc.

Often, however, the transition dynamics of the Markov chain required to run this process exactly are too computationally expensive because of prohibitively large datasets, intractable likelihoods, etc. In such cases it is tempting to instead *approximate* the transition dynamics of the Markov process in question, either deterministically as in the low-rank Gaussian

Received 23 October 2019; revision received 28 January 2021.

* Postal address: Department of Statistical Sciences, 9th Floor, 700 University Ave., Toronto, ON M5G 1Z5, Canada.

** Email address: negrea@utstat.toronto.edu

*** Email address: jeff@math.toronto.edu

© The Author(s) 2021. Published by Cambridge University Press on behalf of Applied Probability Trust.

approximation of [18], or stochastically as in the noisy Metropolis–Hastings procedure of [1]. It is important then to understand whether these approximations will yield stable and reliable results. This paper aims to provide quantitative tools for the analysis of these algorithms. Since the use of approximation for the transition dynamics may be interpreted as a *perturbation* of the transition kernel of the exact MCMC algorithm, we focus on bounds on the convergence of perturbations of Markov chains.

The primary purpose of this paper is to extend existing quantitative bounds on the errors of approximate Markov chains from the uniformly ergodic case in [18] to the geometrically ergodic case (a weaker condition, for which multiple equivalent definitions may be found in [29]). Our work will extend the theoretical results of [18] in the case that the exact chain is reversible, replacing the total variation metric with L_2 distances, and relaxing the uniform contraction condition to $L_2(\pi)$ -geometric ergodicity.

1.1. Geometric ergodicity

When analyzing the performance of exact MCMC algorithms, it is natural to decompose the error in approximation of expectations into a component for the transient phase error of the process and one for the Monte Carlo approximation error. The former may be interpreted as the bias due to not having started the process in the stationary distribution. A Markov chain is *geometrically ergodic* if, from a suitable initial distribution ν , the marginal distribution of the n th iterate of the chain converges to the stationary distribution, with an error that decays as $C(\nu)\rho^n$ for some $\rho \in (0, 1)$ and some constant depending on the initial distribution $C(\nu)$, in some suitable metric on the space of probability measures. The geometric ergodicity condition essentially dictates that the transient phase error of the n th sample decays exponentially quickly in n . The chain is *uniformly* (geometrically) ergodic if C can be chosen independently of the initial distribution. Geometric ergodicity is a desirable property as it ensures that the cumulative transient phase error asymptotically does not dominate the Monte Carlo error, while still being less restrictive than the uniform ergodicity condition, which often fails when the state space is not finite or compact (for example, an AR(1) process is geometrically ergodic but not uniformly ergodic).

When using approximate MCMC methods, one desires the approximation to preserve geometric ergodicity, so that convergence to stationarity is still fast and the transient phase error goes to zero quickly. This is an important issue, especially since Medina-Aguayo *et al.* [25] have shown that intuitive approximations such as Monte Carlo within Metropolis may lead to transient approximating chains.

1.2. Outline of the paper

The outline of this paper is as follows. Section 2 reviews related work. Then Section 3 contains our main theoretical results and their proofs. Theorem 1 therein provides bounds on the distance between stationary distributions, and gives a sufficient condition for the perturbed chain to be geometrically ergodic in $L_2(\pi)$, where π is the stationary distribution of the unperturbed chain. Theorem 2 and Theorem 3 give sufficient conditions for the perturbed chain to be geometrically ergodic according to several other variants of the definition of geometric ergodicity (for different metrics and families of initial distributions), and provide quantitative rates when possible. The remainder of Section 3 establishes bounds on autocorrelations and mean squared error for Monte Carlo estimates of expected values computed with the perturbed chain.

Finally, Section 4 considers noisy and/or approximate Metropolis–Hastings algorithms. It provides sufficient conditions that one can check in order for our results from Section 3 to be

applied. We use this to study Metropolis–Hastings with deterministic approximations to the target density, as well as the Monte Carlo within Metropolis algorithm, as in [24], and provide some examples of how these types of approximations might arise in practice.

2. Related work

This section presents a brief review of related work, discussing convergence of perturbed Markov chains in the uniformly ergodic and geometrically ergodic cases with varying metrics and additional assumptions. The results in the literature require a wide range of assumptions and have a wide range of scopes. The results for uniformly ergodic chains have a simpler aesthetic, in line with what intuition for finite-state-space chains might inspire, as they do not require drift and minorization conditions to state. Our results cover the geometrically ergodic and reversible case, and use properties of reversibility to match the simpler aesthetic found in the literature for the uniformly ergodic case.

In close relation to the present paper, Johndrow *et al.* [18] derive perturbation bounds to assess the robustness of approximate MCMC algorithms. The assumptions upon which their results rely are as follows: the original chain is uniformly contractive in the total variation norm (this implies uniform ergodicity), and the perturbation is sufficiently small (in the operator norm induced by the total variation norm). The main results of their paper are as follows:

- The perturbed kernel is uniformly contractive in the total variation norm.
- The perturbed stationary distribution is close to the original stationary distribution in total variation.
- Explicit bounds are proved on the total variation distance between finite-time approximate sampling distributions and the original stationary distribution.
- Explicit bounds are proved on the total variation difference between the original stationary distribution and the mixture of finite-time approximate sampling distributions.
- Explicit bounds are proved on the mean squared error for integral approximation using the approximate kernel and the true kernel.

The results derived by [18] are applied within the same paper to a wide variety of approximate MCMC problems, including low-rank approximation to Gaussian processes and sub-sampling approximations. In other work, Johndrow and Mattingly [16] use intuitive coupling arguments to establish similar results under the same uniform contractivity assumption.

Further results on perturbations for uniformly ergodic chains may be found in [26]. This work is motivated in part by numerical rounding errors. Various applications of these results may be found in [1]. The only assumption of [26] is that the original chain is uniformly ergodic. The paper is unique in that it makes no assumption regarding the proximity of the original and perturbed kernel, though the level of approximation error does still scale linearly with the total variation distance of the original and perturbed kernels. The main results are explicit bounds on the total variation distance between finite-time sampling distributions, and explicit bounds on the total variation distance between stationary distributions.

The work of Roberts *et al.* [31] (see also [6]) is also motivated by numerical rounding errors. The perturbed kernel is assumed to be derived from the original kernel by a *round-off function*,

which e.g. maps the input to the nearest multiple of 2^{-31} . In such cases, the new state space is at most countable while the old state space may have been uncountable, and so the resulting chains have mutually singular marginal distributions at all finite times and mutually singular stationary distributions (if they have stationary distributions at all). The results of [31] require the analysis of Lyapunov drift conditions and drift functions (which we will avoid by working in an appropriate L_2 space). The key assumptions in [31] are as follows: the original kernel is geometrically ergodic, and V is a Lyapunov drift function for the original kernel; the original and perturbed transition kernels are close in the V -norm; the perturbed kernel is defined via a round-off function with round-off error uniformly sufficiently small; and $\log V$ is uniformly continuous. The main results of the paper include the following:

- If the perturbed kernel is sufficiently close in the V -norm, then geometric ergodicity is preserved.
- If the drift function, V , can be chosen so that $\log V$ is uniformly continuous, and if the round-off errors can be made arbitrarily small, then the kernels can be made arbitrarily close in the V -norm.
- Explicit bounds are proved on the total variation distance between the approximate finite-time sampling distribution and the true stationary distribution.
- Sufficient conditions are given for the approximating stationary distribution to be arbitrarily close in total variation to the true stationary distribution.

Roberts *et al.* [31] also prove results that do not require closeness in the V -norm, or even absolute continuity of the perturbed transitions; in such cases they show that a suitable drift condition on the original chain together with a uniformly small round-off error yields perturbed chains which are geometrically ergodic, and that the stationary measure varies continuously under such perturbations in the topology of weak convergence.

Pillai and Smith [28] provide bounds in terms of the Wasserstein topology (cf. [11]). Their main focus is on approximate MCMC algorithms, especially approximation due to sub-sampling from a large dataset (e.g., when computing the posterior density). Their underlying assumptions are as follows:

- The original and perturbed kernels satisfy a series of *drift-like conditions* with shared parameters.
- The original kernel has finite eccentricity for all states (where eccentricity of a state is defined as the expected distance between the state and a sample from the stationary distribution).
- The *Ricci curvature* of the original kernel has a non-trivial uniform lower bound on a positive measure subset of the state space.
- The transition kernels are close in the Wasserstein metric, uniformly on the mentioned subset.

Their main results under these assumptions are as follows:

- Explicit bounds on the Wasserstein distance between the approximate sampling distribution and the original stationary distribution.

- Explicit bounds on the total variation distance of the original and perturbed stationary distributions and bounds on the mixing times of each chain.
- Explicit bounds on the bias and L_1 error of Monte Carlo approximations.
- Decomposition of the error from approximate MCMC estimation into components from *burn-in*, *asymptotic bias*, and *variance*.
- Rigorous discussion of the trade-off between the above error components.

Rudolf and Schweizer [35] also use the Wasserstein topology. They focus on approximate MCMC algorithms, with applications to auto-regressive processes and stochastic Langevin algorithms for Gibbs random fields. Their results use the following assumptions: the original kernel is Wasserstein ergodic; a Lyapunov drift condition for perturbed kernel is given, with drift function \tilde{V} ; \tilde{V} has finite expectation under the initial distribution; and the perturbation operator is uniformly bounded in a \tilde{V} -normalized Wasserstein norm. Their main results are explicit bounds on the Wasserstein distance and weighted total variation distance between the original and perturbed finite-time sampling distributions, and explicit bounds on the Wasserstein distance between stationary distributions.

Ferré *et al.* [10] build upon the work of Keller and Liverani [20] to provide perturbation results for V -geometrically ergodic Markov chains using a simultaneous drift condition. They show that any perturbation to the transition kernel which shares its drift condition has a stationary distribution, that such a perturbation is also V -geometrically ergodic, and that the perturbed stationary distribution is close to the original one. The assumption of a shared drift condition may be difficult to verify or may not hold in some cases of interest related to approximate or noisy MCMC. Hervé and Ledoux [12] consider finite-rank approximations to a transition kernel. That work gives sufficient conditions for approximations to inherit V -geometric ergodicity; it provides a quantitative relationship between the rates of convergence and bounds the total variation distance between stationary measures. It also provides sufficient conditions for V -geometric ergodicity of a family of finite-rank approximations to a transition kernel to guarantee geometric ergodicity of the kernel, and provides quantitative rates of convergence. In both of these results, as in [10], the results depend on a simultaneous drift condition for the approximations and the original kernel.

Each of the above papers demonstrate bounds on various measures of error from using approximate finite-time sampling distributions and approximate ergodic distributions to calculate expectations of functions. On the other hand, the assumptions underlying the results vary dramatically. The results for uniformly ergodic chains are based on simpler and more intuitive assumptions than those for geometrically ergodic chains. Our work extends these results to geometrically ergodic chains and perturbations while preserving essentially the same level of simplicity in the assumptions. In particular we avoid the need to identify a Lyapunov drift condition, and our assumptions are expressed directly in terms of transition kernels, rather than in terms of a relationship between drift conditions which they satisfy.

3. Perturbation bounds

This section extends the main results from [18] to the $L_2(\pi)$ -geometrically ergodic case, assuming the perturbation $P - P_\epsilon$ has bounded $L_2(\pi)$ operator norm.

3.1. Definitions and notation

Let π be a probability measure on a measurable space (\mathcal{X}, Σ) . We make considerable use of the following norms on signed measures and their corresponding Banach spaces:

$$\begin{aligned} \|\lambda\|_{\text{TV}} &= \sup_{A \in \Sigma} |\lambda(A)|, & \mathcal{M}(\Sigma) &= \{\text{bounded signed measures on } (\mathcal{X}, \Sigma)\}, \\ \|\lambda\|_{L_2(\pi)} &= \left(\int \left(\frac{d\lambda}{d\pi} \right)^2 d\pi \right)^{1/2}, & L_2(\pi) &= \{v \ll \pi : \|v\|_{L_2(\pi)} < \infty\}, \\ \|\cdot\|_{L_{2,0}(\pi)} &= \|\cdot\|_{L_2(\pi)|_{L_{2,0}(\pi)}}, & L_{2,0}(\pi) &= \{v \in L_2(\pi) : v(\mathcal{X}) = 0\}, \\ \|\lambda\|_{L_1(\pi)} &= \int \left| \frac{d\lambda}{d\pi} \right| d\pi, & L_1(\pi) &= \{v \ll \pi : \|v\|_{L_1(\pi)} < \infty\}, \\ \|\lambda\|_{L_\infty(\pi)} &= \text{ess sup}_{X \sim \pi} \frac{d\lambda}{d\pi}(X), & L_\infty(\pi) &= \left\{ v \ll \pi : (\exists b > 0) \left(\left| \frac{dv}{d\pi} \right| < b \quad \pi\text{-a.e.} \right) \right\}. \end{aligned}$$

Note that $L_{2,0}(\pi)$ is a complete subspace of $L_2(\pi)$. Let

$$\mathcal{M}_{+,1} = \{\lambda \in \mathcal{M} : [\forall A \in \Sigma \quad \lambda(A) \geq 0] \text{ and } [\lambda(\mathcal{X}) = 1]\}$$

be the set of probability measures on (\mathcal{X}, Σ) . Note that for any probability measure π , $L_\infty(\pi) \subset L_2(\pi) \subset L_1(\pi) \subset \mathcal{M}(\Sigma)$, though in general they are not complete subspaces of each other when their corresponding norms are not equivalent. For a norm $\|\cdot\|$ on a vector space, we also write $\|\cdot\|$ for the corresponding operator norm on the space of bounded linear operators from V to itself, $\mathcal{B}(V)$.

Definition 1. (Geometric ergodicity) Let P be the kernel of a positive recurrent Markov chain with invariant measure π . Let λ be any measure with $\pi \ll \lambda$. Suppose that $\rho_{\text{TV}}, \rho_1, \rho_2 \in (0, 1)$.

- (i) P is π -a.e.-TV geometrically ergodic with factor ρ_{TV} if there exists $C_{\text{TV}} : \mathcal{X} \rightarrow \mathbb{R}_+$ such that for π -almost every $(\pi$ -a.e.) $x \in \mathcal{X}$ and for all $n \in \mathbb{N}$,

$$\|\delta_x P^n - \pi\|_{\text{TV}} \leq C_{\text{TV}}(x) \rho_{\text{TV}}^n.$$

The *optimal rate* for π -a.e.-TV geometric ergodicity is the infimum over factors for which the above definition holds:

$$\begin{aligned} \rho_{\text{TV}}^* &= \inf \left\{ \rho > 0 \text{ s.t. } \exists C : \mathcal{X} \rightarrow \mathbb{R}_+ \text{ with } \pi(\{x : C(x) < \infty\}) = 1 \text{ and} \right. \\ &\quad \left. \forall n \in \mathbb{N}, \pi\text{-a.e. } x \in \mathcal{X} \quad \|\delta_x P^n - \pi\|_{\text{TV}} \leq C(x) \rho^n \right\}. \end{aligned} \tag{1}$$

- (ii) P is $L_2(\lambda)$ -geometrically ergodic with factor ρ_2 if $P : L_2(\lambda) \rightarrow L_2(\lambda)$ and there exists $C_2 : L_2(\lambda) \cap \mathcal{M}_{+,1} \rightarrow \mathbb{R}_+$ such that for every $v \in L_2(\lambda) \cap \mathcal{M}_{+,1}$ and for all $n \in \mathbb{N}$,

$$\|v P^n - \pi\|_{L_2(\lambda)} \leq C_2(v) \rho_2^n.$$

The *optimal rate* for $L_2(\lambda)$ -geometric ergodicity is the infimum over factors for which the above definition holds:

$$\begin{aligned} \rho_2^* &= \inf \{ \rho > 0 \text{ s.t. } \exists C : L_2(\lambda) \cap \mathcal{M}_{+,1} \rightarrow \mathbb{R}_+ \text{ with} \\ &\quad \forall n \in \mathbb{N}, v \in L_2(\lambda) \cap \mathcal{M}_{+,1} \quad \|v P^n - \pi\|_{L_2(\lambda)} \leq C(v) \rho^n \}. \end{aligned}$$

Remark 1. If P is π -reversible and aperiodic, then P is $L_2(\pi)$ -geometrically ergodic if and only if it is π -a.e.-TV geometrically ergodic, as per [29]. In this case the optimal rate of $L_2(\pi)$ -geometric ergodicity, ρ_2^* , is equal to the spectral radius of $P|_{L_{2,0}(\pi)}$. In this case, the spectrum of P is a subset of $[-\rho_2^*, \rho_2^*] \cup \{1\}$, and P is $L_2(\pi)$ -geometrically ergodic with factor ρ_2^* and $C(\mu) = \|\mu - \pi\|_{L_2(\pi)}$. For more details see Proposition 2 and [29].

We abbreviate *geometric ergodicity* and *geometrically ergodic* as ‘GE’ for brevity going forward.

3.2. Assumptions

We assume throughout that P is the transition kernel for a Markov chain on a countably generated state space \mathcal{X} with σ -algebra Σ , which is reversible with respect to a stationary probability measure π , and is π -irreducible and aperiodic. We call the Markov chain induced by P the ‘original’ chain. The π -reversibility of P makes it natural to work in $L_2(\pi)$, since in this case P is a self-adjoint linear operator on a Hilbert space. This allows us access to the rich, elegant, and mature spectral theory of such operators. See for example [33, Chapter 12] and [9, Chapter 22]. We further assume that P is $L_2(\pi)$ -GE with factor $0 < (1 - \alpha) < 1$. Equivalent definitions of $L_2(\pi)$ -GE are given in Proposition 2. This assumption is weaker than the Doeblin condition used by [18], which implies uniform ergodicity.

Next, we assume that P_ϵ is a second, ‘perturbed’ transition kernel, such that

$$\|P - P_\epsilon\|_{L_2(\pi)} \leq \epsilon$$

for some fixed $\epsilon > 0$, and that

$$P_\epsilon|_{L_2(\pi)} \in \mathcal{B}(L_2(\pi)),$$

i.e. that the perturbed transition kernel maps $L_2(\pi)$ measures to $L_2(\pi)$ measures. The norm condition quantifies the intuition that the perturbation is ‘small’. We assume that P_ϵ is π -irreducible and aperiodic. We demonstrate (in Theorem 1) that under these assumptions P_ϵ has a unique stationary distribution, denoted by π_ϵ , with $\pi_\epsilon \in L_2(\pi)$.

Note that when $\mu \in L_1(\pi)$ we have

$$\|\mu - \pi\|_{TV} = \frac{1}{2} \|\mu - \pi\|_{L_1(\pi)}.$$

On the other hand, $\|\cdot\|_{TV}$ applies to all bounded measures, while $\|\cdot\|_{L_1(\pi)}$ applies only to the subspace of $L_1(\pi)$ measures. Note also that if $\pi \sim \pi_\epsilon$ (the two measures are mutually absolutely continuous), then $L_1(\pi)$ and $L_1(\pi_\epsilon)$ are equal as spaces and their norms are always equal, so in this case we need not distinguish between them.

To summarize, we assume the following.

Assumption 1. (*Assumptions of Section 3.2.*)

- P is a Markov kernel that is
 - π -reversible for a probability measure π ,
 - irreducible and aperiodic,
 - $L_2(\pi)$ -GE with factor $(1 - \alpha)$;
- P_ϵ is a Markov kernel that is
 - irreducible and aperiodic,

- $P_\epsilon : L_2(\pi) \rightarrow L_2(\pi)$, and
- $\|P - P_\epsilon\|_{L_2(\pi)} < \epsilon$.

The assumption that $P_\epsilon : L_2(\pi) \rightarrow L_2(\pi)$ and that $\|P_\epsilon\|_{L_2(\pi)} < \infty$ may seem difficult to verify. However, the following proposition shows us that it is satisfied for P_ϵ constructed based on the Metropolis–Hastings algorithm with suitable *jump kernels*. As long as the jump kernel, J , has $\|J\|_{L_2(\pi)} < \infty$, the assumption will be satisfied. Therefore, this assumption is not excessively restrictive for MCMC applications. The jump kernel J describes the conditional distribution of a new point in the chain proposed from x given that the proposal is accepted; it is related to the proposal kernel, Q , by $\alpha(x)J(x, A) = \int_A a(x, y)Q(x, dy)$, where $a(x, y)$ is the Metropolis–Hastings acceptance ratio and $\alpha(x) = \int_{\mathcal{X}} a(x, y)Q(x, dy)$ is the implied local jump-intensity.

Proposition 1. *If $P_\epsilon(x, \cdot) = (1 - \alpha(x))\delta_x + \alpha(x)J(x, \cdot)$ with $\alpha : \mathcal{X} \rightarrow [0, 1]$ measurable, and $J : L_2(\pi) \rightarrow L_2(\pi)$ and $\|J\|_{L_2(\pi)} < \infty$, then*

$$\|P_\epsilon\|_{L_2(\pi)} \leq 1 + \|J\|_{L_2(\pi)}. \tag{2}$$

Proof of Proposition 1. Consider the operator A on $L_2(\pi)$ given by the formula $[\nu A](C) = \int_C \alpha(x)\nu(dx)$ for all measurable sets C . Its adjoint, A' , is given by the formula $[A'f](x) = \alpha(x)f(x)$ for all $x \in \mathcal{X}$ and $f \in L_2(\pi)$. Since $\alpha : \mathcal{X} \rightarrow [0, 1]$, we have $A' : L_2(\pi) \rightarrow L_2(\pi)$ with $\|A'\|_{L_2(\pi)} \leq 1$. Thus $A : L_2(\pi) \rightarrow L_2(\pi)$ with $\|A\|_{L_2(\pi)} \leq 1$. The same also holds for $I - A$. Now, $P_\epsilon = \tilde{A} + (I - A)J$, so

$$\|P_\epsilon\|_{L_2(\pi)} \leq 1 + \|J\|_{L_2(\pi)}. \quad \square$$

Verifying that $\|P - P_\epsilon\|_{L_2(\pi)}$ is finite and sufficiently small will be the main analytic burden faced when trying to apply our results to more general settings. The development of further tools to determine whether $\|P - P_\epsilon\|_{L_2(\pi)}$ is finite and to bound it quantitatively would be an interesting line of future research.

3.3. Convergence rates and closeness of stationary distributions

Theorem 1. (Geometric ergodicity of the perturbed chain and closeness of the stationary distributions in original norm, $L_2(\pi)$) *Under the assumptions of Section 3.2, if in addition $\epsilon < \alpha$, then $\pi_\epsilon \in L_2(\pi)$,*

$$0 \leq \|\pi - \pi_\epsilon\|_{L_2(\pi)} \leq \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}},$$

P_ϵ is $L_2(\pi)$ -geometrically ergodic with factor $1 - (\alpha - \epsilon)$, and for any initial probability measure $\mu \in L_2(\pi)$,

$$\|\mu P_\epsilon^n - \pi\|_{L_2(\pi)} \leq (1 - (\alpha - \epsilon))^n \|\mu - \pi_\epsilon\|_{L_2(\pi)} + \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}}.$$

The proof of this result is the content of Appendix A.1. We follow the derivation in [18] with minimal structural modification, though the technicalities must be handled differently and additional theoretical machinery is required. We use the fact that the existence of a spectral gap for the restriction of P to $L_{2,0}(\pi)$ yields an inequality of the same form as the uniform contractivity condition, but in the $L_2(\pi)$ -norm as opposed to the total variation norm (cf. Theorem 2.1 of [29]).

Remark 2. Bounds on the differences between measures in $L_2(\pi)$ -norm can be converted into bounds on the total variation distance, since, by Cauchy–Schwarz, for any measure λ and any signed measure $\nu \in L_2(\lambda)$, we have

$$\|\nu\|_{TV} = \frac{1}{2} \|\nu\|_{L_1(\lambda)} \leq \frac{1}{2} \|\nu\|_{L_2(\lambda)}.$$

Thus, for example, under the assumptions of Theorem 1,

$$\|\mu P_\epsilon^n - \pi\|_{TV} \leq \frac{1}{2} \left[(1 - (\alpha - \epsilon))^n \|\mu - \pi_\epsilon\|_{L_2(\pi)} + \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} \right].$$

Similarly, under the assumptions of Theorem 1, we find that P_ϵ is $(L_2(\pi), \|\cdot\|_{TV})$ -GE with factor $1 - (\alpha - \epsilon)$ (see Definition 2 below).

In some situations, such as the computation of mean squared errors in Theorem 5, it may be inconvenient or impossible to use the $L_2(\pi)$ norm when studying some aspects of P_ϵ . The next theorem will allow us to ‘switch’ to other norms which may be more natural for a given task. First, however, we need to introduce one more notion of geometric ergodicity.

Definition 2. ($(V, \|\cdot\|)$ -geometric ergodicity.) Let P be the kernel of a positive recurrent Markov chain with invariant measure π . Let V be a vector space of signed measures on (\mathcal{X}, Σ) containing π , and let $\|\cdot\|$ be a norm on V (for which V may not be complete).

P is $(V, \|\cdot\|)$ -geometrically ergodic with factor ρ if there exists $C : V \cap \mathcal{M}_{+,1} \rightarrow \mathbb{R}_+$ such that for every $\nu \in V \cap \mathcal{M}_{+,1}$ and for all $n \in \mathbb{N}$,

$$\|\nu P^n - \pi\| \leq C(\nu)\rho^n.$$

The *optimal rate* for $(V, \|\cdot\|)$ -geometric ergodicity is the infimum over factors for which the above definition holds:

$$\rho^* = \inf \left\{ \rho > 0 : \exists C : V \cap \mathcal{M}_{+,1} \rightarrow \mathbb{R}_+ \text{ s.t. } \forall n \in \mathbb{N}, \nu \in V \cap \mathcal{M}_{+,1} \quad \|\nu P^n - \pi\| \leq C(\nu)\rho^n \right\}.$$

We will be interested in this definition for the cases that $V = L_\infty(\pi)$ and $\|\cdot\|$ is either $\|\cdot\|_{L_2(\pi)}$ or $\|\cdot\|_{L_1(\pi)}$.

Remark 3. (Relationships between $(L_\infty(\lambda), \|\cdot\|_{L_p(\lambda)})$ -GE, a.e.-TV-GE, and $L_2(\lambda)$ -GE.) Clearly if P is $L_2(\lambda)$ -GE with factor ρ_2 then it is also $(L_\infty(\lambda), \|\cdot\|_{L_2(\lambda)})$ -GE with factor ρ_2 . Conversely, Roberts and Tweedie [32] show that if P is $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE with factor ρ_2 , then it is also a.e.-TV-GE with some factor $\rho_{TV} \in (0, 1)$. However, the factor for a.e.-TV-GE may in fact be worse than the factor for $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE or $(L_\infty(\pi), \|\cdot\|_{L_1(\pi)})$ -GE. Baxendale [3] gives a detailed exposition on the barriers to the comparison of factors for geometric ergodicity given by different equivalent definitions.

In Appendix C we give an example where the optimal rates for $L_2(\pi)$ -GE and $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE are distinct when P is not reversible. If P is π -reversible then the factors for $L_2(\pi)$ -GE, $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE, and $(L_\infty(\pi), \|\cdot\|_{L_1(\pi)})$ -GE must be the same. This result combines a comment and Theorem 3 of [32], both stated but not proved. The formal statement of that result and its proof may be found in Appendix D.

Finally, note that by definition $L_2(\pi)$ -GE is equivalent to $(L_2(\pi), \|\cdot\|_{L_2(\pi)})$ with the same coefficient functions and factors, and that a.e.-TV-GE is equivalent to $(D, \|\cdot\|_{TV})$ -GE where

we can take $D = \text{span}(\{\pi\} \cup \{\delta_x : x \in \mathcal{X} \setminus N, r \in \mathbb{R}\})$ for some π -null set N . The null set N can be taken to be the same for all factors ρ by taking the union over the null sets for factors $\rho \in \mathbb{Q}$ (since a countable union of null sets is still null).

Lemma 1. (Characterization of optimal rates for $(V, \|\cdot\|)$ -GE chains) *If P is $(V, \|\cdot\|)$ -GE with stationary measure π , then the optimal rate for $(V, \|\cdot\|)$ -GE is equal to*

$$\sup_{\mu \in V \cap \mathcal{M}_{+,1}} \limsup_{n \rightarrow \infty} \|\mu P^n - \pi\|^{1/n}. \tag{3}$$

The proof of this result is found in Appendix D.

Remark 4. The quantity $\limsup_{n \rightarrow \infty} \|\mu P^n - \pi\|^{1/n}$ is the local spectral radius of $P - \Pi$ at μ with respect to $\|\cdot\|$, where Π is the rank-1 kernel defined by $\Pi(x, A) = \pi(A)$ for all $x \in \mathcal{X}$ and $A \in \Sigma$.

Lemma 2. ($L_2(\pi)$ -GE, $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE, and $(L_\infty(\pi), \|\cdot\|_{L_1(\pi)})$ -GE are equivalent for π -reversible chains, with equal optimal rates) *Let $\rho \in [0, 1)$. The following are equivalent for a π -reversible Markov chain P :*

- (i) P is $(L_\infty(\pi), \|\cdot\|_{L_1(\pi)})$ -GE with optimal rate ρ .
- (ii) P is $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE with optimal rate ρ .
- (iii) P is $L_2(\pi)$ -GE with optimal rate ρ .
- (iv) The spectral radius of $P|_{L_{2,0}(\pi)}$ is equal to ρ .

Remark 5. Since either of (iii) or (iv) is equivalent to all the conditions listed in [29, Theorem 2.1], indeed all the items listed above are equivalent to all the items listed in their result. We only included (iii) and (iv) here for brevity, since these are the conditions most relevant to the present paper. Moreover, all of these conditions are implied by any of the equivalent conditions for π -a.e.-TV-GE in [29, Proposition 2.1] (though with possibly different optimal rates for each condition therein). The proof of this result is found in Appendix D.

Theorem 1 controls the convergence of the perturbed chain P_ϵ in terms of the ‘original’ norm (from $L_2(\pi)$). We also demonstrate that P_ϵ is geometrically ergodic in the $L_2(\pi_\epsilon)$ norm, as this would also allow us to use the equivalences in [29]. The following two results allow us to transfer the geometric ergodicity of P_ϵ in $L_2(\pi)$ to other notions of geometric ergodicity. Theorem 3 handles the case that the perturbed kernel is reversible, while Theorem 2 handles both the case that the perturbed kernel is reversible and the case that it is non-reversible.

Theorem 2. (Geometric ergodicity of the perturbed chain in the other norms: $L_1(\pi_\epsilon)$, $L_2(\pi_\epsilon)$, total variation) *Under the assumptions of Section 3.2, if $\epsilon < \alpha$, then the following hold:*

- (i) P_ϵ is a.e.-TV-GE with some factor $\rho_{\text{TV}} \in (0, 1)$.
- (ii) P_ϵ is $(L_\infty(\pi_\epsilon), \|\cdot\|_{L_1(\pi_\epsilon)})$ -GE with factor $\rho_1 = (1 - (\alpha - \epsilon))$ and with $C_1(\mu) = \|\mu - \pi_\epsilon\|_{L_2(\pi)}$.
- (iii) If $\pi \in L_\infty(\pi_\epsilon)$ then P_ϵ is $L_2(\pi_\epsilon)$ -GE with factor $\rho_2 = (1 - (\alpha - \epsilon))$ and with

$$C_2(\mu) = \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \|\mu - \pi\|_{L_2(\pi)}.$$

The proof of this result is found in Appendix A.2.

Example 1. For example, consider perturbations of a Gaussian AR(1) process. Let $Z_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$ and let $W_i \stackrel{iid}{\sim} \mu$. Take

$$\begin{aligned} X_{t+1} | X_t &= (1 - \alpha)X_t + Z_{t+1}, \\ X_{t+1}^\epsilon | X_t^\epsilon &= (1 - \alpha)X_t^\epsilon + W_{t+1}. \end{aligned} \tag{4}$$

Then the original chain $\{X_t\}_{t \in \mathbb{N}}$ is not uniformly ergodic, but it is geometrically ergodic. Hence, the results of [1, 18] do not apply. The stationary measure of the exact chain is

$$\pi \equiv \mathcal{N}\left(0, \frac{\sigma^2}{\alpha(2 - \alpha)}\right),$$

it is reversible, and the rate of geometric ergodicity is $(1 - \alpha)$. Note that the perturbed chain, which we will call a μ AR(1) process, may not be reversible, and whether it is geometrically ergodic generally depends on the distribution μ .

Now, letting ϕ_{σ^2} be the $\mathcal{N}(0, \sigma^2)$ density, for any μ with $\frac{d\mu}{d\phi_{\sigma^2}} \in [1 - \epsilon, 1 + \epsilon]$,

$$\begin{aligned} \|P - P_\epsilon\|_{L_2(\pi)}^2 &= \int_{-\infty}^\infty \int_{-\infty}^\infty \left(\frac{\mu(y - (1 - \alpha)x)}{\pi(y)} - \frac{\phi_{\sigma^2}(y - (1 - \alpha)x)}{\pi(y)} \right)^2 \pi(y) dy \pi(x) dx \\ &\leq \int_{-\infty}^\infty \int_{-\infty}^\infty \epsilon^2 \left(\frac{\phi_{\sigma^2}(y - (1 - \alpha)x)}{\pi(y)} dy \right)^2 \pi(y) dy \pi(x) dx \\ &= \epsilon^2 \|P\|_{L_2(\pi)} \\ &= \epsilon^2. \end{aligned} \tag{5}$$

Therefore, when $\epsilon < \alpha$ we can extend the geometric ergodicity of the Gaussian AR process to the $\mu - \text{AR}(1)$ process using Theorem 2. We can also bound the discrepancy of the stationary measure of the perturbed chain from that $\mathcal{N}(0, \frac{\sigma^2}{\alpha(2-\alpha)})$ using Theorem 1. The subsequent results, Corollary 1 and Theorem 4 of this section, may also be applied to this example to bound the discrepancy between the marginal distributions of the μ -AR(1) from an $\mathcal{N}(0, \frac{\sigma^2}{\alpha(2-\alpha)})$ at any time, as well as the approximation error of the time-averaged law of the μ -AR(1) from $\mathcal{N}(0, \frac{\sigma^2}{\alpha(2-\alpha)})$.

Theorem 3. ($L_2(\pi_\epsilon)$ -GE of the perturbed chain, reversible case.) *Under the assumptions of Section 3.2, if $\epsilon < \alpha$ and P_ϵ is π_ϵ -reversible, then P_ϵ is $L_2(\pi_\epsilon)$ -GE with factor $\rho_2 = (1 - \alpha + \epsilon)$ and coefficient function $C(v) = \|v\|_{L_2(\pi_\epsilon)}$. The proof of this result is found in Appendix A.2.*

Corollary 1. (Closeness of stationary distributions in $L_2(\pi_\epsilon)$.) *If $\epsilon < \alpha$ and $\|P - P_\epsilon\|_{L_2(\pi_\epsilon)} \leq \varphi$, then the following hold:*

(i) *If P_ϵ is π_ϵ reversible and $\varphi < \alpha - \epsilon$, then*

$$\|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \leq \frac{\varphi}{\sqrt{(\alpha - \epsilon)^2 - \varphi^2}},$$

and for any $\mu \in L_2(\pi_\epsilon)$,

$$\|\mu P_\epsilon^n - \pi\|_{L_2(\pi_\epsilon)} \leq (1 - (\alpha - \epsilon))^n \|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)} + \frac{\varphi}{\sqrt{(\alpha - \epsilon)^2 - \varphi^2}}.$$

(ii) If $\pi \in L_\infty(\pi_\epsilon)$ and $\varphi < 1$, then

$$\|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \leq \frac{\varphi + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon))}{1 - \varphi},$$

and for any $\mu \in L_\infty(\pi_\epsilon)$,

$$\begin{aligned} \|\mu P_\epsilon^n - \pi\|_{L_2(\pi_\epsilon)} &\leq (1 - (\alpha - \epsilon))^n \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \\ &\quad + \frac{\varphi + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon))}{1 - \varphi}. \end{aligned}$$

The proof of this result is found in Appendix A.2. We turn our attention to bounds on the error of estimation measures of the form $\frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k$, and estimates of the form $\frac{1}{t} \sum_{k=0}^{t-1} f(X_k)$. Firstly, when computing Monte Carlo estimates, the bias is controlled by a time-averaged marginal distribution of the form $\frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k$. This leads us to the following result.

Theorem 4. (Convergence of time-averaged marginal distributions) *Under the assumptions of Section 3.2, suppose $\epsilon < \alpha$ and $\pi_\epsilon \in L_2(\pi)$. Then for any probability distribution $\mu \in L_2(\pi)$,*

$$\left\| \pi - \frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k \right\|_{L_2(\pi)} \leq \frac{1 - (1 - (\alpha - \epsilon))^t}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi)} + \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}}.$$

If additionally $\|P - P_\epsilon\|_{L_2(\pi_\epsilon)} \leq \varphi$, then the following hold:

(i) If P_ϵ is π_ϵ -reversible and $\varphi < \alpha - \epsilon$, then

$$\left\| \pi - \frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k \right\|_{L_2(\pi_\epsilon)} \leq \frac{1 - (1 - (\alpha - \epsilon))^t}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} + \frac{\varphi}{\sqrt{(\alpha - \epsilon)^2 - \varphi^2}}.$$

(ii) If $\pi \in L_\infty(\pi_\epsilon)$ and $\varphi < 1$, and if $\mu \in L_\infty(\pi_\epsilon)$, then

$$\begin{aligned} \left\| \pi - \frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k \right\|_{L_2(\pi_\epsilon)} &\leq \frac{1 - (1 - (\alpha - \epsilon))^t}{t(\alpha - \epsilon)} \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} \\ &\quad + \frac{\varphi + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon))}{1 - \varphi}. \end{aligned}$$

The proof of this result is found in Appendix A.3.1. Relative to the uniform closeness of kernels (in total variation) required in [18], our assumption that the approximating kernel is close in the operator norm induced by $L_2(\pi)$ is non-comparable. This is because our bound is in terms of the L_2 distance, which always upper-bounds the total variation distance (up to a constant factor of 1/2), but our assumption also does not require spatial uniformity, which that of [18] does. Thus, this paper’s assumptions are neither weaker nor stronger than those in [18]. Comparing the above results to the corresponding L_1 result of [18], we see that the transient phase bias part of our L_2 bounds differ from the L_1 transient phase bias bound of [18] only by a factor which is constant in time, but varies with the initial distribution (as is to be expected when moving from uniform ergodicity to geometric ergodicity).

3.4. Mean squared error bounds for Monte Carlo estimates

Suppose that $(X_k^\epsilon)_{k \in \mathbb{N} \cup \{0\}}$ is a realization of the Markov chain with transition kernel P_ϵ and initial distribution μ . The mean squared error of a Monte Carlo estimate of πf made using $(X_k^\epsilon)_{k \leq t}$ is given by

$$MSE_t^\epsilon(\mu, f) = \mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k^\epsilon) \right)^2 \right]. \tag{6}$$

Theorem 5. (Mean squared error of Monte Carlo estimates from the perturbed chain) *Under the assumptions of Section 3.2, if $\epsilon < \alpha$, $X_0^\epsilon \sim \mu$, P_ϵ is π_ϵ -reversible, and $\rho_2 = (1 - (\alpha - \epsilon))$, then for $f \in L'_4(\pi_\epsilon)$,*

(i) *if $f \in L'_2(\pi)$ as well, then*

$$MSE_t^\epsilon(\mu, f) \leq \frac{2\|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)}^2}{(1 - \rho_2)t} + \frac{2^{7/2}\|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)}\|f - \pi_\epsilon f\|_{L'_4(\pi_\epsilon)}^2}{(1 - \rho_2)^2 t^2} + \|f - \pi_\epsilon f\|_{L'_2(\pi)}^2 \left(\frac{\epsilon^2}{\alpha^2 - \epsilon^2} + 2 \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} \frac{1}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi)} \right)$$

and

$$MSE_t^\epsilon(\mu, f) \leq \frac{4\|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)}^2}{(1 - \rho_2)t} + \frac{2^{9/2}\|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)}\|f - \pi_\epsilon f\|_{L'_4(\pi_\epsilon)}^2}{(1 - \rho_2)^2 t^2} + 2\|f - \pi_\epsilon f\|_{L'_2(\pi)}^2 \frac{\epsilon^2}{\alpha^2 - \epsilon^2};$$

(ii) *if $\|P - P_\epsilon\|_{L_2(\pi_\epsilon)} \leq \varphi < (1 - \rho_2)$, then*

$$MSE_t^\epsilon(\mu, f) \leq \frac{2^{7/2}\|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)}\|f - \pi_\epsilon f\|_{L'_4(\pi_\epsilon)}^2}{(1 - \rho_2)^2 t^2} + \|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)}^2 \left(\frac{\varphi^2}{(1 - \rho_2)^2 - \varphi^2} + 2 \frac{1 + \frac{\varphi}{\sqrt{(1 - \rho_2)^2 - \varphi^2}}}{t(1 - \rho_2)} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} \right),$$

and

$$MSE_t^\epsilon(\mu, f) \leq \frac{2^{9/2}\|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)}\|f - \pi_\epsilon f\|_{L'_4(\pi_\epsilon)}^2}{(1 - \rho_2)^2 t^2} + \|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)}^2 \left(\frac{2\varphi^2}{(1 - \rho_2)^2 - \varphi^2} + \frac{4}{t(1 - \rho_2)} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} \right).$$

The proof of this result is found in Appendix A.3.3. Perturbation bounds based upon drift and minorization conditions could provide similar mean squared error bounds for functions in $L_2(\pi_\epsilon)$ with

$$\sup_{x \in \mathcal{X}} \frac{|f|}{\sqrt{V}} < \infty$$

(where V is the function appearing in the drift condition), as in [17]. While that may be a larger class of functions than $L'_4(\pi_\epsilon)$ (depending on what V happens to be), the class $L'_4(\pi_\epsilon)$ is quite rich, making this bound still useful. Moreover, the class of functions to which our mean squared error bounds apply, and the value of the bound itself, depend only on intrinsic features of the Markov chains under consideration. In contrast, bounds based on drift and minorization conditions include extrinsic features—introduced by the user for analytic purposes (such as the drift function, V)—for which many choices might exist, each leading to different function classes and different bounds.

4. Applications to Markov chain Monte Carlo

In this section we apply our theoretical results to some specific variants of MCMC algorithms to obtain guarantees for noisy and/or approximate variants of MCMC algorithms. MCMC is used to generate (correlated) samples approximately from a target distribution for which the (unnormalized) density can be evaluated. The key insight is to construct a (typically reversible) Markov chain for which the stationary distribution is the target distribution. This is possible since the reversibility condition is readily verified locally (without integration).

The most commonly used family of MCMC methods is the Metropolis–Hastings (MH) algorithm. The chain is initialized from some distribution $X_0 \sim \mu_0$. At each step a *proposal* is drawn from some transition kernel, $Y_t \sim Q(X_{t-1}, \cdot)$. Suppose that the kernel $Q(x, \cdot)$ has density $q(\cdot | x)$. The proposal is *accepted* with probability

$$a(Y_t | X_{t-1}) = \min \left(1, \frac{\pi(Y_t)q(X_{t-1} | Y_t)}{\pi(X_{t-1})q(Y_t | X_{t-1})} \right).$$

If the proposal is accepted then $X_t = Y_t$, and if it is *rejected* (not accepted) then $X_t = X_{t-1}$. The combination of proposal and accept/reject steps yields a π -reversible Markov kernel, and reversibility guarantees that the stationary distribution is the target distribution. The user has freedom in selecting the proposal kernel, Q , and some choices lead to better performance than others. The accept/reject step requires evaluating the target density, π , twice on each step.

A large body of research exists guaranteeing that specific MCMC algorithms will be geometrically ergodic (see for example [13,22,30], and many more.). These typically verify geometric ergodicity for a collection of target distributions, π , and for a small family of proposal kernels, Q .

If the target likelihood involves some integral which is computed numerically or by simple Monte Carlo, then the numerical and/or stochastic approximation introduces a *perturbation* to the idealized MCMC scheme. This occurs even in standard and widely used statistical models such as generalized linear mixed effect models (GLMMs), since the random effects are nuisance variables which need to be integrated away, either using Laplace or Gaussian quadrature schemes, or by simple Monte Carlo, in order to evaluate the likelihood. Since the MH algorithm requires evaluation of the density, these each introduce a perturbation in the acceptance ratio, and hence in the actual transition kernel of the MH scheme. We now consider the extent to which our results from Section 3 can be applied to prove geometric ergodicity for certain approximate MCMC algorithms.

4.1. Noisy and approximate MCMC

The noisy (or approximate) Metropolis–Hastings (nMH) algorithm, as found in [1] (see also [25]), was briefly described above. The algorithm is defined exactly the same way as the MH algorithm, except that the *acceptance ratio*, $a(Y_t | X_{t-1})$, is replaced by a (possibly stochastic)

approximation $\widehat{a}(Y_t|X_{t-1}, Z_t)$. Here Z_t denotes some random element providing an additional source of randomness, so that $a(Y_t|X_{t-1}, Z_t)$ is not $\sigma(Y_t, X_{t-1})$ -measurable when the approximation $\widehat{a}(Y_t|X_{t-1}, Z_t)$ is stochastic. In the case of a deterministic approximation, Z_t can be ignored or treated as a constant. The approximation can typically be thought of as replacing the target density in the acceptance ratio with some approximation. This includes most approximate MCMC algorithms which preserve the state space and the Markov property, such as those replacing π with a deterministic approximation or an independent stochastic approximation at each step (as in Monte Carlo within Metropolis). It does not include algorithms which retain the Markov property only for an augmented state space, such as the pseudo-marginal approach of [2].

For our analysis of these algorithms, P will represent the transition kernel for the MH algorithm, while \widehat{P} will represent the kernel for the corresponding nMH chain. The key step in applying our results from Section 3 will be to show the $L_2(\pi)$ closeness of the nMH transition kernel to the MH transition kernel. Again, $\|\cdot\|_{L_2(\pi)}$ is the norm on $L_2(\pi)$ and the corresponding operator norm. We will assume that π and $\{Q(x, \cdot)\}_{x \in \mathcal{X}}$ are all absolutely continuous with respect to the Lebesgue measure and have densities π and $\{q(\cdot|x)\}_{x \in \mathcal{X}}$ respectively. All arguments used would still apply if there were an arbitrary dominating measure in place of the Lebesgue measure. Let $F_{y|x}$ be the regular conditional distribution for Z given $X = x$ and $Y = y$, and let $f_{y|x}$ be its Lebesgue density. Define the *perturbation function* for the nMH algorithm as

$$r(y|x) = \mathbb{E}_{Z \sim F_{y|x}}(a(y|x) - \widehat{a}(y|x, Z)) = \int (a(y|x) - \widehat{a}(y|x, z)) f_{y|x}(z) dz.$$

Theorem 6. (Geometric ergodicity and closeness of stationary distributions for nMH) *Let P be the transition kernel for an MH algorithm with proposal distribution Q , target distribution π , and acceptance ratio $a(\cdot|\cdot)$. Let \widehat{P} be the transition kernel for a corresponding nMH algorithm with approximate/noisy acceptance ratio $\widehat{a}(\cdot|\cdot, \cdot)$. Let $r(\cdot|\cdot)$ be the corresponding perturbation function.*

If $\|Q\|_{L_2(\pi)} < \infty$ and $\sup_{x,y} |r(y|x)| \leq R$, then

$$\|\widehat{P} - P\|_{L_2(\pi)} \leq R(1 + \|Q\|_{L_2(\pi)}). \tag{7}$$

Furthermore, if P is reversible and $L_2(\pi)$ -GE with geometric contraction factor $(1 - \alpha)$, and $\epsilon = R(1 + \|Q\|_{L_2(\pi)}) < \alpha$, then \widehat{P} has a stationary distribution $\widehat{\pi}$, and the assumptions outlined in Section 3.2 hold with $P_\epsilon = \widehat{P}$ and $\pi_\epsilon = \widehat{\pi}$.

Therefore, Theorems 1 to 5 and Corollary 1 can all be applied. In particular, \widehat{P} is $L_2(\pi)$ -GE with factor $1 - (\alpha - R(1 + \|Q\|_{L_2(\pi)}))$, it is a.e.-TV-GE, and

$$\|\widehat{\pi} - \pi\|_{L_2(\pi)} \leq \frac{R(1 + \|Q\|_{L_2(\pi)})}{\sqrt{\alpha^2 - R^2(1 + \|Q\|_{L_2(\pi)})^2}}; \tag{8}$$

furthermore, if \widehat{P} is reversible, then it is $L_2(\widehat{\pi})$ -GE with factor $(1 - (\alpha - R(1 + \|Q\|_{L_2(\pi)})))$.

The above theorem provides an alternative to the analogous result of Corollary 2.3 from [1], relaxing the uniform ergodicity assumption. In particular, it requires that $Q \in \mathcal{B}(L_2(\pi))$ and that $R(1 + \|Q\|_{L_2(\pi)}) < \alpha$. The first of these requirements is not dramatically limiting since the user has control over the choice of Q . The second requirement is also not dramatically limiting, as control over R may be interpreted as limiting the amount of noise in the nMH algorithm, and such control is required regardless in order to ensure the accuracy of approximation in both the geometrically ergodic and uniformly ergodic cases.

4.2. Application to fixed deterministic approximations

Suppose we run a fixed MH algorithm, but replace the target density with one which is close everywhere. Perhaps this alternative density is easier to compute (e.g. replacing an integral with a Laplace approximation as in [19], or replacing a full sample with a coresets for sub-sampled Bayesian inference as in [8]). By construction we would know that the approximate target distribution is close to the ideal target distribution. The question still remains whether geometric ergodicity is preserved. We resolve this question in the case that the approximation has constant relative error.

Corollary 2. *Suppose we can approximate the unnormalized target density, $C\pi$, by $\hat{\pi}$, with a θ -bounded relative error:*

$$\sup_{x \in \mathcal{X}} \left| \log \frac{C\pi(x)}{\hat{\pi}(x)} \right| \leq \theta. \tag{9}$$

If the MH algorithm with proposal kernel Q is $L_2(\pi)$ -GE with factor $(1 - \alpha)$, and if

$$\theta < \frac{\alpha}{2(1 + \|Q\|_{L_2(\pi)})},$$

then the corresponding approximate transition kernel, \hat{P} , is $L_2(\hat{\pi})$ -GE, and

$$\|\hat{\pi} - \pi\|_{L_2(\pi)} \leq \frac{2\theta(1 + \|Q\|_{L_2(\pi)})}{\sqrt{\alpha^2 - 4\theta^2(1 + \|Q\|_{L_2(\pi)})^2}}. \tag{10}$$

Proof. Since the function $x \mapsto 1 \wedge \exp(x)$ is 1-Lipschitz, we have

$$\begin{aligned} |r(y|x)| &= |a(y|x) - \hat{a}(y|x)| \\ &\leq \left| \log \frac{\pi(y)q(x|y)}{\pi(x)q(y|x)} - \log \frac{\hat{\pi}(y)q(x|y)}{\hat{\pi}(x)q(y|x)} \right| \\ &= \left| \log \frac{C\pi(y)}{\hat{\pi}(y)} - \log \frac{C\pi(x)}{\hat{\pi}(x)} \right| \\ &\leq 2\theta. \end{aligned} \tag{11}$$

So \hat{P} will be $L_2(\pi)$ -GE as long as P was geometrically ergodic with some factor $0 \leq (1 - \alpha) < 1$ and

$$\theta < \frac{\alpha}{2(1 + \|Q\|_{L_2(\pi)})}. \tag{12}$$

Moreover, in this case, \hat{P} is reversible. Thus, we can use Theorem 3 to obtain $L_2(\hat{\pi})$ -GE of \hat{P} , with factor $1 - \alpha + 2\theta(1 + \|Q\|_{L_2(\pi)})$. \square

In this scenario, we can also use Theorem 5 to get quantitative bounds for the mean squared error of any Monte Carlo estimates made using \hat{P} , or any of our other results in Theorems 1 to 4 and Corollary 1 as needed.

Example 2. (Independence sampler.) The previous result also immediately gives that if $\frac{d\hat{\pi}}{d\pi}$ is bounded above by $C < \exp(1/4)$ and below by $c > \exp(-1/4)$, then the independence sampler for $\hat{\pi}$ with proposals from π is geometrically ergodic with factor $4\max(\log C, -\log(c))$. This is, however, sub-optimal when compared to [36], which only requires a finite upper bound on $\frac{d\hat{\pi}}{d\pi}$ to establish *uniform ergodicity*.

Example 3. (Laplace approximation for GLMMs.) Generalized linear mixed models (GLMMs) (see [5,23], etc.) are widely used in the modelling of non-normal response variables under repeated or correlated measurements. They are the natural common extension of generalized linear models and linear mixed effects models. They handle dependence between observations by introducing Gaussian latent variables. These *random effects* are nuisance variables for the purpose of inference. In order to perform Bayesian inference for GLMMs, one requires samples from the marginal posterior distribution of the parameters given the data. The marginal posterior, here, is the posterior for the parameters given the observations, in contrast to the joint posterior of the random effects and the parameters given the data.

This can be approached in two ways. One option is to obtain samples for the random effects and parameters jointly given the data, and discard the random effects to get marginal posterior samples for the parameters. The second option is to approximate the likelihood by integrating (numerically) over the random effects, and using the resulting approximate likelihood in the calculations involving the unnormalized posterior for the parameters.

In the second case, when the prior for the parameters is compactly supported, if one had established a result saying that a particular MH procedure for the exact posterior distribution of the parameters would be geometrically ergodic, then one could directly transfer this result to the approximate posterior computed using a Laplace approximation, at least for large enough samples. This is valid since the Laplace approximation has constant relative error on compact sets, and the relative error decreases with sample size (see [37]). Hence, for a large enough sample size, Equation (12) will be satisfied regardless of what the proposal kernel Q was (as long as $\|Q\|_{L_2(\pi)}$ was finite).

Example 4. (Uniform coresets.) In Bayesian inference with large samples, an approach to reducing the computational burden of evaluating the likelihood in the unnormalized posterior for MCMC accept/reject steps is to select a representative subsample of the data and to up-weight the contributions of each of the selected samples in such a way as to best approximate the original likelihood. These up-weighted subsamples are called *coresets*. They naturally give rise to approximate MCMC methods in which the true posterior is replaced by an approximation based upon a coreset. Several methods for coreset construction exist; however, relatively little work has been done to assess their impact upon approximate MCMC methods. We will consider the *uniform coreset* construction of [14] (so named in [8]).

Theorem 3.2 of [8] provides the guarantee that, with probability $(1 - \delta)$, the unnormalized approximate posterior $\hat{C}\hat{\pi}$ based on a uniform coreset of size M will satisfy

$$\sup_{x \in \mathcal{X}} \frac{1}{|\mathcal{L}(x)|} \left| \log \frac{\hat{C}\hat{\pi}(x)}{C\pi(x)} \right| \leq \frac{\sigma}{\sqrt{M}} \left(\frac{3}{2}D + \bar{\eta}\sqrt{2 \log(1/\delta)} \right), \tag{13}$$

where $\sigma = \sum_{n=1}^N \sigma_n$, N is the number of observations,

$$\sigma_n = \sup_{x \in \mathcal{X}} \left| \frac{\mathcal{L}_i(x)}{\mathcal{L}(x)} \right|,$$

$\mathcal{L}_i(x)$ is the log-likelihood of parameter x at the i th observation, $\mathcal{L}(x) = \sum_{i=1}^N \mathcal{L}_i(x)$ is the log-likelihood of the dataset

$$\bar{\eta} = \max_{i,j \in \{1, \dots, N\}} \sup_{x \in \mathcal{X}} \left| \frac{1}{|\mathcal{L}(x)|} \frac{\mathcal{L}_i(x)}{\sigma_i} - \frac{\mathcal{L}_j(x)}{\sigma_j} \right|, \tag{14}$$

and D is the *approximate dimension* of $\{\mathcal{L}_i\}_{i=1}^n$ ([8, Definition 3.1]).

If, in addition to assuming that $\{\sigma_i\}_{i=1}^N$ are all finite as in [8, Section 3], one were to assume that $|\mathcal{L}(x)|$ is bounded as a function of x , then the uniform coresets result would imply the conditions of our Corollary 2, namely that

$$\sup_{x \in \mathcal{X}} \left| \log \frac{C\pi(x)}{\widehat{\pi}(x)} \right| \leq \frac{\sigma \|\mathcal{L}\|_\infty}{\sqrt{M}} \left(\frac{3}{2}D + \bar{\eta} \sqrt{2 \log(1/\delta)} \right), \tag{15}$$

with high probability. Consequently, for any proposal kernel $Q : L_2(\pi) \rightarrow L_2(\pi)$ we should be able to choose M sufficiently large so that with high probability

$$\frac{\sigma \|\mathcal{L}\|_\infty}{\sqrt{M}} \left(\frac{3}{2}D + \bar{\eta} \sqrt{2 \log(1/\delta)} \right) < \frac{\alpha}{2(1 + \|Q\|_{L_2(\pi)})}. \tag{16}$$

Hence the approximating Markov chain will be geometrically ergodic with high probability.

4.3. Application to Monte Carlo within Metropolis

Following [24], we can get bounds for the simple Monte Carlo within Metropolis algorithm (MCwM). This is the special case of nMH where we approximate the likelihood ratio

$$\frac{\pi(y)}{\pi(x)} = \frac{\mathbb{E}\Pi(y, Z)}{\mathbb{E}\Pi(x, Z)}$$

by

$$\left(\frac{\pi(y)}{\pi(x)} \right) = \frac{\sum_{i=1}^N \Pi(y, Z_i)}{\sum_{i=N+1}^{2N} \Pi(x, Z_i)}$$

using a new independent sample taken each time the likelihood is evaluated. In the notation of the previous section,

$$\widehat{a}(y|x, z) = 1 \wedge \frac{q(x|y) \sum_{i=1}^N \Pi(y, z_i)}{q(y|x) \sum_{i=N+1}^{2N} \Pi(x, z_i)}. \tag{17}$$

Let

$$\begin{aligned} W_k(x) &= \frac{1}{k\pi(x)} \sum_{i=1}^k \Pi(x, Z_i), \\ i_k(x)^2 &= \mathbb{E}[W_k(x)^{-2}], \\ s(x) &= \frac{1}{\sqrt{\pi(x)}} \text{StdDev}(\Pi(x, Z_1)). \end{aligned} \tag{18}$$

Lemma 14 of [24] tells us that if there is a $k \in \mathbb{N}$ such that $i_k(x) < \infty$ for all $x \in \mathcal{X}$, then for $N \geq k$,

$$\begin{aligned} |r(y|x)| &\leq a(y|x) \frac{1}{\sqrt{N}} i_k(y) (s(x) + s(y)) \\ &\leq \frac{1}{\sqrt{N}} i_k(y) (s(x) + s(y)). \end{aligned} \tag{19}$$

Corollary 3. *Let P be the MH transition kernel for the target density π and proposal kernel Q . Let \widehat{P}_N be the corresponding MCwM transition kernel when $\pi(\cdot)$ is approximated by $\frac{1}{N} \sum_{i=1}^N \Pi(\cdot, Z_i)$.*

Assume that s and the i_k as defined above are uniformly bounded for some $k \in \mathbb{N}$. Suppose further that

$$N_0 = \max\left(k, \frac{4\|i_k\|_\infty^2 \|s\|_\infty^2 (1 + \|Q\|_{L_2(\pi)})^2}{\alpha^2}\right),$$

and $N \geq \lfloor N_0 \rfloor + 1$.

Then \widehat{P}_N is reversible and $L_2(\pi)$ -GE with factor

$$1 - \alpha + \frac{1}{\sqrt{N/N_0}},$$

and it has a stationary distribution

$$\widehat{\pi}_N(x) \propto \frac{\pi(x)}{N \mathbb{E} \left[\left(\sum_{i=1}^N \Pi(x, Z_i) \right)^{-1} \right]}$$

with

$$\|\pi - \widehat{\pi}_N\|_{L_2(\pi)} \leq \sqrt{\frac{N_0}{N\alpha^2 - N_0}}. \tag{20}$$

Proof. Suppose that $N \geq \lfloor N_0 \rfloor + 1$. From Theorem 1, we know that the perturbed chain \widehat{P}_N is $L_2(\pi)$ -GE with factor $1 - \alpha + \frac{1}{\sqrt{N/N_0}}$ and has a stationary distribution $\widehat{\pi}_N$ with

$$\|\pi - \widehat{\pi}_N\|_{L_2(\pi)} \leq \sqrt{\frac{N_0}{N\alpha^2 - N_0}}. \tag{21}$$

Moreover, by inspection, \widehat{P}_N is reversible with respect to

$$\widehat{\pi}_N(x) \propto \frac{\pi(x)}{N \mathbb{E} \left[\left(\sum_{i=1}^N \Pi(x, Z_i) \right)^{-1} \right]}.$$

Thus, we can use Theorem 3 to obtain $L_2(\widehat{\pi}_N)$ -GE of \widehat{P}_N , with factor $1 - \alpha + \frac{1}{\sqrt{N/N_0}}$. □

Remark 6. A simple scenario under which these i_k and s are uniformly bounded is when the joint density of x and Z is bounded above and below by a multiple of the marginal of x , so that

$$\frac{\Pi(x, z)}{\pi(x)} \in [c, C] \tag{22}$$

for all $(x, z) \in \mathcal{X} \times \mathcal{Z}$. This condition is essentially tight if we wish to take $k = 1$ and take the base measure to be the Lebesgue measure restricted to $U \subset \mathbb{R}^d$; in this case the condition $\|i_k(x)\|_{L_\infty} < \infty$ implies that

$$\int_U \frac{\pi(x)}{\Pi(x, z)} dz = \mathbb{E}_{Z \sim \frac{\Pi(x, \cdot)}{\Pi(x, \cdot)}} \frac{\pi(x)^2}{\Pi(x, \cdot)^2} < \infty \tag{23}$$

for all x . That is, the reciprocal of the conditional density of Z given $X = x$ has a finite integral for each x .

Remark 7. More generally, [24, Lemma 23] tells us that if $\mathbb{E}[W_{k_0}(x)^{-p}] < \infty$ for some $k_0 \in \mathbb{N}$ and $p > 0$, then for $k \geq k_0 \lceil \frac{2}{p} \rceil$, $i_k(x)^2 < \mathbb{E}[W_{k_0}(x)^{-p}]$. Therefore, in order to uniformly bound $i_k(x)$, it is sufficient to bound $\mathbb{E}[W_{k_0}(x)^{-p}]$ uniformly in x for some $k_0 \in \mathbb{N}$, $p > 0$. This is much less restrictive than trying to bound $i_1(x)$. In the case that $p < 1$, $k_0 = 1$, this is much less restrictive than when $p = 2$, $k_0 = 1$; it is equivalent to requiring that tempered versions of the conditional distribution

$$\frac{\Pi(x, \cdot)}{\pi(x)}$$

can be normalized by uniformly bounded normalizing constants. This would be true if, for example, $(Z|X = x) \sim \mathcal{N}(\mu(x), \sigma^2(x))$ with $\sigma^2(x)$ uniformly bounded in x . More generally, using $0 < p < 1$ instead of $p = 2$, whenever the conditional law of Z has uniform exp-poly tails, i.e.

$$\frac{\Pi(x, z)}{\pi(x)} \leq \exp(-C|z - \mu(x)|^\alpha)$$

with $\alpha > 0$, the p -version of the condition would hold.

We could also use Theorem 5 to get quantitative bounds for the mean squared error of any Monte Carlo estimates made using \widehat{P}_N , or any of our other results in Theorems 1 to 4 and Corollary 1 as needed.

Medina-Aguayo *et al.* [24] also consider a case where the assumption that s and the i_k are uniformly bounded is dropped, and instead, the perturbed kernel is restricted to a bounded region. We do not address this case here.

Appendix A. Proofs

A.1 Proof of Theorem 1

The following lemma is contained in the remark after Theorem 2.1 of [29]; we prove it here as well since the proof is so simple.

Lemma 3. (Remark in [29].) *For any probability measure $\mu \in L_2(\pi)$,*

$$\|\mu - \pi\|_{L_2(\pi)}^2 = \|\mu\|_{L_2(\pi)}^2 - 1.$$

Proof.

$$\begin{aligned} 0 \leq \|\mu - \pi\|_{L_2(\pi)}^2 &= \int \left(\frac{d\mu}{d\pi} - 1\right)^2 d\pi = \int \left(\left(\frac{d\mu}{d\pi}\right)^2 - 2\frac{d\mu}{d\pi} + 1\right) d\pi \\ &= \int \left(\frac{d\mu}{d\pi}\right)^2 d\pi - 2 \int d\mu + \int d\pi = \|\mu\|^2 - 1 \|_{L_2(\pi)}. \end{aligned}$$

□

We will make use of the following simplified version of Theorem 2.1 from [29] as well.

Proposition 2. (Equivalent definitions of $L_2(\pi)$ -GE from [29].) *For a reversible Markov chain with kernel P and stationary distribution π on state space \mathcal{X} , the following are equivalent (and ρ is equal in both cases):*

- (i) P is $L_2(\pi)$ -GE with optimal rate ρ and coefficient function $C(\mu) = \|\mu - \pi\|_{L_2(\pi)}$.

(ii) P has $L_{2,0}(\pi)$ -spectral radius and norm both equal to ρ :

$$\sup_{v \in L_{2,0}(\pi) \setminus \{0\}} \frac{\|vP\|_{L_2(\pi)}}{\|v\|_{L_2(\pi)}} = \rho = r(P|_{L_{2,0}(\pi)}),$$

where

$$r(P|_{L_{2,0}(\pi)}) := \sup\{|\rho| : \rho \in \mathbb{C} \text{ and } (P|_{L_{2,0}(\pi)} - \rho I_{L_{2,0}(\pi)}) \text{ is not invertible}\}. \tag{24}$$

Note that while when the kernel is reversible we may take $C(\mu) = \|\mu - \pi\|_{L_2(\pi)}$ in the bound corresponding to $L_2(\pi)$ -GE with optimal rate ρ , this is not true for non-reversible chains. By applying the above theorem in our context we have the following.

Lemma 4. *Under the assumptions of Section 3.2,*

$$\|v_1 P^n - v_2 P^n\|_{L_2(\pi)} \leq (1 - \alpha)^n \|v_1 - v_2\|_{L_2(\pi)}$$

for any probability distributions $v_1, v_2 \in L_2(\pi)$. In particular, taking $v_2 = \pi$,

$$\|v_1 P^n - \pi\|_{L_2(\pi)} \leq (1 - \alpha)^n \|v_1 - \pi\|_{L_2(\pi)} = (1 - \alpha)^n \sqrt{\|v_1\|_{L_2(\pi)}^2 - 1},$$

and applying Cauchy–Schwarz yields

$$\|v_1 P^n - \pi\|_{L_1(\pi)} \leq \|v_1 P^n - \pi\|_{L_2(\pi)} \leq (1 - \alpha)^n \|v_1 - \pi\|_{L_2(\pi)}.$$

We begin with a first result giving sufficient conditions under which the stationary distribution π_ϵ of the perturbed chain is in $L_2(\pi)$.

Lemma 5. *Under the assumptions of Section 3.2, if in addition $\epsilon < \alpha$, then P_ϵ has a unique stationary distribution, $\pi_\epsilon \in L_2(\pi)$, and $\|\pi_\epsilon - \pi\|_{L_2(\pi)} \leq \frac{\epsilon}{\alpha - \epsilon}$.*

Proof. Since P_ϵ is π -irreducible and aperiodic, it has at most one stationary distribution, π_ϵ , with $\pi_\epsilon \ll \pi$ (see for example [9, Corollary 9.2.16]).

Suppose for now that πP_ϵ^n has an $L_2(\pi)$ limit, π_ϵ . Then, using the triangle inequality, the contraction property ($\|P_\epsilon\|_{TV} = 1$), and Cauchy–Schwarz, we have

$$\begin{aligned} \|\pi_\epsilon P_\epsilon - \pi_\epsilon\|_{TV} &\leq \|\pi_\epsilon P_\epsilon - \pi P_\epsilon^n\|_{TV} + \|\pi P_\epsilon^n - \pi_\epsilon\|_{TV} \\ &\leq \|\pi_\epsilon - \pi P_\epsilon^{n-1}\|_{TV} + \|\pi P_\epsilon^n - \pi_\epsilon\|_{TV} \\ &\leq \|\pi_\epsilon - \pi P_\epsilon^{n-1}\|_{L_2(\pi)} + \|\pi P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi)} \xrightarrow{n \rightarrow \infty} 0; \end{aligned}$$

thus we find that π_ϵ must be stationary for P_ϵ .

It remains to verify that $\{\pi P_\epsilon^n\}_{n \in \mathbb{N}}$ is an $L_2(\pi)$ -Cauchy sequence, and thus from completeness it must have an $L_2(\pi)$ -limit. To this end, define $Q_\epsilon = (P_\epsilon - P)$. Let $\mathbf{2}^k = \{0, 1\}^k$ for all $k \in \mathbb{N}$. We will expand $\pi(P + Q_\epsilon)^n$ and use the following facts:

- A. $\forall R \in \mathcal{B}(L_2(\pi)) \ [\pi P^n R = \pi R]$.
- B. $Q_\epsilon : L_2(\pi) \rightarrow L_{2,0}(\pi)$.
- C. $P|_{L_{2,0}(\pi)} \in \mathcal{B}(L_{2,0}(\pi))$ and $\|P|_{L_{2,0}(\pi)}\|_{L_{2,0}(\pi)} \leq (1 - \alpha)$.

Since the operators P and Q_ϵ do not (necessarily) commute, when we expand $(P + Q)^n$ we must have one distinct term per binary sequence of length n . We can then group terms by the number of leading P s, and use (A) to cancel the leading terms.

Let $m, n \in \mathbb{N}$ be arbitrary with $m \leq n$. Then

$$\begin{aligned} & \|\pi P_\epsilon^n - \pi P_\epsilon^m\|_{L_2(\pi)} \\ &= \|\pi(P + Q_\epsilon)^n - \pi(P + Q_\epsilon)^m\|_{L_2(\pi)} \\ &= \left\| \pi \left[\left(\sum_{\mathbf{b} \in 2^n} \prod_{j=1}^n P^{b_j} Q_\epsilon^{1-b_j} \right) - \left(\sum_{\mathbf{b} \in 2^m} \prod_{j=1}^m P^{b_j} Q_\epsilon^{1-b_j} \right) \right] \right\|_{L_2(\pi)} \\ &= \left\| \pi \left[\left(P^n + \sum_{k=0}^{n-1} P^{n-k-1} Q_\epsilon \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k P^{b_j} Q_\epsilon^{1-b_j} \right) \right. \right. \\ &\quad \left. \left. - \left(P^m + \sum_{k=0}^{m-1} P^{m-k-1} Q_\epsilon \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k P^{b_j} Q_\epsilon^{1-b_j} \right) \right] \right\|_{L_2(\pi)} \\ &= \left\| \left(\pi + \sum_{k=0}^{n-1} \pi Q_\epsilon \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k P^{b_j} Q_\epsilon^{1-b_j} \right) - \left(\pi + \sum_{k=0}^{m-1} \pi Q_\epsilon \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k P^{b_j} Q_\epsilon^{1-b_j} \right) \right\|_{L_2(\pi)} \\ &= \left\| \pi \sum_{k=m}^{n-1} Q_\epsilon \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k P^{b_j} Q_\epsilon^{1-b_j} \right\|_{L_2(\pi)} \\ &\leq \epsilon \sum_{k=m}^{n-1} \sum_{\mathbf{b} \in 2^k} \prod_{j=1}^k (1 - \alpha)^{b_j} \epsilon^{1-b_j} \\ &= \epsilon \sum_{k=m}^{n-1} (1 - \alpha + \epsilon)^k \\ &\leq \frac{\epsilon}{\alpha - \epsilon} (1 - \alpha + \epsilon)^m. \end{aligned}$$

Since this upper bound on $\|\pi P_\epsilon^n - \pi P_\epsilon^m\|_{L_2(\pi)}$ decreases to 0 monotonically in $m = \min(m, n)$, the sequence must be $L_2(\pi)$ -Cauchy.

Now, to bound the norm of π_ϵ , we take $m = 0$ and we get that for all $n \in \mathbb{N}$,

$$\|\pi P_\epsilon^n - \pi\|_{L_2(\pi)} \leq \frac{\epsilon}{\alpha - \epsilon}.$$

From the continuity of the norm, it must be the case that $\|\pi_\epsilon - \pi\|_{L_2(\pi)} \leq \frac{\epsilon}{\alpha - \epsilon}$. □

Lemma 6. *Under the assumptions of Section 3.2, if in addition $\epsilon < \alpha$, then*

$$1 \leq \|\pi_\epsilon\|_{L_2(\pi)} \leq \frac{\alpha}{\sqrt{\alpha^2 - \epsilon^2}}$$

and

$$0 \leq \|\pi - \pi_\epsilon\|_{L_2(\pi)} \leq \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}}.$$

Proof. The two lower bounds are immediate from Lemma 3 and the positivity of norms:

$$0 \leq \|\pi - \pi_\epsilon\|_{L_2(\pi)}^2 = \|\pi_\epsilon\|_{L_2(\pi)}^2 - 1.$$

To derive the first upper bound, we apply Lemma 3, our assumptions about the operators P and P_ϵ , and the triangle inequality to $\|\pi - \pi_\epsilon\|_2$:

$$\begin{aligned} \sqrt{\|\pi_\epsilon\|_{L_2(\pi)}^2 - 1} = \|\pi - \pi_\epsilon\|_{L_2(\pi)} &= \|\pi P - \pi_\epsilon P + \pi_\epsilon P - \pi_\epsilon P_\epsilon\|_{L_2(\pi)} \\ &\leq \|\pi P - \pi_\epsilon P\|_{L_2(\pi)} + \|\pi_\epsilon P - \pi_\epsilon P_\epsilon\|_{L_2(\pi)} \\ &\leq (1 - \alpha)\|\pi - \pi_\epsilon\|_{L_2(\pi)} + \epsilon\|\pi_\epsilon\|_{L_2(\pi)} \\ &= (1 - \alpha)\sqrt{\|\pi_\epsilon\|_{L_2(\pi)}^2 - 1} + \epsilon\|\pi_\epsilon\|_{L_2(\pi)}. \end{aligned}$$

Collecting the square roots and squaring both sides yields

$$\alpha^2 \left(\|\pi_\epsilon\|_{L_2(\pi)}^2 - 1 \right) \leq \epsilon^2 \|\pi_\epsilon\|_{L_2(\pi)}^2,$$

which implies that

$$\|\pi_\epsilon\|_{L_2(\pi)}^2 \leq \frac{\alpha^2}{\alpha^2 - \epsilon^2}.$$

Finally, the second upper bound is derived from the first one, again using Lemma 5.1:

$$\|\pi - \pi_\epsilon\|_{L_2(\pi)}^2 = \|\pi_\epsilon\|_{L_2(\pi)}^2 - 1 \leq \frac{\alpha^2}{\alpha^2 - \epsilon^2} - 1 = \frac{\epsilon^2}{\alpha^2 - \epsilon^2}.$$

□

We next observe that our assumptions imply that for small enough perturbations, the perturbed chain P_ϵ is geometrically ergodic in the $L_2(\pi)$ norm.

Lemma 7. *Under the assumptions of Section 3.2, if $\epsilon < \alpha$, then we have that P_ϵ is $L_2(\pi)$ -GE, with factor $\leq 1 - (\alpha - \epsilon)$.*

Proof. Suppose that $v \in L_{2,0}(\pi)$. Then

$$\begin{aligned} \|vP_\epsilon\|_{L_2(\pi)} &\leq \|v(P_\epsilon - P)\|_{L_2(\pi)} + \|vP\|_{L_2(\pi)} \\ &\leq \epsilon\|v\|_{L_2(\pi)} + (1 - \alpha)\|v\|_{L_2(\pi)} \\ &= (1 - \alpha + \epsilon)\|v\|_{L_2(\pi)}. \end{aligned}$$

Thus, for any probability measure $\mu \in L_2(\pi)$, since $\pi_\epsilon \in L_2(\pi)$ we have

$$\begin{aligned} \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi)} &= \|(\mu - \pi_\epsilon)P_\epsilon^n\|_{L_2(\pi)} \\ &\leq (1 - (\alpha - \epsilon))^n \|\mu - \pi_\epsilon\|_{L_2(\pi)}. \end{aligned}$$

□

Combining Lemmas 5 to 7 together with the triangle inequality immediately yields combinedthm.

A.2. Proofs of Theorem 2, Theorem 3, and Corollary 1

Definition 3. Following [29], a subset $S \subset \mathcal{X}$ is called hyper-small for the π -irreducible Markov kernel P with stationary measure π if $\pi(S) > 0$ and there exist $\delta_S > 0$ and $k \in \mathbb{N}$ such that for all $x, y \in \mathcal{X}$,

$$\frac{dP^k(x, \cdot)}{d\pi}(y) \geq \delta_S \mathbf{1}_S(x) \mathbf{1}_S(y),$$

or equivalently, if $P^k(x, A) \geq \delta_S \pi(A)$ for all $x \in S$ and $A \subset S$ measurable.

Lemma 4 of [15] states that on a countably generated state space (as we have assumed herein), every set of positive π -measure contains a hyper-small subset.

Lemma 8. (Existence of hyper-small subsets from [15].) *Suppose that (\mathcal{X}, Σ) is countably generated. Suppose that X is a ϕ -irreducible Markov chain on \mathcal{X} with kernel P for some σ -finite measure ϕ on \mathcal{X} . Then any set $K \subset \mathcal{X}$ with $\phi(K) > 0$ contains a set S_K such that (for some $n_K \in \mathbb{N}$)*

$$\inf_{(x,y) \in S_K \times S_K} \frac{dP^{n_K}(x, \cdot)}{d\pi}(y) = \delta > 0.$$

In the case that a stationary distribution π for P exists, without loss of generality we can take $\phi = \pi$. In this case, it is immediate that any set (S_K, n_K) satisfying Lemma 8 also satisfies Definition 3.

Also of importance to us is the following variant of Proposition 2.1 of [29], which provides a characterization of geometric ergodicity in terms of convergence to a hyper-small set.

Proposition 3. (Equivalent characterizations of π -a.e.-TV-GE from [27,29].) *Suppose that (Ω, Σ) is countably generated, and that X is a ϕ -irreducible Markov chain on \mathcal{X} with kernel P with stationary distribution π . Then the following are equivalent:*

- (i) *There exists $\rho_{TV} \in (0, 1)$ such that P is π -a.e.-TV-GE with factor ρ_{TV} .*
- (ii) *There exist a hyper-small set $S \subset \mathcal{X}$ and constants $\rho_S < 1, C_S \in \mathbb{R}_+$ such that*

$$\left\| \int \frac{\mathbf{1}_S(y)\pi(dy)}{\pi(S)} P^n(y, \cdot) - \pi \right\|_{TV} \leq C_S \rho_S^n \quad \forall n \in \mathbb{N}.$$

- (iii) *There exist a π -a.e. finite, measurable function $V : \mathcal{X} \rightarrow [1, \infty]$ with $\pi(V^2) < \infty$, and $\rho_V \in (0, 1)$, and $C > 0$ such that*

$$2\|\delta_x P^n - \pi\|_{TV} \leq \|\delta_x P^n - \pi\|_V \leq CV(x)\rho_V^n,$$

where $\|\mu\|_V = \sup_{|f| \leq V} |\mu(f)|$.

Proof of Theorem 2.

- (i) Let S be a hyper-small set for P_ϵ (which must exist from Lemma 8, since P_ϵ is π_ϵ -irreducible). Then the measure μ_S defined by

$$\frac{d\mu_S}{d\pi} = \frac{\mathbf{1}_S}{\pi_\epsilon(S)} \frac{d\pi_\epsilon}{d\pi}$$

satisfies (by Hölder’s inequality, and since $\pi_\epsilon \in L_2(\pi)$)

$$\|\mu_S\|_{L_2(\pi)}^2 \leq \|\pi_\epsilon\|_{L_2(\pi)}^2 \pi_\epsilon(S)^{-2} < \infty,$$

and hence $\mu_S \in L_2(\pi)$. Then (by Cauchy–Schwarz again)

$$\left\| \int \frac{\mathbf{1}_S(y)\pi_\epsilon(dy)}{\pi_\epsilon(S)} P_\epsilon^n(y, \cdot) - \pi_\epsilon \right\|_{\text{TV}} \leq \frac{1}{2} \|\mu_S P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi)} \leq \|\mu_S - \pi_\epsilon\|_{L_2(\pi)} (1 - \alpha + \epsilon)^n,$$

which, along with Proposition 3, establishes that P_ϵ is π_ϵ -a.e.-TV-GE with some factor $\rho_{\text{TV}} \in (0, 1)$.

(ii) Suppose that $\mu \in L_\infty(\pi_\epsilon)$. Then $\mu \in L_2(\pi)$ since

$$\frac{d\mu}{d\pi} \leq \|\mu\|_{L_\infty(\pi_\epsilon)} \frac{d\pi_\epsilon}{d\pi}.$$

Since $\mu P_\epsilon^n - \pi_\epsilon \in L_1(\pi_\epsilon) \subset L_1(\pi)$,

$$\|\mu P_\epsilon^n - \pi_\epsilon\|_{L_1(\pi_\epsilon)} = \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_1(\pi)} = 2\|\mu P_\epsilon^n - \pi_\epsilon\|_{\text{TV}}. \tag{25}$$

Applying this equality as well as Cauchy–Schwarz, we get

$$\begin{aligned} \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_1(\pi_\epsilon)} &= \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_1(\pi)} \\ &\leq \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi)} \\ &\leq \|\mu - \pi_\epsilon\|_{L_2(\pi)} (1 - \alpha - \epsilon)^n. \end{aligned} \tag{26}$$

(iii) If $\pi \in L_\infty(\pi_\epsilon)$ and $\mu \in L_2(\pi_\epsilon)$, then

$$\begin{aligned} \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi_\epsilon)}^2 &= \int \left(\frac{d\mu P_\epsilon^n - \pi_\epsilon}{d\pi_\epsilon} \right)^2 d\pi_\epsilon \\ &= \int \left(\frac{d\mu P_\epsilon^n - \pi_\epsilon}{d\pi} \right)^2 \frac{d\pi}{d\pi_\epsilon} d\pi \\ &\leq \|\pi\|_{L_\infty(\pi_\epsilon)} \int \left(\frac{d\mu P_\epsilon^n - \pi_\epsilon}{d\pi} \right)^2 d\pi \\ &= \|\pi\|_{L_\infty(\pi_\epsilon)} \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi)}^2 \\ &\leq \|\pi\|_{L_\infty(\pi_\epsilon)} \|\mu - \pi_\epsilon\|_{L_2(\pi)}^2 (1 - (\alpha - \epsilon))^{2n}. \end{aligned} \tag{27}$$

Proof of Theorem 3. From [4, Lemma 1], since P_ϵ has stationary measure π_ϵ , we have $P_\epsilon : L_2(\pi_\epsilon) \rightarrow L_2(\pi_\epsilon)$. Since P_ϵ is $(L_\infty(\pi_\epsilon), \|\cdot\|_{L_1(\pi_\epsilon)})$ -GE with factor $\rho_1 \leq (1 - (\alpha - \epsilon))$ (as established by Theorem 2) and P_ϵ is reversible, it must also be $L_2(\pi_\epsilon)$ -GE with factor $\rho = \rho_1$ by Lemma 2. □

Proof of Corollary 1. Note that the assumption that $\|P - P_\epsilon\|_{L_2(\pi_\epsilon)} < \varphi$ implies $P - P_\epsilon : L_2(\pi_\epsilon) \rightarrow L_2(\pi_\epsilon)$.

(i) Since P_ϵ is $L_2(\pi_\epsilon)$ -GE with factor $(1 - (\alpha - \epsilon))$ and π_ϵ -reversible, we can reverse the roles of P and P_ϵ , so the result follows by Theorem 1.

(ii) Taking $\mu = \pi$ and $n = 1$ in Theorem(iii), we have

$$\begin{aligned} \|\pi P_\epsilon - \pi_\epsilon\|_{L_2(\pi_\epsilon)}^2 &\leq \|\pi\|_{L_\infty(\pi_\epsilon)} \|\pi - \pi_\epsilon\|_{L_2(\pi)}^2 (1 - (\alpha - \epsilon))^2 \\ &\leq \|\pi\|_{L_\infty(\pi_\epsilon)} \frac{\epsilon^2}{\alpha^2 - \epsilon^2} (1 - (\alpha - \epsilon))^2. \end{aligned} \tag{28}$$

Hence,

$$\begin{aligned} \|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)} &\leq \|\pi P - \pi P_\epsilon\|_{L_2(\pi_\epsilon)} + \|\pi P_\epsilon - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \\ &\leq \varphi \|\pi\|_{L_2(\pi_\epsilon)} + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon)) \\ &= \varphi \sqrt{\|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)}^2 + 1} + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon)) \\ &\leq \varphi (\|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)} + 1) + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon)). \end{aligned} \tag{29}$$

Hence,

$$\|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \leq \frac{\varphi + \|\pi\|_{L_\infty(\pi_\epsilon)}^{1/2} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} (1 - (\alpha - \epsilon))}{1 - \varphi}. \tag{30}$$

Finally,

$$\|\mu P_\epsilon^n - \pi\|_{L_2(\pi_\epsilon)} \leq \|\mu P_\epsilon^n - \pi_\epsilon\|_{L_2(\pi_\epsilon)} + \|\pi_\epsilon - \pi\|_{L_2(\pi_\epsilon)}.$$

The first term is bounded by Theorem 2(iii), and the second term is bounded by (30). \square

A.3. Proofs of Theorem 4 and Theorem 5

A.3.1 *Time-averaging of marginal distributions. Proof of Theorem 4.* The first result of Theorem 4 follows from the triangle inequality and Theorem 1:

$$\begin{aligned} \left\| \pi - \frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k \right\|_{L_2(\pi)} &\leq \frac{1}{t} \sum_{k=0}^{t-1} \left\| \pi - \mu P_\epsilon^k \right\|_{L_2(\pi)} \\ &\leq \frac{1}{t} \sum_{k=0}^{t-1} \left[(1 - (\alpha - \epsilon))^k \|\pi_\epsilon - \mu\|_{L_2(\pi)} + \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} \right] \\ &\leq \frac{1 - (1 - (\alpha - \epsilon))^t}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi)} + \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}}. \end{aligned}$$

The subsequent results follow similarly via Theorems 2 and 3 and Corollary 1. \square

A.3.2 *Covariance bounds.* We turn our attention to the covariance structure of the original and perturbed chains. There is an obvious isometric isomorphism between the space of measures $L_2(\pi)$ and the function space

$$L'_2(\pi) = \{f : \mathcal{X} \rightarrow \mathbb{R} \text{ s.t. } \int f(x)^2 \pi(dx) < \infty\}$$

equipped with the norm $\|f\|_{L'_2(\pi)}^2 = \int f(x)^2 \pi(dx)$, where a measure μ is mapped to its Radon–Nikodym derivative, $\mu \mapsto \frac{d\mu}{d\pi}$. For this reason, we need not distinguish between these spaces, and when dealing with a function $f \in L'_2(\pi)$ we may occasionally abuse notation and treat it as its associated measure. Let X_t and X_t^ϵ denote the original and perturbed chains run from some initial measure $\mu \in L_2(\pi)$.

Corollary 4. *Under the assumptions of Section 3.2,*

(a) *if $X_0 \sim \pi$ (the initial distribution is the stationary distribution), then for $f, g \in L'_2(\pi)$,*

$$\text{Cov}[f(X_t), g(X_s)] \leq (1 - \alpha)^{|t-s|} \|f - \pi f\|_{L'_2(\pi)} \|g - \pi g\|_{L'_2(\pi)} ; \tag{31}$$

(b) *if $\epsilon < \alpha$, and P_ϵ is π_ϵ -reversible, $\rho_2 = (1 - (\alpha - \epsilon))$, and $X_0^\epsilon \sim \pi_\epsilon$, then for $f, g \in L'_2(\pi_\epsilon)$,*

$$\text{Cov}[f(X_t^\epsilon), g(X_s^\epsilon)] \leq \rho_2^{|t-s|} \|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)} \|g - \pi_\epsilon g\|_{L'_2(\pi_\epsilon)} , \tag{32}$$

where for a function $h: \mathcal{X} \rightarrow \mathbb{R}$, πh is the constant function equal to $\int h(s)\pi(ds)$ everywhere.

Proof. The proof of this result follows that of Corollary B.5 in [18]. We only show the proof for the original chain; however, the proof for the perturbed chain is the same, since it is reversible and $L_2(\pi_\epsilon)$ -GE with the appropriate factor, from Theorem 3.

Define the subspace

$$L'_{2,0}(\pi) = \{h \in L'_2(\pi) : \int h(s)\pi(ds) = 0\},$$

and define the operator $F \in \mathcal{B}(L'_{2,0}(\pi))$ by

$$[Ff](x) = \int P(x, dy)f(y) = \mathbb{E}[f(X_1)|X_0 = x].$$

From Lemma 12.6.4 of [21],

$$\sup_{f, g \in L'_2(\pi)} \text{corr}(f(X_0), g(X_t)) = \sup_{\substack{\|f\|_{L'_2(\pi)} = 1 = \|g\|_{L'_2(\pi)} \\ f, g \in L'_{2,0}(\pi)}} \langle f, F^t g \rangle = \|F^t\|_{L'_{2,0}(\pi)}.$$

Consider the canonical isomorphism between $L_2(\pi)$ and $L'_2(\pi)$. The restriction of this isomorphism (on the right) to elements of $L'_{2,0}(\pi)$ yields $L_{2,0}(\pi)$ (on the left)—the signed measures with total measure 0. The image of F under the restricted isomorphism is the adjoint operator of P restricted to $L_{2,0}(\pi)$. Since P is π -reversible, it is self-adjoint in $L_2(\pi)$, so $\|F\|_{L'_{2,0}(\pi)} = \|P\|_{L_{2,0}(\pi)}$. Thus

$$\|F^t\|_{L'_{2,0}(\pi)} \leq \|F\|_{L'_{2,0}(\pi)}^t = \|P\|_{L_{2,0}(\pi)}^t \leq (1 - \alpha)^t.$$

Therefore

$$\text{Cov}(f(X_0), g(X_t)) \leq \|f - \pi f\|_{L'_2(\pi)} \|g - \pi g\|_{L'_2(\pi)} (1 - \alpha)^t.$$

Since *Cov* is symmetric, the shifted and symmetrized result holds for any $f, g \in L'_2(\pi)$:

$$\text{Cov}[f(X_t), g(X_s)] \leq (1 - \alpha)^{|t-s|} \|f - \pi f\|_{L'_2(\pi)} \|g - \pi g\|_{L'_2(\pi)}. \tag{33}$$

We present further bounds for the case that the initial distribution is not the stationary distribution in Corollary 5.

Remark 8. Note in Corollary 4 that

$$\|h - \pi h\|_{L'_2(\pi)} = \sqrt{\|h\|_{L'_2(\pi)}^2 - (\pi h)^2} \leq \|h\|_{L'_2(\pi)}.$$

Also note that

$$\|h\|_{L'_2(\pi)} \leq \|h - \pi(h)\|_{L'_2(\pi)} + |\pi(h)|. \tag{34}$$

Corollary 5. Under the assumptions of Section 3.2,

(a) if $X_0 \sim \mu$, then for $f, g \in L'_4(\pi)$,

$$\begin{aligned} &\text{Cov}(f(X_t), g(X_{t+s})) \\ &\leq (1 - \alpha)^s \|f - \pi f\|_{L'_2(\pi)} \|g - \pi g\|_{L'_2(\pi)} \\ &\quad + 2^{3/2} (1 - \alpha)^{t+s/2} \|\mu - \pi\|_{L_2(\pi)} \|f - \pi f\|_{L'_4(\pi)} \|g - \pi g\|_{L'_4(\pi)} \\ &\quad - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g); \end{aligned}$$

(b) if $\epsilon < \alpha$, and P_ϵ is π_ϵ -reversible, $\rho_2 = (1 - (\alpha - \epsilon))$, and $X_0^\epsilon \sim \mu$, then for $f, g \in L'_4(\pi_\epsilon)$,

$$\begin{aligned} &\text{Cov}(f(X_t^\epsilon), g(X_{t+s}^\epsilon)) \\ &\leq \rho_2^s \|f - \pi_\epsilon f\|_{L'_2(\pi_\epsilon)} \|g - \pi_\epsilon g\|_{L'_2(\pi_\epsilon)} \\ &\quad + 2^{3/2} \rho_2^{t+s/2} \|\mu - \pi\|_{L_2(\pi_\epsilon)} \|f - \pi_\epsilon f\|_{L'_4(\pi_\epsilon)} \|g - \pi_\epsilon g\|_{L'_4(\pi_\epsilon)} \\ &\quad - (\mu P_\epsilon^t f - \pi_\epsilon f) (\mu P_\epsilon^{t+s} g - \pi_\epsilon g). \end{aligned}$$

Proof. This will use the following shorthand notation. Let

$$\begin{aligned} f_0 &= f - \pi f, \\ g_0 &= g - \pi g, \\ \|h\|_\star &= \left(\int (h(x) - \pi h)^2 \pi(dx) \right)^{1/2}, \\ \|h\|_{\star\star} &= \left(\int (h(x) - \pi h)^4 \pi(dx) \right)^{1/4}, \\ C_\mu &= \|\mu - \pi\|_2. \end{aligned}$$

We can interpret $\|\cdot\|_{\star\star}$ as a centred 4-norm. It is certainly bounded above by $\|\cdot\|_4$, the norm on $L'_4(\pi)$. For some results regarding the properties of a Markov transition kernel as an operator on $L'_p(\pi)$ for general p given an L_2 -spectral gap (as is implied by L_2 -GE), please refer to [34].

We only show the proof for the original chain. The result for the perturbed chain has essentially the same proof.

By definition we can express the covariance by the triple integral below. We re-express this integral as a sum of two integrals involving the chain run from stationarity. This will allow us to apply Corollary 4. We have

$$\begin{aligned} & \text{Cov}(f(X_t), g(X_{t+s})) \\ &= \iiint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g)\mu(dx)P^t(x, dy)P^s(y, dz) \\ &= \iiint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g) \left[\frac{d\mu}{d\pi}(x) - 1 \right] \pi(dx)P^t(x, dy)P^s(y, dz) \\ & \quad + \iiint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g)\pi(dx)P^t(x, dy)P^s(y, dz). \end{aligned}$$

We will simplify each of these expressions separately, starting with the second term:

$$\begin{aligned} & \iiint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g)\pi(dx)P^t(x, dy)P^s(y, dz) \\ &= \iint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g)\pi(dy)P^s(y, dz) \\ &= \iint f(y)g(z)\pi(dy)P^s(y, dz) \\ & \quad - (\mu^{P^t}f)(\pi g) - (\pi f)(\mu^{P^{t+s}}g) + (\mu^{P^t}f)(\mu^{P^{t+s}}g) \\ &= \iint f_0(y)g_0(z)\pi(dy)P^s(y, dz) + (\pi f)(\pi g) \\ & \quad - (\mu^{P^t}f)(\pi g) - (\pi f)(\mu^{P^{s+t}}g) + (\mu^{P^t}f)(\mu^{P^{t+s}}g) \\ &= \langle f_0, F^s g_0 \rangle + (\mu^{P^t}f - \pi f)(\mu^{P^{s+t}}g - \pi g). \end{aligned}$$

For the first term we find that

$$\begin{aligned} & \iiint (f(y) - \mu^{P^t}f)(g(z) - \mu^{P^{t+s}}g) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx)P^t(x, dy)P^s(y, dz) \\ &= \iiint f(y)g(z) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx)P^t(x, dy)P^s(y, dz) \\ & \quad - (\mu^{P^t}f) \iint g(z) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx)P^{t+s}(x, dz) \\ & \quad - (\mu^{P^{s+t}}g) \iint f(y) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx)P^t(x, dy) \\ & \quad + (\mu^{P^t}f)(\mu^{P^{s+t}}g) \int \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx) \\ &= \iiint f_0(y)g_0(z) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx)P^t(x, dy)P^s(y, dz) \end{aligned}$$

$$\begin{aligned}
 & - (\mu P^t f - \pi f) \iint g(z) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx) P^{t+s}(x, dz) \\
 & - (\mu P^{s+t} g - \pi g) \iint f(y) \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx) P^t(x, dy) \\
 & - (\pi f)(\pi g) \int \left(\frac{d\mu}{d\pi}(x) - 1 \right) \pi(dx) \\
 & = \left\langle \frac{d\mu}{d\pi} - 1, F^t(f_0 \otimes (F^s g_0)) \right\rangle \\
 & \quad - (\mu P^t f - \pi f) \left\langle \frac{d\mu}{d\pi} - 1, F^{t+s} g \right\rangle - (\mu P^{t+s} g - \pi g) \left\langle \frac{d\mu}{d\pi} - 1, F^t f \right\rangle \\
 & = \left\langle \frac{d\mu}{d\pi} - 1, F^t(f_0 \otimes (F^s g_0)) \right\rangle - 2(\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g),
 \end{aligned}$$

where $f_0 \otimes F^s g_0$ is defined by

$$[f_0 \otimes F^s g_0](y) = f_0(y) \int g_0(z) P^s(y, dz).$$

Putting these together,

$$\begin{aligned}
 & \text{Cov}(f(X_t), g(X_{t+s})) \\
 & = \langle f_0, F^s g_0 \rangle + (\pi f - \mu P^t f)(\pi g - \mu P^{s+t} g) \\
 & \quad + \left\langle \frac{d\mu}{d\pi} - 1, F^t(f_0 \otimes (F^s g_0)) \right\rangle - 2(\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g) \\
 & = \langle f_0, F^s g_0 \rangle + \left\langle \frac{d\mu}{d\pi} - 1, F^t(f_0 \otimes (F^s g_0)) \right\rangle - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g) \\
 & \leq (1 - \alpha)^s \|f\|_* \|g\|_* + (1 - \alpha)^t \|\mu - \pi\|_2 \|f_0 \otimes F^s g_0\|_2 \\
 & \quad - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g) \\
 & \leq (1 - \alpha)^s \|f\|_* \|g\|_* + (1 - \alpha)^t \|\mu - \pi\|_2 \|f_0\|_4 \|F^s g_0\|_4 \\
 & \quad - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g) \\
 & \leq (1 - \alpha)^s \|f\|_* \|g\|_* + (1 - \alpha)^t \|\mu - \pi\|_2 \|f_0\|_4 \|g_0\|_4 \left\| F^s \Big|_{L^4_{4,0}} \right\|_4 \\
 & \quad - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g) \\
 & \leq (1 - \alpha)^s \|f\|_* \|g\|_* + 2^{3/2} (1 - \alpha)^{t+s/2} \|\mu - \pi\|_2 \|f\|_{**} \|g\|_{**} \\
 & \quad - (\mu P^t f - \pi f) (\mu P^{t+s} g - \pi g).
 \end{aligned}$$

The $\langle f_0, F^s g_0 \rangle$ term is bounded using Corollary 4, where we have taken the result in its equivalent form using the $\langle \cdot, \cdot \rangle$ notation and the forward operator F . The term

$$\left\langle \frac{d\mu}{d\pi} - 1, F^t(f_0 \otimes (F^s g_0)) \right\rangle$$

is bounded following the methodology of the proof of [34, Lemma 3.39] (in order, the inequalities are Cauchy–Schwarz, $\|F^s g_0\| \leq \|F^s\| \|g_0\|$ for any norm $\|\cdot\|$, and Proposition 3.17 of [34]). \square

The main motivation in establishing the covariance bounds in Corollaries 4 and 5 is that we will need to sum up covariances in order to establish bounds on the variance component of mean squared error for estimation of $\pi(f)$ via the dependent sample means

$$\frac{1}{t} \sum_{j=0}^{t-1} f(X_j), \quad \frac{1}{t} \sum_{j=0}^{t-1} f(X_j^\epsilon)$$

for an arbitrary starting measure. To this end we will be interested in the following summation result.

Corollary 6. *Under the assumptions of Section 3.2,*

(a) *if $X_0 \sim \mu$, then for $f, g \in L'_4(\pi)$,*

$$\begin{aligned} & \frac{1}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} \text{Cov}(f(X_j), f(X_k)) \\ & \leq \frac{2\|f - \pi f\|_{L'_2(\pi)}^2}{\alpha t} + \frac{2^{7/2} \|\mu - \pi\|_{L_2(\pi)} \|f - \pi f\|_{L'_4(\pi)}^2}{\alpha^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} \mu P^m f - \pi f \right)^2; \end{aligned}$$

(b) *if $\epsilon < \alpha$, and P_ϵ is π_ϵ -reversible, $\rho_2 = (1 - (\alpha - \epsilon))$, and $X_0^\epsilon \sim \mu$, then for $f, g \in L'_4(\pi_\epsilon)$,*

$$\begin{aligned} & \frac{1}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} \text{Cov}(f(X_j^\epsilon), f(X_k^\epsilon)) \\ & \leq \frac{2\|f - \pi f\|_{L'_2(\pi_\epsilon)}^2}{(1 - \rho_2)t} + \frac{2^{7/2} \|\mu - \pi\|_{L_2(\pi_\epsilon)} \|f - \pi f\|_{L'_4(\pi_\epsilon)}^2}{(1 - \rho_2)^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} \mu P_\epsilon^m f - \pi f \right)^2. \end{aligned}$$

Proof. We only show the proof for the original chain. The results for the perturbed chain have essentially the same proof. The proof is largely an exercise in summation of geometric series and meticulous bookkeeping. The first inequality is due to Corollary 5. The second inequality makes use of the fact that $0 < \alpha < 1$. To simplify notation, $C_\mu = \|\mu - \pi\|_{L_2(\pi)}$. We have

$$\begin{aligned} & \frac{1}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} \text{Cov}(f(X_j), f(X_k)) \\ & = \frac{\|f\|_*^2}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} (1 - \alpha)^{|m-n|} - \frac{1}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} (\mu P^m f - \pi f) (\mu P^n f - \pi f) \\ & \quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \sum_{m=0}^{t-1} \sum_{n=0}^{t-1} (1 - \alpha)^{(m+n)/2} \end{aligned}$$

$$\begin{aligned}
 &= \frac{\|f\|_*^2}{t^2} \sum_{m=0}^{t-1} \left(1 + 2 \sum_{s=1}^{t-m-1} (1-\alpha)^s \right) - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2 \\
 &\quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \sum_{m=0}^{t-1} (1-\alpha)^m \left(1 + 2 \sum_{s=1}^{t-m-1} (1-\alpha)^{s/2} \right) \\
 &= \frac{\|f\|_*^2}{t^2} \sum_{m=0}^{t-1} \left(1 + 2 \frac{(1-\alpha) - (1-\alpha)^{t-m}}{\alpha} \right) - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2 \\
 &\quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \sum_{m=0}^{t-1} (1-\alpha)^m \left(1 + 2 \frac{\sqrt{1-\alpha} - \sqrt{1-\alpha}^{t-m}}{1 - \sqrt{1-\alpha}} \right) \\
 &= \frac{\|f\|_*^2}{t^2} \sum_{m=0}^{t-1} \left(\frac{2-\alpha}{\alpha} - \frac{2}{\alpha} (1-\alpha)^{t-m} \right) - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2 \\
 &\quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \sum_{m=0}^{t-1} \left((1-\alpha)^m \frac{1 + \sqrt{1-\alpha}}{1 - \sqrt{1-\alpha}} - 2 \frac{\sqrt{1-\alpha}^{t+m}}{1 - \sqrt{1-\alpha}} \right) \\
 &= \frac{\|f\|_*^2}{t^2} \left(\frac{2-\alpha}{\alpha} t - \frac{2}{\alpha} \frac{(1-\alpha) - (1-\alpha)^{t+1}}{\alpha} \right) - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2 \\
 &\quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \left(\left[\frac{1 + \sqrt{1-\alpha}}{1 - \sqrt{1-\alpha}} \right] \left[\frac{1 - (1-\alpha)^t}{\alpha} \right] \right. \\
 &\quad \left. - \left[\frac{2\sqrt{1-\alpha}^t}{1 - \sqrt{1-\alpha}} \right] \left[\frac{1 - \sqrt{1-\alpha}^t}{1 - \sqrt{1-\alpha}} \right] \right) \\
 &= (2-\alpha) \frac{\|f\|_*^2}{\alpha t} - 2(1-\alpha) \frac{1 - (1-\alpha)^t}{\alpha^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2 \\
 &\quad + \frac{2^{3/2} C_\mu \|f\|_{**}^2}{t^2} \left(\frac{1 + \sqrt{1-\alpha}}{\alpha} \right)^2 (1 - (1-\alpha)^{t/2})^2 \\
 &\leq \frac{2\|f\|_*^2}{\alpha t} + \frac{2^{7/2} C_\mu \|f\|_{**}^2}{\alpha^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} (\mu P^m f - \pi f) \right)^2.
 \end{aligned}$$

□

A.3.3 Mean squared error bounds.

Theorem 7. Under the assumptions of Section 3.2, if $X_0 \sim \mu \in L_2(\pi)$, then

$$\mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k) \right)^2 \right] \leq \frac{2\|f - \pi f\|_2^2}{\alpha t} + \frac{2^{7/2} \|\mu - \pi\|_2 \|f - \pi f\|_4^2}{\alpha^2 t^2}.$$

Proof. The proof proceeds by partitioning the mean squared error via the bias-variance decomposition, then bounding the variance term and noting that our bound for the variance contains an expression which exactly cancels the bias term. We compute that

$$\begin{aligned} & \mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P^k](f) - \frac{1}{t} \sum_{k=0}^{t-1} (f(X_k) - [\mu P^k](f)) \right)^2 \right] \\ &= \left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P^k](f) \right)^2 + \mathbb{E} \left[\left(\frac{1}{t} \sum_{k=0}^{t-1} (f(X_k) - [\mu P^k](f)) \right)^2 \right] \\ &= \left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P^k](f) \right)^2 + \frac{1}{t^2} \sum_{j=0}^{t-1} \sum_{k=0}^{t-1} \text{Cov}(f(X_j), f(X_k)). \end{aligned}$$

The variance term is bounded using Corollary 6:

$$\begin{aligned} & \frac{1}{t^2} \sum_{j=0}^{t-1} \sum_{k=0}^{t-1} \text{Cov}(f(X_j), f(X_k)) \\ & \leq \frac{2\|f - \pi f\|_2^2}{\alpha t} + \frac{2^{7/2} \|\mu - \pi\|_2 \|f - \pi f\|_4^2}{\alpha^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} \mu P^m f - \pi f \right)^2. \end{aligned}$$

Putting these together yields the desired result. □

Remark 9. We note that, as per Remark 8, $\|f - \pi f\| \leq \|f\|_2$. Similarly $\|f - \pi f\|_4 \leq \|f\|_4$. Also, in the case that f is π -essentially bounded, $\|f\|_2 \leq \|f\|_\infty$ and $\|f\|_4 \leq \|f\|_\infty$. These alternative norms may be substituted into the result as necessary in order to make the bounds tractable for a given application.

Remark 10. Comparing our above geometrically ergodic results to the L_1 results of [18] in the uniformly ergodic case, we see that the L_2 and L_1 bounds we establish above differ from the corresponding L_1 bound of [18] only by a factor, which is constant in time, but varies with the initial distribution (as is to be expected when moving from uniform ergodicity to geometric ergodicity). For the mean squared error results, the $\|\cdot\|_*$ -norm in that paper is based on the midrange-centred infinity norm, which as per Remark 9 is an upper bound on what we have.

Proof of Theorem 5. For the first result, we proceed via bias-variance decomposition, as in the corresponding result for the exact chain. However, now the bias under consideration is itself decomposed as the square of a sum of two components. The squared sum is expanded simultaneously with the bias-variance expansion. We compute that

$$\begin{aligned} & \mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k^\epsilon) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\pi(f) - \pi_\epsilon(f) + \frac{1}{t} \sum_{k=0}^{t-1} [\pi_\epsilon - \mu P_\epsilon^k](f) - \frac{1}{t} \sum_{k=0}^{t-1} (f(X_k^\epsilon) - [\mu P_\epsilon^k](f)) \right)^2 \right] \end{aligned}$$

$$\begin{aligned}
 &= ([\pi - \pi_\epsilon](f))^2 + 2([\pi - \pi_\epsilon](f)) \left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) \right) \\
 &+ \left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) \right)^2 + \frac{1}{t^2} \sum_{j=0}^{t-1} \sum_{k=0}^{t-1} \text{Cov}(f(X_j^\epsilon), f(X_k^\epsilon)).
 \end{aligned}$$

We bound the first component of the bias term using versions of Lemma 6:

$$\begin{aligned}
 ([\pi - \pi_\epsilon](f))^2 &= ([\pi - \pi_\epsilon](f - \pi_\epsilon f))^2 \\
 &\leq \begin{cases} \|\pi - \pi_\epsilon\|_{L_2(\pi)}^2 \|f - \pi_\epsilon f\|_{L_2'(\pi)}^2 \\ \|\pi - \pi_\epsilon\|_{L_2(\pi_\epsilon)}^2 \|f - \pi_\epsilon f\|_{L_2'(\pi_\epsilon)}^2 \end{cases} \\
 &\leq \begin{cases} \frac{\epsilon^2}{\alpha^2 - \epsilon^2} \|f - \pi_\epsilon f\|_{L_2'(\pi)}^2 \\ \frac{\varphi^2}{(1 - \rho_2)^2 - \varphi^2} \|f - \pi_\epsilon f\|_{L_2'(\pi_\epsilon)}^2 \end{cases} \quad \text{: given } (*).
 \end{aligned}$$

We bound the variance term using Corollary 6:

$$\begin{aligned}
 &\frac{1}{t^2} \sum_{j=0}^{t-1} \sum_{k=0}^{t-1} \text{Cov}(f(X_j^\epsilon), f(X_k^\epsilon)) \\
 &\leq \frac{2\|f - \pi_\epsilon f\|_{L_2'(\pi_\epsilon)}^2}{(1 - \rho_2)t} + \frac{2^{7/2} \|\mu - \pi_\epsilon\|_{L_2(\pi_\epsilon)} \|f - \pi_\epsilon f\|_{L_4'(\pi_\epsilon)}}{(1 - \rho_2)^2 t^2} - \left(\frac{1}{t} \sum_{m=0}^{t-1} \mu P_\epsilon^m f - \pi_\epsilon f \right)^2.
 \end{aligned}$$

The negative term in this expression exactly cancels out the third bias term in the expansion.

Finally, we bound the second bias term using Lemma 6 and Theorem 4:

$$\begin{aligned}
 &2([\pi - \pi_\epsilon](f)) \left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) \right) \\
 &= 2([\pi - \pi_\epsilon](f - \pi_\epsilon f)) \left(\left[\pi_\epsilon - \frac{1}{t} \sum_{k=0}^{t-1} \mu P_\epsilon^k \right] (f - \pi_\epsilon f) \right) \\
 &\leq 2 \begin{cases} \frac{\varphi}{\sqrt{(1 - \rho_2)^2 - \varphi^2}} \|f - \pi_\epsilon f\|_{L_2'(\pi)} \frac{1 - (1 - (\alpha - \epsilon))^t}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi)} \|f - \pi_\epsilon f\|_{L_2(\pi)} \\ \frac{\varphi^2}{(1 - \rho_2)^2 - \varphi^2} \|f - \pi_\epsilon f\|_{L_2'(\pi_\epsilon)}^2 \frac{1 - \rho_2^t}{t(1 - \rho_2)} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} \|f - \pi_\epsilon f\|_{L_2(\pi_\epsilon)} \end{cases} \quad \text{: given } (*) \\
 &\leq 2 \begin{cases} \frac{\epsilon}{\sqrt{\alpha^2 - \epsilon^2}} \frac{1}{t(\alpha - \epsilon)} \|\pi_\epsilon - \mu\|_{L_2(\pi)} \|f - \pi_\epsilon f\|_{L_2(\pi)}^2 \\ \frac{\varphi}{\sqrt{(1 - \rho_2)^2 - \varphi^2}} \frac{1}{t(1 - \rho_2)} \|\pi_\epsilon - \mu\|_{L_2(\pi_\epsilon)} \|f - \pi_\epsilon f\|_{L_2(\pi_\epsilon)}^2 \end{cases} \quad \text{: given } (*).
 \end{aligned}$$

Putting these together yields the first and third results.

For the second and fourth results we use the fact that for any random variable Z and for any $a, b \in \mathbb{R}$ the following holds:

$$\begin{aligned} \mathbb{E}[(Z - a)^2] &= 2\mathbb{E}[(Z - b)^2] + 2(a - b)^2 - \mathbb{E}[(Z + a - 2b)^2] \\ &\leq 2\mathbb{E}[(Z - b)^2] + 2(a - b)^2. \end{aligned}$$

We have

$$\begin{aligned} &\mathbb{E} \left[\left(\pi(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k^\epsilon) \right)^2 \right] \\ &\leq 2([\pi - \pi_\epsilon](f))^2 + 2\mathbb{E} \left[\left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} f(X_k^\epsilon) \right)^2 \right] \\ &= 2([\pi - \pi_\epsilon](f - \pi_\epsilon f))^2 \\ &\quad + 2\mathbb{E} \left[\left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) - \frac{1}{t} \sum_{k=0}^{t-1} (f(X_k^\epsilon) - [\mu P_\epsilon^k](f)) \right)^2 \right] \\ &= 2([\pi - \pi_\epsilon](f - \pi_\epsilon f))^2 + 2 \left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) \right)^2 \\ &\quad + 2\mathbb{E} \left[\left(\frac{1}{t} \sum_{k=0}^{t-1} (f(X_k^\epsilon) - [\mu P_\epsilon^k](f)) \right)^2 \right] \\ &= 2([\pi - \pi_\epsilon](f - \pi_\epsilon f))^2 + 2 \left(\pi_\epsilon(f) - \frac{1}{t} \sum_{k=0}^{t-1} [\mu P_\epsilon^k](f) \right)^2 \\ &\quad + \frac{2}{t^2} \sum_{j=0}^{t-1} \sum_{k=0}^{t-1} \text{Cov}(f(X_j^\epsilon), f(X_k^\epsilon)). \end{aligned}$$

Applying Corollary 5 to bound the sum of covariances, we find that we are able to exactly cancel the second term in the final expression above. Using the same bound as before for the first expression, we get the final result. □

Appendix B. Proof of Theorem 6

Let

$$\begin{aligned} \gamma(x) &= \mathbb{E}_{y \sim q(y|x)} r(y|x) = \int r(y|x)q(y|x)dy, \\ [v\Gamma](dy) &= v(y)\gamma(y)dy, \\ [vZ](dy) &= \left[\int r(y|x)q(y|x)v(x)dx \right] dy. \end{aligned}$$

Lemma 9. $P - \hat{P} = Z - \Gamma.$

Proof. We first give expressions for the elements of measure for transitions of the original chain. The first formula is the element of measure for transition from an arbitrary, fixed initial point. It is defined for us by the mechanics of the MH algorithm. The second expression is the element of measure for transition from a sample from an initial distribution, ν . It is derived from the first expression by integrating over the sample from ν :

$$\begin{aligned}
 P(x, dx') &= \delta_x(dx') \left[1 - \int (a(y|x)q(y|x)dy) \right] + a(x'|x)q(x'|x)dx', \\
 [\nu P](dx') &= \int \left[\delta_x(dx') \left[1 - \int a(y|x)q(y|x)dy \right] + a(x'|x)q(x'|x)dx' \right] \nu(x)dx \\
 &= \left[\left[1 - \int a(y|x')q(y|x')dy \right] \nu(x') + \int a(x'|x)q(x'|x)\nu(x)dx \right] dx'.
 \end{aligned}$$

The second form of the second expression is an application of Fubini’s theorem. The exchange of the order of integration for the second term in the expression is immediate. For the first term, for an arbitrary non-negative function f ,

$$\int_s \int_t f(s, t)\delta_t(ds)dt = \int_t \int_s f(s, t)\delta_t(ds)dt = \int_t f(t, t)dt = \int_s f(s, s)ds,$$

where the first equality is Fubini’s theorem, the second comes from integrating with respect to s , and the third comes from a change of dummy variable.

Similarly, the elements of measure for transitions from the approximating kernel are expressed below. The first expression, as above, is the element of measure for transition from an arbitrary, fixed initial point. It is defined for us by the mechanics of the nMH algorithm. The second expression is again derived by integrating the first against an initial measure, ν :

$$\begin{aligned}
 \hat{P}(x, dx') &= \delta_x(dx') \left[1 - \iint \hat{a}(y|x, z)q(y|x)f_y(z)dzdy \right] \\
 &\quad + \int \hat{a}(x'|x, z)q(x'|x)f_{x'}(z)dzdx', \\
 [\nu \hat{P}](dx') &= \int \left(\delta_x(dx') \left[1 - \iint \hat{a}(y|x, z)q(y|x)f_y(z)dzdy \right] \right. \\
 &\quad \left. + \int \hat{a}(x'|x, z)q(x'|x)f_{x'}(z)dzdx' \right) \nu(x)dx \\
 &= \left[1 - \iint \hat{a}(y|x', z)q(y|x')f_y(z)dzdy \right] \nu(x')dx' \\
 &\quad + \left[\iint \hat{a}(x'|x, z)q(x'|x)f_{x'}(z)\nu(x)dzdx \right] dx'.
 \end{aligned}$$

The same applications of Fubini’s theorem occur as above.

We may now leverage our notation defined above to simplify the difference of these elements of measure:

$$\begin{aligned} & \left[v(P - \hat{P}) \right] (dx'n) \\ &= \left[\iint \left(\hat{a}(y|x', z) - a(y|x') \right) q(y|x') f_y(z) dz dy \right] v(x') dx' \\ &+ \left[\iint \left(a(x'|x) - \hat{a}(x'nn|x, z) \right) q(x'|x) f_{x'}(z) v(x) dz dx \right] dx'n \\ &= \left[\int r(x'|x) q(x'|x) v(x) dx \right] dx' - \left[\int r(y|x') q(y|x') dy \right] v(x') dx' \\ &= [v(Z - \Gamma)](dx'). \end{aligned}$$

From this one may conclude that $(P - \hat{P} = Z - \Gamma)$ as operators. □

Proof of Theorem 6. It is obvious that if $|r(y|x)| \leq R$ uniformly in $(x, y) \in \mathcal{X}^2$, then

$$(\|\Gamma\|_{L_2(\pi)} \leq R) \tag{35}$$

and

$$(\|Z\|_{L_2(\pi)} \leq R \|Q\|_{L_2(\pi)}). \tag{36}$$

By applying the previous lemma, given the assumptions stated, we have

$$\|P - \hat{P}\|_{L_2(\pi)} \leq R(1 + \|Q\|_{L_2(\pi)}). \tag{37}$$

Appendix C. $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE is distinct from L_2 -GE for non-reversible chains

Let $\mathcal{X} = \mathbb{N} \cup 0$, and let a be a probability mass function on \mathcal{X} . Define transition probabilities by

$$p_{ij} = \begin{cases} a_j & : i = 0, \\ 1 & : i > 0, j = i - 1, \\ 0 & : \text{otherwise.} \end{cases} \tag{38}$$

Let $b_j = \sum_{i=j}^\infty a_i$. It is easy to verify that if $\sum_{j=1}^\infty b_j < \infty$ then

$$\pi_j = \frac{b_j}{\sum_{j=1}^\infty b_j}$$

is the unique stationary probability mass function for $P = [p_{ij}]_{ij \in \mathcal{X}^2}$.

In the special case where $a_j = 2^{-j-1}$, we have $\pi = a$. We continue this example working exclusively with this choice of a . Now,

$$\delta_j P^n = \begin{cases} \pi & : n \geq j + 1, \\ \delta_{n-j} & : n \leq j. \end{cases} \tag{39}$$

Thus, for any initial probability mass function μ ,

$$[\mu P^n]_j = \sum_{i=0}^{n-1} \mu_i \pi_j + \mu_{j+n}. \tag{40}$$

If $\frac{d\mu}{d\pi}(j) = \frac{\mu_j}{\pi_j} \leq \|\mu_{L_\infty(\pi)}\| < \infty$ for all $j \in \mathcal{X}$, then

$$\begin{aligned}
 \|\mu P^n - \pi\|_{L_2(\pi)}^2 &= \sum_{j=0}^{\infty} \pi_j \left(\sum_{i=0}^{n-1} \mu_i + \frac{\mu_{j+n}}{\pi_j} - 1 \right)^2 \\
 &= \sum_{j=0}^{\infty} \pi_j \left(- \sum_{i=n}^{\infty} \mu_i + \frac{\mu_{j+n}}{\pi_{j+n}} \frac{\pi_{j+n}}{\pi_j} \right)^2 \\
 &= \sum_{j=0}^{\infty} \pi_j \left(- \sum_{i=n}^{\infty} \frac{\mu_i}{\pi_i} \pi_i + \frac{\mu_{j+n}}{\pi_{j+n}} \frac{\pi_{j+n}}{\pi_j} \right)^2 \\
 &\leq \sum_{j=0}^{\infty} \pi_j \left(\sum_{i=n}^{\infty} \frac{\mu_i}{\pi_i} \pi_i + \frac{\mu_{j+n}}{\pi_{j+n}} \frac{\pi_{j+n}}{\pi_j} \right)^2 \tag{41} \\
 &\leq \sum_{j=0}^{\infty} 2^{-j-1} \left(\sum_{i=n}^{\infty} \|\mu_{L_\infty(\pi)}\| 2^{-i-1} + \|\mu_{L_\infty(\pi)}\| 2^{-n} \right)^2 \\
 &= \|\mu\|_{L_\infty(\pi)}^2 \sum_{j=0}^{\infty} 2^{-j-1} (2^{-n+1})^2 \\
 &= 4 \|\mu\|_{L_\infty(\pi)}^2 (2^{-n})^2.
 \end{aligned}$$

Hence P is $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE with optimal rate no larger than $1/2$.

For any $\alpha < \sqrt{0.5}$, let $v_j = (1 - \alpha)(\alpha)^j$. Then $v \in L_2(\pi)$, since

$$\begin{aligned}
 \|v\|_{L_2(\pi)}^2 &= \sum_{i=0}^{\infty} 0.5^{i+1} \left(\frac{(1 - \alpha)(\alpha)^i}{0.5^{i+1}} \right)^2 \\
 &= 2(1 - \alpha)^2 \sum_{i=0}^{\infty} (2\alpha^2)^i = \frac{2(1 - \alpha)^2}{1 - 2\alpha^2}.
 \end{aligned}
 \tag{42}$$

Moreover,

$$\begin{aligned}
 \|v P^n - \pi\|_{L_2(\pi)}^2 &= \sum_{j=0}^{\infty} \pi_j \left(\sum_{i=0}^{n-1} v_i + \frac{v_{j+n}}{\pi_j} - 1 \right)^2 \\
 &= \sum_{j=0}^{\infty} 0.5^{j+1} \left(- \sum_{i=n}^{\infty} (1 - \alpha)\alpha^i + (1 - \alpha)\alpha^{j+n}(0.5)^{-j-1} \right)^2 \\
 &= \alpha^{2n} \sum_{j=0}^{\infty} 0.5^{j+1} (-1 + 2(1 - \alpha)(2\alpha)^j)^2 \\
 &= \frac{\alpha^{2n}}{2} \sum_{j=0}^{\infty} (0.5^j - 4(1 - \alpha)\alpha^j + 4(1 - \alpha)^2(2\alpha^2)^j)
 \end{aligned}
 \tag{43}$$

$$\begin{aligned}
 &= \frac{\alpha^{2n}}{2} \left(2 - \frac{4(1-\alpha)}{1-\alpha} + \frac{4(1-\alpha)^2}{1-2\alpha^2} \right) \\
 &= \frac{(2\alpha-1)^2}{1-2\alpha^2} \alpha^{2n}.
 \end{aligned}$$

Thus the convergence rate starting from this initial measure is α .

Since this is true for any $\alpha < 1/\sqrt{2}$, this shows that the $L_2(\pi)$ -GE optimal rate is no smaller than $\sqrt{0.5}$. Hence the $(L_\infty(\pi), \|\cdot\|_{L_2(\pi)})$ -GE and $L_2(\pi)$ -GE optimal rates are different.

Appendix D. Proof of Proof of Lemma 2 and Lemma 1

Proof of Lemma 1. Let

$$\begin{aligned}
 \rho^* &= \left\{ \rho > 0 : \exists C : V \rightarrow \mathbb{R}_+ \text{ s.t. } \forall n \in \mathbb{N}, v \in V \cap \mathcal{M}_{+,1} \quad \|vP^n - \pi\| \leq C(v)\rho^n \right\}, \\
 \hat{\rho} &= \sup_{\mu \in V \cap \mathcal{M}_{+,1}} \limsup_{n \rightarrow \infty} \| \mu P^n - \pi \|^{1/n}.
 \end{aligned} \tag{44}$$

Proof that $\hat{\rho} \leq \rho^*$: Let $\epsilon > 0$. We have

$$\begin{aligned}
 \hat{\rho} &= \sup_{\mu \in V \cap \mathcal{M}_{+,1}} \limsup_{n \rightarrow \infty} \| \mu P^n - \pi \|^{1/n} \\
 &\leq \sup_{\mu \in V \cap \mathcal{M}_{+,1}} \limsup_{n \rightarrow \infty} \| \mu P^n - \pi \|^{1/n} \\
 &\leq \sup_{\mu \in V \cap \mathcal{M}_{+,1}} \limsup_{n \rightarrow \infty} (C_\epsilon(\mu)(\rho^* + \epsilon)^n)^{1/n} \\
 &= \rho^* + \epsilon.
 \end{aligned} \tag{45}$$

Since ϵ is arbitrary, $\hat{\rho} \leq \rho^*$.

Proof that $\hat{\rho} \geq \rho^*$: For all $v \in V \cap \mathcal{M}_{+,1}$, $\limsup_{n \rightarrow \infty} \|vP^n - \pi\|^{1/n} \leq \hat{\rho}$. Let $\epsilon > 0$. Then for all $\mu \in V \cap \mathcal{M}_{+,1}$, $\| \mu P^n - \pi \|^{1/n} > \hat{\rho} + \epsilon$ for at most finitely many $n \in \mathbb{N}$. Let

$$C_\epsilon(\mu) = \max_{n \in \mathbb{N}} \left(1 \vee \frac{\| \mu P^n - \pi \|}{(\rho + \epsilon)^n} \right).$$

Then $C_\epsilon(\mu) < \infty$ since the maximum is over finitely many distinct elements. Therefore $\| \mu P^n - \pi \| \leq C_\epsilon(\mu)(\hat{\rho} + \epsilon)^n$ for all $n \in \mathbb{N}$. This implies that $\hat{\rho} + \epsilon \geq \rho^*$. Since ϵ is arbitrary, $\hat{\rho} \geq \rho^*$. □

Proof of Lemma 2. The equivalence [(iii) \iff (iv)] is proven in [29, Theorem 2.1]. The implication [(iii) \implies (ii)] follows from the inclusion $L_\infty(\pi) \subset L_2(\pi)$, and [(ii) \implies (i)] follows from Cauchy–Schwarz.

Proof that [(ii) \implies (iii)]:

Without loss of generality, assume that ρ is the optimal rate of $(L_\infty(\pi_\epsilon), \|\cdot\|_{L_2(\pi)})$ -GE:

$$\rho = \sup_{v \in L_{\infty,0}(\pi)} \limsup_{t \rightarrow \infty} \|vP^t\|_{L_2(\pi)}^{1/t}. \tag{46}$$

From the proof of [32, Theorem 1], P is π -a.e. geometrically ergodic with some unknown optimal rate. From [29, Theorem 2.1], P is $L_2(\pi)$ -GE with some unknown optimal rate ρ_2 , which is equivalent to the spectral radius of $P|_{L_{2,0}(\pi)}$:

$$\rho_2 = r(P|_{L_{2,0}(\pi)}).$$

It remains to be shown that $\rho_2 \leq \rho$. We will use the spectral measure decomposition of P , as in [29]. Suppose, for a contradiction, that $\rho_2 > \rho$. Let $\bar{\rho} = \frac{\rho + \rho_2}{2}$. Let \mathcal{E} be the spectral measure of P , so that $\mu P^t = \int_{-1}^1 \lambda^t \mu \mathcal{E}(d\lambda)$. If $\rho_2 > \rho$ then either $\mathcal{E}([- \rho_2, -\bar{\rho}]) \neq \mathbf{0}$ or $\mathcal{E}((\bar{\rho}, \rho_2]) \neq \mathbf{0}$. Assume (replacing P by P^2 , ρ by ρ^2 , and ρ_2 by ρ_2^2 if necessary) that $\mathcal{E}((\bar{\rho}, \rho_2]) \neq \mathbf{0}$ and $\mathcal{E}((-1, 0)) = \mathbf{0}$. Then there is some non-zero signed measure ν in the range of $\mathcal{E}((\bar{\rho}, \rho_2])$. Since the spectral projections are orthogonal and $\{1\} \cap (\bar{\rho}, \rho_2] = \emptyset$, we have $\nu \perp \pi$, and hence $\nu(\mathcal{X}) = 0$. Since $L_{\infty,0}(\pi)$ is dense in $L_{2,0}(\pi)$, there is a $\mu \in L_{\infty,0}(\pi)$ with $\|\mu - \nu\|_{L_2(\pi)} < \|\nu\|_{L_2(\pi)}/2$. Then, from the polarization identity,

$$\langle \nu, \mu \rangle_{L_2(\pi)} \geq \frac{3}{8} \|\nu\|_{L_2(\pi)}^2 > 0,$$

and $\mu \neq 0$.

Let $R = \text{range}(\int_{(\bar{\rho}, \rho_2]} \mathcal{E}(d\lambda))$. Then $\text{span}(\nu) \subset R$, so

$$\|proj_R \mu\|_{L_2(\pi)} \geq \|proj_\nu \mu\|_{L_2(\pi)} \geq \frac{3}{8} \|\nu\|_{L_2(\pi)}. \tag{47}$$

Then

$$\begin{aligned} \|\mu P^k\|_{L_2(\pi)}^2 &= \langle \mu P^k, \mu P^k \rangle_{L_2(\pi)} \\ &= \langle \mu, \mu P^{2k} \rangle_{L_2(\pi)} \\ &= \left\langle \mu, \mu \int_{(0, \rho_2]} \lambda^{2k} \mathcal{E}(d\lambda) \right\rangle_{L_2(\pi)} \\ &\geq \left\langle \mu, \mu \int_{(\bar{\rho}, \rho_2]} \lambda^{2k} \mathcal{E}(d\lambda) \right\rangle_{L_2(\pi)} \\ &\geq \left\langle \mu, \mu \int_{(\bar{\rho}, \rho_2]} \bar{\rho}^{2k} \mathcal{E}(d\lambda) \right\rangle_{L_2(\pi)} \\ &= \bar{\rho}^{2k} \|proj_R \mu\|_{L_2(\pi)}^2 \\ &\geq \bar{\rho}^{2k} \frac{9}{64} \|\nu\|_{L_2(\pi)}^2. \end{aligned} \tag{48}$$

Hence $\rho \geq \bar{\rho}$. This contradicts $\rho_2 > \rho$.

Proof that [(i) \implies (ii)]:

Let the optimal rates of $(L_\infty(\pi_\epsilon), \|\cdot\|_{L_1(\pi)})$ -GE and $(L_\infty(\pi_\epsilon), \|\cdot\|_{L_2(\pi)})$ -GE be (respectively)

$$\rho = \sup_{\mu \in L_{\infty,0}(\pi)} \limsup_{n \rightarrow \infty} \|\mu P^n\|_{L_1(\pi)}^{1/n}, \quad \rho_2 = \sup_{\mu \in L_{\infty,0}(\pi)} \limsup_{n \rightarrow \infty} \|\mu P^n\|_{L_2(\pi)}^{1/n}. \tag{49}$$

We want to show that $\rho_2 \leq \rho$.

Let $\epsilon > 0$ be arbitrary. Let $\nu_\epsilon \in L_{\infty,0}(\pi)$ with

$$\limsup_{n \rightarrow \infty} \|\nu_\epsilon P^n\|_{L_2(\pi)}^{1/n} \geq \rho_2 - \epsilon. \tag{50}$$

Then, for some $c(\nu_\epsilon) > 0$, for infinitely many $n \in \mathbb{N}$,

$$\|\nu_\epsilon P^n\|_{L_2(\pi)} \geq (\rho_2 - 2\epsilon)^n. \tag{51}$$

Using the fact that

$$\|\mu\|_{L_1(\pi)} = \sup_{\substack{f \in L'_\infty(\pi) \\ \|f\|_{L'_\infty(\pi)}}} \mu f,$$

and using the self-adjointness of P in $L_2(\pi)$ (since P is reversible), as well as the fact that (a version of) $\frac{dv_\epsilon}{d\pi}$ is some bounded function with

$$\left\| \frac{dv_\epsilon}{d\pi} \right\|_{L'_\infty(\pi)} = \|v_\epsilon\|_{L_\infty(\pi)},$$

we have the following for infinitely many $n \in \mathbb{N}$:

$$\begin{aligned} \|v_\epsilon P^{2n}\|_{L_1(\pi)} &= \sup_{\|f\|_\infty \leq 1} v_\epsilon P^{2n} f \\ &\geq \frac{1}{\|v_\epsilon\|_{L_\infty(\pi)}} v_\epsilon P^{2n} \frac{dv_\epsilon}{d\pi} \\ &= \frac{1}{\|v_\epsilon\|_{L_\infty(\pi)}} \langle v_\epsilon P^{2n}, v_\epsilon \rangle \\ &= \frac{1}{\|v_\epsilon\|_{L_\infty(\pi)}} \langle v_\epsilon P^n, v_\epsilon P^n \rangle \\ &= \frac{1}{\|v_\epsilon\|_{L_\infty(\pi)}} \|v_\epsilon P^n\|_{L_2(\pi)}^2 \\ &\geq \frac{1}{\|v_\epsilon\|_{L_\infty(\pi)}} (\rho_2 - 2\epsilon)^{2n}. \end{aligned} \tag{52}$$

Thus $\rho_2 - 2\epsilon \leq \rho$. Since ϵ was arbitrary, we find that $\rho_2 \leq \rho$. □

Acknowledgements

We thank Daniel Rudolf for very helpful comments on the first version of our preprint. We also thank Gareth O. Roberts, Peter Rosenthal, and Don Hadwin for helpful discussions.

References

- [1] ALQUIER, P., FRIEL, N., EVERITT, R. AND BOLAND, A. (2016). Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Statist. Comput.* **26**, 29–47.
- [2] ANDRIEU, C. AND ROBERTS, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.* **37**, 697–725.
- [3] BAXENDALE, P. H. (2005). Renewal theory and computable convergence rates for geometrically ergodic Markov chains. *Ann. Appl. Prob.* **15**, 700–738.
- [4] BAXTER, J. R. AND ROSENTHAL, J. S. (1995). Rates of convergence for everywhere-positive Markov chains. *Statist. Prob. Lett.* **22**, 333–338.
- [5] BRESLOW, N. E. AND CLAYTON, D. G. (1993). Approximate inference in generalized linear mixed models. *J. Amer. Statist. Assoc.* **88**, 9–25.
- [6] BREYER, L., ROBERTS, G. AND ROSENTHAL, J. A note on geometric ergodicity and floating-point roundoff error. *Statist. Prob. Lett.* **53**, 123–127.
- [7] BROOKS, S., GELMAND, A., JONES, G. L. AND MENG, X.-L. (eds) (2011). *Handbook of Markov Chain Monte Carlo*. Chapman and Hall, Boca Raton, FL.
- [8] CAMPBELL, T. AND BRODERICK, T. (2019). Automated scalable Bayesian inference via Hilbert coresets. *J. Mach. Learning Res.* **20**, 551–588.
- [9] DOUC, R., MOULINES, E., PRIOURET, P. AND SOULIER, P. (2018). *Markov Chains*. Springer, Cham.

- [10] FERRÉ, D., HERVÉ, L. AND LEDOUX, J. (2013). Regular perturbation of V-geometrically ergodic Markov chains. *J. Appl. Prob.* **50**, 184–194.
- [11] GIBBS, A. (2004). Convergence in the Wasserstein metric for Markov chain Monte Carlo algorithms with applications to image restoration. *Stoch. Models* **20**, 473–492.
- [12] HERVÉ, L. AND LEDOUX, J. (2014). Approximating Markov chains and V-geometric ergodicity via weak perturbation theory. *Stoch. Process. Appl.* **124**, 613–638.
- [13] HOBERT, J. P. AND GEYER, C. J. (1998). Geometric ergodicity of Gibbs and block Gibbs samplers for a hierarchical random effects model. *J. Multivariate Anal.* **67**, 414–430.
- [14] HUGGINS, J., CAMPBELL, T. AND BRODERICK, T. (2016). Coresets for scalable Bayesian logistic regression. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, eds D. LEE *et al.* et al., NIPS, Barcelona, pp. 4080–4088.
- [15] JAIN, N. AND JAMISON, B. (1967). Contributions to Doebelin’s theory of Markov processes. *Prob. Theory Relat. Fields* **8**, 19–40.
- [16] JOHNDROW, J. E. AND MATTINGLY, J. C. (2017). *Coupling and decoupling to bound an approximating Markov chain*. Preprint. Available at <https://arxiv.org/abs/1706.02040>.
- [17] JOHNDROW, J. E. AND MATTINGLY, J. C. (2017). *Error bounds for approximations of Markov chains used in Bayesian sampling*. Preprint. Available at <https://arxiv.org/abs/1711.05382>.
- [18] JOHNDROW, J. E., MATTINGLY, J. C., MUKHERJEE, S. AND DUNSON, D. (2015). *Approximations of Markov chains and high-dimensional Bayesian inference*. Preprint. Available at <https://arxiv.org/abs/1508.03387v1>.
- [19] KASS, R. E., TIERNEY, L. AND KADANE, J. B. (1990). The validity of posterior expansions based on Laplace’s method. In *Bayesian and Likelihood Methods in Statistics and Econometrics*, Elsevier, Amsterdam, pp. 473–488.
- [20] KELLER, G. AND LIVERANI, C. (1999). Stability of the spectrum for transfer operators. *Ann. Scuola Norm. Sci.* **28**, 141–152.
- [21] LIU, J. S. (2008). *Monte Carlo Strategies in Scientific Computing*. Springer, New York.
- [22] LIVINGSTONE, S., BETANCOURT, M., BYRNE, S. AND GIROLAMI, M. (2019). On the geometric ergodicity of Hamiltonian Monte Carlo. *Bernoulli* **25**, 3109–3138.
- [23] MCCULLOCH, C. E. AND NEUHAUS, J. M. (2005). Generalized linear mixed models. In *Encyclopedia of Biostatistics*, eds P. ARMITAGE AND T. COLTON, JOHN WILEY, NEW YORK. Available at <https://doi.org/10.1002/0470011815.b2a10021>.
- [24] MEDINA-AGUAYO, F., RUDOLF, D. AND SCHWEIZER, N. (2019). Perturbation bounds for Monte Carlo within Metropolis via restricted approximations. *Stoch. Process. Appl.* **130**, 2200–2227.
- [25] MEDINA-AGUAYO, F. J., LEE, A. AND ROBERTS, G. O. (2016). Stability of noisy Metropolis–Hastings. *Statist. Comput.* **26**, 1187–1211.
- [26] MITROPHANOV, A. Y. (2005). Sensitivity and convergence of uniformly ergodic Markov chains. *J. Appl. Prob.* **42**, 1003–1014.
- [27] NUMMELIN, E. AND TWEEDIE, R. L. (1978). Geometric ergodicity and R-positivity for general Markov chains. *Ann. Prob.* **6**, 404–420.
- [28] PILLAI, N. S. AND SMITH, A. (2014). *Ergodicity of approximate MCMC chains with applications to large data sets*. Preprint. Available at <https://arxiv.org/abs/1405.0182>.
- [29] ROBERTS, G. O. AND ROSENTHAL, J. S. (1997). Geometric ergodicity and hybrid Markov chains. *Electron. Commun. Probab.* **2**, 13–25.
- [30] ROBERTS, G. O. AND ROSENTHAL, J. S. (2004). General state space Markov chains and MCMC algorithms. *Prob. Surveys* **1**, 20–71.
- [31] ROBERTS, G. O., ROSENTHAL, J. S. AND SCHWARTZ, P. O. (1998). Convergence properties of perturbed Markov chains. *J. Appl. Prob.* **35**, 1–11.
- [32] ROBERTS, G. O. AND TWEEDIE, R. L. (2001). Geometric L^2 and L^1 convergence are equivalent for reversible Markov chains. *J. Appl. Prob.* **38**, 37–41.
- [33] RUDIN, W. (1991). *Functional Analysis*, 2nd edn. McGraw-Hill, New York.
- [34] RUDOLF, D. (2011). Explicit error bounds for Markov chain Monte Carlo. *Dissertationes Math.* **485**, 93 pp. Available at <https://arxiv.org/abs/1108.3201>.
- [35] RUDOLF, D. AND SCHWEIZER, N. (2015). *Perturbation theory for Markov chains via Wasserstein distance*. Preprint. Available at <https://arxiv.org/abs/1503.04123>.
- [36] SMITH, R. L. AND TIERNEY, L. (1996). Exact transition probabilities for the independence Metropolis sampler. Available at <http://www.rls.sites.oasis.unc.edu/postscript/rs/exact.pdf>.
- [37] TIERNEY, L. AND KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81**, 82–86.