

Exact sampling for intractable probability distributions via a Bernoulli factory

James M. Flegal

*Department of Statistics
University of California, Riverside
e-mail: jflegal@ucr.edu*

and

Radu Herbei

*Department of Statistics
The Ohio State University
e-mail: herbei@stat.osu.edu*

Abstract: Many applications in the field of statistics require Markov chain Monte Carlo methods. Determining appropriate starting values and run lengths can be both analytically and empirically challenging. A desire to overcome these problems has led to the development of exact, or perfect, sampling algorithms which convert a Markov chain into an algorithm that produces i.i.d. samples from the stationary distribution. Unfortunately, very few of these algorithms have been developed for the distributions that arise in statistical applications, which typically have uncountable support. Here we study an exact sampling algorithm using a geometrically ergodic Markov chain on a general state space. Our work provides a significant reduction to the number of input draws necessary for the Bernoulli factory, which enables exact sampling via a rejection sampling approach. We illustrate the algorithm on a univariate Metropolis-Hastings sampler and a bivariate Gibbs sampler, which provide a proof of concept and insight into hyper-parameter selection. Finally, we illustrate the algorithm on a Bayesian version of the one-way random effects model with data from a styrene exposure study.

AMS 2000 subject classifications: Primary 60J22.

Keywords and phrases: Markov chain, Monte Carlo, perfect sampling, Bernoulli factory, geometric ergodicity.

Received December 2010.

1. Introduction

Suppose we want to explore a probability distribution π defined on X . Further suppose π is intractable in the sense that direct (i.i.d.) sampling is unavailable. In this setting, the Markov chain Monte Carlo (MCMC) method can be a useful tool since it is often straightforward to construct and simulate an ergodic Markov chain that has π as its stationary distribution (Chen et al., 2000; Liu, 2001; Robert and Casella, 1999). The two main drawbacks of MCMC relative to direct sampling from π are (i) the difficulty in ascertaining how long the Markov

chain needs to be run before it gets “close” to π (see, e.g., Jones and Hobert, 2001), and (ii) the difficulty in deriving and calculating asymptotically valid standard errors for the ergodic averages that are used to approximate intractable expectations under π (see, e.g., Flegal et al., 2008).

A desire to overcome these problems has led to the development of clever techniques using a Markov chain to create an algorithm that produces i.i.d. draws from the stationary distribution (e.g., Craiu and Meng, 2011; Green and Murdoch, 1999; Huber, 2004; Propp and Wilson, 1996; Wilson, 2000). Unfortunately, very few of these so-called perfect sampling algorithms have been developed for the distributions that arise in realistic statistical applications, which typically have uncountable support. Asmussen et al. (1992) and Blanchet and Meng (2005) provide one such algorithm applicable to Markov chains on general state spaces. The main assumption necessary is that the chain satisfies a one-step minorization condition. As we describe later, under this condition the stationary distribution admits a *mixture representation*, suggesting the following two-step sampling approach: sample the discrete distribution corresponding to the mixture weights, then sample the selected mixture component.

This approach has never been successfully implemented, however, it has been used to obtain approximate draws from π , see for example Blanchet and Thomas (2007), Hobert et al. (2006) and Hobert and Robert (2004). The difficult part is drawing from the discrete distribution corresponding to the mixture weights, which is done via a rejection sampling approach. In this paper, we provide solutions to a number of practical problems and illustrate the algorithm on three examples. This requires overcoming two challenges: (i) obtaining a dominating proposal distribution and (ii) generating a Bernoulli variate to decide whether a proposed draw is accepted or not. While (ii) might seem trivial, the challenge is that we are unable to (exactly) compute the success probability for this Bernoulli variate.

A solution to (i) requires identification of a bounding (proposal) probability mass function for the target mass function. Previously, Blanchet and Meng (2005) proposed an upper bound on moments associated with the target mass function. Blanchet and Thomas (2007) used output from a preliminary run of the Markov chain to construct an approximate upper bound. We provide an explicit bound for Markov chains satisfying a geometric drift condition using results from Roberts and Tweedie (1999).

A solution to (ii) will determine whether to accept a proposed draw. The decision is made by generating a Bernoulli random variable with success probability that involves the ratio between the target and proposal mass functions. In our case, the target mass function is unknown, apparently making this step impossible. However, one can still generate such a Bernoulli variate, using a so-called *Bernoulli factory* (Keane and O’Brien, 1994). Briefly, a Bernoulli factory is an algorithm that outputs a Bernoulli variate with success probability $f(p)$, from i.i.d. Bernoulli(p) variates, when f is known but p is unknown.

The rejection sampling approach requires a Bernoulli factory algorithm for $f(p) = ap$ and $a \in (1, \infty)$. Nacu and Peres (2005) and Latuszynski et al. (2011) provide an algorithm when $a = 2$, but their algorithms are computationally

demanding and scale poorly for $a \in (1, \infty)$. For example when $p \in (0, .4)$, one requires at least 65,536 Bernoulli(p) random variables to generate one Bernoulli($2p$) variate. In this paper, we provide an algorithm for any $a \in (1, \infty)$ that reduces the computational time substantially. For example when $p \in (0, .4)$, we can obtain a Bernoulli($2p$) variate with only 256 Bernoulli(p) random variables. This is an important reduction because the Bernoulli factory accounts for much of the computational time in the exact sampling algorithm. Section 3 contains a full description the Bernoulli factory and our modification.

Our solutions to (i) and (ii) yield an exact sampling algorithm for π . The algorithm is suitable even for intractable distributions on general state spaces: that is, for distributions that typically arise in statistical applications. This is an important extension, since very few existing algorithms apply to general state spaces, but it is limited in the sense that one must be able to establish a drift and associated minorization condition for the underlying Markov chain. The current algorithm can be computationally demanding, however we have successfully implemented it in three examples.

Our first example considers a univariate Metropolis-Hastings sampler for which we obtain 1000 i.i.d. draws. The second example considers a slightly more complicated bivariate Gibbs sampler where we again obtain 1000 i.i.d. draws. These two examples could be considered toy examples in the sense that i.i.d. observations are available for each. However, they provide insights into the performance and hyper-parameter selection of the algorithm.

Our final example considers a Bayesian version of the classical one-way random effects model that is widely used to analyze data. We illustrate the exact sampling algorithm, using data from a styrene exposure study, to obtain 20 i.i.d. draws. This is the first successful implementation of an exact sampling algorithm for a model of this type. Our analysis considers a balanced design and requires development of a suitable drift condition, which improves upon the existing drift constants of [Tan and Hobert \(2009\)](#).

These examples give hope for exact sampling algorithms for general state space Markov chains in more complicated settings. Even if we are unable to obtain multiple draws in these settings, a single exact draw will alleviate the need for burn-in entirely.

The rest of this paper is organized as follows. Section 2 provides the mixture representation of π , details the rejection sampling approach and bounds the tail probabilities of the proposal distribution. Section 3 introduces the Bernoulli factory and proposes a new target function that speeds up the algorithm significantly. Section 4 gives the full exact sampling algorithm. Sections 5 and 6 implement the algorithm for two toy examples and a Bayesian version of a one-way random effects model, respectively. Finally, Section 7 discusses our implementation and provides some general recommendations to practitioners.

2. Exact sampling via a mixture distribution

Suppose we want to explore the intractable probability measure $\pi(dx)$ defined on the measurable space $(X, \mathcal{B}(X))$. Let $P : X \times \mathcal{B}(X) \rightarrow [0, 1]$ be a Markov

transition function and let $X = \{X_n\}_{n=0}^\infty$ denote the corresponding Markov chain. Then for $x \in \mathsf{X}$ and a measurable set A ,

$$P(x, A) = \Pr(X_{n+1} \in A | X_n = x) .$$

Assume that π is an invariant measure for the chain; i.e., $\pi(A) = \int_{\mathsf{X}} P(x, A) \pi(dx)$ for all measurable A . Assume further that X satisfies the usual regularity conditions, which are irreducibility, aperiodicity and positive Harris recurrence. For definitions, see [Meyn and Tweedie \(1993\)](#) and [Roberts and Rosenthal \(2004\)](#). Finally, assume we are able to simulate the chain; that is, given $X_n = x$, we have the ability to draw from $P(x, \cdot)$.

The exact sampling algorithm considered here utilizes a mixture representation for π ([Asmussen et al., 1992](#); [Hobert et al., 2006](#); [Hobert and Robert, 2004](#)), however, we must first develop the split chain. The main assumption necessary is that X satisfies a one-step minorization condition, i.e. there exists a function $s : \mathsf{X} \rightarrow [0, 1]$ satisfying $\int_{\mathsf{X}} s(x) \pi(dx) > 0$ and some probability measure $Q(dy)$ on $(\mathsf{X}, \mathcal{B}(\mathsf{X}))$ such that,

$$P(x, A) \geq s(x) Q(A) \text{ for all } x \in \mathsf{X} \text{ and } A \in \mathcal{B}(\mathsf{X}) . \quad (1)$$

Given X satisfies the one-step minorization condition at (1), then P can be decomposed as

$$P(x, dy) = s(x) Q(dy) + (1 - s(x)) R(x, dy) , \quad (2)$$

where

$$R(x, dy) = \frac{P(x, dy) - s(x) Q(dy)}{1 - s(x)} ,$$

and define $R(x, dy) = 0$ if $s(x) = 1$. It is helpful to think of (2) as a mixture of two Markov transition functions with probabilities $s(x)$ and $1 - s(x)$. Equation (2) shows that it is possible to simulate X_{n+1} given $X_n = x$ as follows: Flip a coin (independently) that comes up heads with probability $s(x)$. If the coin is a head, take $X_{n+1} \sim Q(\cdot)$; if it's a tail, take $X_{n+1} \sim R(x, \cdot)$.

This decomposition has several important applications in MCMC. Indeed, it can be used to perform regenerative simulation ([Hobert et al., 2002](#); [Mykland et al., 1995](#)) and to derive computable bounds on the convergence rate of X ([Lund and Tweedie, 1996](#); [Roberts and Tweedie, 1999](#); [Rosenthal, 1995, 2002](#)).

Now consider a new Markov chain that actually includes the coin flips mentioned above. Let $X' = \{(X_n, \delta_n)\}_{n=0}^\infty$ be a Markov chain with state space $\mathsf{X} \times \{0, 1\}$. If the current state is $(X_n, \delta_n) = (x, \delta)$, then the next state, (X_{n+1}, δ_{n+1}) , is drawn as follows. If $\delta = 1$, then $X_{n+1} \sim Q(\cdot)$; while if $\delta = 0$, $X_{n+1} \sim R(x, \cdot)$. Then, conditional on $X_{n+1} = x'$, $\delta_{n+1} \sim \text{Bernoulli}(s(x'))$. This chain is called the split chain. Equation (2) implies that, marginally, the sequence of X_n values in the split chain has the same overall probability law as the original Markov chain X ([Nummelin, 1984](#), Chapter 4). Note that, if $\delta_n = 1$, then the distribution of (X_{n+1}, δ_{n+1}) does not depend on x .

Remark 1. We can avoid drawing from $Q(\cdot)$ entirely by changing the order slightly. Given x is the current state, we can simply generate $X_{i+1} \sim P(x, \cdot)$ in the usual manner and then generate $\delta_i | X_i, X_{i+1}$ with

$$\Pr(\delta_i = 1 | X_i, X_{i+1}) = \frac{s(X_i)q(X_{i+1})}{k(X_{i+1}|X_i)}, \quad (3)$$

where $q(\cdot)$ and $k(\cdot|x)$ are the densities corresponding to $Q(\cdot)$ and P (Nummelin, 1984, p. 62).

2.1. A mixture representation of π

The reason for introducing the split chain is that it possesses an accessible atom, $X \times \{1\}$. Indeed, each time the set $X \times \{1\}$ is entered, the split chain stochastically restarts itself (because the next X_n has distribution Q). Let $\mathbb{N} = \{1, 2, 3, \dots\}$ and $(X_0, \delta_0) \in X \times \{1\}$, then define the first return time to the atom as

$$\tau = \min \{n \in \mathbb{N} : (X_n, \delta_n) \in X \times \{1\}\}.$$

Our assumptions about P imply that $E(\tau) < \infty$ and hence the sequence $\{p_n\}_{n=1}^{\infty}$ defined by

$$p_n = \frac{\Pr(\tau \geq n)}{E(\tau)}$$

is nonnegative, nonincreasing, and sums to one. Let T denote a discrete random variable on \mathbb{N} with $\Pr(T = n) = p_n$; also let Q_n be the conditional distribution of X_n given that the split chain does not return to the atom before time n . Thus for any $n \in \mathbb{N}$ and any measurable A , $Q_n(A) = \Pr(X_n \in A | \tau \geq n)$. (Note that $Q_1 \equiv Q$.) Then π can be written as the following mixture of the Q_n values:

$$\pi(dx) = \sum_{n=1}^{\infty} p_n Q_n(dx). \quad (4)$$

Remark 2. The representation at (4) can be obtained from results in Asmussen et al. (1992) by applying their methods to the split chain. Alternatively, Hobert and Robert (2004) obtain the representation when $s(x)$ has the specific form $\varepsilon I_C(x)$ and Hobert et al. (2006) obtain the representation with the more general minorization shown here.

The representation at (4) offers an alternative sampling scheme for π . First, make a random draw from the set $\{Q_1, Q_2, Q_3, \dots\}$ according to the probabilities p_1, p_2, p_3, \dots and then make an independent random draw from the chosen Q_n .

Sampling algorithm for π :

1. Draw T such that $\Pr(T = n) = p_n$ for $n = 1, 2, 3, \dots$, call the result t .
 2. Make a draw from $Q_t(\cdot)$.
-

Drawing from Q_n is simple even when $n \geq 2$. Indeed, just repeatedly simulate n iterations of the split chain, and accept X_n the first time that $\delta_1 = \dots = \delta_{n-1} = 0$. The challenging part of this recipe is drawing from the set $\{Q_1, Q_2, Q_3, \dots\}$, i.e. simulating a random variable T , since the p_n values are not computable. [Hobert et al. \(2006\)](#) and [Hobert and Robert \(2004\)](#) approximate the p_n values, which yields approximate draws from T and thus π . In this paper, we obtain exact draws from T that result in exact draws from π .

Remark 3. Let $K|T = n$ be the number of simulations of the split chain before we get a draw from Q_n , conditional on $T = n$ from Step 1 above. Then $K|T = n$ is geometric with mean $E(K|T = n) = 1/P(\tau \geq n)$. Unfortunately, [Blanchet and Meng \(2005\)](#) show $E(K) = \infty$, and justifiably argue (4) should not be used for multiple replications. This presents a major challenge in the applicability of our algorithm and others that can be similarly expressed, some of which are discussed in the next section.

2.2. Rejection sampler for T

There is one case where simulating T is simple ([Hobert et al., 2006](#); [Hobert and Robert, 2004](#)). Suppose that in the minorization condition (1), $s(x) \equiv \varepsilon > 0$ for all $x \in X$ (implying the Markov chain is uniformly ergodic) and consider the procedure for simulating the split chain with this constant s . In particular, note that the coin flip determining whether δ_n is 0 or 1 does not depend on x , and it follows that the number of steps until the first return to the accessible atom has a geometric distribution. Indeed, $\Pr(\tau \geq n) = (1 - \varepsilon)^{n-1}$. Hence, $E(\tau) = 1/\varepsilon$ and

$$p_n = \Pr(T = n) = \frac{\Pr(\tau \geq n)}{E(\tau)} = \varepsilon(1 - \varepsilon)^{n-1},$$

so T also has a geometric distribution. Therefore it is easy to make exact draws from π .

[Hobert and Robert \(2004\)](#) show this exact sampling algorithm is equivalent to [Murdoch and Green's \(1998\)](#) Multigamma Coupler and to [Wilson's \(2000\)](#) Read-Once algorithm. It is interesting that (4) can be used to reconstruct perfect sampling algorithms based on coupling from the past despite the fact that its derivation involves no backward simulation arguments. Of course, this exact sampling algorithm will be useless from a practical standpoint if ε is too small.

Unfortunately, in statistical inference problems, the MCMC algorithms are usually driven by Markov chains that are not uniformly ergodic and, hence, cannot satisfy (1) with a constant s . Moreover, there is no efficient method to simulate T where s is non-constant. (When s is non-constant, the distribution of τ is complex and its mass function is not available in closed form. Hence, the mass function of T is also unknown, which precludes direct simulation of T .) Therefore we must resort to indirect methods of simulating T .

Fortunately, simulating τ is trivial—indeed, one can simply run the split chain and count the number of steps until it returns to $X \times \{1\}$. Because this provides an unlimited supply of i.i.d. copies of τ , we can use a rejection sampling

approach (Asmussen et al., 1992; Blanchet and Meng, 2005) to simulate T from the i.i.d. sequence τ_1, τ_2, \dots (where $\tau_1 \stackrel{d}{=} \tau$).

Suppose there exists a function $d : \mathbb{N} \rightarrow [0, 1]$ such that $\sum_{n=1}^{\infty} d(n) = D < \infty$ and $P(\tau \geq n) \leq Md(n)$ where M is a finite, positive constant. Consider a rejection sampler with candidate mass function $d(\cdot)/D$. Thus

$$\frac{\Pr(T = n)}{d(n)/D} = \frac{\Pr(\tau \geq n)/\mathbb{E}(\tau)}{d(n)/D} = \frac{D}{\mathbb{E}(\tau)} \frac{\Pr(\tau \geq n)}{d(n)} \leq \frac{D}{\mathbb{E}(\tau)} M,$$

which justifies the following rejection sampler.

Rejection sampler for simulating T :

1. Draw $T \sim \frac{d(\cdot)}{D}$. Call the result n and let $a = 1/[Md(n)]$.
 2. Draw an independent Bernoulli random variable, B , with success probability $a \Pr(\tau \geq n)$. If $B = 1$, accept $T = n$; if $B = 0$, return to Step 1.
-

Unfortunately, the standard method of simulating B (by computing $a \Pr(\tau \geq n)$ and comparing it to an independent $\text{Uniform}(0, 1)$ random variable) is not available to us because the mass function of τ is unavailable in closed form. However, we may draw B without knowing the value of $\Pr(\tau \geq n)$ using a supply of i.i.d. copies of τ . This is the basis of our exact sampling approach.

Suppose $a \in (0, 1]$ and let $p = \Pr(\tau \geq n)$, then there exists a simple solution to generate $B \sim \text{Bernoulli}(ap)$, which Fill (1998) calls “engineering a coin flip”. Indeed simulate a single τ and define

$$W = \begin{cases} 1 & \text{if } \tau \geq n \\ 0 & \text{if } \tau < n \end{cases},$$

hence $W \sim \text{Bernoulli}(p)$. If we independently simulate $V \sim \text{Bernoulli}(a)$ as usual and set $B = VW$, then

$$\Pr(B = 1) = \Pr([V = 1] \cap [W = 1]) = \Pr(V = 1) \Pr(W = 1) = ap.$$

That is, $B \sim \text{Bernoulli}(ap)$, obtained by simulating a single τ and a single Bernoulli V .

When $a \in (1, \infty)$, we will obtain $B \sim \text{Bernoulli}(ap)$ via the Bernoulli factory described in Section 3. For now assume such a simulation is possible, then what remains to establish is a computable tail probability bound, i.e. the sequence $d(n)$ and the constant M .

2.3. Tail probability bound

Blanchet and Meng (2005) bound the moments of τ , however they do not explicitly determine computable values M and $d(n)$. Fortunately, $\Pr(\tau \geq n)$ can be bounded above by a known constant times a known geometric mass function if X satisfies a geometric drift and associated one-step minorization conditions.

We will say a drift condition holds if there exists some function $V : \mathsf{X} \mapsto [1, \infty)$, some $0 < \lambda < 1$ and some $b < \infty$, such that

$$E[V(X_{i+1})|X_i = x] \leq \lambda V(x) + I_{(x \in C)} b \quad \text{for all } x \in \mathsf{X}. \quad (5)$$

In addition, we require the associated one-step minorization condition as follows; assume that $s(x)$ is bounded below by ε on C and that

$$\Pr(x, A) \geq \varepsilon Q(A) \quad \text{for all } x \in C, \quad A \in \mathcal{B}(\mathsf{X}). \quad (6)$$

[Hobert and Robert \(2004\)](#) provide the following bound on $\Pr(\tau \geq n)$ based results in [Roberts and Tweedie \(1999\)](#). Define $A = \sup_{x \in C} E[V(X_{i+1})|X_i = x]$, $J = (A - \varepsilon)/\lambda$, and

$$\beta^* = \begin{cases} \lambda^{-1} & \text{if } J < 1, \\ \exp \left\{ \frac{\log \lambda \log(1 - \varepsilon)}{\log J - \log(1 - \varepsilon)} \right\} \leq \lambda^{-1} & \text{if } J \geq 1. \end{cases}$$

Then letting $\phi(\beta) = \log \beta / \log \lambda^{-1}$, if $\beta \in (1, \beta^*)$, we have

$$\begin{aligned} \Pr(\tau \geq n) &\leq \beta \left[\frac{b}{\varepsilon(1 - \lambda)} \right]^{\phi(\beta)} \left[\frac{1 - \beta(1 - \varepsilon)}{1 - (1 - \varepsilon)(J/(1 - \varepsilon))^{\phi(\beta)}} \right] \beta^{-n} \\ &= Md(n), \end{aligned} \quad (7)$$

where $d(n) = \beta^{-n}$ and

$$M = \beta \left[\frac{b}{\varepsilon(1 - \lambda)} \right]^{\phi(\beta)} \left[\frac{1 - \beta(1 - \varepsilon)}{1 - (1 - \varepsilon)(J/(1 - \varepsilon))^{\phi(\beta)}} \right].$$

Note $\sum_{n=1}^{\infty} d(n) = \sum_{n=1}^{\infty} \beta^{-n} = \frac{1}{\beta - 1} = D < \infty$ since $\beta \in (1, \beta^*)$. Having established the inequality in (7), we next detail how to generate $B \sim \text{Bernoulli}(ap)$ when $a > 1$.

3. Bernoulli factory

Given a sequence $W = \{W_n\}_{n \geq 1}$ of i.i.d. $\text{Bernoulli}(p)$ random variables, where p is unknown, a Bernoulli factory is an algorithm that simulates a random variable $B \sim \text{Bernoulli}(f(p))$, where f is a known function. For the exact sampling algorithm, we require a Bernoulli factory where $f(p) = ap$. This idea arose in [Asmussen et al. \(1992\)](#) when proposing an exact sampling algorithm for general regenerative processes.

Consider $f : S \mapsto [0, 1]$, where $S \subset (0, 1)$. [Keane and O'Brien \(1994\)](#) show it is possible to simulate a random variable $B \sim \text{Bernoulli}(f(p))$ for all $p \in S$ if and only if f is constant, or f is continuous and satisfies, for some $n \geq 1$,

$$\min\{f(p), 1 - f(p)\} \geq \min\{p, 1 - p\}^n \quad \forall p \in S. \quad (8)$$

While [Keane and O'Brien \(1994\)](#) develop the necessary and sufficient conditions on f , they do not provide a detailed description of an algorithm. [Nacu and Peres](#)

(2005) suggest a constructive algorithm via Bernstein polynomials for fast simulation, i.e. the number of input Bernoulli(p) variates needed for the algorithm has exponentially bounded tails. However, we find no practical implementation since it requires dealing with sets of exponential size. Our approach is based on the recent work of Latuszynski et al. (2011), which avoids keeping track of large sets by introducing a single auxiliary random variable.

The general approach, in the formulation of Latuszynski et al. (2011), is to construct two random approximations to $f(p)$, denoted U_n and L_n , which depend on W_1, W_2, \dots, W_n and satisfy

$$1 \geq U_n = U_n(W_1, \dots, W_n) \geq U_{n+1} \geq L_{n+1} \geq L_n = L_n(W_1, \dots, W_n) \geq 0 \text{ a.s.} \quad (9)$$

The random variables U_n and L_n approximate $f(p)$ in the sense that $E(U_n) \searrow f(p)$ and $E(L_n) \nearrow f(p)$ as $n \rightarrow \infty$. The decision to continue sampling or output a zero or a one in the Bernoulli factory is made using an auxiliary Uniform(0, 1) variable.

Remark 4. The almost sure monotonicity requirement in (9) is typically difficult to attain and thus Latuszynski et al. (2011) relax it by using super/sub-martingales instead.

3.1. Modified target function

For the rejection sampling approach to simulating T , we have the ability to simulate W by setting $W_i = I(\tau_i \geq n)$ for $i \geq 1$. We require a single $B \sim \text{Bernoulli}(ap)$, where $a = 1/[Md(n)]$ is a known constant such that $a > 0$. The outcome B determines if we accept or reject the proposed value. For $a \in (0, 1]$ we use the simple solution in Section 2 and for $a \in (1, \infty)$ we use the Bernoulli factory.

Unfortunately, the function $f(p) = ap$ on $(0, 1/a)$ does not satisfy (8) and cannot be simulated via the Bernoulli factory. However, when restricted to $f(p) = \min\{ap, 1 - \omega\}$ for $\omega > 0$, such a simulation is possible.

Nacu and Peres (2005) and Latuszynski et al. (2011) provide a detailed algorithm for $a = 2$ and $0 < \omega < 1/4$. Their construction requires a minimum of 65,536 input variables (see Table 1) before the requirement $U_n \leq 1$ at (9) is met. This is due to the fact that $f(p) = \min\{2p, 1 - \omega\}$ is not differentiable and the Bernstein polynomials can approximate general Lipschitz functions at a rate of $1/\sqrt{n}$ (see part (i) of Lemma 6 from Nacu and Peres, 2005). However, when the target function f is twice differentiable, the rate increases to $1/n$ (see part (ii) of the same Lemma).

This suggests the number of Bernoulli(p) input variates required may decrease significantly by using a twice differentiable f . With this in mind, we propose extending ap smoothly from $[0, \frac{1-\omega}{a}]$ to $[0, 1]$. Fix $\delta < \omega$ and consider the following function

$$F : \left[0, 1 - \frac{1-\omega}{a}\right] \rightarrow [0, \delta] \quad F(p) = \delta \int_0^{ap/\delta} e^{-t^2} dt,$$

which is bounded by δ and twice differentiable, with $F'(p) = a \exp\{-a^2 p^2 / \delta^2\}$ and $F''(p) = -2p \frac{a^3}{\delta^2} \exp\{-a^2 p^2 / \delta^2\} \leq 0$. Standard calculus also gives $|F''(p)| \leq a^2 \frac{\sqrt{2}}{\delta \sqrt{e}}$. (F is related to the Gauss error function (erf), though it can be simply calculated from a standard normal distribution function.)

Then define our target function f as

$$f(p) = \begin{cases} ap & \text{if } p \in \left[0, \frac{1-\omega}{a}\right) \\ (1-\omega) + F\left(p - \frac{1-\omega}{a}\right) & \text{if } p \in \left[\frac{1-\omega}{a}, 1\right] \end{cases}. \quad (10)$$

In other words, we have extended ap such that f defined at (10) is twice differentiable with $|f''| \leq C \equiv a^2 \frac{\sqrt{2}}{\delta \sqrt{e}}$. Define $a(n, k) = f(k/n)$ and $b(n, k) = a(n, k) + C/(2n)$ using f at (10), then we can state Algorithm 4 of [Latuszynski et al. \(2011\)](#) with our modification.

Algorithm I.

1. Simulate $G_0 \sim \text{Uniform}(0, 1)$.
2. Compute $m = \min\{m \in \mathbb{N} : b(2^m, 2^m) \leq 1\}$. Set $n = 2^m$, $\tilde{L}_{\frac{n}{2}} = 0$ and $\tilde{U}_{\frac{n}{2}} = 1$.
3. Compute $H_n = \sum_{i=1}^n W_i$, $L_n = a(n, H_n)$ and $U_n = b(n, H_n)$.
4. Compute

$$L_n^* = \sum_{i=0}^{H_n} \frac{\binom{n - \frac{n}{2}}{H_n - i} \binom{\frac{n}{2}}{i}}{\binom{n}{H_n}} a\left(\frac{n}{2}, i\right) \quad \text{and} \quad U_n^* = \sum_{i=0}^{H_n} \frac{\binom{n - \frac{n}{2}}{H_n - i} \binom{\frac{n}{2}}{i}}{\binom{n}{H_n}} b\left(\frac{n}{2}, i\right).$$

5. Compute

$$\tilde{L}_n = \tilde{L}_{\frac{n}{2}} + \frac{L_n - L_n^*}{U_n^* - L_n^*} \left(\tilde{U}_{\frac{n}{2}} - \tilde{L}_{\frac{n}{2}}\right) \quad \text{and} \quad \tilde{U}_n = \tilde{U}_{\frac{n}{2}} - \frac{U_n^* - U_n}{U_n^* - L_n^*} \left(\tilde{U}_{\frac{n}{2}} - \tilde{L}_{\frac{n}{2}}\right).$$

6. If $G_0 \leq \tilde{L}_n$ set $B = 1$; if $G_0 \geq \tilde{U}_n$ set $B = 0$.
7. If $\tilde{L}_n < G_0 < \tilde{U}_n$, set $n = 2n$, return to step 3.
8. Output B .

Theorem 1. *Suppose a , δ and ω are constants such that $0 < a < \infty$ and $0 < \delta < \omega < 1$. Further suppose $f(p)$ as defined at (10), $ap \in [0, 1 - \omega]$ and $W = \{W_n\}_{n \geq 1}$ are i.i.d. Bernoulli(p). Then Algorithm I outputs $B \sim \text{Bernoulli}(ap)$. Moreover the probability that it needs $N > n$ iterations equals $a^2 / n\delta\sqrt{2e}$.*

Proof. See Appendix A. □

Remark 5. Note that the probability Algorithm I needs $N > n$ iterations is independent of the unknown value of p . Hence the number of Bernoulli(p) variates required will also be independent of p .

TABLE 1
 Comparison of count of input Bernoulli(p) variates to implement the Bernoulli factory.

a	2 (N&P)	2	5	10	20
Minimum	65536	256	2048	8192	32768
Mean Count	65536	562.9	2439.8	10373	43771
S.D. Count	0	21046	7287.6	54836	3.908e5

Algorithm 1 provides a constructive algorithm for $a \in (1, \infty)$ and reduces the number of input variables by a factor of over 100, which we demonstrate in the following example. This is a critical improvement since the Bernoulli factory accounts for most of the computational demands of the exact sampling algorithm.

3.2. Bernoulli factory example

Consider generating 10,000 Bernoulli(ap) variates for various values of a while setting $p = 0.01$, $\omega = 1/5$ and $\delta = 1/6$. Table 1 displays the minimum number of Bernoulli(p) variates required, along with the observed mean and standard deviation for the count of Bernoulli(p) variates used to generate 10,000 Bernoulli(ap) variates. We can see for $a = 2$ the minimum and observed mean have been reduced substantially when comparing our target function with the [Nacu and Peres \(2005\)](#) target (N&P). The reduction in observed input Bernoulli(p) variates represents a 120 times reduction in computational time. The N&P implementation always stops the simulation at the minimum, and hence the standard deviation of the count is 0. Table 1 also shows the input variates required increases as a increases. Simulations for other p values (not shown) provide very similar results.

4. Exact sampling algorithm

The Bernoulli factory algorithm, Theorem 1, requires $ap < 1 - \omega$. To this end, let $\kappa > 1$ such that $1/\kappa < 1 - \omega$ and from (7)

$$\Pr(\tau \geq n) \leq Md(n) < Md(n)\kappa.$$

Then letting $a = 1/\lceil Md(n)\kappa \rceil$ we have $a\Pr(\tau \geq n) \leq 1/\kappa < 1 - \omega < 1$ for all n . The following algorithm results in exact draws from π .

Exact sampling algorithm for π :

1. Draw $T^* \sim \text{Geometric}(1 - 1/\beta)$, i.e. $\Pr(T^* = k) = (1/\beta)^{k-1}(1 - 1/\beta)$ for $k = 1, 2, \dots$, and call the result n . Set $a = 1/\lceil Md(n)\kappa \rceil$.
 2. If $a \leq 1$, draw a single τ random variable. Let $W = I(\tau \geq n)$ and independently draw $V \sim \text{Bernoulli}(a)$. Set $B = VW$.
 3. If $a > 1$, use the Bernoulli factory to obtain B , a single Bernoulli random variable with success probability $p = a\Pr(\tau \geq n)$.
 4. If $B = 1$, accept $T^* = n$; if $B = 0$, return to Step 1.
 5. Make a draw from $Q_n(\cdot)$.
-

The algorithm requires selection of ε , C , λ , and κ from a range of possible values (depending on the drift and minorization). Further, $\beta \in (1, \beta^*)$ must be selected depending on the previously selected parameters. Each selection impacts the algorithm performance and we suggest investigation of different settings for a given example. Our examples in Sections 5 and 6 discuss hyperparameter selection and provide further recommendations.

5. Toy examples

This section contains two toy examples in the sense that we can obtain i.i.d. samples for each, hence there is no practical reason for considering MCMC algorithms. The purpose is to gain insights into the exact sampler and study its performance.

5.1. Metropolis-Hastings example

This section illustrates the exact sampling algorithm for a Metropolis-Hastings sampler. Suppose that $\mathsf{X} = [0, \infty)$ and $\pi(dx) = f_X(x) dx$ where $f_X(x) = e^{-x} I(x \geq 0)$. Consider the function $V(x) = e^{cx}$ for some $c > 0$ and suppose we use a Metropolis sampler with a symmetric proposal density $g(\cdot|x)$, which is supported on $[x - \gamma, x + \gamma]$, $\gamma > 0$. Then, for $x > \gamma$,

$$\begin{aligned} E[V(X_{i+1})|X_i = x] &= \int_{x-\gamma}^x V(y)g(y|x)dy + \int_x^{x+\gamma} V(y)g(y|x)dy \frac{f(y)}{f(x)} \\ &\quad + \int_x^{x+\gamma} V(x)g(y|x)dy \left(1 - \frac{f(y)}{f(x)}\right) \\ &= \int_x^{x+\gamma} \left(V(2x-y) + V(y) \frac{f(y)}{f(x)} + V(x) \left(1 - \frac{f(y)}{f(x)}\right) \right) g(y|x)dy \\ &= V(x) \int_x^{x+\gamma} \left(e^{-c(y-x)} + e^{(c-1)(y-x)} + 1 - e^{-(y-x)} \right) g(y|x)dy. \end{aligned}$$

Selecting g to be the uniform density $g(y|x) = \frac{1}{2\gamma} I_{y \in [x-\gamma, x+\gamma]}$, we get

$$E[V(X_{i+1})|X_i = x] = V(x) \frac{1}{2\gamma} \int_0^\gamma \left(e^{-cz} + e^{(c-1)z} + 1 - e^{-z} \right) dz. \quad (11)$$

When $x \in [0, \gamma]$,

$$\begin{aligned} E[V(X_{i+1})|X_i = x] &= \int_0^x V(y)g(y|x)dy + \int_x^{x+\gamma} V(y)g(y|x)dy \frac{f(y)}{f(x)} \\ &\quad + \int_x^{x+\gamma} V(x)g(y|x)dy \left(1 - \frac{f(y)}{f(x)}\right) \\ &= \int_0^x e^{cy} \frac{1}{2\gamma} dy + \frac{e^{cx}}{2\gamma} \int_x^{x+1} e^{(c-1)(y-x)} dy \\ &\quad + \frac{e^{cx}}{2\gamma} \int_x^{x+1} \left(1 - e^{-(y-x)}\right) dy. \end{aligned} \quad (12)$$

When combined, (11) and (12) obtain a drift condition for the Metropolis-Hastings sampler considered. However, the selection of the constants c and γ is crucial to obtaining a reasonable computation time. The constants β and M , which are described in Section 2.3, also depend on c and γ . Based on our example in Section 3, our strategy is to maximize β , which in turn results in small values for $a = 1/[Md(n)\kappa]$. Figure 1(a) displays a contour map of β^* for $c < 0.3$ and $\gamma < 10$. Based on this plot, we select $c = 0.028$ and $\gamma = 4$, resulting in $\beta = 1.0243$ (as seen below). Evaluating the integrals in (11) and (12) gives the drift condition

$$E[V(X_{i+1})|X_i = x] \leq \lambda V(x) + I_{(x \in C)} b$$

where $\lambda = 0.977$, $b = 0.1$ and the small set $C = [0, 4]$. The interval $[0, \gamma]$ is indeed a small set, since, when $x \in [0, \gamma]$,

$$\begin{aligned} P(x, dy) &\geq g(y|x)dy \min \left\{ 1, \frac{f(y)}{f(x)} \right\} = \frac{1}{2\gamma} I_{y \in [0, x+\gamma]} dy \min \left\{ 1, \frac{f(y)}{f(x)} \right\} \\ &\geq \frac{1}{2\gamma} I_{y \in [0, \gamma]} e^{-y} dy \\ &= \frac{1 - e^{-\gamma}}{2\gamma} \left(\frac{1}{1 - e^{-\gamma}} e^{-y} I_{y \in [0, \gamma]} dy \right). \end{aligned}$$

This establishes the necessary minorization condition $P(x, dy) \geq s(x)Q(dy)$ where

$$Q(dy) = q(y)dy = \frac{1}{1 - e^{-\gamma}} e^{-y} I_{y \in [0, \gamma]} dy \text{ and } s(x) = \frac{1 - e^{-\gamma}}{2\gamma} I_{x \in [0, \gamma]}.$$

The remaining numerical elements required for the geometric probability bound are: $A = \sup_{x \in C} E[V(X_{i+1})|X_i = x] = 1.09197$ and $J = (A - \varepsilon)/\lambda = 0.99283 < 1$. Then $\beta = 1/\lambda = 1.0243$. Finally, the Bernoulli factory hyper-parameters are $\kappa = 5/4$ and $\delta = 1/6$ resulting in $\omega = 0.2$.

Following Mykland et al. (1995), we can now simulate the split chain as follows: (1) draw X_{n+1} from $g(\cdot|X_n)$ and accept it with probability $\min\{1, \frac{f(X_{n+1})}{f(X_n)}\}$; (2) if the candidate in step (1) is rejected, set $\delta_n = 0$, otherwise, generate δ_n as a Bernoulli variate with success probability given by

$$Pr(\delta_n = 1) = \frac{s(X_n)q(X_{n+1})}{g(X_{n+1}|X_n) \min \left\{ 1, \frac{f(X_{n+1})}{f(X_n)} \right\}}.$$

Using the exact sampling algorithm, we generated 1000 i.i.d. Exp(1) random variates. Figure 1(b) shows a Q-Q plot of the observed draws versus the theoretical quantiles of the Exp(1) distribution. During our simulations none of the proposed T^* values which required the use of the Bernoulli factory were accepted. This is due to the fact that our Metropolis-Hastings sampler regenerates very fast, roughly in about 20 moves. Thus for large T^* , when the Bernoulli factory is necessary, the probability $Pr(\tau > T^*)$ is negligible and the Bernoulli factory outputs a zero. Improvements via modified drift and minorization or hyper-parameter selection may improve this situation.

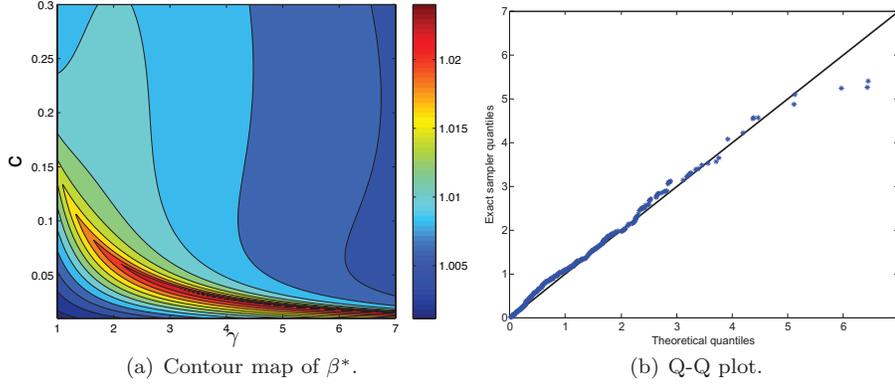


FIG 1. Plots for Metropolis-Hastings example.

5.2. Gibbs example

Suppose $Y_i|\mu, \theta \sim N(\mu, \theta)$ independently for $i = 1, \dots, m$ where $m \geq 3$ and assume the standard invariant prior $\nu(\mu, \theta) \propto \theta^{-1/2}$. The resulting posterior density is

$$\pi(\mu, \theta|y) \propto \theta^{-(m+1)/2} \exp\left\{-\frac{m}{2\theta}(s^2 + (\bar{y} - \mu)^2)\right\} \quad (13)$$

where s^2 is the usual biased sample variance. It is easy to see the full conditional densities, $f(\mu|\theta, y)$ and $f(\theta|\mu, y)$, are given by $\mu|\theta, y \sim N(\bar{y}, \theta/m)$ and $\theta|\mu, y \sim \text{IG}((m-1)/2, m[s^2 + (\bar{y} - \mu)^2]/2)$, hence a Gibbs sampler is appropriate. (We say $W \sim \text{IG}(\alpha, \beta)$ if its density is proportional to $w^{-(\alpha+1)}e^{-\beta/w}I(w > 0)$.) We consider the Gibbs sampler that updates θ then μ ; that is, letting $x' = (\theta', \mu')$ denote the current state and $x = (\theta, \mu)$ denote the future state, the transition looks like $(\theta', \mu') \rightarrow (\theta, \mu') \rightarrow (\theta, \mu)$. This Markov chain then has state space $X = \mathbb{R}^+ \times \mathbb{R}$ and transition density

$$k(\theta, \mu|\theta', \mu') = f(\theta|\mu', y)f(\mu|\theta, y). \quad (14)$$

Appendix B provides a drift and minorization condition using a small set in the form

$$C = \{(\theta, \mu) \in \mathbb{R}^+ \times \mathbb{R} : V(\mu, \theta) \leq d\},$$

where $V(\theta, \mu) = 1 + (\mu - \bar{y})^2$.

Suppose $\bar{y} = 1$, $s^2 = 4$, and $m = 11$. We set $\lambda = 0.5$ and $d = b/(\lambda - (m-3)^{-1}) = 11/3$ resulting in $\varepsilon = 0.5750034$ and $\beta^* = 1.3958$. Given the allowable range $\beta \in (1, \beta^*)$, we select $\beta = 1.35$, which turns out to be extremely important for implementation. Selection of β very close to β^* or close to 1 seems to cause problems for the Bernoulli factory because of large constant multipliers $a = 1/[Md(n)\kappa]$ given typical proposed $T^* = n$ (which depend on β). Experimentation has shown us values somewhat close to β^* seem to provide the best results. Finally, the Bernoulli factory hyper-parameters are $\kappa = 5/4$ and $\delta = 1/6$ resulting in $\omega = 0.2$.

TABLE 2

Summary constants for Gibbs sampler example with $\beta = 1.35$. The row min refers to the minimum number of observed τ s to implement the Bernoulli factory and $p(n) = P(T^* = n)$.

n	1	2	3	4	5	6	7	8	9	10
$p(n)$	0.259	0.192	0.142	0.105	0.078	0.058	0.043	0.032	0.023	0.017
a	0.08	0.11	0.14	0.19	0.26	0.35	0.47	0.64	0.86	1.16
min	-	-	-	-	-	-	-	-	-	128
n	11	12	13	14	15	16	17	18	19	20
$p(n)$	0.013	0.010	0.007	0.005	0.004	0.003	0.002	0.002	0.001	0.001
a	1.57	2.12	2.87	3.87	5.22	7.05	9.52	12.85	17.35	23.42
min	128	256	512	1024	2048	4096	8192	8192	16384	32768

Table 2 summarizes the resulting constants given the hyper-parameter choices. Notice the Bernoulli factory will be necessary for proposed values greater than or equal to 10, that is with probability 0.067. The constants for values greater than 20 are not listed. However, these values are of interest since the minimum number of observed τ values becomes extremely large. While values in this range are uncommon, $P(T^* > 20) \approx 0.002$, they occur with enough frequency to slow down the algorithm substantially.

The exact sampling algorithm was used to generate 1000 exact draws from the posterior density at (13). Generating the 1000 draws required approximately 35 hours of computational time, about 2 minutes per draw. A total of 27,665 T^* values were proposed of which 69 (0.25%) were greater than 20. Implementing the Bernoulli factory required $1.52e9$ τ values, or $1.52e6$ τ values per exact draw. However, most of the computational time, and necessary τ values, were used for a small number of proposed T^* values. Similar to the Metropolis-Hastings example, none of the accepted T^* values were from the Bernoulli factory.

The largest proposed T^* value was 37 with $a \approx 3848$ requiring $1.07e9$ τ values (of which 0 were ≥ 37) to implement the Bernoulli factory. This value alone accounted for about 70% of the total number of τ values, and hence about 70% of the total computational time. It should be noted that the largest accepted T^* value was 9, so proposals unlikely to be accepted account for most of the computational demands. Removing only the largest proposal, the remaining 999 draws required approximately 10 hours of computation, or about 35 seconds per draw.

Alternatively for this example, we can sequentially sample from (13) to obtain i.i.d. draws (Flegal et al., 2008). Figure 2 compares the 1000 exact draws to 1000 i.i.d. draws from a sequential sampler using a Q-Q plot for both μ and θ . We can see from these plots that the exact sampling algorithm is indeed working well.

6. Bayesian random effects model

This section considers a Bayesian version of the one-way random effects model given by

$$Y_{ij} = \phi_i + \zeta_{ij}, \quad i = 1, \dots, q, \quad j = 1, \dots, m_i$$

where the random effects ϕ_i are i.i.d. $N(\mu, \sigma_\phi^2)$ and independently the errors ζ_{ij} are i.i.d. $N(0, \sigma_e^2)$. Thus $(\mu, \sigma_\phi^2, \sigma_e^2)$ is the unknown parameter.

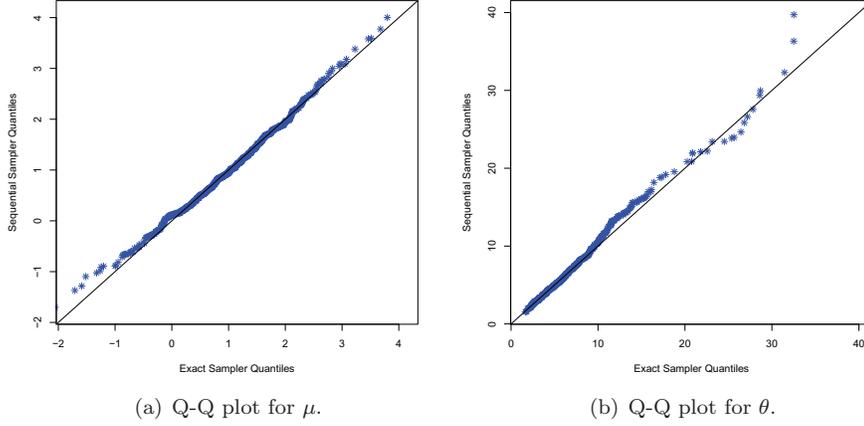


FIG 2. Q-Q plots comparing 1000 exact draws to 1000 i.i.d. draws from a sequential sampler.

Bayesian analysis using this model requires specifying a prior distribution, for which we consider the family of inverse gamma priors

$$\pi(\mu, \sigma_\phi^2, \sigma_e^2) \propto (\sigma_\phi^2)^{-(\alpha_1+1)} e^{-\beta_2/\sigma_\phi^2} (\sigma_e^2)^{-(\alpha_2+1)} e^{-\beta_2/\sigma_e^2}$$

where α_1 , α_2 , β_1 and β_2 are hyper-parameters. If we let $y = \{y_{ij}\}$ and $\phi = \{\phi_i\}$ denote the vectors of observed data and random effects respectively, then the posterior density is as follows

$$\pi(\phi, \mu, \sigma_\phi^2, \sigma_e^2) \propto f(y|\phi, \mu, \sigma_\phi^2, \sigma_e^2) f(\phi|\mu, \sigma_\phi^2, \sigma_e^2) \pi(\mu, \sigma_\phi^2, \sigma_e^2), \quad (15)$$

where

$$f(y|\phi, \mu, \sigma_\phi^2, \sigma_e^2) = \prod_{i=1}^q \prod_{j=1}^{m_i} (2\pi\sigma_e^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2\sigma_e^2} (y_{ij} - \phi_i)^2\right\}$$

and

$$f(\phi|\mu, \sigma_\phi^2, \sigma_e^2) = \prod_{i=1}^q (2\pi\sigma_\phi^2)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2\sigma_\phi^2} (\phi_i - \mu)^2\right\}.$$

For ease of exposition, we will suppress the dependency on the data y and define the usual summary statistics: $\bar{y}_i = m_i^{-1} \sum_j y_{ij}$, $M = \sum_i m_i$, $\bar{y} = M^{-1} \sum_i \sum_j y_{ij}$, $SST = \sum_i m_i (\bar{y}_i - \bar{y})^2$ and $SSE = \sum_i \sum_j (y_{ij} - \bar{y}_i)^2$.

We consider a block Gibbs sampler that updates $\theta = (\sigma_\phi^2, \sigma_e^2)$ then $\xi = (\mu, \phi)$, that is $(\theta', \xi') \rightarrow (\theta, \xi') \rightarrow (\theta, \xi)$. The necessary full conditionals can be obtained via manipulation of (15). That is, $f(\theta|\xi')$ is the product of two inverse gammas such that

$$\sigma_\phi^2|\xi' \sim IG\left(\frac{q}{2} + \alpha_1, \frac{w_1(\xi')}{2} + \beta_1\right)$$

and

$$\sigma_e^2|\xi' \sim IG\left(\frac{M}{2} + \alpha_2, \frac{w_2(\xi') + SSE}{2} + \beta_2\right)$$

where $w_1(\xi) = \sum_{i=1}^q (\phi_i - \mu)^2$ and $w_2(\xi) = \sum m_i (\phi_i - \bar{y}_i)^2$. Further, $f(\xi|\theta)$ is multivariate normal density whose parameters are given in [Tan and Hobert \(2009\)](#). This Markov chain then has state space $\mathsf{X} = \mathbb{R}^{q+1} \times \mathbb{R}_+^2$ and transition density

$$k(\xi, \theta | \xi', \theta') = f(\theta|\xi')f(\xi|\theta) = f(\sigma_\phi^2|\xi')f(\sigma_e^2|\xi')f(\xi|\theta).$$

Implementation of the exact sampling algorithm requires a drift and associated minorization condition as at (5) and (6). [Hobert and Geyer \(1998\)](#), [Jones and Hobert \(2001, 2004\)](#) and [Tan and Hobert \(2009\)](#) analyze variations of the proposed block Gibbs sampler, however none obtain sufficient constants for a practical implementation of our algorithm. To this end, the following theorem improves upon the drift constants of [Tan and Hobert \(2009\)](#) for a balanced design while using a simplified version of their drift function.

Theorem 2. *Let $m_i = m$ for all $i = 1, \dots, q$ and let $\Delta_2 = 1 - 1/[q(m+1)] + \max\{q(m+1)/m^2, 1/m\}$. Further let $V : \mathbb{R}^{q+1} \times \mathbb{R}_+^2 \rightarrow [1, \infty)$ such that $V(\xi, \theta) = K + \delta_1 w_1(\xi) + \delta_2 w_2(\xi)$ where $K \geq 1$, $\delta_1 > 0$ and $\delta_2 > 0$ and define*

$$\lambda^* = \max\left\{\frac{1}{q + 2\alpha_1 - 2}, \frac{\delta_1 \Delta_2 / \delta_2 + q + 1}{M + 2\alpha_2 - 2}\right\}.$$

Then there exists $K \geq 1$, $\delta_1 > 0$ and $\delta_2 > 0$ such that $\lambda^ < 1$ and (5) holds. That is, for any $\lambda \in (\lambda^*, 1)$,*

$$E(V(\xi, \theta)|\xi', \theta') \leq \lambda V(\xi', \theta') + b 1_{(\xi', \theta') \in C} \quad (16)$$

where

$$\begin{aligned} b &= K(1 - \lambda) + \frac{2\delta_1\beta_1}{q + 2\alpha_1 - 2} + \frac{(\delta_1\Delta_2 + \delta_2(q+1))(SSE + 2\beta_2)}{M + 2\alpha_2 - 2} \\ &\quad + (\delta_1 + m\delta_2) \sum_{i=1}^q (\bar{y}_i - \bar{y})^2 \end{aligned}$$

and

$$C = \{(\xi, \theta) \in \mathbb{R}^{q+1} \times \mathbb{R}_+^2 : V(\xi, \theta) \leq d\}$$

where $d = b/(\lambda - \lambda^*)$.

Proof. See [Appendix C](#). □

The drift condition still holds if we increase the small set to

$$C = \{(\xi, \theta) \in \mathbb{R}^{q+1} \times \mathbb{R}_+^2 : K + \delta_1 w_1(\xi) \leq d, K + \delta_2 w_2(\xi) \leq d\},$$

for which [Appendix C](#) provides the associated minorization condition.

6.1. Styrene exposure dataset

We will implement the exact sampling algorithm using the styrene exposure dataset from Lyles et al. (1997) analyzed previously by Jones and Hobert (2001) and Tan and Hobert (2009). The data, summarized in Table 3, is from a balanced design such that $m_i = m = 3$ for $i = 1, \dots, q$, $q = 13$, and $M = mq = 39$.

We consider prior hyper-parameter values $\alpha_1 = \alpha_2 = 0.1$ and $\beta_1 = \beta_2 = 10$. The drift function at (16) requires specification of drift parameters $K = 50$, $\Delta_2 = 6.7585$, $\delta_1 = 1$, and $\delta_2 = 1$ which results in $\lambda^* = 0.5580$. We then choose $\lambda = 0.97$, $b = 37.88927$ and $d = 91.96992$, resulting in $\varepsilon = 0.01269784$.

Using these settings $\beta^* = 1.000092$ and we choose $\beta = 1.000083$ resulting in $M = 10.19413$. We again used Bernoulli factory hyper-parameters of $\kappa = 5/4$ and $\delta = 1/6$ resulting in $\omega = 0.2$. In this case, the Bernoulli factory is necessary for proposals greater than 30,706, approximately 8% of proposed values. It is extremely likely the output from the Bernoulli factory will be zero since a sample of 1000 i.i.d. τ values yielded only a maximum of 1278.

The exact sampling algorithm was run until we obtained a 20 i.i.d. draws from the posterior at (15) which took 31,887 proposed T^* values and 2.61e8 τ values for the Bernoulli factory. The accepted T^* values and i.i.d. θ values are listed in Table 4. Notice the maximum accepted T^* was 287, which is well below

TABLE 3
Styrene exposure data summary statistics.

Worker	1	2	3	4	5	6	7
\bar{y}_i	3.302	4.587	5.052	5.089	4.498	5.186	4.915
Worker	8	9	10	11	12	13	
\bar{y}_i	4.876	5.262	5.009	5.602	4.336	4.813	
$\bar{y} = 4.089$, SST = 11.430, SSE = 14.711							

TABLE 4
List of 20 i.i.d. θ draws from the posterior at (15) with the accepted T^* values.

Draw	T^*	σ_ϕ^2	σ_e^2
1	145	2.477	1.2624
2	18	1.234	1.7698
3	286	3.058	1.8791
4	40	2.607	1.2079
5	76	2.177	2.0603
6	287	6.513	1.2870
7	39	5.961	1.5295
8	103	1.642	1.2093
9	194	2.150	1.8129
10	195	2.112	1.4871
11	101	1.101	1.7166
12	2	1.659	1.0805
13	5	5.544	1.2856
14	9	1.505	1.6600
15	1	2.105	1.4137
16	150	2.681	0.7317
17	63	3.131	1.1506
18	64	3.514	1.8245
19	52	2.119	1.1743
20	62	3.571	1.5009

our observed maximum of 1278 from 1000 i.i.d. τ values. Hence, drawing from Q_n was easy and almost all of the simulation time was used for the Bernoulli factory. Obtaining the 20 i.i.d. draws required 24 days of computational time utilizing six processors in parallel (equating to 144 days on a single processor).

7. Discussion

This paper describes an exact sampling algorithm using a geometrically ergodic Markov chain on a general state space. The algorithm is applicable for any Markov chain where one can establish a drift and associated minorization with computable constants. The limitation of the method is that the simulation time may be prohibitive.

Blanchet and Thomas (2007) implement an approximate version using the Bayesian probit regression example from van Dyk and Meng (2001) with regeneration settings provided by Roy and Hobert (2007). This example is ill-suited using the proposed algorithm because of computational limitations related to the Bernoulli factory and in obtaining a practical ε . Specifically, we found (in simpler examples) obtaining a single draw from π sometimes required millions of i.i.d. τ variates. Unfortunately, even using non-constant $s(x)$, the probit example requires about 14,000 Markov chain draws per τ (Flegal and Jones, 2010). Hence obtaining a single draw from π would require an obscene number of draws from X . Implementation for more complicated Markov chains, such as this, likely requires further improvements, or a lot of patience.

Careful analysis of the Markov chain sampler is necessary to find useful drift and minorization constants. Most research establishing drift and minorization is undertaken to prove geometric ergodicity, in which case the obtained constants are of secondary importance. However, performance of the exact sampling algorithm is heavily dependent on these constants. Improving them may be enough to obtain exact samples in many settings.

Alternatively, the speed of the overall algorithm would improve if one could find a bound using non-constant $s(x)$ or a sharper bound with ε . The current bound at (7) could potentially be modified upon by only considering specific models, or specific classes of models.

Finally, one could obtain further improvements to the Bernoulli factory since it requires most of the necessary τ variates. Our work has already obtained a 100 times reduction in computational time. However there may be further improvements available for the Bernstein polynomial coefficients, modifications to Algorithm 4 of Latuszynski et al. (2011) or an entirely different method to estimate f . Hyper-parameter settings also impact performance and could be investigated further.

Acknowledgments

This work was started during the AMS/NSF sponsored Mathematics Research Communities conference on “Modern Markov Chains and their Statistical Applications,” which was held at the Snowbird Ski and Summer Resort in Snowbird,

Utah between June 28 and July 2, 2009. We are especially grateful to Jim Hobert who proposed this research problem. We would also like to thank fellow conference participants Adam Guetz, Xia Hua, Wai Liu, Yufei Liu, and Vivek Roy. Finally, we are grateful to the anonymous referees, anonymous associate editor, Galin Jones, Ioannis Kosmidis and Krzysztof Latuszyński for their constructive comments and helpful discussions, which resulted in many improvements.

Appendix A: Proof of Theorem 1

Proof. By construction f is a smooth function $f : [0, 1] \rightarrow [0, 1 - \varepsilon]$ for some $0 < \varepsilon < \omega - \delta$. Proposition 3.1 of [Latuszynski et al. \(2011\)](#) and Lemma 6 of [Nacu and Peres \(2005\)](#) prove existence of an algorithm that simulates $f(p)$ if

- (i) f has second derivative f'' which is continuous and
- (ii) the coefficients a and b satisfy

$$a(2n, k) \binom{2n}{k} \geq \sum_{i=0}^k a(n, i) \binom{n}{i} \binom{n}{k-i}, \quad (17)$$

$$b(2n, k) \binom{2n}{k} \leq \sum_{i=0}^k b(n, i) \binom{n}{i} \binom{n}{k-i}. \quad (18)$$

Condition (i) is clearly satisfied by construction, so it remains to check condition (ii). Since the coefficients a and b are defined through f , inequalities (17) and (18) will be checked using the properties of f .

Recall $b(n, k) = a(n, k) + C/(2n)$, the inequalities (17) and (18) above can be re-expressed ([Nacu and Peres, 2005](#)) as

$$a(2n, k) \geq E(a(n, X)) \text{ and } b(2n, k) \leq E(b(n, X)),$$

where X is a hypergeometric random variable, with parameters $(2n, k, n)$. Using the definition of a , and the fact that f is concave, the first inequality is a direct application of Jensen's inequality. The second part is a straight forward application of Lemma 6 from [Nacu and Peres \(2005\)](#) and the properties of the hypergeometric distribution.

Finally, the probability the algorithm needs $N > n$ follows directly from definitions of the coefficients a and b and Theorem 2.5 of [Latuszynski et al. \(2011\)](#). \square

Appendix B: Toy Gibbs drift and minorization

B.1. Drift condition

Let $X = \{X_n\}_{n \geq 0}$ be the Markov chain corresponding to the Gibbs transition kernel given in (14). Recall $\mathbf{X} = \mathbb{R}^+ \times \mathbb{R}$, $x' = (\theta', \mu')$ denotes the current state

and $x = (\theta, \mu)$ denotes the future state. Jones and Hobert (2001) establish a minorization and Rosenthal (1995)–type drift condition for $m \geq 5$, and hence prove the associated Markov chain is geometrically ergodic. Using their argument we show a Roberts-and-Tweedie-type drift condition at (5) (Roberts and Tweedie, 1999, 2001) holds using the function $V(x) = V(\theta, \mu) = 1 + (\mu - \bar{y})^2$. Conditional independence (based on the update order, see Jones and Hobert, 2001) yields

$$\begin{aligned} E[V(X_{i+1})|X_i = x'] &= E[V(\theta, \mu)|\theta', \mu'] \\ &= E[V(\theta, \mu)|\mu'] \\ &= E\{E[V(\theta, \mu)|\theta]|\mu'\} . \end{aligned}$$

Since $\mu|\theta, y \sim N(\bar{y}, \theta/m)$, the inner expectation is

$$E[V(\theta, \mu)|\theta] = E\left[\left(1 + (\mu - \bar{y})^2\right)|\theta\right] = 1 + \text{Var}(\mu|\theta) = 1 + \frac{\theta}{m} .$$

Then since $\theta|\mu', y \sim \text{IG}((m-1)/2, m[s^2 + (\bar{y} - \mu')^2]/2)$,

$$E[\theta|\mu'] = \frac{m[s^2 + (\bar{y} - \mu')^2]}{m-3} ,$$

and hence

$$E[V(X_{i+1})|X_i = x'] = \frac{1 + (\mu' - \bar{y})^2}{m-3} + \frac{s^2 + m - 4}{m-3} . \quad (19)$$

Let $\lambda \in ((m-3)^{-1}, 1)$, $b = (s^2 + m - 4)/(m-3)$, $d \geq b/(\lambda - (m-3)^{-1})$ and

$$C = \{(\theta, \mu) \in \mathbb{R}^+ \times \mathbb{R} : V(\mu, \theta) \leq d\} ,$$

then the drift condition at (5) is satisfied, that is

$$E[V(X_{i+1})|X_i = x'] \leq \lambda V(x') + I_{(x' \in C)} b \text{ for all } x' \in X .$$

It is easy to see from (19) that

$$A = \sup_{x' \in C} E[V(X_{i+1})|X_i = x] = \frac{\sup_{x' \in C} V(x')}{m-3} + b = \frac{d}{m-3} + b .$$

B.2. Minorization condition

Now we establish the associated minorization condition at (6) using a similar argument to Jones and Hobert (2001). Let $C_\mu = \{\mu \in \mathbb{R} : 1 + (\mu - \bar{y})^2 \leq d\}$, then for any $\mu' \in C_\mu$

$$k(\theta, \mu|\theta', \mu') = f(\theta|\mu')f(\mu|\theta) \geq f(\mu|\theta) \inf_{\mu \in C_\mu} f(\theta|\mu) .$$

Recall $f(\theta|\mu)$ is an IG density, thus $g(\theta) := \inf_{\mu \in C_\mu} f(\theta|\mu)$ can be written in closed form (Jones and Hobert, 2004; Rosenthal, 1996; Tan and Hobert, 2009),

$$\inf_{\mu \in C_\mu} f(\theta|\mu) = \begin{cases} IG\left(\frac{m-1}{2}, \frac{m(s^2+d-1)}{2}; \theta\right) & \text{if } \theta < \theta^* \\ IG\left(\frac{m-1}{2}, \frac{ms^2}{2}; \theta\right) & \text{if } \theta \geq \theta^* \end{cases}$$

where $\theta^* = m(d-1) \left[(m-1) \log\left(1 + \frac{d-1}{s^2}\right) \right]^{-1}$ and $IG(\alpha, \beta; x)$ is the inverse gamma density evaluated at x . If we further define

$$\varepsilon = \int_{\mathbb{R} \times \mathbb{R}_+} f(\mu|\theta) \inf_{\mu \in C_\mu} f(\theta|\mu) d\mu d\theta = \int_{\mathbb{R}_+} \inf_{\mu \in C_\mu} f(\theta|\mu) d\theta$$

and density $q(\theta, \mu) = \varepsilon^{-1} g(\theta) f(\mu|\theta)$, then

$$k(\theta, \mu|\theta', \mu') \geq \varepsilon q(\theta, \mu) .$$

Letting $Q(\cdot)$ be the probability measure associated with the density q , then the minorization condition from (6) holds, that is for any set A and any $(\theta', \mu') \in C$

$$P(x, A) \geq \varepsilon Q(A) \text{ for all } A \in \mathcal{B}(X) .$$

Notice the minorization condition holds for any $d > 0$.

Simulating the split chain requires evaluation of (3),

$$\begin{aligned} Pr(\delta' = 1|\mu', \theta', \mu, \theta) &= \frac{\varepsilon q(\theta, \mu)}{k(\mu, \theta|\mu', \theta')} \\ &= \frac{\varepsilon \varepsilon^{-1} g(\theta) f(\mu|\theta)}{f(\theta|\mu') f(\mu|\theta)} \\ &= \frac{g(\theta)}{f(\theta|\mu')} \\ &= \frac{\left[\frac{m(s^2 + I_{\{\theta < \theta^*\}}(d-1))}{2} \right]^{(m-1)/2}}{\left[\frac{m(s^2 + (\bar{y} - \mu')^2)}{2} \right]^{(m-1)/2}} \\ &= \frac{\frac{\theta^{-(m-1)/2-1}}{\Gamma((m-1)/2)} \exp\left\{-\left[\frac{m(s^2 + I_{\{\theta < \theta^*\}}(d-1))}{2\theta}\right]\right\}}{\frac{\theta^{-(m-1)/2-1}}{\Gamma((m-1)/2)} \exp\left\{-\left[\frac{m(s^2 + (\bar{y} - \mu')^2)}{2\theta}\right]\right\}} \\ &= \left[\frac{s^2 + I_{\{\theta < \theta^*\}}(d-1)}{s^2 + (\bar{y} - \mu')^2} \right]^{(m-1)/2} \\ &\quad \cdot \exp\left\{-\frac{m I_{\{\theta < \theta^*\}}(d-1)}{2\theta} + \frac{m(\bar{y} - \mu')^2}{2\theta}\right\} . \end{aligned}$$

Notice that $Pr(\delta'|\mu', \theta', \mu, \theta)$ is free of ε , θ' , and μ .

Appendix C: One-way random effects drift and minorization

C.1. Proof of Theorem 2

Proof. Notice

$$E(V(\xi, \theta)|\xi', \theta') = E\{E(V(\xi, \theta)|\theta)|\xi', \theta'\} \quad (20)$$

where the inner expectation becomes

$$E(V(\xi, \theta)|\theta) = K + \delta_1 E(w_1(\xi)|\theta) + \delta_2 E(w_2(\xi)|\theta).$$

Note that

$$E(w_1(\xi)|\theta) = E\left[\sum_{i=1}^q (\phi_i - \mu)^2|\theta\right] = \sum_{i=1}^q \left\{ \text{Var}(\phi_i - \mu|\theta) + \left(E(\phi_i - \mu|\theta)\right)^2 \right\}.$$

For general designs (balanced or unbalanced) [Tan and Hobert \(2009\)](#) prove

$$\sum_{i=1}^q \text{Var}(\phi_i - \mu|\theta) \leq \Delta_1 \sigma_\phi^2 + \Delta_2 \sigma_e^2$$

where,

$$\Delta_1 = \min \left\{ q \left(\sum_{i=1}^q \frac{m_i}{m_i + 1} \right)^{-1}, \frac{q \cdot \max\{m_1, \dots, m_q\}}{M} \right\}$$

and

$$\Delta_2 = \sum_{i=1}^q \frac{1}{m_i} - \sum_{i=1}^q \frac{1}{M(1 + m_i)} + \max \left\{ q \left(\sum_{i=1}^q \frac{m_i}{m_i + 1} \right)^{-1}, \frac{q}{M} \right\}.$$

Simplifying Δ_1 and Δ_2 under the balanced design, we obtain $\Delta_1 = 1$ and $\Delta_2 = 1 - 1/[q(m + 1)] + \max\{q(m + 1)/m^2, 1/m\}$.

From the full conditionals ([Tan and Hobert, 2009](#))

$$E(\phi_i - \mu|\theta) = \frac{m\sigma_\phi^2}{\sigma_e^2 + m\sigma_\phi^2}(\bar{y}_k - \bar{y}),$$

then it follows that (with $\Delta_1 = 1$)

$$\begin{aligned} E(w_1(\xi)|\theta) &\leq \sigma_\phi^2 + \Delta_2 \sigma_e^2 + \left(\frac{m\sigma_\phi^2}{\sigma_e^2 + m\sigma_\phi^2} \right)^2 \sum_{i=1}^q (\bar{y}_i - \bar{y})^2 \\ &\leq \sigma_\phi^2 + \Delta_2 \sigma_e^2 + \sum_{i=1}^q (\bar{y}_i - \bar{y})^2. \end{aligned}$$

Similarly,

$$E(w_2(\xi)|\theta) \leq (q+1)\sigma_e^2 + m \sum_{i=1}^q (\bar{y}_i - \bar{y})^2$$

To complete the calculation in (20), recall $\sigma_\phi^2|\xi'$ and $\sigma_e^2|\xi'$, thus

$$E(\sigma_\phi^2|\xi') = \frac{w_1(\xi') + 2\beta_1}{q + 2\alpha_1 - 2}$$

and

$$E(\sigma_e^2|\xi') = \frac{w_2(\xi') + SSE + 2\beta_2}{M + 2\alpha_2 - 2}.$$

It follows that

$$\begin{aligned} E(V(\xi, \theta)|\xi', \theta') &\leq K + \frac{1}{q + 2\alpha_1 - 2} \left(\delta_1 w_1(\xi') \right) + \frac{\delta_1 \Delta_2 / \delta_2 + q + 1}{M + 2\alpha_2 - 2} \left(\delta_2 w_2(\xi') \right) \\ &+ \frac{2\delta_1 \beta_1}{q + 2\alpha_1 - 2} + \frac{(\delta_1 \Delta_2 + \delta_2 (q + 1))(SSE + 2\beta_2)}{M + 2\alpha_2 - 2} + \\ &+ (\delta_1 + m\delta_2) \sum_{i=1}^q (\bar{y}_i - \bar{y})^2. \end{aligned} \quad (21)$$

Note there exists $\delta_1 > 0$ and $\delta_2 > 0$ such that $\lambda^* < 1$, hence the drift condition at (5) is satisfied. \square

For the exact sampling algorithm, we can see from (21) and the definitions of C and b that

$$A = \sup_{x' \in C} E[V(X_{i+1})|X_i = x] \leq \frac{(d-K)}{q + 2\alpha_1 - 2} + \frac{\delta_1 \Delta_2 / \delta_2 + q + 1}{M + 2\alpha_2 - 2} (d-K) + b + \lambda^*.$$

C.2. Minorization condition

Next we show the associated minorization condition holds. The argument will be similar to the toy Gibbs example from Appendix B. Let $C_\xi = \{\xi \in \mathbb{R}^{q+1} : 1 + \delta_1 w_1(\xi) \leq d, 1 + \delta_2 w_2(\xi) \leq d\}$, then the associated minorization condition holds if we can find an ε and $q(\xi, \theta)$ such that for any $\xi' \in C_\xi$

$$\begin{aligned} k(\xi, \theta | \xi', \theta') &= f(\sigma_\phi^2|\xi') f(\sigma_e^2|\xi') f(\xi|\theta) \\ &\geq f(\xi|\theta) \inf_{\xi \in C_\xi} f(\sigma_\phi^2|\xi) \inf_{\xi \in C_\xi} f(\sigma_e^2|\xi) \\ &= f(\xi|\theta) g_1(\sigma_\phi^2) g_2(\sigma_e^2) \\ &= \varepsilon q(\xi, \theta). \end{aligned} \quad (22)$$

The two infimums can be found analytically as before:

$$g_1(\sigma_\phi^2) = \inf_{\xi \in C_\xi} f(\sigma_\phi^2|\xi) = \begin{cases} IG\left(\frac{q}{2} + \alpha_1, \frac{d-K}{2\delta_1} + \beta_1; \sigma_\phi^2\right) & \text{if } \sigma_\phi^2 \leq \sigma_\phi^* \\ IG\left(\frac{q}{2} + \alpha_1, \beta_1; \sigma_\phi^2\right) & \text{if } \sigma_\phi^2 > \sigma_\phi^* \end{cases}$$

and

$$g_2(\sigma_e^2) = \inf_{\xi \in C_\xi} f(\sigma_e^2 | \xi) = \begin{cases} IG\left(\frac{M}{2} + \alpha_2, \frac{(d-K)/\delta_2 + SSE}{2} + \beta_2\right) & \text{if } \sigma_e^2 \leq \sigma_e^* \\ IG\left(\frac{M}{2} + \alpha_2, \frac{SSE}{2} + \beta_2\right) & \text{if } \sigma_e^2 > \sigma_e^* \end{cases}.$$

The points σ_ϕ^* and σ_e^* are the intersection points of the two inverse gamma densities determined by

$$\sigma^* = \frac{b_1 - b_2}{a(\log b_1 - \log b_2)}$$

where a , b_1 and b_2 are the parameters of the two inverse gamma distributions.

We can define

$$\begin{aligned} \varepsilon &= \int_{\mathbb{R}^{q+1} \times \mathbb{R}_+^2} f(\xi | \theta) g(\theta) d\xi d\theta = \int_{\mathbb{R}_+^2} g(\theta) d\theta \\ &= \int_{\mathbb{R}_+} g_1(\sigma_\phi^2) d\sigma_\phi^2 \int_{\mathbb{R}_+} g_2(\sigma_e^2) d\sigma_e^2 \end{aligned}$$

and density $q(\xi, \theta) = \varepsilon^{-1} g_1(\sigma_\phi^2) g_2(\sigma_e^2) f(\xi | \theta)$, then (22) holds. Since this minorization condition holds for any $d > 0$, this establishes the associated minorization condition.

Simulating the split chain requires evaluation of (3) similar to the calculation in Appendix B,

$$\begin{aligned} Pr(\delta' | \xi', \theta', \xi, \theta) &= \frac{\varepsilon q(\theta, \xi)}{k(\xi, \theta | \xi', \theta')} \\ &= \frac{g_1(\sigma_\phi^2)}{f(\sigma_\phi^2 | \xi')} \frac{g_2(\sigma_e^2)}{f(\sigma_e^2 | \xi')} \\ &= \left[\frac{I_{\{\sigma_\phi^2 < \sigma_\phi^*\}}(d-K)/\delta_1 + 2\beta_1}{w_1(\xi') + 2\beta_1} \right]^{\frac{M}{2} + \alpha_1} \\ &\quad \cdot \exp \left\{ -\frac{I_{\{\sigma_\phi^2 < \sigma_\phi^*\}}(d-K)/\delta_1}{2\sigma_\phi^2} + \frac{w_1(\xi')}{2\sigma_\phi^2} \right\} \\ &\quad \cdot \left[\frac{I_{\{\sigma_e^2 < \sigma_e^*\}}(d-K)/\delta_2 + SSE + 2\beta_2}{w_2(\xi') + SSE + 2\beta_2} \right]^{\frac{M}{2} + \alpha_2} \\ &\quad \cdot \exp \left\{ -\frac{I_{\{\sigma_e^2 < \sigma_e^*\}}(d-K)/\delta_2}{2\sigma_e^2} + \frac{w_2(\xi')}{2\sigma_e^2} \right\}. \end{aligned}$$

References

ASMUSSEN, S., GLYNN, P. W., AND THORISSON, H. (1992). Stationarity detection in the initial transient problem. *ACM Trans. Model. Comput. Simul.*, 2:130–157.

- BLANCHET, J. AND MENG, X.-L. (2005). Exact sampling, regeneration and minorization conditions. Technical report, Columbia University. Available at: <http://www.columbia.edu/~jb2814/papers/JSMsent.pdf>.
- BLANCHET, J. AND THOMAS, A. C. (2007). Exact simulation and error-controlled sampling via regeneration and a Bernoulli factory. Technical report, Harvard University. Available at: <http://www.acthomas.ca/papers/factory-sampling.pdf>.
- CHEN, M.-H., SHAO, Q.-M., AND IBRAHIM, J. G. (2000). *Monte Carlo Methods in Bayesian Computation*. Springer-Verlag Inc. [MR1742311](#)
- CRAIU, R. V. AND MENG, X.-L. (2011). Perfection with research: Exact MCMC sampling. In Brooks, S., Gelman, A., Jones, G., and Meng, X., editors, *Handbook of Markov Chain Monte Carlo*, pages 199–225. Chapman & Hall/CRC Press.
- FILL, J. A. (1998). An interruptible algorithm for perfect sampling via Markov chains. *Annals of Applied Probability*, 8:131–162. [MR1620346](#)
- FLEGAL, J. M., HARAN, M., AND JONES, G. L. (2008). Markov chain Monte Carlo: Can we trust the third significant figure? *Statistical Science*, 23:250–260. [MR2516823](#)
- FLEGAL, J. M. AND JONES, G. L. (2010). Batch means and spectral variance estimators in Markov chain Monte Carlo. *The Annals of Statistics*, 38:1034–1070. [MR2604704](#)
- GREEN, P. J. AND MURDOCH, D. J. (1999). Exact sampling for Bayesian inference: Towards general purpose algorithms. In Berger, J. O., Bernardo, J. M., Dawid, A. P., Lindley, D. V., and Smith, A. F. M., editors, *Bayesian Statistics 6*, pages 302–321, Oxford. Oxford University Press. [MR1723502](#)
- HOBERT, J. P. AND GEYER, C. J. (1998). Geometric ergodicity of Gibbs and block Gibbs samplers for a hierarchical random effects model. *Journal of Multivariate Analysis*, 67:414–430. [MR1659196](#)
- HOBERT, J. P., JONES, G. L., PRESNELL, B., AND ROSENTHAL, J. S. (2002). On the applicability of regenerative simulation in Markov chain Monte Carlo. *Biometrika*, 89:731–743. [MR1946508](#)
- HOBERT, J. P., JONES, G. L., AND ROBERT, C. P. (2006). Using a Markov chain to construct a tractable approximation of an intractable probability distribution. *Scandinavian Journal of Statistics*, 33(1):37–51. [MR2255108](#)
- HOBERT, J. P. AND ROBERT, C. P. (2004). A mixture representation of π with applications in Markov chain Monte Carlo and perfect sampling. *Ann. Appl. Probab.*, 14:1295–1305. [MR2071424](#)
- HUBER, M. (2004). Perfect sampling using bounding chains. *The Annals of Applied Probability*, 14(2):734–753. [MR2052900](#)
- JONES, G. L. AND HOBERT, J. P. (2001). Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statistical Science*, 16:312–334. [MR1888447](#)
- JONES, G. L. AND HOBERT, J. P. (2004). Sufficient burn-in for Gibbs samplers for a hierarchical random effects model. *The Annals of Statistics*, 32:784–817. [MR2060178](#)

- KEANE, M. S. AND O'BRIEN, G. (1994). A Bernoulli factory. *ACM Transactions on Modeling and Computer Simulation*, 2:213–219.
- LATUSZYNSKI, K., KOSMIDIS, I., PAPASPILIOPOULOS, O., AND ROBERTS, G. O. (2011). Simulating events of unknown probabilities via reverse time martingales. *Random Structures & Algorithms*, 38(4):441–452. [MR2829311](#)
- LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer, New York. [MR1842342](#)
- LUND, R. B. AND TWEEDIE, R. L. (1996). Geometric convergence rates for stochastically ordered Markov chains. *Mathematics of Operations Research*, 20:182–194. [MR1385873](#)
- LYLES, R. H., KUPPER, L. L., AND RAPPAPORT, S. M. (1997). Assessing regulatory compliance of occupational exposures via the balanced one-way random effects ANOVA model. *Journal of Agricultural, Biological, and Environmental Statistics*, 2:64–86. [MR1812251](#)
- MEYN, S. P. AND TWEEDIE, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, London. [MR1287609](#)
- MURDOCH, D. J. AND GREEN, P. J. (1998). Exact sampling from a continuous state space. *Scandinavian Journal of Statistics*, 25:483–502. [MR1650023](#)
- MYKLAND, P., TIERNEY, L., AND YU, B. (1995). Regeneration in Markov chain samplers. *Journal of the American Statistical Association*, 90:233–241. [MR1325131](#)
- NACU, Ş. AND PERES, Y. (2005). Fast simulation of new coins from old. *Annals of Applied Probability*, 15:93–115. [MR2115037](#)
- NUMMELIN, E. (1984). *General Irreducible Markov Chains and Non-negative Operators*. Cambridge University Press, London. [MR0776608](#)
- PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9:223–252. [MR1611693](#)
- ROBERT, C. P. AND CASELLA, G. (1999). *Monte Carlo Statistical Methods*. Springer, New York. [MR1707311](#)
- ROBERTS, G. O. AND ROSENTHAL, J. S. (2004). General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71. [MR2095565](#)
- ROBERTS, G. O. AND TWEEDIE, R. L. (1999). Bounds on regeneration times and convergence rates for Markov chains. *Stochastic Processes and their Applications*, 80:211–229. Corrigendum (2001) 91:337–338.
- ROBERTS, G. O. AND TWEEDIE, R. L. (2001). Corrigendum to “Bounds on regeneration times and convergence rates for Markov chains”. *Stochastic Processes and their Applications*, 91:337–338. [MR1807678](#)
- ROSENTHAL, J. S. (1995). Minorization conditions and convergence rates for Markov chain Monte Carlo. *Journal of the American Statistical Association*, 90:558–566. [MR1340509](#)
- ROSENTHAL, J. S. (1996). Analysis of the Gibbs sampler for a model related to James-Stein estimators. *Statistics and Computing*, 6:269–275.
- ROSENTHAL, J. S. (2002). Quantitative convergence rates of Markov chains: A simple account. *Electronic Communications of Probability*, 7:123–128. [MR1917546](#)

- ROY, V. AND HOBERT, J. P. (2007). Convergence rates and asymptotic standard errors for Markov chain Monte Carlo algorithms for Bayesian probit regression. *Journal of the Royal Statistical Society, Series B*, 69(4):607–623. [MR2370071](#)
- TAN, A. AND HOBERT, J. P. (2009). Block Gibbs sampling for Bayesian random effects models with improper priors: Convergence and regeneration. *Journal of Computational and Graphical Statistics*, 18:861–878. [MR2598033](#)
- VAN DYK, D. A. AND MENG, X.-L. (2001). The art of data augmentation (with discussion). *Journal of Computational and Graphical Statistics*, 10:1–50. [MR1936358](#)
- WILSON, D. B. (2000). How to couple from the past using a read-once source of randomness. *Random Structures and Algorithms*, 16:85–113. [MR1728354](#)