

Unbiased Markov chain Monte Carlo with couplings

Pierre E. Jacob^{*}, John O’Leary[†], Yves F. Atchadé[‡]

May 22, 2018

Abstract

Markov chain Monte Carlo (MCMC) methods provide consistent approximations of integrals as the number of iterations goes to infinity. MCMC estimators are generally biased after any fixed number of iterations, which complicates both parallel computation and the construction of confidence intervals. We propose to remove this bias by using couplings of Markov chains together with a telescopic sum argument of Glynn & Rhee (2014). The resulting unbiased estimators can be computed in parallel, with confidence intervals following directly from the Central Limit Theorem for i.i.d. variables. We discuss practical couplings for popular algorithms such as Metropolis-Hastings, Gibbs samplers, and Hamiltonian Monte Carlo. We establish the theoretical validity of the proposed estimators and study their efficiency relative to the underlying MCMC algorithms. Finally, we illustrate the performance and limitations of the method on toy examples, a variable selection problem, and an approximation of the cut distribution arising in Bayesian inference for models made of multiple modules.

1 Context

Markov chain Monte Carlo (MCMC) methods constitute a popular class of algorithms to approximate high-dimensional integrals such as those arising in statistics and many other fields [Liu, 2008, Robert and Casella, 1999, Brooks et al., 2011, Green et al., 2015]. These iterative methods provide estimators that are consistent in the limit of the number of iterations but potentially biased for any fixed number of iterations, as the Markov chains are rarely started at stationarity. This “burn-in” bias limits the potential gains from running independent chains in parallel [Rosenthal, 2000]. Consequently, efforts have focused on exploiting parallel processors within each iteration [Tjelmeland, 2004, Brockwell, 2006, Lee et al., 2010, Jacob et al., 2011, Calderhead, 2014, Goudie et al., 2017, Yang et al., 2017], or on the design of parallel chains targeting different distributions [Altekar et al., 2004, Wang et al., 2015, Srivastava et al., 2015]. Nevertheless, MCMC estimators are ultimately justified by asymptotics in the number of iterations. This imposes a severe limitation on the scalability of MCMC methods on modern computing hardware, with increasingly many processors and stagnating clock speeds.

We propose a general construction of unbiased estimators of integrals with respect to a target probability distribution using MCMC kernels. Thanks to the lack of bias, estimators can be generated independently in parallel and averaged over, thus achieving the standard Monte Carlo convergence rate as the number of parallel replicates goes to infinity. Confidence intervals can be constructed via the standard Central Limit Theorem (CLT) for i.i.d. variables, asymptotically valid in the number of parallel replicates, in contrast with confidence intervals for the standard MCMC approach. Indeed these

^{*}Department of Statistics, Harvard University, Cambridge, USA. Email: pjacob@fas.harvard.edu

[†]Department of Statistics, Harvard University, Cambridge, USA. Email: joleary@g.harvard.edu

[‡]Department of Statistics, University of Michigan, Ann Arbor, USA. Email: yvesa@umich.edu

are justified asymptotically in the number of iterations [e.g. [Flegal et al., 2008](#), [Gong and Flegal, 2016](#), [Atchadé, 2016](#), [Vats et al., 2018](#)], although they might also provide useful guidance in the non-asymptotic regime.

Our contribution follows the path-breaking work of [Glynn and Rhee \[2014\]](#), which demonstrates the unbiased estimation of integrals with respect to an invariant distribution using couplings. Their construction is illustrated on Markov chains represented by iterated random functions, and leverages the contraction properties of such functions. [Glynn and Rhee \[2014\]](#) also consider Harris recurrent chains for which an explicit minorization condition holds. Previously, [McLeish \[2011\]](#) employed similar debiasing techniques to obtain “nearly unbiased” estimators from a single MCMC chain. More recently [Jacob et al. \[2017a\]](#) remove the bias from particle Gibbs samplers [[Andrieu et al., 2010](#)] targeting the smoothing distribution in state-space models, by coupling chains such that they meet exactly in finite time without analytical knowledge on the underlying Markov kernels. The present article brings this type of Rhee & Glynn estimators to generic MCMC algorithms, along with new unbiased estimators with reduced variance. The proposed construction involves couplings of MCMC chains, which we provide for various algorithms, including Metropolis–Hastings, Gibbs and Hamiltonian Monte Carlo samplers.

Couplings of MCMC algorithms have been used to study their convergence properties, from both theoretical and practical points of view [e.g. [Reutter and Johnson, 1995](#), [Johnson, 1996](#), [Rosenthal, 1997](#), [Johnson, 1998](#), [Neal, 1999](#), [Roberts and Rosenthal, 2004](#), [Johnson, 2004](#), [Johndrow and Mattingly, 2017](#)]. Couplings of Markov chains also underpin perfect samplers [[Propp and Wilson, 1996](#), [Murdoch and Green, 1998](#), [Casella et al., 2001](#), [Flegal and Herbei, 2012](#), [Lee et al., 2014](#), [Huber, 2016](#)]. A notable difference of the proposed approach is that only two chains have to be coupled for the proposed estimator to be unbiased, without further assumptions on the state space or on the target distribution. Thus the approach applies more broadly than perfect samplers [see [Glynn, 2016](#)], while yielding unbiased estimators rather than exact samples. Couplings of pairs of Markov chains also formed the basis of the approach of [Neal \[1999\]](#), with a similar motivation for parallel computation.

In [Section 2](#), we introduce the estimators and a coupling of random walk Metropolis–Hastings chains as an illustration. In [Section 3](#), we establish properties of these estimators under certain assumptions. In [Section 4](#), we propose couplings of popular MCMC algorithms, using maximal couplings and common random number strategies. In [Section 5](#), we demonstrate the applicability of our approach with examples including a bimodal distribution and a classic Gibbs sampler for nuclear pump failure data. We then consider more challenging tasks including variable selection in high dimension and the approximation of the cut distribution that arises in inference for models made of modules. [[Liu et al., 2009](#), [Plummer, 2014](#), [Jacob et al., 2017b](#)]. We summarize and discuss our findings in [Section 6](#). Scripts in R [[R Core Team, 2015](#)] are available online¹, and supplementary materials with extra numerical illustrations are available on the first author’s webpage.

2 Unbiased estimation from coupled chains

2.1 Basic “Rhee-Glynn” estimator

Given a target probability distribution π on a Polish space \mathcal{X} and a measurable real-valued test function h integrable with respect to π , we want to estimate the expectation $\mathbb{E}_\pi[h(X)] = \int h(x)\pi(dx)$. Let P denote a Markov transition kernel on \mathcal{X} that leaves π invariant, and let π_0 be some initial probability distribution on \mathcal{X} . Our estimators are based on a coupled pair of Markov chains $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$,

¹ Link: github.com/pierrejacob/debiasedmcmc.

which marginally start from π_0 and evolve according to P . More specifically, let \bar{P} be a transition kernel on the joint space $\mathcal{X} \times \mathcal{X}$ such that $\bar{P}((x, y), A \times \mathcal{X}) = P(x, A)$ and $\bar{P}((x, y), \mathcal{X} \times A) = P(y, A)$ for any $x, y \in \mathcal{X}$ and measurable set A . We then construct the coupled Markov chain $(X_t, Y_t)_{t \geq 0}$ as follows. We draw (X_0, Y_0) such that $X_0 \sim \pi_0$, and $Y_0 \sim \pi_0$. Given (X_0, Y_0) , we draw $X_1 \sim P(X_0, \cdot)$. For any $t \geq 1$, given $X_0, (X_1, Y_0), \dots, (X_t, Y_{t-1})$, we draw $(X_{t+1}, Y_t) \sim \bar{P}((X_t, Y_{t-1}), \cdot)$. We consider the following assumptions.

Assumption 2.1. *As $t \rightarrow \infty$, $\mathbb{E}[h(X_t)] \rightarrow \mathbb{E}_\pi[h(X)]$. Furthermore, there exists an $\eta > 0$ and $D < \infty$ such that $\mathbb{E}[|h(X_t)|^{2+\eta}] \leq D$ for all $t \geq 0$.*

Assumption 2.2. *The chains are such that the meeting time $\tau := \inf\{t \geq 1 : X_t = Y_{t-1}\}$ satisfies $\mathbb{P}(\tau > t) \leq C \delta^t$ for all $t \geq 0$, for some constants $C < \infty$ and $\delta \in (0, 1)$.*

Assumption 2.3. *The chains stay together after meeting, i.e. $X_t = Y_{t-1}$ for all $t \geq \tau$.*

By construction, each of the two marginal chains $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ has initial distribution π_0 and transition kernel P . Assumption 2.1 requires these chains to result in a uniformly bounded $(2 + \eta)$ -moment of h ; more discussion on moments of Markov chains can be found in Tweedie [1983]. Since X_0 and Y_0 may be drawn from any coupling of π_0 with itself, it is possible to set $X_0 = Y_0$. However, X_1 is then generated from $P(X_0, \cdot)$, so that $X_1 \neq Y_0$ in general. Thus one cannot force the meeting time to be small by setting $X_0 = Y_0$. Assumption 2.2 puts a condition on the coupling operated by \bar{P} , and would not be satisfied for an independent coupling. Coupled kernels that satisfy Assumption 2.2 can be designed using e.g. common random numbers and maximal couplings. We present a simple case in Section 2.2 and further examples in Section 4. We stress that the state space is not assumed to be discrete, and that the constants D and η of Assumption 2.1 and C and δ of Assumption 2.2 do not need to be known to implement the proposed approach. Assumption 2.3 typically holds by design; coupled chains satisfying this assumption are termed “faithful” in Rosenthal [1997].

Under these assumptions we introduce the following motivation for an unbiased estimator of $\mathbb{E}_\pi[h(X)]$, following Glynn and Rhee [2014]. We begin by writing $\mathbb{E}_\pi[h(X)]$ as $\lim_{t \rightarrow \infty} \mathbb{E}[h(X_t)]$. Then for any fixed $k \geq 0$,

$$\begin{aligned} \mathbb{E}_\pi[h(X)] &= \mathbb{E}[h(X_k)] + \sum_{t=k+1}^{\infty} (\mathbb{E}[h(X_t)] - \mathbb{E}[h(X_{t-1})]) && \text{expanding the limit as a telescoping sum,} \\ &= \mathbb{E}[h(X_k)] + \sum_{t=k+1}^{\infty} (\mathbb{E}[h(X_t)] - \mathbb{E}[h(Y_{t-1})]) && \text{since the chains have the same marginals,} \\ &= \mathbb{E}[h(X_k)] + \sum_{t=k+1}^{\infty} (h(X_t) - h(Y_{t-1})) && \text{swapping the expectations and limit,} \\ &= \mathbb{E}[h(X_k)] + \sum_{t=k+1}^{\tau-1} (h(X_t) - h(Y_{t-1})) && \text{since the terms corresponding to } t \geq \tau \text{ are null.} \end{aligned}$$

We note that the sum in the last equation is zero if $k+1 > \tau-1$. The heuristic argument above suggests that the estimator $H_k(X, Y) = h(X_k) + \sum_{t=k+1}^{\tau-1} (h(X_t) - h(Y_{t-1}))$ should have expectation $\mathbb{E}_\pi[h(X)]$.

This estimator requires τ calls to \bar{P} and $\max(1, k+1-\tau)$ calls to P ; thus under Assumption 2.2 its cost has a finite expectation. In Section 3 we establish the validity of the estimator under the three conditions above; this formally justifies the swap of expectation and limit. The estimator can be viewed as a debiased version of $h(X_k)$, where the term $\sum_{t=k+1}^{\tau-1} (h(X_t) - h(Y_{t-1}))$ acts as bias correction. Thanks to this unbiasedness property, we can sample $R \in \mathbb{N}$ independent copies of $H_k(X, Y)$ in parallel and

average the results to estimate $\mathbb{E}_\pi[h(X)]$. Unbiasedness is guaranteed for any choice of $k \geq 0$, but both cost and variance of $H_k(X, Y)$ are sensitive to k .

Before presenting examples and enhancements to the estimator above, we discuss the relationship between our approach and existing work. There is a rich literature applying forward couplings to study Markov chains convergence [Johnson, 1996, 1998, Thorisson, 2000, Lindvall, 2002, Rosenthal, 2002, Johnson, 2004, Douc et al., 2004], and to obtain new algorithms such as perfect samplers [Huber, 2016] and the methods of Neal [1999] and Neal and Pinto [2001]. Our approach is closely related to Glynn and Rhee [2014], who employ pairs of Markov chains to obtain unbiased estimators. The present work combines similar arguments with couplings of MCMC algorithms and proposes further improvements to remove bias at a reduced loss of efficiency.

Indeed Glynn and Rhee [2014] did not apply their methodology to the MCMC setting. They consider chains associated with contractive iterated random functions [see also Diaconis and Freedman, 1999], and Harris recurrent chains with an explicit minorization condition. A minorization condition refers to a small set \mathcal{C} , $\lambda > 0$, an integer $m \geq 1$, and a probability measure ν such that, for all $x \in \mathcal{C}$ and measurable set A , $P^m(x, A) \geq \lambda\nu(A)$. It is explicit if the set, constant and probability measure are known by the user. Finding explicit small sets that are practically useful is a challenging technical task, even for MCMC experts. If available, explicit minorization conditions could also be employed to identify regeneration times, leading to unbiased estimators amenable to parallel computation in the framework of Mykland et al. [1995] and Brockwell and Kadane [2005]. By contrast Johnson [1996, 1998], Neal [1999] more explicitly address the question of coupling MCMC algorithms such that pairs of chains meet exactly, without analytical knowledge on the target distribution. The present article focuses on the use of these couplings in the framework of Glynn and Rhee [2014].

2.2 Coupled Metropolis–Hastings example

Before further examination of our estimator and its properties, we present a coupling of Metropolis–Hastings (MH) chains that will typically satisfy Assumptions 2.1–2.3 in realistic settings; this coupling was proposed in Johnson [1998] as part of a method to diagnose convergence. We postpone discussion of couplings for other MCMC algorithms to Section 4. We recall that each iteration t of the MH algorithm [Hastings, 1970] begins by drawing a proposal X^* from a Markov kernel $q(X_t, \cdot)$, where X_t is the current state. The next state is set to $X_{t+1} = X^*$ if $U \leq \pi(X^*)q(X^*, X_t)/\pi(X_t)q(X_t, X^*)$, where U denotes a uniform random variable on $[0, 1]$, and $X_{t+1} = X_t$ otherwise.

We define a pair of chains so that each proceeds marginally according to the MH algorithm and jointly so that the chains will meet exactly after a random number of steps. We suppose the pair of chains are in states X_t and Y_{t-1} , and consider how to generate X_{t+1} and Y_t so that $\{X_{t+1} = Y_t\}$ might occur.

If $X_t \neq Y_{t-1}$, the event $\{X_{t+1} = Y_t\}$ cannot occur if both chains reject their respective proposals, X^* and Y^* . Meeting will occur if these proposals are identical and if both are accepted. Marginally, the proposals follow $X^*|X_t \sim q(X_t, \cdot)$ and $Y^*|Y_{t-1} \sim q(Y_{t-1}, \cdot)$. If $q(x, x^*)$ can be evaluated for all x, x^* , then one can sample from the maximal coupling between the two proposal distributions, which is the coupling of $q(X_t, \cdot)$ and $q(Y_{t-1}, \cdot)$ maximizing the probability of the event $\{X^* = Y^*\}$. How to sample from maximal couplings of continuous distributions is well-known [Thorisson, 2000] and described in Section 4.1 for completeness. One can accept or reject the two proposals using a common uniform random variable U . The chains will stay together after they meet: at each step after meeting, the proposals will be identical with probability one, and jointly accepted or rejected with a common uniform variable. This coupling does not require explicit minorization conditions, nor contractive properties of a random

Algorithm 1 Unbiased “time-averaged” estimator $H_{k:m}(X, Y)$ of $\mathbb{E}_\pi[h(X)]$.

1. Draw X_0, Y_0 from an initial distribution π_0 and draw $X_1 \sim P(X_0, \cdot)$.
 2. Set $t = 1$. While $t < \max(m, \tau)$, where $\tau = \inf\{t \geq 1 : X_t = Y_{t-1}\}$,
 - draw $(X_{t+1}, Y_t) \sim \bar{P}((X_t, Y_{t-1}), \cdot)$,
 - set $t \leftarrow t + 1$.
 3. For each $\ell \in \{k, \dots, m\}$, compute $H_\ell(X, Y) = h(X_\ell) + \sum_{t=\ell+1}^{\tau-1} (h(X_t) - h(Y_{t-1}))$.
- Return $H_{k:m}(X, Y) = (m - k + 1)^{-1} \sum_{\ell=k}^m H_\ell(X, Y)$; or compute $H_{k:m}(X, Y)$ with (2.1).
-

function representation of the chain.

2.3 Time-averaged estimator

To motivate our next estimator, we note that we can compute $H_k(X, Y)$ for several values of k from the same realization of the coupled chains, and that the average of these is unbiased as well. For any fixed integer m with $m \geq k$, we can run coupled chains for $\max(m, \tau)$ iterations, compute the estimator $H_\ell(X, Y)$ for each $\ell \in \{k, \dots, m\}$, and take the average $H_{k:m}(X, Y) = (m - k + 1)^{-1} \sum_{\ell=k}^m H_\ell(X, Y)$, as we summarize in Algorithm 1. We refer to $H_{k:m}(X, Y)$ as the *time-averaged estimator*; the estimator $H_k(X, Y)$ is retrieved when $m = k$. Alternatively we could average the estimators $H_\ell(X, Y)$ using weights $w_\ell \in \mathbb{R}$ for $\ell \in \{k, \dots, m\}$, to obtain $\sum_{\ell=k}^m w_\ell H_\ell(X, Y)$. This will be unbiased if $\sum_{\ell=k}^m w_\ell = 1$. The computation of weights to minimize the variance of $\sum_{\ell=k}^m w_\ell H_\ell(X, Y)$ for a given test function h is an open question.

Rearranging terms in $(m - k + 1)^{-1} \sum_{\ell=k}^m H_\ell(X, Y)$, we can write the time-averaged estimator as

$$H_{k:m}(X, Y) = \frac{1}{m - k + 1} \sum_{\ell=k}^m h(X_\ell) + \sum_{\ell=k}^{\tau-1} \min\left(1, \frac{\ell - k + 1}{m - k + 1}\right) (h(X_{\ell+1}) - h(Y_\ell)). \quad (2.1)$$

The term $(m - k + 1)^{-1} \sum_{\ell=k}^m h(X_\ell)$ corresponds to a standard MCMC average with m total iterations and a burn-in period of $k - 1$ iterations. We can interpret the other term as a bias correction. If $\tau \leq k + 1$, then the correction term equals zero. This provides some intuition about the choice of k and m : large k values lead to the bias correction being equal to zero with large probability, and large values of m result in $H_{k:m}(X, Y)$ being similar to an estimator obtained from a long MCMC run. Thus we expect the variance of $H_{k:m}(X, Y)$ to be similar to that of MCMC for appropriate choices of k and m .

The estimator $H_{k:m}(X, Y)$ requires τ calls to \bar{P} and $\max(1, m + 1 - \tau)$ calls to P , which is comparable to m calls to P when m is large. Thus both the variance and the cost of $H_{k:m}(X, Y)$ will approach those of MCMC estimators for large values of k and m . This motivates the use of the estimator $H_{k:m}(X, Y)$ with $m > k$, since the time-averaged estimator allows us to limit the loss of efficiency associated with the removal of the burn-in bias. We discuss the choice of k and m in further detail in Section 3 and in the experiments.

2.4 Practical considerations

Once we have run the first two steps of Algorithm 1, we can store X_k and (X_t, Y_{t-1}) for $k + 1 \leq t \leq m$ for later use: the test function h does not have to be specified at run-time.

One typically resorts to thinning the output of an MCMC sampler if the test function of interest is

unknown at run-time, if the memory cost of storing long chains is prohibitive, or if the cost of evaluating the test function of interest is significant compared to the cost of each MCMC iteration [e.g. Owen, 2017]. This is also possible in the proposed framework: one can consider a variation of Algorithm 1 where each call to the Markov kernels P and \bar{P} would be replaced by multiple calls.

Algorithm 1 terminates after τ calls to \bar{P} and $\max(1, m+1-\tau)$ calls to P . For the proposed couplings, calls to \bar{P} are approximately twice as expensive as calls to P . Therefore, the cost of $H_{k:m}(X, Y)$ is comparable to $2\tau + \max(1, m+1-\tau)$ iterations of the underlying MCMC algorithm. This cost is random and will generally depend on the specific coupling underlying the estimator.

As in regular Monte Carlo estimation, the use of a fixed computation budget yielding a random number of complete estimator calculations requires care. The naive approach – to take the average of completed estimators and discard ongoing calculations – can produce biased results [Glynn and Heidelberger, 1990]. Still, unbiased estimation is possible, as in Corollary 7 of the aforementioned article.

In addition to estimating integrals, it is often of interest to visualize the target distribution. We use our estimator to construct histograms for the marginal distributions of π by targeting $\mathbb{E}_\pi[\mathbf{1}(X(i) \in A)]$ for various intervals A , where $X(i)$ denotes the i -th component of X . We can also obtain confidence intervals for these histogram probabilities by computing the variance of the estimators of $\mathbb{E}_\pi[\mathbf{1}(X(i) \in A)]$. Such histograms are presented in Section 5 with 95% confidence intervals as grey vertical boxes and point estimates as black vertical bars. Note that the proposed estimators can take values outside the range of the test function h , so that the proposed histograms may include negative values as probability estimates; see Jacob and Thiery [2015] on the possibility of non-negative unbiased estimators.

2.5 Signed measure estimator

We can formulate the proposed estimation procedure in terms of a signed measure $\hat{\pi}$ defined by

$$\hat{\pi}(\cdot) = \frac{1}{m-k+1} \sum_{\ell=k}^m \delta_{X_\ell}(\cdot) + \sum_{\ell=k}^{\tau-1} \min\left(1, \frac{\ell-k+1}{m-k+1}\right) (\delta_{X_{\ell+1}}(\cdot) - \delta_{Y_\ell}(\cdot)), \quad (2.2)$$

obtained by replacing test function evaluations by delta masses in (2.1), as in Section 4 of Glynn and Rhee [2014]. The measure $\hat{\pi}$ is of the form $\hat{\pi}(\cdot) = \sum_{\ell=1}^N \omega_\ell \delta_{Z_\ell}(\cdot)$ with $\sum_{\ell=1}^N \omega_\ell = 1$ and where the atoms (Z_ℓ) are values among (X_t) and (Y_t) . Some of the weights (ω_ℓ) might be negative, making $\hat{\pi}$ a signed empirical measure. The unbiasedness property states $\mathbb{E}[\sum_{\ell=1}^N \omega_\ell h(Z_\ell)] = \mathbb{E}_\pi[h(X)]$.

One can consider the convergence behavior of $\hat{\pi}^R(\cdot) = R^{-1} \sum_{r=1}^R \hat{\pi}^{(r)}(\cdot)$ towards π , where $(\hat{\pi}^{(r)})$ are independent replications of $\hat{\pi}$. Glynn and Rhee [2014] obtain a Glivenko–Cantelli result for a similar measure related to their estimator. In the current setting, assume for simplicity that π is univariate or else consider only one of its marginals. We redefine the weights and the atoms to write $\hat{\pi}^R(\cdot) = \sum_{\ell=1}^{N_R} \omega_\ell \delta_{Z_\ell}(\cdot)$. Introduce the function $s \mapsto \hat{F}^R(s) = \sum_{\ell=1}^{N_R} \omega_\ell \mathbf{1}(Z_\ell \leq s)$ on \mathbb{R} . Proposition 3.2 states that \hat{F}^R converges to F uniformly with probability one, where F is the cumulative distribution function of π .

The function $s \mapsto \hat{F}^R(s)$ is not monotonically increasing because of negative weights among (ω_ℓ) . Therefore, for any $q \in (0, 1)$ there might more than one index ℓ such that $\sum_{i=1}^{\ell-1} \omega_i \leq q$ and $\sum_{i=1}^{\ell} \omega_i > q$; the quantile estimate might be defined as Z_ℓ for any such ℓ . The convergence of \hat{F}^R to F indicates that all such estimates are expected to converge to the q -th quantile of π . Therefore the signed measure representation leads to a way of estimating quantiles of the target distribution in a consistent way as $R \rightarrow \infty$. The construction of confidence intervals for these quantiles, perhaps by bootstrapping the R independent copies, stands as an interesting area for future research.

Another route to estimate quantiles of π would be to project $\hat{\pi}^R$, or some of its marginals, onto the

space of probability measures. For instance, one could search for the vector $(\bar{\omega}_\ell)$ in the N_R -simplex $\{\bar{\omega}_\ell, \ell = 1, \dots, N_R : \bar{\omega}_\ell \geq 0, \sum_{\ell=1}^{N_R} \bar{\omega}_\ell = 1\}$ such that $\bar{\pi}^R(\cdot) = \sum_{\ell=1}^{N_R} \bar{\omega}_\ell \delta_{Z_\ell}(\cdot)$ is closest to $\hat{\pi}^R(\cdot)$, in some sense. That sense could be a generalization of the Wasserstein metric to signed measures [Mainini, 2012]. Another option would be to estimate F using isotonic regression [Chatterjee et al., 2015], considering $\hat{F}^R(s)$ for various values s as noisy measurements of $F(s)$; this amounts to another projection of $\hat{\pi}^R(\cdot)$ onto probability measures. One could hope that as $\hat{\pi}^R$ approaches π , the projection $\bar{\pi}^R$ would also converge to π , preserving consistency in $R \rightarrow \infty$. In that case, $(\bar{\omega}_\ell, Z_\ell)_{\ell=1}^{N_R}$ are weighted samples which can be used to approximate quantiles or plot histograms approximating π . Another appeal of $\bar{\pi}^R$ is that weighted averages $\sum_{i=1}^{N_R} \bar{\omega}_\ell h(Z_\ell)$ are guaranteed to take values in the convex hull of the range of h .

3 Theoretical properties and guidance

We state our main result for the estimator $H_k(X, Y)$, which extends directly to $H_{k:m}(X, Y)$.

Proposition 3.1. *Under Assumptions 2.1-2.3, for all $k \geq 0$, the estimator $H_k(X, Y)$ has expectation $\mathbb{E}_\pi[h(X)]$, a finite variance, and a finite expected computing time.*

From the proof of Proposition 3.1, it is clear that Assumption 2.2 could be weakened: geometric tails of the meeting time are sufficient but not necessary. The main consequence of Proposition 3.1 is that an average of R independent copies of $H_{k:m}(X, Y)$ converges to $\mathbb{E}_\pi[h(X)]$ as $R \rightarrow \infty$, and that a central limit theorem holds.

Concerning the signed measure estimator of (2.2), following Glynn and Rhee [2014] we provide Proposition 3.2, which applies to univariate target distributions or to marginals of the target.

Proposition 3.2. *Under Assumptions 2.2-2.3, for all $m \geq k \geq 0$, and assuming that $(X_t)_{t \geq 0}$ converges to π in total variation, introduce the function $s \mapsto \hat{F}^R(s) = \sum_{\ell=1}^{N_R} \omega_\ell \mathbb{1}(Z_\ell \leq s)$, where $(\omega_\ell, Z_\ell)_{\ell=1}^{N_R}$ are weighted atoms obtained from R independent copies of $\hat{\pi}$ in (2.2). Denote by F the cumulative distribution function of π . Then $\sup_{s \in \mathbb{R}} |\hat{F}^R(s) - F(s)| \xrightarrow{R \rightarrow \infty} 0$ almost surely.*

In Section 3.1, we discuss the variance and efficiency of $H_{k:m}(X, Y)$, and the effect of k and m . In Section 3.2, we investigate the verification of Assumption 2.2 using drift conditions.

3.1 Variance and efficiency

Estimators $H_{k:m}^{(r)}(X, Y)$, for $r = 1, \dots, R$, can be generated in parallel and averaged. More estimators can be produced in a given computing budget if each estimator is cheaper to produce. The trade-off can be understood in the framework of Glynn and Whitt [1992], also used in Rhee and Glynn [2012] and Glynn and Rhee [2014], by defining the asymptotic inefficiency as the product of the variance and expected cost of the estimator. Indeed, the product of expected cost and variance is equal to the asymptotic variance of $R^{-1} \sum_{r=1}^R H_{k:m}^{(r)}(X, Y)$ as the computational budget, as opposed to the number of estimators R , goes to infinity [Glynn and Whitt, 1992]. Of primary interest is the comparison of this asymptotic inefficiency with the asymptotic variance of standard MCMC estimators.

We start by writing the time-averaged estimator of (2.1) as

$$H_{k:m}(X, Y) = \text{MCMC}_{k:m} + \text{BC}_{k:m},$$

where $\text{MCMC}_{k:m}$ is the MCMC average $(m - k + 1)^{-1} \sum_{\ell=k}^m h(X_\ell)$ and $\text{BC}_{k:m}$ is the bias correction

term. The variance of $H_{k:m}(X, Y)$ can be written

$$\mathbb{V}[H_{k:m}(X, Y)] = \mathbb{E}[(\text{MCMC}_{k:m} - \mathbb{E}_\pi[h(X)])^2] + 2\mathbb{E}[(\text{MCMC}_{k:m} - \mathbb{E}_\pi[h(X)])\text{BC}_{k:m}] + \mathbb{E}[\text{BC}_{k:m}^2].$$

Defining the mean squared error of the MCMC estimator as $\text{MSE}_{k:m} = \mathbb{E}[(\text{MCMC}_{k:m} - \mathbb{E}_\pi[h(X)])^2]$, Cauchy-Schwarz inequality yields

$$\mathbb{V}[H_{k:m}(X, Y)] \leq \text{MSE}_{k:m} + 2\sqrt{\text{MSE}_{k:m}}\sqrt{\mathbb{E}[\text{BC}_{k:m}^2]} + \mathbb{E}[\text{BC}_{k:m}^2]. \quad (3.1)$$

To bound $\mathbb{E}[\text{BC}_{k:m}^2]$ we introduce a geometric drift condition on the Markov kernel P .

Assumption 3.1. *The Markov kernel P is π -invariant, φ -irreducible and aperiodic, and there exists a measurable function $V : \mathcal{X} \rightarrow [1, \infty)$, $\lambda \in (0, 1)$, $b < \infty$ and a small set \mathcal{C} such that for all $x \in \mathcal{X}$,*

$$\int P(x, dy)V(y) \leq \lambda V(x) + b\mathbf{1}(x \in \mathcal{C}).$$

We refer the reader to [Meyn and Tweedie \[2009\]](#) for the definitions and general concepts of Markov chains, in particular Chapter 5 for aperiodicity, φ -irreducibility and small sets, and Chapter 15 for geometric drift conditions; see also [Roberts and Rosenthal \[2004\]](#). Geometric drift conditions are known to hold for various MCMC algorithms [e.g. [Roberts and Tweedie, 1996a](#), [Janner and Roberts, 2007](#), [Atchade, 2006](#), [Khare and Hobert, 2013](#), [Choi and Hobert, 2013](#), [Pal and Khare, 2014](#)]. Assumption 3.1 often plays a central role in establishing geometric ergodicity [e.g. Theorem 9 in [Roberts and Rosenthal, 2004](#)]. We show next that this assumption enables an informative bound on $\mathbb{E}[\text{BC}_{k:m}^2]$.

Proposition 3.3. *Suppose that Assumptions 2.2-2.3 and 3.1 hold, with a function V for which $\int V(x)\pi_0(dx)$ is finite. If the function h is such that $\sup_{x \in \mathcal{X}} |h(x)|/V(x)^\beta < \infty$ for some $\beta \in [0, 1/2)$, then for all $m \geq k \geq 0$, we have*

$$\mathbb{E}[\text{BC}_{k:m}^2] \leq \frac{C_{\delta, \beta} \delta_\beta^k}{(m - k + 1)^2},$$

for some constants $C_{\delta, \beta} < +\infty$, and $\delta_\beta = \delta^{1-2\beta} \in (0, 1)$, with $\delta \in (0, 1)$ as in Assumption 2.2.

Using Proposition 3.3, (3.1) becomes

$$\mathbb{V}[H_{k:m}(X, Y)] \leq \text{MSE}_{k:m} + 2\sqrt{\text{MSE}_{k:m}} \frac{\sqrt{C_{\delta, \beta} \delta_\beta^k}}{m - k + 1} + \frac{C_{\delta, \beta} \delta_\beta^k}{(m - k + 1)^2}. \quad (3.2)$$

The variance of $H_{k:m}(X, Y)$ is thus bounded by the mean squared error of an MCMC estimator, and additive terms that vanish geometrically when k increases, and polynomially as $m - k$ increases.

In order to facilitate the comparison between the efficiency of $H_{k:m}(X, Y)$ and that of MCMC estimators, we add simplifying assumptions. First, the right-most terms of (3.2) decrease geometrically with k , at a rate driven by $\delta_\beta = \delta^{1-2\beta}$ where δ is as in Assumption 2.2. This motivates a choice of k depending on the distribution of the meeting time τ . In practice, we can sample independent realizations of the meeting time and choose k such that $\mathbb{P}(\tau > k)$ is small; i.e. we choose k as a large quantile of the meeting times.

Secondly, as k increases and for $m \geq k$, we expect $(m - k + 1)\text{MSE}_{k:m}$ to converge to $\mathbb{V}[(m - k + 1)^{-1/2} \sum_{t=k}^m h(X_t)]$, where X_k would be distributed according to π . Denote this variance by $V_{k,m}$. The limit of $V_{k,m}$ as $m \rightarrow \infty$ is the asymptotic variance of the MCMC estimator, denoted by V_∞ . We do the simplifying assumption that k is large enough for $(m - k + 1)\text{MSE}_{k:m}$ to be approximately $V_{k,m}$.

Furthermore, we approximate the cost of $H_{k:m}(X, Y)$ by the cost of m calls to P . Dropping the third term on the right-hand side of (3.2), which is of smaller magnitude than the second term, we obtain the approximate inequality

$$\mathbb{E}[2\tau + \max(1, m + 1 - \tau)]\mathbb{V}[H_{k:m}(X, Y)] \lesssim \frac{m}{m - k + 1}V_{k,m} + 2\frac{m\sqrt{V_{k,m}C_{\delta,\beta}\delta_{\beta}^k}}{(m - k + 1)^{3/2}}.$$

In order for the left-hand side to be comparable to the asymptotic variance of MCMC, we can choose m such that $m/(m - k + 1) \approx 1$, e.g. by defining m as a large multiple of k . The second term on the right-hand side is negligible compared to the first as either k or m increases. This informal series of approximations suggests that we can retrieve an asymptotic efficiency comparable to the underlying MCMC estimators with appropriate choices of k and m . In other words, the bias of MCMC can be removed at the cost of an increased variance, which can in turn be reduced by choosing large enough values of k and m . Large values of k and m are to be traded against the desired level of parallelism: one might prefer to keep m small, yielding a suboptimal efficiency for $H_{k:m}(X, Y)$, but enabling more independent copies to be generated in a given computing time.

Thus we propose to choose k such that $\mathbb{P}(\tau > k)$ is small, and m as a large multiple of k , for the asymptotic inefficiency to be comparable to that of the underlying MCMC algorithm; more precise recommendations would depend on the target, on the budget constraint and on the degree of parallelism of available hardware.

3.2 Verifying Assumption 2.2

We discuss how Assumption 3.1 can be used to verify Assumption 2.2. Informally, Assumption 3.1 guarantees that the bivariate chain $\{(X_t, Y_{t-1}), t \geq 1\}$ visits $\mathcal{C} \times \mathcal{C}$ infinitely often, where \mathcal{C} is a small set. Therefore, if there is a positive probability of the event $\{X_{t+1} = Y_t\}$ for every t such that $(X_t, Y_{t-1}) \in \mathcal{C} \times \mathcal{C}$, then we expect Assumption 2.2 to hold. The next result formalizes that intuition. The proof is based a modification of an argument by Douc et al. [2004]. For convenience, we introduce $\mathcal{D} = \{(x, y) \in \mathcal{X} \times \mathcal{X} : x = y\}$. Hence Assumption 2.3 reads $\bar{P}((x, x), \mathcal{D}) = 1$, for all $x \in \mathcal{X}$.

Proposition 3.4. *Suppose that P satisfies Assumption 3.1 with a small set \mathcal{C} of the form $\mathcal{C} = \{x : V(x) \leq L\}$ where $\lambda + \frac{b}{1+L} < 1$. Suppose also that there exists $\epsilon \in (0, 1)$ such that*

$$\inf_{x, y \in \mathcal{C}} \bar{P}((x, y), \mathcal{D}) \geq \epsilon. \quad (3.3)$$

Then there exists a finite constant C' , and $\kappa \in (0, 1)$, such that for all $n \geq 1$,

$$\mathbb{P}(\tau > n) \leq C'\pi_0(V)\kappa^n,$$

where $\pi_0(V) = \int V(x)\pi_0(dx)$. Hence Assumption 2.2 holds as long as $\pi_0(V) < \infty$.

4 Couplings of MCMC algorithms

To compute our unbiased estimators, we must couple Markov chains in a way that satisfies Assumptions 2.2-2.3. We discuss such couplings for Metropolis–Hastings, Gibbs, and Hamiltonian Monte Carlo. These couplings are widely applicable since they do not require extensive analytical knowledge of the target

distribution; however they are not optimal in general, and we expect case-specific constructions to yield more efficient estimators. We begin in Section 4.1 by reviewing maximal couplings.

4.1 Sampling from a maximal coupling

The maximal coupling between two distributions p and q on a space \mathcal{X} is the distribution of a pair of random variables (X, Y) that maximizes $\mathbb{P}(X = Y)$, subject to the marginal constraints $X \sim p$ and $Y \sim q$. A procedure to sample from a maximal coupling is described in Algorithm 2. Here $\mathcal{U}([a, b])$ refers to the uniform distribution on the interval $[a, b]$ for $a < b$. We write p and q for both these distributions and their probability density functions with respect to a common dominating measure. Algorithm 2 is well-known and described e.g. in Section 4.5 of Chapter 1 of [Thorisson \[2000\]](#); in [Johnson \[1998\]](#) it is termed γ -coupling.

We justify Algorithm 2 and compute its cost. Denote by (X, Y) the output of the algorithm. First, X follows p from step 1. To prove that Y follows q , we introduce a measurable set A and check that $\mathbb{P}(Y \in A) = \int_A q(y)dy$. We write $\mathbb{P}(Y \in A) = \mathbb{P}(Y \in A, \text{step 1}) + \mathbb{P}(Y \in A, \text{step 2})$, where the events $\{\text{step 1}\}$ and $\{\text{step 2}\}$ refer to the algorithm terminating at step 1 or 2. We compute

$$\mathbb{P}(Y \in A, \text{step 1}) = \int_A \int_0^{+\infty} \mathbb{1}(w \leq q(x)) \frac{\mathbb{1}(0 \leq w \leq p(x))}{p(x)} p(x) dw dx = \int_A \min(p(x), q(x)) dx,$$

from which we deduce that $\mathbb{P}(\text{step 1}) = \int_{\mathcal{X}} \min(p(x), q(x)) dx$. For $\mathbb{P}(Y \in A, \text{step 2})$ to be equal to $\int_A (q(x) - \min(p(x), q(x))) dx$, we need

$$\int_A (q(x) - \min(p(x), q(x))) dx = \mathbb{P}(Y \in A | \text{step 2}) \left(1 - \int_{\mathcal{X}} \min(p(x), q(x)) dx \right),$$

and we conclude that the distribution of Y given $\{\text{step 2}\}$ should have a density equal to $\tilde{q}(x) = (q(x) - \min(p(x), q(x))) / (1 - \int \min(p(x'), q(x')) dx')$ for all x . Step 2 is a standard rejection sampler using q as a proposal distribution to target \tilde{q} , which concludes the proof that $Y \sim q$. We now confirm that Algorithm 2 indeed maximizes the probability of $\{X = Y\}$. Under the algorithm,

$$\mathbb{P}(X = Y) = \int_{\mathcal{X}} \min(p(x), q(x)) dx = \frac{1}{2} \int_{\mathcal{X}} (p(x) + q(x) - |p(x) - q(x)|) dx = 1 - d_{\text{TV}}(p, q),$$

where $d_{\text{TV}}(p, q) = 1/2 \int_{\mathcal{X}} |p(x) - q(x)| dx$ is the total variation distance. By the coupling inequality [[Lindvall, 2002](#)], this proves that the algorithm implements a maximal coupling.

To address the cost of Algorithm 2, observe that the probability of acceptance in step 2 is given by

$$\mathbb{P}(W^* \geq p(Y^*)) = 1 - \int_{\mathcal{X}} \min(p(y), q(y)) dy.$$

Step 1 costs one draw from p , one evaluation from p and one from q . Each attempt in the rejection sampler of step 2 costs one draw from q , one evaluation from p and one from q . We refer to the cost of one draw and two evaluations by “one unit”, for simplicity. Then, there is a Geometric number of attempts in step 2, with mean $(1 - \int_{\mathcal{X}} \min(p(y), q(y)) dy)^{-1}$, and step 2 occurs with probability $1 - \int_{\mathcal{X}} \min(p(y), q(y)) dy$. Therefore the expected cost is of two units, for all distributions p and q . To summarize, the expected cost of the algorithm does not depend on total variation distance between p and q , and the probability of $\{X = Y\}$ is precisely one minus that distance.

Alternative couplings described in [Johnson \[1996\]](#), [Neal \[1999\]](#) include the following strategy, for univariate p and q . Let F_p^- and F_q^- be the quantile functions associated with p and q , and let U

Algorithm 2 Sampling from a maximal coupling of p and q .

1. Sample $X \sim p$ and $W|X \sim \mathcal{U}([0, p(X)])$. If $W \leq q(X)$, output (X, X) .
 2. Otherwise, sample $Y^* \sim q$ and $W^*|Y^* \sim \mathcal{U}([0, q(Y^*)])$ until $W^* > p(Y^*)$, and output (X, Y^*) .
-

denote a uniform random variable on $[0, 1]$. Then $X = F_p^-(U)$ and $Y = F_q^-(U)$ computed with the same realization of U constitute an optimal transport coupling of p and q , also called an “increasing rearrangement” [Villani, 2008, Chapter 1]. Such couplings minimize the expected distance between X and Y , which could be useful in the present context in combination with maximal couplings; this is left as an avenue of future research. Note that in multivariate settings, optimal transport of Normal distributions can be implemented following e.g. Knott and Smith [1984]; however, sampling from the optimal transport between arbitrary distributions is a challenging task.

4.2 Metropolis–Hastings

A couplings of MH chains due to Johnson [1998] was described in Section 2.2; the coupled kernel $\bar{P}((X_t, Y_{t-1}), \cdot)$ is summarized in the following procedure.

1. Sample $(X^*, Y^*)|(X_t, Y_{t-1})$ from a maximal coupling of $q(X_t, \cdot)$ and $q(Y_{t-1}, \cdot)$.
2. Sample $U \sim \mathcal{U}([0, 1])$.
3. If $U \leq \min(1, \pi(X^*)q(X^*, X_t)/\pi(X_t)q(X_t, X^*))$, then $X_{t+1} = X^*$, otherwise $X_{t+1} = X_t$.
4. If $U \leq \min(1, \pi(Y^*)q(Y^*, Y_{t-1})/\pi(Y_{t-1})q(Y_{t-1}, Y^*))$, then $Y_t = Y^*$, otherwise $Y_t = Y_{t-1}$.

Here we address the verification of Assumptions 2.1-2.3. Assumption 2.1 can be verified for MH chains under conditions on the target and the proposal [Nummelin, 2002, Roberts and Rosenthal, 2004]. In some settings the explicit drift function given in Theorem 3.2 of Roberts and Tweedie [1996b] may be used to verify Assumption 2.2 as in Section 3.2. In certain settings, the probability of coupling at the next step given that the chains are in X_t and Y_{t-1} can be controlled as follows. First, the probability of proposing the same value X^* depends on the total variation distance between $q(X_t, \cdot)$ and $q(Y_{t-1}, \cdot)$, which is typically strictly positive if X_t and Y_{t-1} are in bounded subsets of \mathcal{X} . Furthermore, the probability of accepting X^* is often lower-bounded away from zero on bounded subsets of \mathcal{X} , for instance when $\pi(x) > 0$ for all $x \in \mathcal{X}$.

In high dimension, the probability of proposing the same value X^* is low unless X_t is close to Y_{t-1} . It might therefore be preferable to use a series of updates on low-dimensional components of the chain states as in a Metropolis-within-Gibbs strategy (see Section 4.4), or to combine maximal couplings with optimal transport couplings mentioned in the previous section. Scalability with respect to dimension is investigated in Section 4.5.

The optimal choice of proposal distribution for a single MH chain might not be optimal in the proposed coupling construction. For instance, in the case of Normal random walk proposals with variance Σ , larger variances lead to smaller total variation distances between $q(X_t, \cdot)$ and $q(Y_{t-1}, \cdot)$ and thus larger coupling probabilities for the proposals. However meeting events only occur if proposals are accepted, which is unlikely if Σ is too large. This trade-off could lead to optimal choices of Σ that are different than the optimal choices known for the marginal chains [Roberts et al., 1997], which deserves further investigation.

Among extensions of the Metropolis–Hastings algorithm, Metropolis adjusted Langevin algorithms [e.g. Roberts and Tweedie, 1996a] are such that the proposal distribution given X_t is a Normal with

mean $X_t + h\nabla \log \pi(X_t)/2$ and variance $h\Sigma$, for some tuning parameter $h > 0$ and covariance matrix Σ . These Normal proposal distributions could be maximally coupled as well. Another important extension consists in adapting the proposal distribution during the run of the chains [Andrieu and Thoms, 2008, Atchadé et al., 2011]; it is unclear whether such strategies could be used in the proposed framework.

4.3 Hamiltonian Monte Carlo

Hamiltonian or Hybrid Monte Carlo [HMC, Duane et al., 1987, Neal, 1993, 2011, Durmus et al., 2017, Betancourt et al., 2017] is a popular MCMC algorithm using gradients of the target density, in which each iteration t is defined as follows. The state X_t is treated as the initial position $q(0)$ of a particle under a potential energy function given by $-\log \pi$. The initial momentum $p(0)$ of the particle is drawn at random, typically from a Normal distribution [see Livingstone et al., 2017]. One can numerically approximate the solution of Hamiltonian dynamics defined by:

$$\frac{d}{dt}q(t) = p(t), \quad \frac{d}{dt}p(t) = \nabla \log \pi(q(t)),$$

over a time interval $[0, T]$, where T denotes the trajectory length. For instance, one might use a leap-frog integrator [Hairer et al., 2005] with L steps and a step-size of ϵ so that $T = \epsilon L$. Finally, a Metropolis–Hastings step sets X_{t+1} either to $q(T)$ or X_t .

The use of common random numbers for the initial velocities and the uniform variables of the acceptance steps leads to pairs of chains converging to one another, under conditions on the target distribution such as strict log-concavity. This is used in Mangoubi and Smith [2017] to quantify the mixing properties of HMC. In Heng and Jacob [2017], such coupled HMC steps are combined with coupled random walk MH steps that produce exact meeting times, and the verification of Assumptions 2.1-2.3 is discussed.

4.4 Gibbs sampling

Gibbs sampling consists in updating components of a Markov chain by alternately sampling from conditional distributions of the target [Chapter 10 of Robert and Casella, 1999]. In Bayesian statistics, these conditional distributions sometimes belong to a standard family such as Normal, Gamma, or Inverse Gamma. If all conditional distributions are standard, then the Markov kernel of the Gibbs sampler is itself tractable, and a maximal coupling can be implemented following Section 4.1. However, in many cases at least one of the conditional updates is intractable and requires a Metropolis step. We therefore focus on maximal couplings of each conditional update, using either full conditional distributions or Metropolis updates. Controlling the probability of meeting at the next step over a set, as required for the application of Proposition 3.4, can be done on a case-by-case basis. Drifts conditions for Gibbs samplers also tend to rely on case-by-case arguments [see e.g. Rosenthal, 1996].

In generic state space models, the conditional particle filter is an MCMC algorithm targeting the distribution of the latent process given the observations. It is a Gibbs sampler on an extended state space [Andrieu et al., 2010]. Couplings of such Gibbs samplers are the focus of Jacob et al. [2017a], where a combination of common random numbers and maximal couplings leads to pairs of chains that satisfy Assumptions 2.1-2.3.

4.5 Scaling with the dimension

We compare the scaling behavior of MH, Gibbs, and HMC couplings as the dimension d of the target distribution increases.




Figure 1: Scaling of $\mathbb{E}[\tau]$ with the dimension of the target $\mathcal{N}(0, V)$, where $V_{i,j} = 0.5^{|i-j|}$, as described in Section 4.5. For the Normal random walk Metropolis–Hastings algorithm (1a) “scaling 1” corresponds to a proposal covariance matrix $\Sigma = V/d$, and “scaling 2” corresponds to $\Sigma = V$. For MH-within-Gibbs (1b), the different lines correspond to different number of MH steps performed to update each component. For HMC (1c), the lines correspond to different trajectory lengths.

Consider a d -dimensional Normal target distribution $\mathcal{N}(0, V)$, where V is a $d \times d$ covariance matrix with entry (i, j) equal to $0.5^{|i-j|}$. In an MH algorithm on the joint target, we consider Gaussian random walks, where the covariance matrix Σ of the proposals is set to V/d (“scaling 1” in Figure 1). The division by d follows the recommendations of Roberts et al. [1997]. We consider another strategy where Σ is set to V (“scaling 2”). Our second algorithm is an MH-within-Gibbs approach where each univariate component is updated with a Metropolis step using Normal proposals with unit variance. We consider performing 1, 2 or 5 of such steps for each component under a systematic scan of the components. Each iteration refers to a complete scan of all components. Finally, we consider a mixture of MH and HMC kernels. At each iteration we perform an HMC step with 90% probability, and otherwise we perform an MH step with a Normal proposal distribution with variance 10^{-5} times the identity matrix. The HMC kernel uses a leap-frog integrator with 20 sub-steps, and we try different trajectory lengths. For all strategies, we initialize the chains from the target distribution.

For a range of values of the dimension d , we run the coupled chains until they meet. We present the average meeting time over $R = 100$ independent repetitions in Figure 1 to visualize the relationship between $\mathbb{E}[\tau]$ and dimension for various MCMC algorithms. Figure 1a illustrates that the coupling of MH on the joint space fails quickly as the dimension d increases. Note the logarithmic scale on the y -axis. We obtain worse performance with $\Sigma = V/d$ than with $\Sigma = V$, even though the marginal chains mix more quickly with $\Sigma = V/d$. The MH-within-Gibbs approach scales more favorably with dimension, as indicated in Figure 1b. Performing multiple MH steps per component further decreases the average meeting time. Finally, we present results for HMC in Figure 1c for three different trajectory lengths. In this setting, the scaling of HMC with respect to the dimension is qualitatively similar to that of the Gibbs sampler and is sensitive to the choice of trajectory length.

These experiments suggest that the proposed methodology can be implemented in realistic dimensions. In particular, strategies that leverage the dependence structure of the target or gradient information can result in short meeting times even in high dimension.

5 Illustrations

5.1 Bimodal target

We use a bimodal target distribution and a random walk Metropolis–Hastings algorithm to illustrate various aspects of the proposed method and highlight challenging situations.

We consider a mixture of univariate Normals with density $\pi(x) = 0.5 \cdot \mathcal{N}(x; -4, 1) + 0.5 \cdot \mathcal{N}(x; +4, 1)$, which we sample from using random walk Metropolis–Hastings with Normal proposal distributions of variance $\sigma_q^2 = 9$. This enables regular jumps between the modes of π . We set the initial distribution π_0 to $\mathcal{N}(10, 10^2)$, so that chains are more likely to start near the mode at $+4$ than the mode at -4 . Over 1,000 independent runs, we find that the meeting time τ has an average of 20 and a 99% quantile of 105.

We consider the task of estimating $\int \mathbf{1}(x > 3)\pi(dx) \approx 0.421$. First, we consider the choice of k and m . Over 1,000 independent experiments, we approximate the expected cost $\mathbb{E}[2\tau + \max(1, m - \tau + 1)]$, the variance $\mathbb{V}[H_{k:m}(X, Y)]$, and compute the inefficiency as the product of the two. We then divide the inefficiency by the asymptotic variance of the MCMC estimator, denoted by V_∞ , which we obtain from 10^6 iterations and a burn-in period of 10^4 using the R package CODA [Plummer et al., 2006].

We present the results of this test in Table 1. First, we see that the inefficiency is sensitive to the choice of k and m : simply setting k and m to one would be highly inefficient. Secondly, we see that when k and m are large enough we can retrieve an inefficiency comparable to that of the underlying MCMC algorithm. A relative inefficiency close to 1 indicates that the estimator variance is similar to that of MCMC. The ideal choice of k and m will depend on tradeoffs between inefficiency, the desired level of parallelism, and the number of processors available. Whether it is preferable to run coupled chains with $k = 200$, $m = 2,000$ for an inefficiency of 1.3 or $k = 200$, $m = 4,000$ for an inefficiency of 1.2 is likely to depend on the context. We present a histogram of the target distribution, obtained using $k = 200$, $m = 4,000$, in Figure 2a.

k	m	Cost	Variance	Inefficiency / V_∞
1	$1 \times k$	37	4.1e+02	1867.4
1	$10 \times k$	39	3.6e+02	1693.5
1	$20 \times k$	45	3.0e+02	1615.3
100	$1 \times k$	119	9.0e+00	129.8
100	$10 \times k$	1019	2.3e-02	2.9
100	$20 \times k$	2019	7.9e-03	1.9
200	$1 \times k$	219	2.4e-01	6.4
200	$10 \times k$	2019	5.3e-03	1.3
200	$20 \times k$	4019	2.4e-03	1.2

Table 1: Cost, variance, and inefficiency divided by MCMC asymptotic variance V_∞ for various choices of k and m , for the test function $h : x \mapsto \mathbf{1}(x > 3)$ in the mixture target example of Section 5.1.

Next, we consider a more challenging case by setting $\sigma_q^2 = 1^2$, and we use again $\pi_0 = \mathcal{N}(10, 10^2)$. These values make it difficult for the chains to jump between the modes of π . Over $R = 1,000$ runs we find an average meeting time of 769, with a 99% quantile of 9,186. When the chains start in different modes, the meeting times are often dramatically larger than when the chains start by the same mode. One can still recover reasonable estimates of the target distribution, but k and m have to be set to larger values. With $k = 20,000$ and $m = 30,000$, we obtain the 95% confidence interval $[0.397, 0.430]$ for $\int \mathbf{1}(x > 3)\pi(dx) \approx 0.421$. We show a histogram of π in Figure 2b.

Finally, we consider a third case in which σ_q^2 is again set to one, but π_0 is set to $\mathcal{N}(10, 1)$. This initialization makes it unlikely for a chain to start near the mode at -4 . The pair of chains typically




Figure 2: Histograms of the mixture target distribution of Section 5.1, obtained with the proposed unbiased estimators, based on a Normal random walk Metropolis–Hastings algorithm, with a proposal variance σ_q^2 and an initial distribution π_0 , over $R = 1,000$ experiments.

converge to the right-most mode and meet in a small number of iterations. Over $R = 1,000$ replications, we find an average meeting time of 9 and a 99% quantile of 35. A 95% confidence interval on $\int \mathbf{1}(x > 3)\pi(dx)$ obtained from the estimators with $k = 50$, $m = 500$ is $[0.799, 0.816]$, far from the true value of 0.421. The associated histogram of π is shown in Figure 2c.

Sampling 9,000 additional estimators yields a 95% confidence interval $[-0.353, 1.595]$, again using $k = 50$, $m = 500$. Among these extra 9,000 values, a few correspond to cases where one chain jumped to the left-most mode before meeting the other. This resulted in large meeting times, and thus in a much larger empirical variance. Upon noticing a large empirical variance one can then decide to use larger values of k and m ; the challenging situation is when the empirical variance is small even though the number of replicates is seemingly large. We conclude that although our estimators are unbiased and are consistent in the limit as $R \rightarrow \infty$, poor performance of the underlying Markov chains can still produce misleading results for any finite R .

5.2 Gibbs sampler for nuclear pump failure data

Next we consider a classic Gibbs sampler for a model of pump failure counts, used e.g. in [Murdoch and Green \[1998\]](#) to illustrate the implementation of perfect samplers for continuous distributions. We refer to the latter article for the case-specific calculations associated with the implementation of perfect samplers. Here we compare the proposed method with the regeneration approach of [Mykland et al. \[1995\]](#), which was illustrated on the same example and which was motivated by the same practical concerns: choosing the number of iterations to discard as burn-in, constructing confidence intervals, and using parallel processors.

The data consist of operating times $(t_n)_{n=1}^K$ and failure counts $(s_n)_{n=1}^K$ for $K = 10$ pumps at the Farley-1 nuclear power station, as first described in [Gaver and O’Muircheartaigh \[1987\]](#). The model specifies $s_n \sim \text{Poisson}(\lambda_n t_n)$ and $\lambda_n \sim \text{Gamma}(\alpha, \beta)$, where $\alpha = 1.802$, $\beta \sim \text{Gamma}(\gamma, \delta)$, $\gamma = 0.01$, and $\delta = 1$. The Gibbs sampler for this model consists of the following update steps:

$$\begin{aligned} \lambda_n \mid \text{rest} &\sim \text{Gamma}(\alpha + s_n, \beta + t_n) \quad \text{for } n = 1, \dots, K, \\ \beta \mid \text{rest} &\sim \text{Gamma}(\gamma + 10\alpha, \delta + \sum_{n=1}^K \lambda_n). \end{aligned}$$

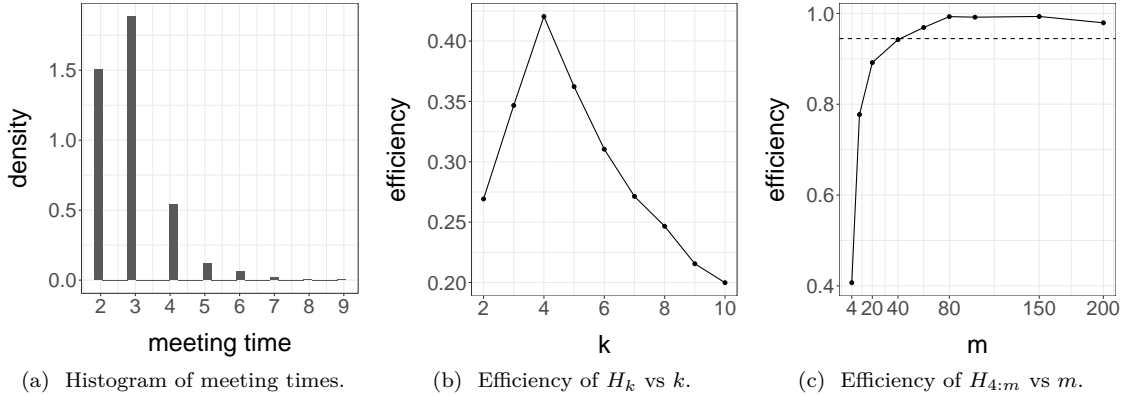


Figure 3: Gibbs sampling in the pump failure example of Section 5.2. Histogram of the meeting times in 3a. Efficiency of $H_k(X, Y)$ as a function of k in 3b and of $H_{k:m}(X, Y)$ as a function of m for $k = 4$ in 3c. The test function is $h : (\lambda_1, \dots, \lambda_K, \beta) \mapsto \beta$.

Here the Gamma (α, β) distribution refers to the distribution with density $x \mapsto \Gamma(\alpha)^{-1} \beta^\alpha x^{\alpha-1} \exp(-\beta x)$. We initialize all parameter values to 1. To form our estimator we apply maximal couplings at each conditional update of the Gibbs sampler, as described in Section 4.4.

We begin by drawing 1,000 meeting times to obtain the histogram in Figure 3a. Following the guidelines of Section 3.1, we set $k = 7$, corresponding to the 99% quantile of τ and $m = 10k = 70$. We then generate 10,000 independent estimates for the test function $h(\lambda_1, \dots, \lambda_K, \beta) = \beta$. Figure 3b shows the efficiency of $H_k(X, Y)$, defined as $(\mathbb{E}[\max(k, \tau)] \cdot \mathbb{V}[H_k(X, Y)])^{-1}$, for a range of k values. The choice of $k = 7$ appears somewhat conservative relative to the efficiency-maximizing value of $k = 4$. Figure 3c shows the efficiency of $H_{4:m}(X, Y)$ as a function of m . The horizontal dashed line represents the inefficiency associated with $k = 7$ and $m = 70$, and illustrates that the efficiency obtained by following the heuristics is close to the maximum that we observe.

It is natural to compare our estimator with the regenerative approach of Mykland et al. [1995], which also provides a way of parallelizing MCMC and of constructing confidence intervals. In that paper the authors show how to use detailed knowledge of a Markov chain to construct regeneration times – random times between which the chain forms independent and identically distributed “tours”. The authors define a consistent estimator for arbitrary test functions, whose asymptotic variance takes a particularly simple form. The estimator is obtained by aggregating over these independent tours. The authors give a set of preferred tuning parameters, which we adopt for our test below.

Applying the regeneration approach to 1,000 Gibbs sampler runs of 5,000 iterations each, we observe on average 1,996 complete tours with an average length of 2.50 iterations. These values agree with the count of 1,967 tours of average length 2.56 reported in Mykland et al. [1995]. We also observe a posterior mean estimate for β of 2.47 with a variance of 1.89×10^{-4} over the 1,000 independent runs, which implies an efficiency value of $(5,000 \cdot 1.89 \times 10^{-4})^{-1} = 1.06$. This exceeds the efficiency of 0.94 achieved by our estimator with our heuristic choice of $k = 7$ and $m = 70$. On the other hand, the regeneration approach requires more extensive analytical work with the underlying Markov chain; we refer to Mykland et al. [1995] for a detailed description. For reference, the underlying Gibbs sampler achieves an efficiency of 1.08, based on a long run with 5×10^5 iterations and a burn-in of 10^3 .

5.3 Variable selection

We consider a variable selection setting following [Yang et al. \[2016\]](#) to illustrate the proposed method on a high-dimensional discrete state space.

For integers p and n (potentially with $p > n$) let $Y \in \mathbb{R}^n$ represent a response variable depending on covariates $X_1, \dots, X_p \in \mathbb{R}^n$. We consider the task of inferring a binary vector $\gamma \in \{0, 1\}^p$ representing which covariates to select as predictors of Y , with the convention that X_i is selected if $\gamma_i = 1$. For any γ , we write $|\gamma| = \sum_{i=1}^p \gamma_i$ for the number of selected covariates and X_γ for the $n \times |\gamma|$ matrix of covariates chosen by γ . Inference on γ relies on a linear regression model relating Y to X_γ ,

$$Y = X_\gamma \beta_\gamma + w, \quad \text{where } w \sim \mathcal{N}(0, \sigma^2 I_n).$$


We assume a prior on γ of $\pi(\gamma) \propto p^{-\kappa|\gamma|} \mathbf{1}(|\gamma| \leq s_0)$. This distribution puts mass only on vectors γ with fewer than s_0 ones, imposing a degree of sparsity. Conditional on γ , we assume a Normal prior for the regression coefficient vector $\beta_\gamma \in \mathbb{R}^{|\gamma|}$ with zero mean and variance $g\sigma^2(X'_\gamma X_\gamma)^{-1}$. We give the precision σ^{-2} an improper prior $\pi(\sigma^{-2}) \propto 1/\sigma^{-2}$. This leads to the marginal likelihood

$$\pi(Y|X, \gamma) \propto \frac{(1+g)^{-|\gamma|/2}}{(1+g(1-R_\gamma^2))^{n/2}}, \quad \text{where } R_\gamma^2 = \frac{Y'X_\gamma(X'_\gamma X_\gamma)^{-1}X'_\gamma Y}{Y'Y}.$$

To approximate the distribution $\pi(\gamma|X, Y)$, [Yang et al. \[2016\]](#) rely on an MCMC algorithm whose kernel P is a mixture of two Metropolis kernels. The first component $P_1(\gamma, \cdot)$ selects a coordinate $i \in \{1, \dots, p\}$ uniformly at random and flips γ_i to $1 - \gamma_i$. The resulting vector γ^* is then accepted with probability $1 \wedge \pi(\gamma^*|X, Y)/\pi(\gamma|X, Y)$. Sampling a vector γ' from the second kernel $P_2(\gamma, \cdot)$ proceeds as follows. If $|\gamma|$ equals zero or p , then γ' is set to γ . Otherwise, coordinates i_0, i_1 are drawn uniformly among $\{j : \gamma_j = 0\}$ and $\{j : \gamma_j = 1\}$, respectively. The proposal γ^* is such that $\gamma_{i_0}^* = \gamma_{i_1}$ and $\gamma_{i_1}^* = \gamma_{i_0}$, and $\gamma_j^* = \gamma_j$ for the other components. Then γ' is set to γ^* with probability $1 \wedge \pi(\gamma^*|X, Y)/\pi(\gamma|X, Y)$, and to γ otherwise. Finally the MCMC kernel $P(\gamma, \cdot)$ targets $\pi(\gamma|X, Y)$ by sampling from $P_1(\gamma, \cdot)$ or from $P_2(\gamma, \cdot)$ with equal probability. Note that each MCMC iteration can only benefit from parallel processors to a limited extent, since $|\gamma|$ is always less than s_0 , itself chosen to be a small value in most settings.

To sample a pair of states $(\gamma', \tilde{\gamma}')$ given $(\gamma, \tilde{\gamma})$, we consider the following coupled version of the MCMC algorithm described above. First, we use a common uniform random variable to decide whether to sample from a coupling of P_1 to itself, \bar{P}_1 , or a coupling of P_2 to itself, \bar{P}_2 . The coupled kernel $\bar{P}_1((\gamma, \tilde{\gamma}), \cdot)$ proposes flipping the same coordinate for both vectors γ and $\tilde{\gamma}$ and then uses a common uniform random variable in the acceptance step. For the coupled kernel $\bar{P}_2((\gamma, \tilde{\gamma}), \cdot)$, we need to select two pairs of indices, (i_0, \tilde{i}_0) and (i_1, \tilde{i}_1) . We obtain the first pair by sampling from a maximal coupling of the discrete uniform distributions on $\{j : \gamma_j = 0\}$ and $\{j : \tilde{\gamma}_j = 0\}$. This yields indices (i_0, \tilde{i}_0) with the greatest possible probability that $i_0 = \tilde{i}_0$. We use the same approach to sample a pair (i_1, \tilde{i}_1) to maximize the probability that $i_1 = \tilde{i}_1$. Finally we use a common uniform variable to accept or reject the proposals. If either vector γ or $\tilde{\gamma}$ has no zeros or no ones, then it is kept unchanged.

We recall that one can sample from a maximal coupling of two discrete probability distributions $q = (q_1, \dots, q_N)$ and $\tilde{q} = (\tilde{q}_1, \dots, \tilde{q}_N)$ as follows. First, let $c = (c_1, \dots, c_N)$ be the distribution with probabilities $c_n = (q_n \wedge \tilde{q}_n)/\alpha$ for $\alpha = \sum_{n=1}^N q_n \wedge \tilde{q}_n$ and define residual distributions q' and \tilde{q}' with probabilities $q'_n = (q_n - \alpha c_n)/(1 - \alpha)$ and $\tilde{q}'_n = (\tilde{q}_n - \alpha c_n)/(1 - \alpha)$. Then with probability α , draw $i \sim c$ and output (i, i) . Otherwise draw $i \sim q'$ and $\tilde{i} \sim \tilde{q}'$ and output (i, \tilde{i}) . The resulting pair follows a maximal coupling of q and \tilde{q} , and the procedure involves $\mathcal{O}(N)$ operations for N the size of the state



(a) Meeting times divided by p , and deciles.

(b) Inclusion probabilities of the first 20 variables, in the case $p = 1,000$.

Figure 4: Left: meeting times divided p for $p \in \{100, 250, 500, 750, 1000\}$ and $n = 500$, in the variable selection setting of Section 5.3, over $R = 1,000$ experiments. The dashed lines indicate the 9 empirical deciles of the meeting times divided by p . Right: posterior probabilities of inclusion for the first 20 variables, in a setting with $p = 1,000$. The dotted line represents an MCMC estimate taken as ground truth, while the error bars represent 95% confidence intervals obtained with $R = 1,000$ unbiased estimators.

space.

We now consider an experiment similar to those of Yang et al. [2016]. We generate covariates X_{ij} as independent standard Normal variables, define

$$\beta^* = \text{SNR} \sqrt{\frac{\sigma_0^2 \log(p)}{n}} (2, -3, 2, 2, -3, 3, -2, 3, -2, 3, 0, \dots, 0)' \in \mathbb{R}^p,$$

with $\sigma_0^2 = 1$ and $\text{SNR} = 1$, and generate Y given X and β^* from the model, with $n = 500$ and $\sigma^2 = 1$. We set $s_0 = 100$, $g = p^3$, and $\kappa = 0.1$. A sample from the initial distribution π_0 is drawn by defining a vector of p zeros, sampling s_0 coordinates uniformly at random in $\{1, \dots, p\}$, without replacement, and setting the corresponding entries to 1 with probability 0.5.

For values of p between 100 and 1000, we run coupled chains until they meet, $R = 1,000$ times independently. The meeting times divided by p are shown in Figure 4a, along with dashed lines connecting deciles. From the graph we see that these scaled meeting times are approximately constant, suggesting that meeting times increase linearly with p . This is in line with the findings of Yang et al. [2016], where mixing times are shown to increase linearly with p .

In the case $p = 1,000$, we choose $k = 50,000$ and $m = 100,000$ and generate $R = 1,000$ independent unbiased estimators of the inclusion probabilities $\mathbb{P}(\gamma_i = 1 | X, Y)$ for all $i \in \{1, \dots, p\}$. We compare the estimates with an MCMC estimate obtained with 10^6 iterations and a burn-in of 50,000. The MCMC estimates for the first 20 variables are represented by a dotted line in Figure 4b, and 95% confidence intervals obtained from the unbiased estimators are shown as error bars. The probabilities of inclusion of the other variables are all close to zero. We observe an agreement between the proposed estimators and the MCMC estimates, considered as ground truth.

Assessing the convergence of MCMC can be difficult in high-dimensional discrete state spaces. We might not know whether the chain sticks to a small number of states because they are more probable than all other states or because the chain is stuck in a local mode. One can plot the target probabilities along the iterations of the chain and repeat over multiple chains to check that they all converge to the same range of values [see Figure 1 in Yang et al., 2016]. These considerations assume a different significance

in the proposed framework, where confidence intervals are constructed using the Central Limit Theorem for independent variables. We note that the results of this section are sensitive to the prior parameter κ and on the degree of correlations between the columns of X ; we refer to [Yang et al. \[2016\]](#) for a further discussion, and to [Johnson \[2004\]](#) for a similar Bayesian variable selection setting where couplings of chains are used to assess convergence.

5.4 Cut distribution

Our proposed estimator can be used to approximate the cut distribution, which poses a significant challenge for existing MCMC methods [[Plummer, 2014](#), [Jacob et al., 2017b](#)].

Consider two models, one with parameters θ_1 and data Y_1 and another with parameters θ_2 and data Y_2 , where we might allow the likelihood Y_2 to depend on both θ_1 and θ_2 . For instance the first model could be a regression with data Y_1 and coefficients θ_1 , and the second model could be another regression whose covariates are the residuals, coefficients, or fitted values of the first regression [[Pagan, 1984](#), [Murphy and Topel, 2002](#)]. In principle one could introduce an encompassing model and conduct joint inference on θ_1 and θ_2 via the posterior distribution. However, misspecification of either model would then lead to misspecification of the ensemble and misleading quantification of uncertainties, as noted in several studies [e.g. [Liu et al., 2009](#), [Plummer, 2014](#), [Lunn et al., 2009](#), [McCandless et al., 2010](#), [Zigler, 2016](#), [Blangiardo et al., 2011](#)].

The cut distribution [[Spiegelhalter et al., 2003](#), [Plummer, 2014](#)] allows the propagation of the uncertainty about θ_1 to inference on θ_2 while preventing misspecification in the second model from affecting estimation in the first. The cut distribution is defined as

$$\pi_{\text{cut}}(\theta_1, \theta_2) = \pi_1(\theta_1)\pi_2(\theta_2|\theta_1),$$

where $\pi_1(\theta_1)$ refers to the distribution of θ_1 given Y_1 in the first model alone, and $\pi_2(\theta_2|\theta_1)$ refers to the distribution of θ_2 given Y_2 and θ_1 in the second model. Often, the density $\pi_2(\theta_2|\theta_1)$ can only be evaluated up to a constant in θ_2 , which may vary as a function of θ_1 . This makes the cut distribution intractable for standard MCMC algorithms [[Plummer, 2014](#)]. However the cut distribution has been used in practice, as in the references above, and can be justified with decision-theoretic arguments [[Jacob et al., 2017b](#)].

A naive approach consists of first running an MCMC algorithm targeting $\pi_1(\theta_1)$, obtaining a sample $(\theta_1^n)_{n=1}^{N_1}$, perhaps after discarding a burn-in period and thinning the chain. Then, for each θ_1^n , one can run an MCMC algorithm targeting $\pi_2(\theta_2|\theta_1^n)$, yielding N_2 samples. One might again discard some burn-in and thin the chains, or just keep the final state of each chain. The resulting joint samples approximate the cut distribution. However, the validity of this approach relies on a double limit in N_1 and N_2 . Diagnosing convergence may also be difficult given the number of chains in the second stage, each of which targets a different distribution $\pi_2(\theta_2|\theta_1^n)$.

If one could sample $\theta_1 \sim \pi_1$ and $\theta_2|\theta_1 \sim \pi_2(\cdot|\theta_1)$, then the pair (θ_1, θ_2) would follow the cut distribution. The same two-stage rationale can be applied in the proposed framework. Consider a test function $(\theta_1, \theta_2) \mapsto h(\theta_1, \theta_2)$. Writing \mathbb{E}_{cut} for expectations with respect to π_{cut} , the law of iterated expectations yields

$$\mathbb{E}_{\text{cut}}[h(\theta_1, \theta_2)] = \int \left(\int h(\theta_1, \theta_2)\pi_2(d\theta_2|\theta_1) \right) \pi_1(d\theta_1) = \int \bar{h}(\theta_1)\pi_1(d\theta_1).$$

Here $\bar{h}(\theta_1) = \int h(\theta_1, \theta_2)\pi_2(d\theta_2|\theta_1)$. In the proposed framework, one can make an unbiased estimator of $\bar{h}(\theta_1)$ for all θ_1 , and plug the estimators into an unbiased estimator of $\int h(\theta_1)\pi_1(d\theta_1)$. This is even

clearer using the signed measure representation of Section 2.5: one can obtain a signed measure $\hat{\pi}_1(\cdot) = \sum_{\ell=1}^N \omega_\ell \delta_{\theta_{1,\ell}}(\cdot)$ approximating π_1 , and then obtain an unbiased estimator of $\bar{h}(\theta_{1,\ell})$ for all ℓ , denoted by \bar{H}_ℓ . Then the weighted average $\sum_{\ell=1}^N \omega_\ell \bar{H}_\ell$ is an unbiased estimator of $\mathbb{E}_{\text{cut}}[h(\theta_1, \theta_2)]$ by the law of iterated expectations. Such estimators can be generated independently in parallel, and their average provides a consistent approximation of an expectation with respect to the cut distribution.

We consider the example described in [Plummer \[2014\]](#), inspired by an investigation of the international correlation between human papillomavirus (HPV) prevalence and cervical cancer incidence [Maucort-Boulch et al. \[2008\]](#). The first module concerns HPV prevalence, with data independently collected in 13 countries. The parameter $\theta_1 = (\theta_{1,1}, \dots, \theta_{1,13})$ receives a Beta(1, 1) prior distribution independently for each component. The data (Y_1, \dots, Y_{13}) consist of 13 pairs of integers. The first represents the number of women infected with high-risk HPV, and the second represents population sizes. The likelihood specifies a Binomial model for Y_i , independently for each component i . The first posterior is given by a product of Beta distributions.

The second module concerns the relationship between HPV prevalence and cancer incidence, and posits a Poisson regression. The parameters are $\theta_2 = (\theta_{2,1}, \theta_{2,2})$ and receive a Normal prior with zero mean and variance 10^3 . The likelihood in this module is given by

$$Z_{1,i} \sim \text{Poisson}(\exp(\theta_{2,1} + \theta_{1,i}\theta_{2,2} + Z_{2,i})) \quad \text{for } i \in \{1, \dots, 13\},$$

where the data $(Z_{1,i}, Z_{2,i})_{i=1}^{13}$ are pairs of integers. The first component represents numbers of cancer cases, while the second is a number of woman-years of follow-up. The Poisson regression model might be misspecified, which motivates departures from the joint model [\[Plummer, 2014\]](#).

In our setting, we can draw directly from the first posterior, denoted by $\pi_1(\theta_1)$. We can thus obtain a sample $(\theta_1^n)_{n=1}^N$ approximating $\pi_1(\theta_1)$. For each θ_1^n , we consider a Metropolis–Hastings algorithm targeting $\pi_2(\theta_2|\theta_1^n)$, using a Normal random walk proposal with variance Σ . We set the initial distribution π_0 to be standard Normal, and the variance Σ to be unit diagonal matrix. From 1,000 coupled chains, we observe a 95% quantile of the meeting times at 381. We set $k = 381$ and $m = 5k$, and estimate the posterior mean and variance using $R = 1,000$ independent estimators. We use these estimated moments to propose a new Normal distribution for π_0 , and we set Σ to the estimated posterior variance. We run $R = 10,000$ coupled chains with these new settings, and observe the meeting times shown in the histogram of [Figure 5a](#). The average meeting time is 15, with a maximum of 110. Applying the 95% quantile heuristic, we set $k = 49$ and $m = 10k$. The two histograms of [Figures 5b](#) and [5c](#) show that the marginal distributions are approximated accurately. The overlaid red curves indicate the “ground truth”, obtained by running 1,000 steps of MCMC targeting $\pi_2(\theta_2|\theta_1)$ for each of 10,000 independent draws of θ_1 from π_1 and keeping the final state of each chain.

6 Discussion

With the telescopic sum argument of [Glynn and Rhee \[2014\]](#) and couplings of MCMC algorithms, we have introduced unbiased estimators of integrals with respect to the target distribution, with a variance that can be controlled with tuning parameters (k and m). In practice, we propose to choose k as a large quantile of the meeting times. Then m can be set as a large multiple of k . Improving on these simple guidelines stands as a subject for future research. In numerical experiments we have argued that the proposed estimators yield a practical way of parallelizing MCMC computations, applicable to a wide range of settings; the supplementary materials contain further experiments. We stress that




Figure 5: Meeting times and histograms of $\theta_{2,1}$ and $\theta_{2,2}$ in the epidemiological example of Section 5.4, produced with $R = 10,000$ estimators, $k = 49$ and $m = 490$.

coupling pairs of Markov chains does not improve their marginal mixing properties, and that misleading confidence intervals are obtained for any fixed computational budget if the initial distribution and the MCMC kernels are poorly designed, as in Section 5.1. Misleading results would also be obtained if an impatient user stops the coupled chains before they meet.

We have considered different couplings of MCMC algorithms, based on maximal couplings and common random numbers. These allow us to compute unbiased estimators without requiring further analytical knowledge about the target distribution or about the MCMC kernels. More sophisticated couplings are also possible, and include optimal transport couplings to bring two chains together quickly and, subsequently, facilitate exact meeting using a maximal coupling.

Regarding convergence diagnostics, the proposed framework yields the following representation for the total variation between π_k and π , where π_k denotes the marginal distribution of X_k :

$$\begin{aligned} d_{\text{TV}}(\pi_k, \pi) &= \frac{1}{2} \sup_{h: |h| \leq 1} |\mathbb{E}[h(X_k)] - \mathbb{E}_\pi[h(X)]| \\ &= \frac{1}{2} \sup_{h: |h| \leq 1} \left| \mathbb{E} \left[\sum_{t=k+1}^{\tau-1} (h(X_t) - h(Y_{t-1})) \right] \right|, \end{aligned}$$

Here the supremum ranges over all bounded measurable functions under the stated assumptions. The equality above has several implications. For instance, by the triangle inequality we obtain $d_{\text{TV}}(\pi_k, \pi) \leq \min(1, \mathbb{E}[\max(0, (\tau - k + 1))])$, and we can approximate $\mathbb{E}[\max(0, (\tau - k + 1))]$ by Monte Carlo for a range of k values. On the other hand, considering any particular function h yields a lower bound on $d_{\text{TV}}(\pi_k, \pi)$.

Thanks to the potential for complete parallelization, the proposed framework can facilitate consideration of MCMC kernels that might be too expensive for serial implementation. For instance, one can improve MH-within-Gibbs samplers by performing more MH steps per component update (as in Section 4.5), HMC by using smaller step-sizes in the integrator [Heng and Jacob, 2017], and particle MCMC by using more particles at each step [Andrieu et al., 2010, Jacob et al., 2017a]. We expect the optimal tuning of MCMC kernels to differ in the proposed framework than when used marginally.

There are also settings in which the unbiased property itself might be appealing. For instance, if MCMC is used to approximate a gradient in a stochastic gradient algorithm, or an expectation in an Expectation-Maximization algorithm, the unbiasedness property could have benefits for the optimizer [see e.g. Doucet and Tadic, 2017]. Moreover, in the context of Bayesian inference with misspecified

models made of multiple components [Liu et al., 2009, Jacob et al., 2017b], we have demonstrated that the unbiased property leads to convenient approximations of the cut distribution [Plummer, 2014].

The design of generic and efficient MCMC kernels is a topic of active ongoing research [e.g. Andrieu et al., 2009, Murray et al., 2010, Goodman et al., 2010, Pollock et al., 2016, Vanetti et al., 2017, Titsias and Yau, 2017, and references therein]. These kernels can be used in the proposed framework as long as appropriate couplings can be found. The particular case of conditional particle filter is explored in Jacob et al. [2017a] and the case of Hamiltonian Monte Carlo in Heng and Jacob [2017]; both require specific couplings and associated arguments to establish the validity of the resulting estimators.

Acknowledgement.

The authors are grateful to Jeremy Heng and Luc Vincent-Genod for useful discussions. The authors gratefully acknowledge support by the National Science Foundation through grants DMS-1712872 (Pierre E. Jacob), and DMS-1513040 (Yves F. Atchadé).

References

- G. Altekar, S. Dwarkadas, J. P. Huelsenbeck, and F. Ronquist. Parallel Metropolis coupled Markov chain Monte Carlo for Bayesian phylogenetic inference. *Bioinformatics*, 20(3):407–415, 2004. 1
- C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373, 2008. 12
- C. Andrieu, G. O. Roberts, et al. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009. 22
- C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010. 2, 12, 21
- Y. F. Atchade. An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift. *Methodology and Computing in applied Probability*, 8:235–254, 2006. 8
- Y. F. Atchadé. Markov chain Monte Carlo confidence intervals. *Bernoulli*, 22(3):1808–1838, 2016. 2
- Y. F. Atchadé, G. Fort, E. Moulines, and P. Priouret. Adaptive Markov chain Monte Carlo : theory and methods. In D. Barber, A. T. Cemgil, and S. Chiappa, editors, *Bayesian Time Series Models*, pages 33–52. Cambridge Univ. Press, 2011. 12
- M. Betancourt, S. Byrne, S. Livingstone, and M. Girolami. The geometric foundations of Hamiltonian Monte Carlo. *Bernoulli*, 23(4A):2257–2298, 2017. 12
- M. Blangiardo, A. Hansell, and S. Richardson. A Bayesian model of time activity data to investigate health effect of air pollution in time series studies. *Atmospheric Environment*, 45(2):379–386, 2011. 19
- A. E. Brockwell. Parallel Markov chain Monte Carlo simulation by pre-fetching. *Journal of Computational and Graphical Statistics*, 15(1):246–261, 2006. 1
- A. E. Brockwell and J. B. Kadane. Identification of regeneration times in mcmc simulation, with application to adaptive schemes. *Journal of Computational and Graphical Statistics*, 14(2):436–458, 2005. 4

- S. P. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov chain Monte Carlo*. CRC press, 2011. 1
- B. Calderhead. A general construction for parallelizing Metropolis–Hastings algorithms. *Proceedings of the National Academy of Sciences*, 111(49):17408–17413, 2014. 1
- G. Casella, M. Lavine, and C. P. Robert. Explaining the perfect sampler. *The American Statistician*, 55(4):299–305, 2001. 2
- S. Chatterjee, A. Guntuboyina, B. Sen, et al. On risk bounds in isotonic and other shape restricted regression problems. *The Annals of Statistics*, 43(4):1774–1800, 2015. 7
- H. M. Choi and J. P. Hobert. The Pólya–Gamma Gibbs sampler for Bayesian logistic regression is uniformly ergodic. *Electronic Journal of Statistics*, 7:2054–2064, 2013. 8
- P. Diaconis and D. Freedman. Iterated random functions. *SIAM review*, 41(1):45–76, 1999. 4
- R. Douc, E. Moulines, and J. S. Rosenthal. Quantitative bounds on convergence of time-inhomogeneous markoc chains. *Annals of Applied Probability*, 14(4):1643–1665, 2004. 4, 9, 33
- A. Doucet and V. Tadic. Asymptotic bias of stochastic gradient search. 2017. 21
- S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987. 12
- A. Durmus, E. Moulines, and E. Saksman. On the convergence of Hamiltonian Monte Carlo. *arXiv preprint arXiv:1705.00166*, 2017. 12
- J. M. Flegal and R. Herbei. Exact sampling for intractable probability distributions via a Bernoulli factory. *Electronic Journal of Statistics*, 6:10–37, 2012. 2
- J. M. Flegal, M. Haran, and G. L. Jones. Markov chain Monte Carlo: can we trust the third significant figure? *Statistical Science*, pages 250–260, 2008. 2
- D. P. Gaver and I. G. O’Muircheartaigh. Robust empirical Bayes analyses of event rates. *Technometrics*, 29(1):1–15, 1987. 15
- P. W. Glynn. Exact simulation versus exact estimation. In *Winter Simulation Conference (WSC), 2016*, pages 193–205. IEEE, 2016. 2
- P. W. Glynn and P. Heidelberger. Bias properties of budget constraint simulations. *Operations Research*, 38(5):801–814, 1990. 6
- P. W. Glynn and C.-H. Rhee. Exact estimation for Markov chain equilibrium expectations. *Journal of Applied Probability*, 51(A):377–389, 2014. 2, 3, 4, 6, 7, 20, 28
- P. W. Glynn and W. Whitt. The asymptotic efficiency of simulation estimators. *Operations Research*, 40(3):505–520, 1992. 7
- L. Gong and J. M. Flegal. A practical sequential stopping rule for high-dimensional Markov chain Monte Carlo. *Journal of Computational and Graphical Statistics*, 25(3):684–700, 2016. doi: 10.1080/10618600.2015.1044092. URL <http://dx.doi.org/10.1080/10618600.2015.1044092>. 2

- J. Goodman, J. Weare, et al. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010. [22](#)
- R. J. Goudie, R. M. Turner, D. De Angelis, and A. Thomas. Massively parallel MCMC for Bayesian hierarchical models. *arXiv preprint arXiv:1704.03216*, 2017. [1](#)
- P. J. Green, K. Łatuszyński, M. Pereyra, and C. P. Robert. Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statistics and Computing*, 25(4):835–862, 2015. [1](#)
- E. Hairer, G. Wanner, and C. Lubich. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer-Verlag, New York, 2005. [12](#)
- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970. [4](#)
- J. Heng and P. E. Jacob. Unbiased Hamiltonian Monte Carlo with couplings. *arXiv preprint arXiv:1709.00404*, 2017. [12](#), [21](#), [22](#)
- M. Huber. *Perfect simulation*, volume 148. CRC Press, 2016. [2](#), [4](#)
- P. E. Jacob and A. H. Thiery. On non-negative unbiased estimators. *The Annals of Statistics*, 43(2):769–784, 2015. [6](#)
- P. E. Jacob, C. P. Robert, and M. H. Smith. Using parallel computation to improve independent Metropolis–Hastings based estimation. *Journal of Computational and Graphical Statistics*, 20(3):616–635, 2011. [1](#)
- P. E. Jacob, F. Lindsten, and T. B. Schön. Smoothing with couplings of conditional particle filters. *arXiv preprint arXiv:1701.02002*, 2017a. [2](#), [12](#), [21](#), [22](#)
- P. E. Jacob, L. M. Murray, C. C. Holmes, and C. P. Robert. Better together? Statistical learning in models made of modules. *arXiv preprint arXiv:1708.08719*, 2017b. [2](#), [19](#), [22](#)
- S. F. Jarner and G. O. Roberts. Convergence of heavy-tailed Monte Carlo Markov chain algorithms. *Scandinavian Journal of Statistics*, 34(4):781–815, 2007. [8](#)
- J. E. Johndrow and J. C. Mattingly. Coupling and decoupling to bound an approximating Markov chain. *arXiv preprint arXiv:1706.02040*, 2017. [2](#)
- V. E. Johnson. Studying convergence of Markov chain Monte Carlo algorithms using coupled sample paths. *Journal of the American Statistical Association*, 91(433):154–166, 1996. [2](#), [4](#), [10](#)
- V. E. Johnson. A coupling-regeneration scheme for diagnosing convergence in Markov chain Monte Carlo algorithms. *Journal of the American Statistical Association*, 93(441):238–248, 1998. [2](#), [4](#), [10](#), [11](#)
- V. E. Johnson. On numerical aspects of Bayesian model selection in high and ultrahigh-dimensional settings. *Bayesian Analysis*, 8 (4):741–758, 2004. [2](#), [4](#), [19](#)
- K. Khare and J. P. Hobert. Geometric ergodicity of the Bayesian lasso. *Electronic Journal of Statistics*, 7:2150–2163, 2013. doi: 10.1214/13-EJS841. URL <http://dx.doi.org/10.1214/13-EJS841>. [8](#)
- M. Knott and C. S. Smith. On the optimal mapping of distributions. *Journal of Optimization Theory and Applications*, 43(1):39–49, 1984. [11](#)

- A. Lee, C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *Journal of computational and graphical statistics*, 19(4):769–789, 2010. [1](#)
- A. Lee, A. Doucet, and K. Łatuszyński. Perfect simulation using atomic regeneration with application to sequential Monte Carlo. *arXiv preprint arXiv:1407.5770*, 2014. [2](#)
- T. Lindvall. *Lectures on the coupling method*. Dover Books on Mathematics, 2002. [4](#), [10](#)
- F. Liu, M. Bayarri, and J. Berger. Modularization in Bayesian analysis, with emphasis on analysis of computer models. *Bayesian Analysis*, 4(1):119–150, 2009. [2](#), [19](#), [22](#)
- J. S. Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008. [1](#)
- S. Livingstone, M. F. Faulkner, and G. O. Roberts. Kinetic energy choice in Hamiltonian/hybrid Monte Carlo. *arXiv preprint arXiv:1706.02649*, 2017. [12](#)
- D. Lunn, N. Best, D. Spiegelhalter, G. Graham, and B. Neuenschwander. Combining MCMC with sequential PKPD modelling. *Journal of Pharmacokinetics and Pharmacodynamics*, 36(1):19–38, 2009. [19](#)
- E. Mainini. A description of transport cost for signed measures. *Journal of Mathematical Sciences*, 181(6):837–855, 2012. [7](#)
- O. Mangoubi and A. Smith. Rapid mixing of Hamiltonian Monte Carlo on strongly log-concave distributions. *arXiv preprint arXiv:1708.07114*, 2017. [12](#)
- D. Maucort-Boulch, S. Franceschi, and M. Plummer. International correlation between human papillomavirus prevalence and cervical cancer incidence. *Cancer Epidemiology Biomarkers & Prevention*, 17(3):717–720, 2008. [20](#)
- L. C. McCandless, I. J. Douglas, S. J. Evans, and L. Smeeth. Cutting feedback in Bayesian regression adjustment for the propensity score. *The International Journal of Biostatistics*, 6(2), 2010. [19](#)
- D. McLeish. A general method for debiasing a Monte Carlo estimator. *Monte Carlo methods and applications*, 17(4):301–315, 2011. [2](#)
- S. Meyn and R. Tweedie. *Markov chains and stochastic stability*. 2 edition, 2009. [8](#), [30](#)
- D. J. Murdoch and P. J. Green. Exact sampling from a continuous state space. *Scandinavian Journal of Statistics*, 25(3):483–502, 1998. [2](#), [15](#)
- K. M. Murphy and R. H. Topel. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 20(1):88–97, 2002. [19](#)
- I. Murray, R. Prescott Adams, and D. J. MacKay. Elliptical slice sampling. 2010. [22](#)
- P. Mykland, L. Tierney, and B. Yu. Regeneration in Markov chain samplers. *Journal of the American Statistical Association*, 90(429):233–241, 1995. [4](#), [15](#), [16](#)
- R. M. Neal. Bayesian learning via stochastic dynamics. *Advances in neural information processing systems*, pages 475–475, 1993. [12](#)

- R. M. Neal. Circularly-coupled Markov chain sampling. Technical report, Department of Statistics, University of Toronto, 1999. [2](#), [4](#), [10](#)
- R. M. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov chain Monte Carlo*, 2(11), 2011. [12](#)
- R. M. Neal and R. L. Pinto. Improving Markov chain Monte Carlo estimators by coupling to an approximating chain. 2001. [4](#)
- E. Nummelin. MC's for MCMC'ists. *International Statistical Review / Revue Internationale de Statistique*, 70(2):pp. 215–240, 2002. ISSN 03067734. [11](#)
- A. B. Owen. Statistically efficient thinning of a Markov chain sampler. *Journal of Computational and Graphical Statistics*, (just-accepted), 2017. [6](#)
- A. Pagan. Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, pages 221–247, 1984. [19](#)
- S. Pal and K. Khare. Geometric ergodicity for Bayesian shrinkage models. *Electron. J. Statist.*, 8(1): 604–645, 2014. doi: 10.1214/14-EJS896. URL <http://dx.doi.org/10.1214/14-EJS896>. [8](#)
- M. Plummer. Cuts in Bayesian graphical models. *Statistics and Computing*, pages 1–7, 2014. [2](#), [19](#), [20](#), [22](#)
- M. Plummer, N. Best, K. Cowles, and K. Vines. CODA: convergence diagnosis and output analysis for MCMC. *R news*, 6(1):7–11, 2006. [14](#)
- M. Pollock, P. Fearnhead, A. M. Johansen, and G. O. Roberts. The scalable Langevin exact algorithm: Bayesian inference for big data. *arXiv preprint arXiv:1609.03436*, 2016. [22](#)
- J. G. Propp and D. B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random structures and Algorithms*, 9(1-2):223–252, 1996. [2](#)
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015. URL <https://www.R-project.org>. [2](#)
- A. Reutter and V. E. Johnson. *General strategies for assessing convergence of MCMC algorithms using coupled sample paths*. Institute of Statistics & Decision Sciences, Duke University, 1995. [2](#)
- C.-H. Rhee and P. W. Glynn. A new approach to unbiased estimation for SDE's. In *Proceedings of the Winter Simulation Conference*, page 17. Winter Simulation Conference, 2012. [7](#)
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer-Verlag, New York, first edition, 1999. [1](#), [12](#)
- G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71, 2004. [2](#), [8](#), [11](#)
- G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996a. [8](#), [11](#)
- G. O. Roberts and R. L. Tweedie. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110, 1996b. [11](#)

- G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120, 1997. [11](#), [13](#)
- J. S. Rosenthal. Analysis of the Gibbs sampler for a model related to James-Stein estimators. *Statistics and Computing*, 6(3):269–275, 1996. [12](#)
- J. S. Rosenthal. Faithful couplings of Markov chains: now equals forever. *Advances in Applied Mathematics*, 18(3):372 – 381, 1997. ISSN 0196-8858. [2](#), [3](#)
- J. S. Rosenthal. Parallel computing and Monte Carlo algorithms. *Far east journal of theoretical statistics*, 4(2):207–236, 2000. [1](#)
- J. S. Rosenthal. Quantitative convergence rates of Markov chains: A simple account. *Electronic Communications in Probability*, 7:123–128, 2002. [4](#)
- D. Spiegelhalter, A. Thomas, N. Best, and D. Lunn. WinBUGS user manual, 2003. [19](#)
- S. Srivastava, V. Cevher, Q. Dinh, and D. Dunson. WASP: Scalable Bayes via barycenters of subset posteriors. In *Artificial Intelligence and Statistics*, pages 912–920, 2015. [1](#)
- H. Thorisson. *Coupling, stationarity, and regeneration*, volume 14. Springer New York, 2000. [4](#), [10](#)
- M. K. Titsias and C. Yau. The Hamming ball sampler. *Journal of the American Statistical Association*, pages 1–14, 2017. [22](#)
- H. Tjelmeland. Using all Metropolis–Hastings proposals to estimate mean values. Technical report, Department of Mathematical Sciences, Norwegian University of Science and Technology, 2004. [1](#)
- R. Tweedie. The existence of moments for stationary Markov chains. *Journal of Applied Probability*, 20(1):191–196, 1983. [3](#)
- P. Vanetti, A. Bouchard-Côté, G. Deligiannidis, and A. Doucet. Piecewise deterministic Markov chain Monte Carlo. *arXiv preprint arXiv:1707.05296*, 2017. [22](#)
- D. Vats, J. M. Flegal, G. L. Jones, et al. Strong consistency of multivariate spectral variance estimators in Markov chain Monte Carlo. *Bernoulli*, 24(3):1860–1909, 2018. [2](#)
- M. Vihola. Unbiased estimators and multilevel Monte Carlo. *arXiv preprint arXiv:1512.01022*, 2015. [28](#)
- C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008. [11](#)
- X. Wang, F. Guo, K. A. Heller, and D. B. Dunson. Parallelizing MCMC with random partition trees. In *Advances in Neural Information Processing Systems*, pages 451–459, 2015. [1](#)
- S. Yang, Y. Chen, E. Bernton, and J. S. Liu. On parallelizable Markov chain Monte Carlo algorithms with waste-recycling. *Statistics and Computing*, pages 1–9, 2017. [1](#)
- Y. Yang, M. J. Wainwright, and M. I. Jordan. On the computational complexity of high-dimensional Bayesian variable selection. *The Annals of Statistics*, 44(6):2497–2532, 2016. [17](#), [18](#), [19](#)
- C. M. Zigler. The central role of Bayes theorem for joint estimation of causal effects and propensity scores. *The American Statistician*, 70(1):47–54, 2016. [19](#)

A Proofs

A.1 Proof of Proposition 3.1

We present the proof for $H_0(X, Y)$, a similar reasoning holds for $H_k(X, Y)$. We follow the same arguments as in Glynn and Rhee [2014], Vihola [2015]. To study $H_0(X, Y) = h(X_0) + \sum_{t=1}^{\tau-1} (h(X_t) - h(Y_{t-1}))$, we introduce $\Delta_0 = h(X_0)$ and $\Delta_t = h(X_t) - h(Y_{t-1})$ for all $t \geq 1$, and define $H_0^n(X, Y) = \sum_{t=0}^n \Delta_t$. For simplicity we assume that $\Delta_t \in \mathbb{R}$, which corresponds to studying the component-wise behavior of $H_0(X, Y)$.

We have $\mathbb{E}[\tau] < \infty$ from Assumption 2.2, so that the computing time of $H_0(X, Y)$ has a finite expectation. Together with Assumption 2.3, this implies that $H_0^n(X, Y) \rightarrow H_0(X, Y)$ almost surely as $n \rightarrow \infty$. We now show that $H_0^n(X, Y)$ is a Cauchy sequence in L_2 , the complete space of random variables with finite two moments, that is, $\sup_{n' \geq n} \mathbb{E}[(H_0^{n'}(X, Y) - H_0^n(X, Y))^2] \rightarrow 0$ as $n \rightarrow \infty$. For $n' \geq n$, we compute

$$\mathbb{E}[(H_0^{n'}(X, Y) - H_0^n(X, Y))^2] = \sum_{s=n+1}^{n'} \sum_{t=n+1}^{n'} \mathbb{E}[\Delta_s \Delta_t],$$

and consider each term $\mathbb{E}[\Delta_s \Delta_t]$ for $(s, t) \in \{n+1, \dots, n'\}^2$. The Cauchy-Schwarz inequality implies $\mathbb{E}[\Delta_s \Delta_t] \leq (\mathbb{E}[\Delta_s^2] \cdot \mathbb{E}[\Delta_t^2])^{1/2}$, and noting that $\mathbb{E}[\Delta_t^2] = \mathbb{E}[\Delta_t^2 \cdot \mathbf{1}(\tau > t)]$, we can apply Hölder's inequality with $p = 1 + \eta/2$, $q = (2 + \eta)/\eta$ for any $\eta > 0$ to obtain

$$\mathbb{E}[\Delta_t^2 \cdot \mathbf{1}(\tau > t)] \leq \mathbb{E}[|\Delta_t|^{2+\eta}]^{1/(1+\eta/2)} \cdot \mathbb{E}[\mathbf{1}(\tau > t)]^{\eta/(2+\eta)} \leq \mathbb{E}[|\Delta_t|^{2+\eta}]^{1/(1+\eta/2)} \cdot (C\delta^t)^{\eta/(2+\eta)}.$$

Here we have used Assumption 2.2 to bound $\mathbb{E}[\mathbf{1}(\tau > t)]$. We can also use Assumption 2.1 together with Minkowski's inequality to bound $\mathbb{E}[|\Delta_t|^{2+\eta}]^{1/(1+\eta/2)}$ by a constant \tilde{C} , for all $t \geq 0$. Defining $\tilde{\delta} = \delta^{\eta/(2+\eta)} \in (0, 1)$ then gives the bound $\mathbb{E}[\Delta_t^2] \leq \tilde{C}\tilde{\delta}^t$ for all $t \geq 0$. This implies $\mathbb{E}[(H_0^{n'}(X, Y) - H_0^n(X, Y))^2] \leq \tilde{C}\tilde{\delta}^n$ for some other constant \tilde{C} , and thus $(H_0^n(X, Y))$ is Cauchy in L_2 . This proves that its limit $H_0(X, Y)$ has finite first and second moments. Assumption 2.1 implies that $\lim_{n \rightarrow \infty} \mathbb{E}[H_0^n(X, Y)] = \mathbb{E}_\pi[h(X)]$, by a telescopic sum argument, so we conclude that $\mathbb{E}[H_0(X, Y)] = \mathbb{E}_\pi[h(X)]$. We can also obtain an explicit representation of $\mathbb{E}[H_0(X, Y)^2]$ as the limit of $\mathbb{E}[H_0^n(X, Y)^2]$ when $n \rightarrow \infty$.

A.2 Proof of Proposition 3.2

We adopt a similar strategy to that of Glynn and Rhee [2014], Section 4. For the $H_k(X, Y)$ case, the unbiased estimator of $F(s) = \mathbb{P}_\pi(X \leq s)$ over R samples is of the form

$$\hat{F}^R(s) = \frac{1}{R} \sum_{r=1}^R \left(\mathbf{1}(X_k^{(r)} \leq s) + \sum_{\ell=k+1}^{\tau^{(r)}} (\mathbf{1}(X_\ell^{(r)} \leq s) - \mathbf{1}(Y_{\ell-1}^{(r)} \leq s)) \right).$$

Here $\tau^{(r)}$ denotes the meeting time of the r -th independent run. We want to show that as $R \rightarrow \infty$, $\sup_s |F(s) - \hat{F}^R(s)| \xrightarrow{a.s.} 0$. Define $G_{X,k}^R(s) = R^{-1} \sum_{r=1}^R \mathbf{1}(X_k^{(r)} \leq s)$. For $\ell > k$, define

$$G_{X,\ell}^R(s) = \frac{1}{R} \sum_{r=1}^R \mathbf{1}(X_\ell^{(r)} \leq s) \cdot \mathbf{1}(\ell \leq \tau^{(r)}), \quad G_{Y,\ell}^R(s) = \frac{1}{R} \sum_{r=1}^R \mathbf{1}(Y_{\ell-1}^{(r)} \leq s) \cdot \mathbf{1}(\ell \leq \tau^{(r)}).$$

Then $\hat{F}^R(s) = G_{X,k}^R(s) + \sum_{\ell=k+1}^{\infty} (G_{X,\ell}^R(s) - G_{Y,\ell}^R(s))$. By the standard Glivenko-Cantelli theorem, as $\mathbb{R} \rightarrow \infty$ we have

$$\begin{aligned} \sup_s |G_{X,k}^R(s) - \mathbb{P}(X_k \leq s)| &\xrightarrow{a.s.} 0, & \sup_s |G_{X,\ell}^R(s) - \mathbb{E}[\mathbf{1}(X_\ell \leq s) \cdot \mathbf{1}(\ell \leq \tau)]| &\xrightarrow{a.s.} 0, \\ \sup_s |G_{Y,\ell}^R(s) - \mathbb{E}[\mathbf{1}(Y_{\ell-1} \leq s) \cdot \mathbf{1}(\ell \leq \tau)]| &\xrightarrow{a.s.} 0, \end{aligned}$$

for each $\ell > k$. Next, we observe that, for all s, ℓ ,

$$\mathbb{E}[(\mathbf{1}(X_\ell \leq s) - \mathbf{1}(Y_{\ell-1} \leq s)) \cdot \mathbf{1}(\tau \geq \ell)] = \mathbb{P}(X_\ell \leq s) - \mathbb{P}(X_{\ell-1} \leq s).$$

This holds for a simple reason in our setting. For any $h(\cdot)$ and any ℓ ,

$$\begin{aligned} \mathbb{E}[h(X_\ell) - h(Y_{\ell-1})] &= \mathbb{E}[(h(X_\ell) - h(Y_{\ell-1}))\mathbf{1}(\tau > \ell)] + \mathbb{E}[(h(X_\ell) - h(Y_{\ell-1}))\mathbf{1}(\tau \leq \ell)] \\ &= \mathbb{E}[(h(X_\ell) - h(Y_{\ell-1}))\mathbf{1}(\tau > \ell)] \end{aligned}$$

since if $\tau \leq \ell$ then $X_\ell = Y_{\ell-1}$ by Assumption 2.3. Applying this result with $h(\cdot) = \mathbf{1}(\cdot \leq s)$ yields the desired statement.

The above implies that for any finite $i \geq k$ we have

$$\begin{aligned} &\left| G_{X,k}^R(s) - \mathbb{P}(X_k \leq s) + \sum_{\ell=k+1}^i \left(G_{X,\ell}^R(s) - G_{Y,\ell}^R(s) - \mathbb{E}[(\mathbf{1}(X_\ell \leq s) - \mathbf{1}(Y_{\ell-1} \leq s)) \cdot \mathbf{1}(\ell \leq \tau)] \right) \right| \\ &= \left| G_{X,k}^R(s) - \mathbb{P}(X_k \leq s) + \sum_{\ell=k+1}^i \left(G_{X,\ell}^R(s) - G_{Y,\ell}^R(s) - (\mathbb{P}(X_\ell \leq s) - \mathbb{P}(X_{\ell-1} \leq s)) \right) \right| \\ &= \left| \left(G_{X,k}^R(s) + \sum_{\ell=k+1}^i (G_{X,\ell}^R(s) - G_{Y,\ell}^R(s)) \right) - \mathbb{P}(X_i \leq s) \right| \end{aligned}$$

Hence

$$\sup_s \left| \left(G_{X,k}^R(s) + \sum_{\ell=k+1}^i (G_{X,\ell}^R(s) - G_{Y,\ell}^R(s)) \right) - \mathbb{P}(X_i \leq s) \right| \rightarrow 0.$$

We have assumed that $(X_t)_{t \geq 0}$ converges to π in total variation, which implies $\sup_s |\mathbb{P}(X_i \leq s) - F(s)| \rightarrow 0$ as $i \rightarrow \infty$. Also, we note that for all s ,

$$\left| \sum_{\ell > i} G_{X,\ell}^R(s) - G_{Y,\ell}^R(s) \right| \leq \sum_{\ell > i} \frac{1}{R} \sum_{r=1}^R \mathbf{1}(\ell \leq \tau^{(r)}) \rightarrow \sum_{\ell > i} \mathbb{P}(\ell \leq \tau)$$

almost surely by the strong law of large numbers. Assumption 2.2 implies that this quantity goes to 0 as $i \rightarrow \infty$.

Combining these observations with the result obtained for finite i , we conclude that $\sup_s |\hat{F}^R(s) - F(s)| \rightarrow 0$ as $R \rightarrow \infty$, almost surely. The reasoning holds for the function \hat{F}^R corresponding to $H_{k:m}(X, Y)$ instead of $H_k(X, Y)$; such a function is simply the average of a finite number of functions associated with $H_\ell(X, Y)$ for $\ell \in \{k, \dots, m\}$.

A.3 Proof of Proposition 3.3

Throughout the proof C denotes a generic finite constant whose actual value may change from one appearance to the next. We will use the usual Markov chain notation. In particular if $f : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is

a measurable function then $\bar{P}[f](x, y) := \int_{\mathcal{X} \times \mathcal{X}} \bar{P}((x, y), dz) f(z)$. Note that from the construction of \bar{P} , if f depends only on x (resp. y), that is $f(x, y) = f(x)$ (resp. $f(x, y) = f(y)$), then $\bar{P}[f](x, y) = Pf(x)$ (resp. $\bar{P}[f](x, y) = Pf(y)$).

For $k \geq 1$, we consider the general problem of bounding $\mathbb{E}[S_k^2]$, where S_k is of the form

$$S_k = \mathbf{1}(\tau > k) \sum_{t=k}^{\tau-1} b_t (h(X_t) - h(Y_{t-1})),$$

for some arbitrary bounded sequence $(b_t)_{t \geq 0}$. Fix an integer $N \geq k$, and set

$$S_k^{(N)} = \mathbf{1}(\tau > k) \sum_{t=k}^{N \wedge \tau - 1} b_t (h(X_t) - h(Y_{t-1})).$$

The same argument as in the proof of Proposition 3.1 can be applied here and shows that $S_k^{(N)} \in L^2$, is a Cauchy sequence in L_2 that converges to S_k , as $N \rightarrow \infty$, so that

$$\mathbb{E}[S_k^2] = \lim_{N \rightarrow \infty} \mathbb{E}[(S_k^{(N)})^2].$$

Since $|h|_{V^\beta} := \sup_x |h(x)|/V^\beta(x) < \infty$, and under Assumption 3.1, there exists a measurable function $g : \mathcal{X} \rightarrow \mathbb{R}$ such that $|g|_{V^\beta} < \infty$, and $g - Pg = h - \pi(h)$. To see this, first note that the drift condition (3.1) implies that for any $\beta \in (0, 1/2)$, we have $PV^\beta(x) \leq \lambda^\beta V^\beta(x) + b^\beta \mathbf{1}(x \in \mathcal{C})$, for all $x \in \mathcal{X}$. Indeed by Jensen's inequality $PV^\beta(x) \leq (PV(x))^\beta \leq (\lambda V(x) + b \mathbf{1}(x \in \mathcal{C}))^\beta \leq \lambda^\beta V^\beta(x) + b^\beta \mathbf{1}(x \in \mathcal{C})$, using the fact that for all $x, y \geq 0$ and $\alpha \in [0, 1]$, $(x + y)^\alpha \leq x^\alpha + y^\alpha$. The drift condition in V^β together with the fact that P is ϕ -irreducible and aperiodic implies that there exists $\rho_\beta \in (0, 1)$, $C_\beta < \infty$ such that for all $x \in \mathcal{X}$, $n \geq 0$,

$$\|P^n(x, \cdot) - \pi\|_{V^\beta} \leq C_\beta V^\beta(x) \rho_\beta^n, \quad (\text{A.1})$$

where for a function $W : \mathcal{X} \rightarrow [1, \infty)$, the W -norm between two probability measures μ, ν is defined as

$$\|\mu - \nu\|_W := \sup_{f \text{ meas.: } |f|_W \leq 1} |\mu(f) - \nu(f)|,$$

and $|f|_W := \sup_x |f(x)|/W(x)$. This result can be found in Theorem 15.0.1 of [Meyn and Tweedie \[2009\]](#). It follows from (A.1) that $\sum_{j \geq 0} |P^j(h - \pi(h))(x)| \leq |h|_{V^\beta} \sum_{j \geq 0} \|P^j(x, \cdot) - \pi\|_{V^\beta} \leq \frac{C_\beta |h|_{V^\beta} V^\beta(x)}{1 - \rho_\beta} < \infty$. Hence the function

$$g(x) = \sum_{j \geq 0} P^j(h - \pi(h))(x), \quad x \in \mathcal{X},$$

is well-defined and measurable (as a limit of a sequence of measurable functions) and satisfies $|g(x)| \leq C_\beta |h|_{V^\beta} V^\beta(x)/(1 - \rho_\beta)$. And since PV^β is finite everywhere, by Lebesgue's dominated convergence we deduce that Pg is finite everywhere as well and

$$Pg(x) = \int g(y) P(x, dy) = \sum_{j \geq 0} \int (P^j(h - \pi(h))(y)) P(x, dy) = \sum_{j \geq 1} P^j(h - \pi(h))(x).$$

Hence $g - Pg = h - \pi(h)$, as claimed.

Hence, with $Z_t := (X_t, Y_{t-1})$, and $\bar{g}(x, y) = g(x) - g(y)$, we have

$$h(X_t) - h(Y_{t-1}) = \bar{g}(Z_t) - \bar{P}[\bar{g}](Z_t).$$

Using this and a telescoping sum argument, we write

$$\begin{aligned}
S_k^{(N)} &= \sum_{t=k}^{N-1} b_t (h(X_t) - h(Y_{t-1})) \mathbf{1}(\tau > t) \\
&= \sum_{t=k}^{N-1} b_t (\bar{g}(Z_t) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t) \\
&= \sum_{t=k}^{N-1} b_t (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t) \\
&\quad + \sum_{t=k}^{N-1} (b_t \bar{g}(Z_t) \mathbf{1}(\tau > t) - b_{t+1} \bar{g}(Z_{t+1}) \mathbf{1}(\tau > t+1)) \\
&\quad + \sum_{t=k}^{N-1} (b_{t+1} \mathbf{1}(\tau > t+1) - b_t \mathbf{1}(\tau > t)) \bar{g}(Z_{t+1}).
\end{aligned}$$

Since $\bar{g}(Z_{t+1}) = 0$ on $\{\tau = t+1\}$, the last term in the above display reduces to $\sum_{t=k}^{N-1} (b_{t+1} - b_t) \bar{g}(Z_{t+1}) \mathbf{1}(\tau > t+1)$, and we obtain

$$\begin{aligned}
S_k^{(N)} &= \sum_{t=k}^{N-1} b_t (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t) \\
&\quad + b_k \bar{g}(Z_k) \mathbf{1}(\tau > k) - b_N \bar{g}(Z_N) \mathbf{1}(\tau > N) + \sum_{t=k}^{N-1} (b_{t+1} - b_t) \bar{g}(Z_{t+1}) \mathbf{1}(\tau > t+1). \quad (\text{A.2})
\end{aligned}$$

Let \mathcal{F}_t denote the sigma-algebra generated by the variables $X_0, (X_1, Y_0), \dots, (X_t, Y_{t-1})$. Note that $\{\tau > t\}$ belongs to \mathcal{F}_t . Hence

$$\begin{aligned}
\mathbb{E} [b_t (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t) | \mathcal{F}_t] &= b_t \mathbf{1}(\tau > t) \mathbb{E} [\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t) | \mathcal{F}_t] \\
&= b_t \mathbf{1}(\tau > t) (\bar{P}[\bar{g}](Z_t) - \bar{P}[\bar{g}](Z_t)) = 0.
\end{aligned}$$

In other words, $\left\{ \left(\sum_{t=k}^j b_t (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t), \mathcal{F}_j \right), k \leq j \leq N-1 \right\}$ is a martingale. The orthogonality of the martingale increments gives

$$\mathbb{E} \left[\left(\sum_{t=k}^{N-1} b_t (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t) \right)^2 \right] = \sum_{t=k}^{N-1} b_t^2 \mathbb{E} \left[(\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t))^2 \mathbf{1}(\tau > t) \right].$$

We use this together with (A.2), the convexity of the squared norm, and Minkowski's inequality to conclude that

$$\begin{aligned}
\mathbb{E} \left[(S_k^{(N)})^2 \right] &\leq 4 \sum_{t=k}^{N-1} b_t^2 \mathbb{E} \left[(\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t))^2 \mathbf{1}(\tau > t) \right] + 4b_k^2 \mathbb{E} [\bar{g}^2(Z_k) \mathbf{1}(\tau > k)] \\
&\quad + 4b_N^2 \mathbb{E} [\bar{g}^2(Z_N) \mathbf{1}(\tau > N)] + 4 \left[\sum_{t=k}^{N-1} |b_{t+1} - b_t| \mathbb{E}^{1/2} (\bar{g}^2(Z_{t+1}) \mathbf{1}(\tau > t+1)) \right]^2. \quad (\text{A.3})
\end{aligned}$$

Assumption 3.1 together with $\pi_0(V) < \infty$, implies that

$$\sup_{n \geq 0} \mathbb{E}[V(X_n)] \leq C, \quad (\text{A.4})$$

for some finite constant C . Indeed, $\mathbb{E}[V(X_n)] = \int \pi_0(dx) P^n V(x)$, and a repeated application of the drift condition (3.1) implies that $P^n V(x) \leq \lambda^n V(x) + \frac{b}{1-\lambda}$, for all $x \in \mathcal{X}$. For any $t \geq 0$, and for $1 < p = \frac{1}{2\beta}$, and $\frac{1}{p} + \frac{1}{q} = 1$, we have

$$\begin{aligned}
\mathbb{E} \left[(V^\beta(X_t) + V^\beta(Y_{t-1}))^2 \mathbf{1}(\tau > t) \right] &\leq \mathbb{E}^{1/p} \left[(V^\beta(X_t) + V^\beta(Y_{t-1}))^{2p} \right] \mathbb{P}^{1/q}(\tau > t) \quad (\text{H\"older}) \\
&\leq \left\{ \mathbb{E}^{1/(2p)} [V^{2p\beta}(X_t)] + \mathbb{E}^{1/(2p)} [V^{2p\beta}(Y_{t-1})] \right\}^2 \\
&\hspace{15em} \times \mathbb{P}^{1/q}(\tau > t) \quad (\text{Minkowski}) \\
&\leq C \mathbb{P}^{1/q}(\tau > t) \hspace{10em} (\text{by (A.4)}) \\
&\leq C \delta_\beta^t \hspace{15em} (\text{by Assumption 2.2}),
\end{aligned}$$

where $\delta_\beta = \delta^{1/q}$ with δ as in Assumption 2.2. Note that all the expectations on the right hand side of (A.3) are of the form $\mathbb{E} [\bar{g}^2(Z_t) \mathbf{1}(\tau > t)]$, except for the term $\mathbb{E} \left[(\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t))^2 \mathbf{1}(\tau > t) \right]$. However by the martingale difference property, $\mathbb{E} [\bar{P}[\bar{g}](Z_t) (\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t)) \mathbf{1}(\tau > t)] = 0$, so that for all $t \geq k$,

$$\begin{aligned}
\mathbb{E} \left[(\bar{g}(Z_{t+1}) - \bar{P}[\bar{g}](Z_t))^2 \mathbf{1}(\tau > t) \right] &= \mathbb{E} [\mathbf{1}(\tau > t) \bar{g}(Z_{t+1})^2] - \mathbb{E} [\bar{g}(Z_{t+1}) \bar{P}[\bar{g}](Z_t) \mathbf{1}(\tau > t)] \\
&= \mathbb{E} [\mathbf{1}(\tau > t) \bar{g}(Z_{t+1})^2] - \mathbb{E} \left[(\bar{P}[\bar{g}](Z_t))^2 \mathbf{1}(\tau > t) \right] \quad (\text{by conditioning on } \mathcal{F}_t) \\
&\leq \mathbb{E} [\mathbf{1}(\tau > t) \bar{g}(Z_{t+1})^2] \\
&\leq |g|_{V^\beta}^2 \mathbb{E} \left[\mathbf{1}(\tau > t) (V^\beta(X_{t+1}) + V^\beta(Y_t))^2 \right] \\
&\leq C \delta_\beta^t,
\end{aligned}$$

using the same arguments as above. On the other hand, for all $t \geq k$, the term $\mathbb{E} [\bar{g}^2(Z_t) \mathbf{1}_{\{\tau > t\}}]$ satisfies

$$\mathbb{E} [\bar{g}^2(Z_t) \mathbf{1}(\tau > t)] \leq |g|_{V^\beta}^2 \mathbb{E} \left[(V^\beta(X_t) + V^\beta(Y_{t-1}))^2 \mathbf{1}(\tau > t) \right] \leq C \delta_\beta^t,$$

as seen above. In conclusion, all the expectations appearing in (A.3) are upper bounded by some constant times terms of the form δ_β^t . We conclude that

$$\mathbb{E} \left[(S_k^{(N)})^2 \right] \leq C \left(b_k^2 \delta_\beta^k + b_N^2 \delta_\beta^N + \sum_{t=k}^{N-1} b_k^2 \delta_\beta^t + \left[\sum_{t=k}^{N-1} |b_{t+1} - b_t| \delta_\beta^t \right]^2 \right).$$

Letting $N \rightarrow \infty$, we conclude that

$$\mathbb{E}[S_k^2] \leq C \left(b_k^2 \delta_\beta^k + \sum_{j \geq k} b_j^2 \delta_\beta^j + \left[\sum_{j \geq k} |b_{j+1} - b_j| \delta_\beta^j \right]^2 \right).$$

In the particular case of $\eta_{k:m}$, we have $\eta_{k:m} = \mathbf{1}(\tau > k) \sum_{t=k}^{\tau-1} \min \left(1, \frac{t+1-k}{m+1-k} \right) (h(X_{t+1}) - h(Y_t))$. Hence $b_k = 0$, $b_t = (t-k)/(m-k+1)$ if $k < t \leq m+1$, $b_t = 1$ if $t > m+1$. We then obtain the bound of Proposition 3.3.

A.4 Proof of Proposition 3.4

Here Z_n is defined as (X_n, Y_{n-1}) for all $n \geq 1$. Assumption (3.3) implies that for $(x, y) \in \mathcal{C} \times \mathcal{C}$, \bar{P} can be written as a mixture

$$\bar{P}((x, y), dz) = \epsilon_{x,y} \nu_{x,y}(dz) + (1 - \epsilon_{x,y}) R((x, y), dz),$$

where $\epsilon_{x,y} \geq \epsilon$, $\nu_{x,y}(dz)$ is a restriction of $\bar{P}((x, y), dz)$ on \mathcal{D} (that is for any measurable subset A of \mathcal{D} , $\bar{P}((x, y), A) = P((x, y), A \cap \mathcal{D}) / P((x, y), \mathcal{D})$), and $R((x, y), dz)$ is the restriction of $\bar{P}((x, y), dz)$ on $(\mathcal{X} \times \mathcal{X}) \setminus \mathcal{D}$. This means that whenever $(x, y) \in \mathcal{C} \times \mathcal{C}$ one can sample from $\bar{P}((x, y), \cdot)$ by drawing independently a Bernoulli random variable J , with probability of success $\epsilon_{x,y}$. Then if $J = 1$, we draw from $\nu_{x,y}$, if $J = 0$, we draw from $R((x, y), \cdot)$. From this decomposition, the proof of the proposition follows the same lines as in Douc et al. [2004], and we give the details only for completeness. We cannot directly invoke their result since their assumptions do not seem to apply to our setting.

Set $\bar{V}(x, y) = \frac{1}{2}(V(x) + V(y))$. First we show that the bivariate kernel satisfies a geometric drift towards $\mathcal{C} \times \mathcal{C}$. That is, there exists $\alpha \in (0, 1)$ such that

$$\bar{P}\bar{V}(x, y) \leq \alpha \bar{V}(x, y), \quad (x, y) \notin \mathcal{C} \times \mathcal{C}. \quad (\text{A.5})$$

Indeed for $(x, y) \notin \mathcal{C} \times \mathcal{C}$, since $V \geq 1$, and $\mathcal{C} = \{V \leq L\}$, $\bar{V}(x, y) \geq (1 + L)/2$. In other words, $\frac{1}{2} \leq \bar{V}(x, y)/(1 + L)$. Therefore,

$$\bar{P}\bar{V}(x, y) = \frac{1}{2}(PV(x) + PV(y)) \leq \lambda \bar{V}(x, y) + \frac{b}{2} \leq \lambda \bar{V}(x, y) + \frac{b}{1+L} \bar{V}(x, y) \leq \alpha \bar{V}(x, y),$$

with $\alpha = \lambda + \frac{b}{1+L} < 1$. We set

$$B = \max \left(1, \frac{1}{\alpha} \sup_{(x,y) \in \mathcal{C} \times \mathcal{C}} \frac{\bar{P}[\bar{V} \mathbf{1}_{\mathcal{D}^c}](x, y)}{\bar{V}(x, y)} \right) \leq \frac{\lambda + b}{\alpha}.$$

In this section $\mathbf{1}_{\mathcal{S}}(\cdot)$ refers to the indicator function on the set \mathcal{S} . Let N_n denote the number of visits to $\mathcal{C} \times \mathcal{C}$ by time n . Then

$$\mathbb{P}(\tau > n) = \mathbb{P}(\tau > n, N_{n-1} \geq j) + \mathbb{P}(\tau > n, N_{n-1} < j).$$

The event $\{\tau > n, N_{n-1} \geq j\}$ implies that no success occurred within at least j independent Bernoulli random variables each with probability of success at least ϵ . Hence

$$\mathbb{P}(\tau > n, N_{n-1} \geq j) \leq (1 - \epsilon)^j.$$

For the second term, we have (since $B \geq 1$, and the chains stay together after meeting via Assumption 2.3),

$$\mathbb{P}(\tau > n, N_{n-1} \leq j - 1) \leq \mathbb{P}(Z_n \notin \mathcal{D}, B^{-N_{n-1}} \geq B^{-(j-1)}) = \mathbb{P}(\mathbf{1}_{\mathcal{D}^c}(Z_n) B^{-N_{n-1}} \geq B^{-(j-1)}).$$

Then use Markov's inequality to conclude that

$$\begin{aligned} \mathbb{P}(\tau > n, N_{n-1} \leq j-1) &\leq B^{j-1} \mathbb{E} [\mathbf{1}_{\mathcal{D}^c}(Z_n) B^{-N_{n-1}}] \\ &\leq B^{j-1} \mathbb{E} [\mathbf{1}_{\mathcal{D}^c}(Z_n) B^{-N_{n-1}} \bar{V}(Z_n)] = \alpha^n B^{j-1} \mathbb{E}[M_n], \end{aligned}$$

where $M_n = \mathbf{1}_{\mathcal{D}^c}(Z_n) \alpha^{-n} B^{-N_{n-1}} \bar{V}(Z_n)$ (set $N_0 = 0$ so that M_1 is well-defined). The result follows by noting that $\{M_n, \mathcal{F}_n\}$ is a super-martingale, where $\mathcal{F}_n = \sigma(Z_1, \dots, Z_n)$, so that $\mathbb{E}[M_n] \leq \mathbb{E}[M_1] \leq \pi_0(V) + \pi_0(PV) \leq (1 + \lambda)\pi_0(V) + b$, which implies that

$$\mathbb{P}(\tau > n) \leq (1 - \epsilon)^j + \alpha^n B^{j-1} ((1 + \lambda)\pi_0(V) + b).$$

Since $\alpha < 1$, there exists an integer $k_0 \geq 1$ such that $\alpha B^{\frac{1}{k_0}} < 1$. In that case for $n \geq k_0$ one can take $j = \lceil n/k_0 \rceil$, to get

$$\mathbb{P}(\tau > n) \leq \left\{ (1 - \epsilon)^{\frac{1}{k_0}} \right\}^n + ((1 + \lambda)\pi_0(V) + b) \left\{ \alpha B^{\frac{1}{k_0}} \right\}^n,$$

as claimed.

The argument that $\{M_n, \mathcal{F}_n\}$ is a super-martingale is as follows. We need to show that for all $n \geq 1$, $\mathbb{E}[M_{n+1} | \mathcal{F}_n] \leq M_n$. Note that $\mathbb{E}[M_{n+1} | \mathcal{F}_n] = 0 \leq M_n$ on $Z_n \in \mathcal{D}$. So it is enough to assume that $Z_n \notin \mathcal{D}$. Now, suppose also that $Z_n \notin \mathcal{C} \times \mathcal{C}$. Then $N_n = N_{n-1}$, and

$$\begin{aligned} \mathbb{E}[M_{n+1} | \mathcal{F}_n] &= \alpha^{-n-1} \mathbb{E} [B^{-N_{n-1}} \mathbf{1}_{\mathcal{D}^c}(Z_{n+1}) \bar{V}(Z_{n+1}) | \mathcal{F}_n], \\ &= \alpha^{-n-1} B^{-N_{n-1}} \mathbb{E} [\mathbf{1}_{\mathcal{D}^c}(Z_{n+1}) \bar{V}(Z_{n+1}) | Z_n], \\ &\leq \alpha^{-n-1} B^{-N_{n-1}} \mathbb{E} [\bar{V}(Z_{n+1}) | Z_n], \\ &\leq \alpha^{-n} B^{-N_{n-1}} \bar{V}(Z_n), \\ &= M_n. \end{aligned}$$

Suppose now that $Z_n \in \mathcal{C} \times \mathcal{C}$. Then $N_n = N_{n-1} + 1$. Hence

$$\begin{aligned} \mathbb{E}[M_{n+1} | \mathcal{F}_n] &= \alpha^{-n-1} B^{-N_{n-1}-1} \mathbb{E} [\mathbf{1}_{\mathcal{D}^c}(Z_{n+1}) \bar{V}(Z_{n+1}) | \mathcal{F}_n], \\ &= \alpha^{-n} B^{-N_{n-1}} \bar{V}(Z_n) \frac{1}{\alpha B} \frac{\bar{P}[\mathbf{1}_{\mathcal{D}^c} \bar{V}](Z_n)}{\bar{V}(Z_n)}, \\ &\leq \alpha^{-n} B^{-N_{n-1}} \bar{V}(Z_n) = M_n. \end{aligned}$$