

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

# Sampling Conditionally on a Rare Event via Generalized Splitting

Zdravko I. Botev

UNSW Sydney, Australia, botev@unsw.edu.au, <http://web.maths.unsw.edu.au/~zdravkobotev/>

Pierre L'Ecuyer

Université de Montréal, Canada, and Inria–Rennes, France [lecuyer@iro.umontreal.ca](mailto:lecuyer@iro.umontreal.ca), <https://www.iro.umontreal.ca/~lecuyer/>

We propose and analyze a generalized splitting method to sample approximately from a distribution conditional on the occurrence of a rare event. This has important applications in a variety of contexts in operations research, engineering, and computational statistics. The method uses independent trials starting from a single particle. We exploit this independence to obtain asymptotic and non-asymptotic bounds on the total variation error of the sampler. Our main finding is that the approximation error depends crucially on the relative variability of the number of points produced by the splitting algorithm in one run, and that this relative variability can be readily estimated via simulation. We illustrate the relevance of the proposed method on an application in which one needs to sample (approximately) from an intractable posterior density in Bayesian inference.

*Key words:* conditional distribution; Monte Carlo splitting; Markov chain Monte Carlo; rare-event

*History:*

---

## 1. Introduction

We consider the problem of generating samples from a conditional distribution when the conditioning is on the occurrence of an event that has a small probability. We have a random variable  $\mathbf{X}$  defined over a probability space  $(\mathbb{R}^d, \mathcal{B}, \mathbb{P})$ , where  $\mathcal{B}$  can be taken as the Borel sigma-field, and  $\mathbf{X}$  has a probability density function (pdf)  $f$ . We assume it is easy to sample exactly from the density  $f$ . The rare event on which we condition can be written in the form  $B = \{S(\mathbf{x}) \geq \gamma\} \in \mathcal{B}$  for an appropriately chosen measurable function  $S: \mathbb{R}^d \rightarrow \mathbb{R}$  called the *importance function*. The conditional pdf is then

$$q(\mathbf{x}) = \frac{f(\mathbf{x})\mathbb{I}\{S(\mathbf{x}) \geq \gamma\}}{\ell(\gamma)}, \quad \mathbf{x} = (x_1, \dots, x_d)^\top, \quad (1)$$

where  $\mathbb{I}$  is the indicator function, and

$$\ell = \ell(\gamma) = \mathbb{P}(S(\mathbf{X}) \geq \gamma) \quad (2)$$

is the appropriate (unknown) normalizing constant, which we assume is so small that estimating it via the naive acceptance-rejection method (simulate  $\mathbf{X} \sim f$  until  $S(\mathbf{X}) \geq \gamma$ ) is impractical.

Sampling from a distribution conditional on a rare event has many applications. For example, suppose we want to generate  $\mathbf{X}$  from an arbitrary density proportional to  $p(\mathbf{x})$  for  $\mathbf{x} \in \mathbb{R}^d$ , for some known function  $p$ , and that it is too hard to generate samples directly from this density. Since  $p$  is known, we may be able to find a density  $f$  such that  $\sup_{\mathbf{x}} p(\mathbf{x})/f(\mathbf{x}) < \gamma$  for some constant  $\gamma < \infty$ . Then to generate  $\mathbf{X}$ , it suffices to generate a pair of independent random variables  $\mathbf{X} \sim f$  and  $U \sim \text{U}(0,1)$  conditional on the event  $p(\mathbf{X})/f(\mathbf{X}) \geq \gamma U$ , which is frequently a rare event (Kroese et al. 2011)[Section 14.5]. This fits our framework by taking  $S((\mathbf{x}, u)) = p(\mathbf{x})/(f(\mathbf{x})u)$ .

Another application is Bayesian Lasso regression (Park and Casella 2008), in which inference requires repeated simulation of a vector  $\beta$  of model parameters, conditional on the regularization constraint  $\|\beta\|_1 < \gamma$ . We give a detailed example of this in Section 5.

A third type of application occurs in the setting where we want to estimate the probability  $\ell$  of the rare event and to understand under which circumstances the rare event is likely to occur. A popular method to estimate  $\ell$  is importance sampling, and the optimal way to do it is to sample under a density  $f$  proportional to the original density conditional on the rare event, and then adjust the estimator using a likelihood ratio (Tuffin et al. 2014, Botev and Ridder 2014, Botev et al. 2011). This also fits our framework. In this context, it can be very useful to sample from the conditional density to get insight on how the rare event occurs. For instance, in a network with unreliable links, one may want to sample random configurations of all the links conditionally on a failure of the network, to better understand what (typically) makes the network fail (Botev et al. 2014, 2012).

The sampling methods examined in this paper are based on the *generalized splitting* (GS) algorithm of Botev and Kroese (2012) for drawing a collection of random vectors whose distribution converges to a target distribution with pdf of the form (1). To apply GS, we first select an increasing sequence of levels  $-\infty = \gamma_0 < \gamma_1 < \dots < \gamma_r = \gamma$  for the importance function  $S$ . This can be done in pilot runs via a run (Botev and Kroese 2012).

The algorithm uses a branching process that favors states  $\mathbf{X}$  having a large value of  $S(\mathbf{X})$  by resampling them conditional on staying above the current threshold, thus “splitting” those states into new copies, and then discarding those that do not reach the next level. At the end, the states that have reached the last level  $\gamma$  are retained. This process is replicated several times independently and all the retained states are collected to form an empirical version of the target conditional distribution. There are many ways of choosing the total number of replications (or *trials*). For example, one can fix them in advance to a constant  $n$ , or one can repeat the procedure until  $n$  trials have provided at least one retained state each, or until the total number of retained states is more than  $t$ , or until a certain computing budget (CPU time) has elapsed. In the latter case, one can either complete the current trial, or discard it, or just take the states retained so far from that trial.

There is a large variety of splitting-type or interacting particle algorithms to sample the state of a Markov chain approximately from its steady-state distribution conditional on a rare event; see for example Glasserman et al. (1999), Cérou et al. (2005), Cérou et al. (2012), L'Ecuyer et al. (2007, 2009), Andrieu et al. (2010), Bréhier et al. (2016), and the references given there. The analysis of these algorithms consists in most cases in proving their unbiasedness when estimating the expectation of a random variable that can be nonzero only when the rare event occurs (estimating the probability of the rare event is a special case of this), and sometimes showing their asymptotic efficiency when the probability of the rare event decreases toward zero (Dean and Dupuis 2009).

In this paper, we are interested in the different problem of bounding the difference between the exact conditional distribution and the distribution obtained by picking a state from the sample returned by the splitting algorithm. We do this for some variants of the GS method of Botev and Kroese (2012). L'Ecuyer et al. (2018) proved that this method provides an unbiased estimator of the expected value of a cost function, but also showed that a state picked at random from the set of retained states at the last level does not follow the true conditional distribution in general.

On the other hand, the distance between the two distributions converges to zero when the number of replicates increases toward infinity. The aim of the present paper is to study how this convergence occurs and to establish explicit *non-asymptotic* (risk) bounds on the total variation (TV) error between the two distributions, their mean absolute value, and

the expectation of the TV error in the case when it is a random variable. Our bounds are expressed in terms of simple mathematical expectations that can be estimated easily from the simulation output.

We provide convergence results for two versions of the GS algorithm. In both, we assume that whenever a trial returns no state from the rare event set (an empty trial), we discard it and try again. In the first version, we run GS until we have  $n$  non-empty trials, for some fixed  $n > 0$ . We prove that the TV distance between the true conditional distribution and the distribution of a state picked at random from the retained states from this GS version is bounded by  $c_1/n$  where  $c_1$  is an unknown constant that can be estimated from the simulation output. In the second version, we run GS until the total number of retained states exceeds  $t$ , for some fixed positive integer  $t$ . For this version, we show that the convergence rate is of the form  $c_2(t)t^{-3/2} = \mathcal{O}(t^{-3/2})$ , where the quantity  $c_2(t)$  is bounded uniformly in  $t$  and can be estimated from the simulation output. The derivation of these bounds is made possible thanks to the fact that GS produces independent trials, each one starting from a single particle, and this permits us to use results from renewal theory for our analysis.

Typically, approximate simulation from the target pdf (1) is accomplished using Markov chain Monte Carlo (MCMC) (for example, Jones and Hobert (2001), Taimre et al. (2019)). While MCMC sampling can be simple to implement, it still poses the challenge of analyzing its output and deciding how close the sampling or empirical distribution is to the desired target distribution (Jones and Hobert 2001). The reason for this difficulty is that MCMC generates a sequence of *dependent* random vectors  $\mathbf{Y}_1, \mathbf{Y}_2, \dots$ . Typical graphical diagnostic tools like autocorrelation plots are heuristics, which do not easily provide precise qualitative measure of how close the simulated random variables follow the target distribution. Also we have to choose from an infinite number of possible one-dimensional plots. In contrast, our bounds on the TV error present a more rigorous and theoretically justified convergence assessment than the autocorrelation plots typically used in MCMC.

The rest of the paper is organized as follows. In Section 2, we recall the GS algorithm used in this paper. In Section 3, we define our versions of GS used for sampling conditional on the rare event. In Section 4, we state our main new results on the convergence of the distance between the empirical and true conditional distributions, and bounds on this distance. The proofs are given in the Online Supplement. In Section 5, we show how our methodology

can be applied in a practical setting, namely to sample approximately from the posterior density of the Bayesian Lasso. In this example, we show how the non-asymptotic risk bounds can be used to assess convergence and to estimate the error committed when using GS to sample from the conditional distribution. We also compare the simulation accuracy of GS with that of the sequential Monte Carlo method (C erou et al. 2012).

## 2. Background on Generalized Splitting

We recall the GS method for estimating the rare-event probability  $\ell$  in (2). This method is a simple generalization of the classical multilevel splitting technique for rare-event simulation (Kahn and Harris 1951, Glasserman et al. 1999, Garvels et al. 2002, L'Ecuyer et al. 2009). Our background material here is similar to the one given in L'Ecuyer et al. (2018).

The idea of GS is to define a discrete-time Markov chain with state  $\mathbf{Y}$ , which evolves via a branching-type random mechanism that pushes it toward a state corresponding to  $\{\mathbf{Y} \in B\} \equiv \{S(\mathbf{Y}) \geq \gamma\}$  in (1). To estimate the (rare-event) probability (2) via GS, we first need to choose:

1. an integer  $s \geq 2$ , called the *splitting factor*, and
2. an integer  $\tau > 0$  and real numbers  $-\infty = \gamma_0 < \gamma_1 < \dots < \gamma_\tau = \gamma$  for which

$$\rho_l := \mathbb{P}(S(\mathbf{Y}) \geq \gamma_l \mid S(\mathbf{Y}) \geq \gamma_{l-1}) \approx 1/s$$

for  $l = 1, \dots, \tau$  (except for  $\rho_\tau$ , which can be larger than  $1/s$ ). These  $\gamma_l$ 's represent the  $\tau$  levels of the splitting algorithm. In Section 5 we give particular choices of  $s$  and  $\tau$  that are relevant to our examples.

For each level  $\gamma_l$  we construct a Markov chain whose stationary density is equal to the density of  $\mathbf{Y}$  conditional on  $S(\mathbf{Y}) \geq \gamma_l$  (a truncated density), given by

$$q_l(\mathbf{y}) := f(\mathbf{y}) \frac{\mathbb{I}\{S(\mathbf{y}) \geq \gamma_l\}}{\mathbb{P}(S(\mathbf{Y}) \geq \gamma_l)}. \quad (3)$$

Note that  $q_0 = f$  and  $q_\tau = q$ . We denote by  $\kappa_l \equiv \kappa_l(\cdot \mid \cdot)$  the transition kernel of this Markov chain:  $\kappa_l(d\mathbf{y} \mid \mathbf{x})$  represents the probability that the next state is in  $d\mathbf{y}$  when the current state is  $\mathbf{x}$ . There are many ways of constructing this Markov chain and  $\kappa_l$ . A practical example using Gibbs sampling will be given in Section 5.

---

**Algorithm 1** GS Sampler 1

---

**Require:**  $s, \tau, \gamma_1, \dots, \gamma_\tau$ generate a vector  $\mathbf{Y}$  from its unconditional density  $f$ **if**  $S(\mathbf{Y}) < \gamma_1$  **then**    **return**  $\mathcal{X}_\tau = \emptyset$  and  $\mathring{M} = 0$ **else**     $\mathcal{X}_1 \leftarrow \{\mathbf{Y}\}$     {this state  $\mathbf{Y}$  has reached at least the first level}    **for**  $l = 2$  **to**  $\tau$  **do**         $\mathcal{X}_l \leftarrow \emptyset$     {list of states that have reached the level  $\gamma_l$ }        **for all**  $\mathbf{Y} \in \mathcal{X}_{l-1}$  **do**            set  $\mathbf{Y}_0 \leftarrow \mathbf{Y}$     {we will simulate  $s$  steps from this state}            **for**  $j = 1$  **to**  $s$  **do**                sample  $\mathbf{Y}_j$  from  $\kappa_{l-1}(d\mathbf{y} | \mathbf{Y}_{j-1})$                 **if**  $S(\mathbf{Y}_j) \geq \gamma_l$  **then**                    add  $\mathbf{Y}_j$  to  $\mathcal{X}_l$     {this state has reached the next level}**return** the list  $\mathring{\mathcal{Y}} = \mathcal{X}_\tau$  of retained states and its cardinality  $\mathring{M}$ .

---

The original GS algorithm is summarized in Algorithm 1, and is also given in L'Ecuyer et al. (2018). The algorithm returns a list  $\mathring{\mathcal{Y}}$  of *retained states* that belong to  $B = \{\mathbf{y} : S(\mathbf{y}) \geq \gamma\}$ , as well as the size of this list. This list is a multiset, in the sense that it may contain the same state more than once. The list  $\mathring{\mathcal{Y}}$  can be empty and its cardinality  $\mathring{M} = 0$ . The o-ring symbol in the notation is a reminder that the size of the set can be zero. In the remainder of this article, we define  $\mathcal{Y}$  and  $M$  as the versions of  $\mathring{\mathcal{Y}}$  and  $\mathring{M}$ , conditional on  $\mathring{M} \geq 1$ .

Let  $\mathcal{A}$  denote a  $\sigma$ -algebra of Borel measurable subsets of  $\mathbb{R}^d$ . For some of our results,  $\mathcal{A}$  will be a more restricted class than the Borel subsets of  $\mathbb{R}^d$ . Algorithm 1 can be used to estimate  $\mathbb{P}(\mathbf{Y} \in A)$  for any  $A \in \mathcal{A}$  via the unbiased estimator:

$$\widehat{\mathbb{P}}(A) = \mathring{H}(A)/s^{\tau-1}, \quad (4)$$

where  $\mathring{H}(A) = |\mathring{\mathcal{Y}} \cap A|$  is the number of states  $\mathbf{Y} \in \mathring{\mathcal{Y}}$  that belong to  $A$ . In practice, one will replicate this algorithm several times and take the average. The unbiasedness is implied by the following lemma, proved in L'Ecuyer et al. (2018).

LEMMA 1 (L'Ecuyer et al. (2018)). *For any measurable function  $h: \mathbb{R}^d \mapsto \mathbb{R}$  and any measurable subset  $A \subseteq B$ , we have*

$$\mathbb{E}_{\text{GS}} \left[ \sum_{\mathbf{Y} \in \mathring{\mathcal{Y}}} h(\mathbf{Y}) \mathbb{I}\{\mathbf{Y} \in A\} \right] = s^{\tau-1} \mathbb{E}[h(\mathbf{Y}) \mathbb{I}\{\mathbf{Y} \in A\}], \quad (5)$$

where the expectation on the left-hand-side is with respect to  $\mathring{\mathcal{Y}}$  from Algorithm 1 and the expectation on the right-hand-side is with respect to the original density  $f$ .

By taking  $h$  as the identity function in (5), we obtain that  $\widehat{\mathbb{P}}(A)$  in (4) is unbiased for  $\mathbb{P}(A)$ :

$$\mathbb{E}_{\text{GS}}[\widehat{H}(A)] = \mathbb{E}_{\text{GS}} \left[ \sum_{\mathbf{Y} \in \mathring{\mathcal{Y}}} \mathbb{I}\{\mathbf{Y} \in A\} \right] = s^{\tau-1} \mathbb{P}(\mathbf{Y} \in A),$$

and therefore

$$\mathbb{E}_{\text{GS}}[\widehat{\mathbb{P}}(A)] = \mathbb{E}_{\text{GS}}[\widehat{H}(A)/s^{\tau-1}] = \mathbb{P}(\mathbf{Y} \in A).$$

Moreover, since  $\mathbb{E}_{\text{GS}}[\mathring{M}] = s^{\tau-1} \mathbb{P}(\mathbf{Y} \in B)$  and  $A \subseteq B$ , we have that

$$\frac{\mathbb{E}_{\text{GS}}[\mathring{H}(A)]}{\mathbb{E}_{\text{GS}}[\mathring{M}]} = \mathbb{P}(\mathbf{Y} \in A \mid \mathbf{Y} \in B) = \frac{\mathbb{P}(\mathbf{Y} \in A)}{\mathbb{P}(\mathbf{Y} \in B)}.$$

### 3. Sampling Conditionally on a Rare Event

When estimating an expectation as in (5), an empty list  $\mathring{\mathcal{Y}}$  poses no problem: the unbiased estimator just takes the value 0 in that case. But for our purpose of sampling from a conditional distribution, we insist that there are no empty sets of retained states. To make sure that the set of retained states is non-empty, we modify the original GS so that each trial returns at least one state. Whenever a GS run returns an empty list, we simply discard it and try again. This gives Algorithm 2.

---

#### Algorithm 2 GS Sampler 2

---

**Require:**  $s, \tau, \gamma_1, \dots, \gamma_\tau$

**repeat**

    run Algorithm 1

**until**  $\mathring{M} > 0$

**return** the list  $\mathcal{Y} = \mathcal{X}_\tau$  of retained states and its cardinality  $M = |\mathcal{Y}|$ .

---

Does this algorithm still provide an unbiased estimator? An important observation is that if we replace  $\mathring{\mathcal{Y}}$  by  $\mathcal{Y}$  in (5), the equality is no longer true. That is, we get a *biased*

estimator of the expectation on the right. However, our main goal here is not to estimate this expectation, but to sample approximately from the conditional distribution, and we will analyze methods that use Algorithm 2 for this purpose. As mentioned earlier, there are several ways of doing it. In this paper, we examine the following two versions: (a) run a fixed number  $n$  of iid replicates of Algorithm 2 and (b) perform replicates until there are more than  $t$  retained states in total. These two approaches are detailed in Algorithms 3 and 4, respectively. In both cases, at the end we collect all the retained states in a multiset  $\mathcal{Y}_\cup$ . For the first version the cardinality of the returned set  $\mathcal{Y}_\cup$  is at least  $n$ , whereas in the second case it is at least  $t$  and  $N(t)$  is the (random) number of calls to Algorithm 2. We summarize these two versions as follows.

---

**Algorithm 3** Sampling an empirical distribution from  $n$  iid non-empty GS replications

---

**Require:**  $s, \tau, \gamma_1, \dots, \gamma_\tau$  and  $n$

**for**  $i = 1, \dots, n$  **do**

run Algorithm 2 to obtain the list  $\mathcal{Y}_i$  of size  $M_i$

**return** the empirical distribution  $\widehat{\mathbb{Q}}_n$  of the states in the set  $\mathcal{Y}_\cup := \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_n$

---



---

**Algorithm 4** Sampling an empirical distribution with more than  $t$  retained states

---

**Require:**  $s, \tau, \gamma_1, \dots, \gamma_\tau$  and  $t$

$i \leftarrow 0$  and  $T_0 \leftarrow 0$

**repeat**

$i \leftarrow i + 1$

run Algorithm 2 to obtain the list  $\mathcal{Y}_i$  and its cardinality  $M_i$

$T_i \leftarrow T_{i-1} + M_i$

**until**  $T_i > t$

**return**  $N(t) \leftarrow i$  and the empirical distribution  $\widehat{\mathbb{Q}}_{N(t)}$  of the set of states  $\mathcal{Y}_\cup := \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_{N(t)}$

---

Note that  $\widehat{\mathbb{Q}}_n$  (or  $\widehat{\mathbb{Q}}_{N(t)}$ ) is a random distribution; it is the distribution *conditional* on  $\mathcal{Y}_\cup$ . The *unconditional* distribution of a state obtained by generating  $\mathcal{Y}_\cup$  and then selecting one state randomly from  $\mathcal{Y}_\cup$  is also of interest: this is the (a priori) distribution of a state sampled from  $\widehat{\mathbb{Q}}_n$  (or  $\widehat{\mathbb{Q}}_{N(t)}$ ), but before we run the GS algorithm to construct  $\mathcal{Y}_\cup$ . We will denote these two unconditional distributions by

$$\mathbb{Q}_n(A) := \mathbb{E}[\widehat{\mathbb{Q}}_n(A)] \quad (\text{for Algorithm 3})$$

and

$$\bar{Q}_t(A) := \mathbb{E}[\widehat{Q}_{N(t)}(A)] \quad (\text{for Algorithm 4})$$

for all  $A \in \mathcal{A}$ , where the expectation is with respect to the realization of  $\mathcal{Y}_\cup$ . We saw earlier that

$$\frac{\mathbb{E}[\mathring{H}(A)]}{\mathbb{E}[\mathring{M}]} = \mathbb{P}(\mathbf{Y} \in A \mid \mathbf{Y} \in B).$$

Now let  $H(A) = |\mathcal{Y} \cap A|$  be the number of states returned by Algorithm 2 that belong to  $A$ . We have  $\mathbb{E}[H(A)] = \mathbb{E}[\mathring{H}(A) \mid \mathring{M} > 0] = \mathbb{E}[\mathring{H}(A)] / \mathbb{P}(\mathring{M} > 0)$ . Likewise,  $\mathbb{E}[M] = \mathbb{E}[\mathring{M}] / \mathbb{P}(\mathring{M} > 0)$ . Therefore,

$$\frac{\mathbb{E}[H(A)]}{\mathbb{E}[M]} = \frac{\mathbb{E}[\mathring{H}(A)]}{\mathbb{E}[\mathring{M}]} = \mathbb{P}(\mathbf{Y} \in A \mid \mathbf{Y} \in B).$$

We also know that  $\widehat{Q}_n(A)$  and  $\widehat{Q}_{N(t)}(A)$  converge with probability one to  $\mathbb{E}[H(A)] / \mathbb{E}[M]$  when  $n \rightarrow \infty$  and when  $t \rightarrow \infty$ , respectively, from the strong law of large numbers applied to the numerator and the denominator. Thus, they converge almost surely to the desired conditional probability  $Q(A) := \mathbb{P}(\mathbf{Y} \in A \mid \mathbf{Y} \in B)$ .

#### 4. Convergence Analysis

We now analyze the convergence of the empirical distribution of the retained states,  $\widehat{Q}_n$  (or  $\widehat{Q}_{N(t)}$ ), as well as its expected (i.e., unconditional on  $\mathcal{Y}_\cup$ ) version  $Q_n$  (or  $\bar{Q}_t$ ), to the true conditional distribution  $Q$ . The aim is to obtain non-asymptotic or risk bounds on the distance between  $Q$  and the empirical distribution, and its expected (unconditional) version. For a given class  $\mathcal{A}$  of measurable sets, we consider the three error criteria:

1. The TV error between the expected (unconditional) distribution  $Q_n$  and  $Q$ , that is:

$$\sup_{A \in \mathcal{A}} |Q_n(A) - Q(A)|.$$

This error measures the size of the “bias” of  $\widehat{Q}_n$  as an estimator of the true  $Q$ .

2. The worst-case mean absolute error of the conditional distribution  $\widehat{Q}_n$ , defined as:

$$\sup_{A \in \mathcal{A}} \mathbb{E} \left| \widehat{Q}_n(A) - Q(A) \right|.$$

3. The (random) TV error,  $\sup_{A \in \mathcal{A}} \left| \widehat{Q}_n(A) - Q(A) \right|$ , of the conditional distribution  $\widehat{Q}_n$ , and its expected value:

$$\mathbb{E} \sup_{A \in \mathcal{A}} \left| \widehat{Q}_n(A) - Q(A) \right|.$$

By permuting the positions of the expectation, absolute value function, and the supremum ( $\sup |\mathbb{E}[\cdot]| \rightarrow \sup \mathbb{E}|\cdot| \rightarrow \mathbb{E} \sup |\cdot|$ ), we find that the three error criteria dominate each other as follows:

$$\overbrace{\mathbb{E} \sup_{A \subseteq \mathcal{A}} |\widehat{Q}_n(A) - Q(A)|}^{\text{expected TV error}} \geq \sup_{A \subseteq \mathcal{A}} \overbrace{\mathbb{E} |\widehat{Q}_n(A) - Q(A)|}^{\text{mean absolute error}} \geq \overbrace{\sup_{A \subseteq \mathcal{A}} |\mathbb{E}[\widehat{Q}_n(A)] - Q(A)|}^{\text{TV error}}.$$

In other words, the expected TV error is the *most stringent* of these three errors. In fact, the (expected) TV error of the empirical distribution is so stringent that it does not converge, unless the class of sets  $\mathcal{A}$  is restricted. To ensure convergence, in Section 4.2 we will take  $\mathcal{A}$  to be a restricted class of subsets. In contrast, the TV and mean absolute errors do not require any restrictions on the class  $\mathcal{A}$  and for these error criteria we simply take  $\mathcal{A}$  to be the class of all Borel subsets of  $\mathbb{R}^d$ .

#### 4.1. Convergence of Total Variation and Mean Absolute Errors

Let  $m := \mathbb{E}[M]$  and  $\text{Var}(M)$  denote the expectation and variance of  $M$ , which is the output of either Algorithm 3, or Algorithm 4. In this section, we state theorems giving non-asymptotic bounds on the TV error and the (worst-case) mean absolute error. The proofs of the following results are in the Online Supplement.

**THEOREM 1 (Sampling via  $n$  iid runs of GS).** *The TV error is bounded as*

$$\sup_A |\mathbb{Q}_n(A) - Q(A)| \leq c_1 n^{-1}$$

where  $c_1 := \left( \text{Var}(M) + \sqrt{\text{Var}(M)\mathbb{E}[M^2]} \right) m^{-2}$ . *The worst-case mean absolute error is bounded as*

$$\sup_A \mathbb{E} |\widehat{Q}_n(A) - Q(A)| \leq \tilde{c}_1(n) n^{-1/2}$$

where  $\tilde{c}_1(n) := \left( \sqrt{\mathbb{E}M^2} + \sqrt{3\mathbb{E}M^4/n} \right) m^{-1}$  is bounded uniformly in  $n$ .

The terms  $c_1$  and  $\tilde{c}_1(n)$  in these bounds can be estimated from the simulation output.

**THEOREM 2 (Sampling until GS returns  $t$  states).** *In this case, the TV error is bounded as*

$$\sup_A |\overline{Q}_t(A) - Q(A)| \leq c_2(t) (t/m)^{-3/2},$$

where  $c_2(t) := \sqrt{(4/3)\mathbb{E}[M^3]\mathbb{E}[M^2]} (m + \mathbb{E}[M^2]/t) m^{-3}$  is bounded uniformly in  $t$ . *The worst-case mean absolute error is bounded as*

$$\sup_A \mathbb{E} |\widehat{Q}_{N(t)}(A) - Q(A)| \leq \tilde{c}_2(t) (t/m)^{-1/2},$$

where  $\tilde{c}_2(t) := (\mathbb{E}M^2)^{1/2}m^{-1} + (\mathbb{E}M^2)m^{-3/2}t^{-1/2}$  is also uniformly bounded in  $t$ .

Again, the terms  $c_2(t)$  and  $\tilde{c}_2(t)$  can be estimated easily by simulation: it suffices to estimate  $\mathbb{E}M^2$  and  $\mathbb{E}M^3$  by their empirical versions. The constant  $m$  in  $(t/m)^{-3/2}$  could be absorbed into  $c_2(t)$ , but we choose not to do this, because we want to be able to compare  $\overline{\mathbb{Q}}_t$  and  $\mathbb{Q}_n$  on a common scale, where  $n$  (the simulation effort of Algorithm 3) and  $t/m$  (the average simulation effort of Algorithm 4 for large  $t$ ) are the same. The key point to notice is that we get a better rate for the bound for  $\overline{\mathbb{Q}}_t$  than for  $\mathbb{Q}_n$ .

In the next result, we obtain an improved convergence rate of  $\mathcal{O}(1/t^2)$ , but at the price of introducing in the bound an  $\mathcal{O}(\exp(-\omega t))$  term (for some  $\omega > 0$ ) which is hard to estimate. This term converges exponentially fast in  $t$ , so it is asymptotically negligible when  $t \rightarrow \infty$ , but it is not necessarily negligible for a given (finite)  $t$ . So we have an asymptotically better bound that we cannot easily estimate. In practical settings, we may prefer the  $\mathcal{O}((t/m)^{-3/2})$  bound from Theorem 2 that we can more easily estimate to the  $\mathcal{O}((t/m)^{-2})$  bound that we cannot completely estimate.

**THEOREM 3 (Sampling until GS returns more than  $t$  states; asymptotic version).**

We have

$$\sup_A |\overline{\mathbb{Q}}_t(A) - \mathbb{Q}(A)| \leq c_3(t/m)^{-2} + \mathcal{O}(\exp(-\omega t)),$$

where  $\omega > 0$  is a (typically unknown) constant and

$$c_3 := \frac{\mathbb{E}[M^2|M-1-2r|]}{2m^3}$$

with  $r := (\mathbb{E}M^2 + m)/(2m)$ .

This result does not include a statement about the mean absolute error, because the bounds of the mean absolute errors in Theorems 1 and 2 already converge at the optimal asymptotic rate, and thus cannot be improved.

#### 4.2. Convergence of the Empirical Conditional Distribution $\widehat{\mathbb{Q}}_n$

We now examine the convergence of the TV error between the empirical distribution  $\widehat{\mathbb{Q}}_n$  and  $\mathbb{Q}$  when  $n \rightarrow \infty$ . This distribution is random, and any realization is discrete with finite support, so obviously it cannot converge to  $\mathbb{Q}$  in TV with  $\mathcal{A}$  taken as all the Borel sets, because by taking  $A$  as the finite set  $\mathcal{Y}_\cup$ , we get  $\widehat{\mathbb{Q}}_n(A) = 1$  for any  $n$ , but  $\mathbb{Q}(A) = 0$  (assuming that  $\mathbb{Q}$  has a density). Thus, as mentioned previously, we necessarily have to restrict the class  $\mathcal{A}$ . We start by giving conditions for TV convergence with probability 1 under the following restrictions on the class  $\mathcal{A}$ .

ASSUMPTION 1. *Suppose that one of the following two conditions holds:*

1.  $\mathcal{A}$  is a class with finite Vapnik-Chervonenkis (VC) dimension, or
2.  $\mathcal{A}$  is the class of all convex sets in  $\mathbb{R}^d$ , and the transition kernel in Algorithm 1 has a probability density  $\kappa_l(\mathbf{y} | \mathbf{x})$ .

THEOREM 4 (**Almost-Sure TV Convergence**). *Under Assumption 1, we have almost sure TV convergence:*

$$\sup_{A \subseteq \mathcal{A}} |\widehat{\mathbb{Q}}_n(A) - \mathbb{Q}(A)| \longrightarrow 0 \quad \text{with probability 1 when } n \rightarrow \infty.$$

The notion of VC dimension is discussed for example by Vapnik (2013). Roughly speaking, it measures the flexibility of a class of subsets to correctly classify data defined over  $\mathbb{R}^d$ , and in our context it measures the complexity of the class of sets  $\mathcal{A}$ . Sets with higher VC dimension are more complex.

Note that the class of convex sets has an infinite VC dimension, which is why the second option in Assumption 1 requires the extra regularity condition on the transition kernel. This condition will be satisfied if  $\kappa_l$  is the transition kernel of a Gibbs sampler, but will not be satisfied for the kernel of a Metropolis-Hastings sampler (Kroese et al. 2011, Equation 6.3, Page 226). Note that the condition does not require that we have a closed form (simple) formula for the transition density  $\kappa_l(\mathbf{y} | \mathbf{x})$ . It only requires that it exists.

Our next result (proof in Online Supplement) provides bounds on the expected TV error of the empirical distribution, where  $\mathcal{A}$  is a class of sets with a finite VC dimension.

THEOREM 5 (**Bound on Expected TV for Empirical Distribution**). *Suppose the class  $\mathcal{A}$  has finite VC dimension  $v$ . Then, the expected TV error made by using the empirical distribution  $\widehat{\mathbb{Q}}_n$  as an approximation of  $\mathbb{Q}$  is bounded as follows:*

$$\mathbb{E} \sup_{A \in \mathcal{A}} |\widehat{\mathbb{Q}}_n(A) - \mathbb{Q}(A)| \leq \frac{\sqrt{\text{Var}(M)}}{m\sqrt{n}} + \frac{2\sqrt{v \ln(2n) \mathbb{E}[M^2 \ln M]}}{m\sqrt{n}} \psi_1(v, n),$$

where

$$\psi_1 = \psi_1(v, n) := \sqrt{\frac{(\ln(2) + v + v \ln(2n/v)) \mathbb{E}[M^2]}{v \ln(2n) \mathbb{E}[M^2 \ln M]} + \frac{1}{\ln(2n)}} < \infty$$

is bounded uniformly in  $(v, n, \tau)$ .

As an example, let  $[\mathbf{a}, \mathbf{b}] = \{\mathbf{y} \in \mathbb{R}^d : \mathbf{a} \leq \mathbf{y} \leq \mathbf{b}\}$  represent a rectangle in  $\mathbb{R}^d$ , and suppose  $\mathcal{A}$  is the class of all rectangles in  $\mathbb{R}^d$ . Then  $v = 2d$  (Sauer 1972). If  $\mathbf{a} = -\infty$ , that is,  $\mathcal{A}$

as the class of one-sided intervals of the form  $[-\infty, \mathbf{b}]$ , then  $v = d + 1$ . In this case, the previous theorem can provide a bound on the expected value of the Kolmogorov-Smirnov statistic:

$$\text{ks}(n) := \mathbb{E} \left[ \sup_{\mathbf{x} \in \mathbb{R}^d} \left| \widehat{\mathbb{Q}}_n(\mathbf{X} \leq \mathbf{x}) - \mathbb{Q}(\mathbf{X} \leq \mathbf{x}) \right| \right]. \quad (6)$$

We will use this type of error bound in Section 5.2 when we assess the quality of our approximate sampling from a Bayesian posterior.

Using the metric entropy of the class  $\mathcal{A}$ , it is also possible to obtain a bound without the logarithmic growth term  $\ln(n) \ln(M)$  in Theorem 5, and to get an expected TV bound that depends solely on the relative second moment of  $M$ .

**THEOREM 6 (Second Bound on Expected TV for Empirical Distribution).**

Let  $\tau$  be the number of levels in Algorithm 1 with splitting factor  $s$  and suppose that  $\mathcal{A}$  has VC dimension  $v < \infty$ . Then the empirical distribution  $\widehat{\mathbb{Q}}_n$  satisfies:

$$\mathbb{E} \sup_{A \in \mathcal{A}} |\widehat{\mathbb{Q}}_n(A) - \mathbb{Q}(A)| \leq \frac{\sqrt{\text{Var}(M)}}{m\sqrt{n}} + \frac{(s+1)4\sqrt{v} \mathbb{E}[M^2]}{m\sqrt{n}} \psi_2(\tau, v, n, s),$$

where

$$\psi_2 = \psi_2(\tau, v, n, s) := \sum_{k=1}^{\lceil \tau + \log_s \sqrt{n} \rceil} \frac{1}{s^k} \left( \frac{\ln 2}{2nv} + \frac{1 + \ln(v+1)}{v} + 1 + \ln(2s^{2k}) \right)^{1/2} < \infty$$

is bounded uniformly in  $(\tau, v, n, s)$ .

Unfortunately, as we shall see in Section 5.2, the constant  $\psi_2$  in this bound is much larger than  $\psi_1$  in Theorem 5. As a result,  $n$  has to be impractically large for the above bound to beat the simpler bound in Theorem 5. Nevertheless, the result is still of theoretical interest as it shows that the rate of convergence in expectation of the TV distance can be improved from  $\mathcal{O}(\ln(n)/\sqrt{n})$  to the canonical rate of  $\mathcal{O}(1/\sqrt{n})$ . In addition, the term  $\mathbb{E}[M^2 \ln M]/m^2$  in Theorem 5 does not appear in Theorem 6.

**REMARK 1 (SIMPLIFICATIONS DUE TO EXISTENCE OF A DENSITY).** If the transition density  $\kappa_l(\mathbf{x}|\mathbf{y})$  is available in closed form and easily evaluated, we can do much better by dropping the restrictions that the class  $\mathcal{A}$  has a finite VC dimension. Instead, if  $\kappa(\mathbf{x}|\mathbf{y}) \equiv \kappa_\tau(\mathbf{x}|\mathbf{y})$  is a transition density with stationary pdf  $q$ , then we can define the empirical density:

$$\hat{q}_n(\mathbf{x}) := \frac{1}{T_n} \sum_{\mathbf{Y} \in \mathcal{Y}_U} \kappa(\mathbf{x}|\mathbf{Y}),$$

so that we can use Sheffé's identity (Devroye and Lugosi 2001, Theorem 5.1) to simplify the uniform deviation over the class  $\mathcal{B}$  of Borel measurable sets:

$$2 \sup_{A \in \mathcal{B}} \left| \int_A \hat{q}_n(\mathbf{x}) d\mathbf{x} - \mathbb{Q}(A) \right| = \int |\hat{q}_n(\mathbf{x}) - q(\mathbf{x})| d\mathbf{x}$$

Therefore, the bound on the expected TV distance simplifies as follows:

$$\begin{aligned} 2\mathbb{E} \left[ \sup_{A \in \mathcal{B}} \left| \int_A \hat{q}_n(\mathbf{x}) d\mathbf{x} - \mathbb{Q}(A) \right| \right] &\leq \mathbb{E} \int \left| \frac{1}{n\bar{M}_n} \sum_{\mathbf{Y} \in \mathcal{Y}_U} \kappa(\mathbf{x}|\mathbf{Y}) - \frac{1}{nm} \sum_{\mathbf{Y} \in \mathcal{Y}_U} \kappa(\mathbf{x}|\mathbf{Y}) \right| d\mathbf{x} \\ &\quad + \mathbb{E} \int \left| \frac{1}{nm} \sum_{\mathbf{Y} \in \mathcal{Y}_U} \kappa(\mathbf{x}|\mathbf{Y}) - q(\mathbf{x}) \right| d\mathbf{x} \\ &\leq \frac{1}{m\sqrt{n}} \left( \sqrt{\text{Var}(M)} + \sqrt{\int \text{Var}(\kappa(\mathbf{x}|\mathbf{Y})) d\mathbf{x}} \right). \end{aligned}$$

Thus, provided the integrated variance  $\int \text{Var}(\kappa(\mathbf{x}|\mathbf{Y})) d\mathbf{x}$  can be estimated easily, this bound can be used as a simpler alternative to Theorem 5. We do not pursue this possibility further in this article.

## 5. Numerical Example: Bayesian Lasso

In this section we consider an application of the splitting sampler in Algorithm 2 to the problem of posterior simulation in Bayesian inference. We estimate the bounds in Theorems 1 to 6 in order to assess the convergence of Algorithms 3 and 4. This convergence assessment can be used to either assess whether any Bayesian credible intervals are reliably estimated from the simulation output, or to rank the performance of implementations that use different Markov chain kernels  $\kappa_l$  (the Markov chain that yields the smallest TV error will be the preferred one).

### 5.1. Approximate Posterior Simulations via Splitting

One of the simplest and most widely used linear regression models for data  $\mathbf{y} = (y_1, \dots, y_{n'})^\top$  is the *Bayesian Lasso* (Park and Casella 2008), in which the point-estimator of the regression coefficient  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_d)^\top \in \mathbb{R}^d$  is defined as the minimizer of the constrained least squares problem:

$$\min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2, \quad \text{subject to } \|\boldsymbol{\beta}\|_1 \leq \gamma,$$

where: (1)  $\mathbf{X}$  is a matrix with  $d$  columns (predictors); (2) the term  $\|\boldsymbol{\beta}\|_1 = |\beta_1| + \dots + |\beta_d|$  is the *least absolute shrinkage and selection operator* (Lasso); and (3)  $\gamma \geq 0$  is the Lasso

regularization parameter. In a Bayesian context, inference for  $(\boldsymbol{\beta}, \sigma^2)$  requires repeated drawing of pairs  $(\boldsymbol{\beta}, \sigma)$  from the posterior pdf:

$$q(\boldsymbol{\beta}, \sigma | \mathbf{y}, \gamma) \propto \phi(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}; \sigma^2 \mathbf{I}) \sigma^{-2} \mathbb{I}\{\|\boldsymbol{\beta}\|_1 \leq \gamma\} \times d! / (2\gamma)^d, \quad (7)$$

where: a)  $\phi(\mathbf{x}; \Sigma)$  denotes the multivariate normal pdf with mean zero and covariance matrix  $\Sigma$  evaluated at  $\mathbf{x}$ ; b) the factor  $\sigma^{-2}$  results from using an uninformative prior for the scale  $\sigma$ , and c)  $\mathbb{I}\{\|\boldsymbol{\beta}\|_1 \leq \gamma\} \times d! / (2\gamma)^d$  is the prior of  $\boldsymbol{\beta}$ , uniform over the feasible set. Note that, unlike the more common Laplace prior used in the Bayesian Lasso (Park and Casella 2008), here the prior enforces the constraint on  $\boldsymbol{\beta}$  directly.

To sample a new state  $(\boldsymbol{\beta}_k, \sigma_k)$  during the course of splitting, we need to simulate from a transition density  $\kappa_l((\sigma_k, \boldsymbol{\beta}_k) | (\sigma_{k-1}, \boldsymbol{\beta}_{k-1}))$ , which is stationary with respect to the density (3). We simulate a move from  $(\sigma_{k-1}, \boldsymbol{\beta}_{k-1})$  to  $(\sigma_k, \boldsymbol{\beta}_k)$  as follows. Given  $\boldsymbol{\beta}_{k-1}$ , we sample

$$(1/\sigma_k^2) \sim \text{Gamma}((n' + 1)/2, \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_{k-1}\|_2^2/2),$$

which is the gamma distribution with mean  $(n' + 1)/\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_{k-1}\|_2^2$  and shape parameter  $(n' + 1)/2$ . Given  $(\sigma_k, \boldsymbol{\beta}_{k-1})$ , we simulate  $\boldsymbol{\beta}_k$  via a ‘‘hit-and-run’’ Gibbs sampler (Kroese et al. 2011, Page 240). In other words, the new state is  $\boldsymbol{\beta}_k = \boldsymbol{\beta}_{k-1} + \lambda \mathbf{d}$ , where  $\mathbf{d}$  is a point uniformly distributed on the surface of the  $d$ -dimensional unit hyper-sphere, and the scalar  $\lambda$  is simulated according to:

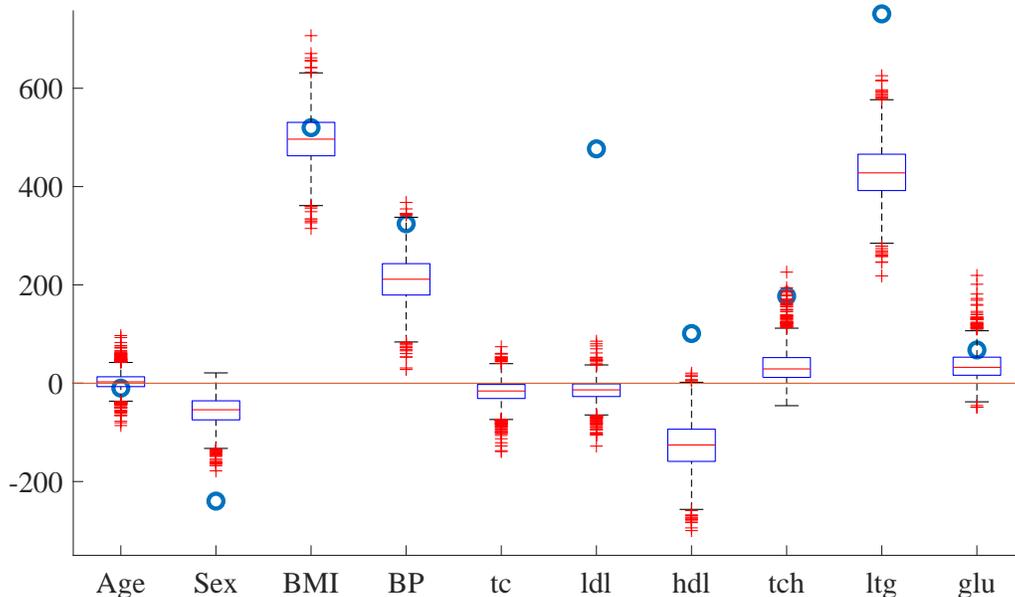
$$(\lambda | \mathbf{d}, \sigma_k, \boldsymbol{\beta}_{k-1}) \sim \varphi(\lambda | \mathbf{d}, \sigma_k, \boldsymbol{\beta}_{k-1}) := \frac{q(\boldsymbol{\beta}_{k-1} + \lambda \mathbf{d}, \sigma_k | \mathbf{y}, \gamma)}{\int q(\boldsymbol{\beta}_k + \xi \mathbf{d}, \sigma_k | \mathbf{y}, \gamma) d\xi}.$$

The conditional pdf  $\varphi(\lambda | \mathbf{d}, \sigma, \boldsymbol{\beta}_{k-1})$  is a univariate truncated normal, which can be simulated easily (Botev and L'Ecuyer 2017).

As a concrete illustration we use the ‘‘diabetes dataset’’ (Park and Casella 2008), consisting of  $n' = 442$  patients. For each patient, we have a record of  $d = 10$  predictor variables (age, sex, body mass index, and 7 blood serum measurements, so that  $\mathbf{X}$  is a matrix of size  $442 \times 10$ ), and a response variable, which measures the severity of nascent diabetes. We fix  $\gamma = 1200$ , which corresponds to the Lasso regularization parameter value used by Park and Casella (2008).

To simulate from the Bayesian posterior (7) we ran Algorithm 2 with splitting factor  $s = 100$  and  $n = 10^4$  using the following  $\tau = 4$  levels:  $(\gamma_1, \gamma_2, \gamma_3, \gamma_4) = (1907, 1368, 1230, 1200)$

to obtain the multiset  $\mathcal{Y}_U$ . The first three levels were chosen so that  $\rho_j \approx 0.01$  for  $j = 1, \dots, 3$ . The values for  $\tau, \gamma_1, \dots, \gamma_\tau$  were selected by running the adaptive pilot algorithm in (Botev et al. 2012)[Algorithm 4]. The marginal empirical distribution of each coefficient  $\beta_j$  is illustrated in Figure 1 as a boxplot.

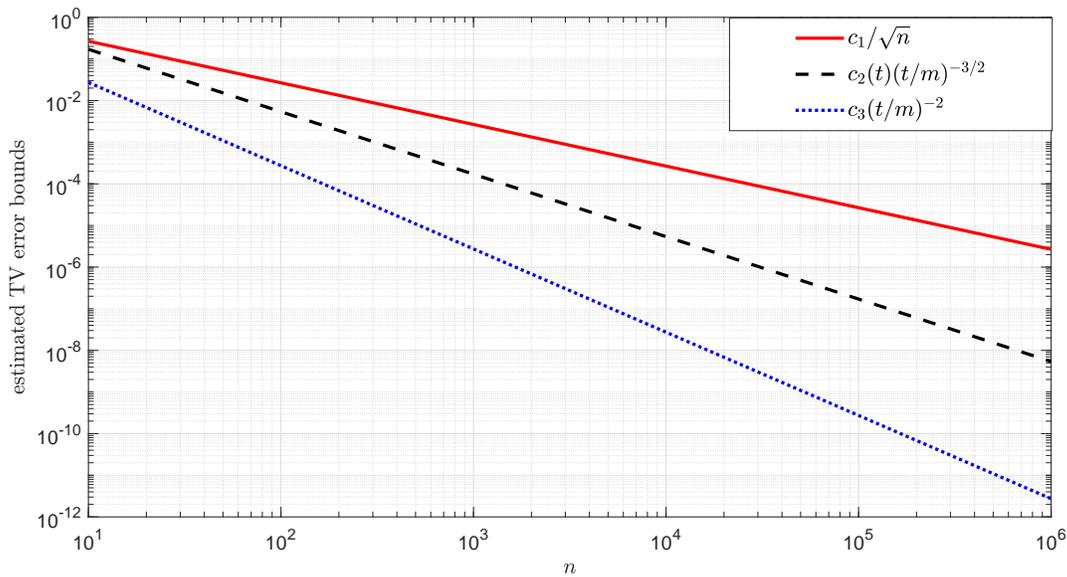


**Figure 1** Empirical marginal distributions of the ten coefficients  $\beta_j$  corresponding to the ten predictors, sampled approximately from (7). For comparison, the unconstrained (ordinary) least squares solution for each  $\beta_j$  is displayed as a circle.

## 5.2. Convergence Assessment via Theoretical Bounds

Using the output of Algorithm 2 from the previous section we calculated point estimates of the unknown terms,  $c_1, \tilde{c}_1(n), c_2(t), \tilde{c}_2(t), c_3, \psi_1, \psi_2$ , appearing in Theorems 1 through 6. Note that all the unknown terms depend on moments of  $M$ . For example, some of the point-estimates of the moments of  $M$  are  $(\mathbb{E}M, \mathbb{E}M^2) \approx (5.9, 71)$ . Figure 2 shows the estimates of  $c_1/n, c_2(t)(t/m)^{-3/2}$ , and  $c_3(t/m)^{-2}$ , which bound the TV error (see Theorems 1 to 3), on a common scale with  $t = n \times 5.9$  (since  $m = \mathbb{E}[M] \approx 5.9$ ).

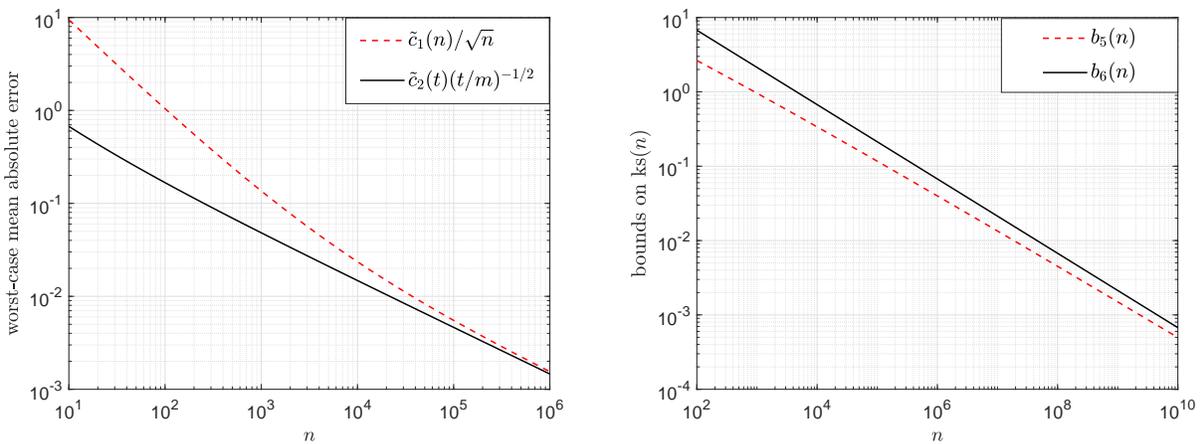
There is one major take-home message from Figure 2, namely, that Algorithm 4 (sampling to exceed  $t$  states) simulates more closely (in terms of TV error) from the target distribution  $\mathbb{Q}$  than Algorithm 3 ( $n$  iid non-empty replications). Of course, the downside



**Figure 2** Comparison of three bounds on the TV error,  $c_1/\sqrt{n}$ ,  $c_2(t) \times (t/m)^{-3/2}$ , and  $c_3 \times (t/m)^{-2}$ , where  $t = n \times m$ .

of using Algorithm 4 is that the number of trials,  $N(t)$ , is random (with expectation  $t/m$  for large  $t$ ).

In addition, reading off from Figure 2 we can see that if we run Algorithm 4 with  $t > 5.9 \times 10^3$ , then the TV error between  $\overline{\mathbb{Q}}_t$  and  $\mathbb{Q}$  is estimated as less than  $10^{-3}$  using the non-asymptotic bound  $c_2(t)(t/m)^{-3/2}$  and as less than  $10^{-5}$  using the asymptotic bound  $c_3(t/m)^{-2}$  (it is asymptotic, because we ignored the asymptotically negligible  $\mathcal{O}(\exp(-\omega t))$  term in Theorem 3).



**Figure 3** Left: estimates of the worst-case mean absolute error; Right: estimates of the expected TV error.

As for the mean absolute error, the left pane of Figure 3 shows the estimated bounds  $\tilde{c}_1(n)n^{-1/2}$  and  $\tilde{c}_2(t)(t/m)^{-1/2}$  given in Theorems 1 and 2, respectively, using  $t = n \times m$ .

It is clear that the bound  $\tilde{c}_2(t)(t/m)^{-1/2}$  is always smaller. Note that both bounds are asymptotically equivalent to first order — as  $n$  becomes larger, the two bounds converge to each other. Based on the mean absolute error, in this example we again conclude that Algorithm 4 (sample more than  $t = m \times n$  states) is a better performing sampler than Algorithm 3 ( $n$  iid non-empty runs).

Next, we apply the results of Theorems 5 and 6 to bound the expectation of the Kolmogorov-Smirnov statistic,  $\text{ks}(n)$ , given in (6). Let  $b_5(n)$  and  $b_6(n)$  be the upper bounds on (6) derived in Theorems 5 and 6, respectively (here  $v = d + 1 = 11$ ). The right pane of Figure 3 shows the estimated bounds on the value of  $\text{ks}(n)$ . There are a number of observations to be made.

First, we can see that for the range of the plot,  $b_5(n) = \mathcal{O}(\ln(n)/\sqrt{n})$  yields a better risk bound than  $b_6(n) = \mathcal{O}(1/\sqrt{n})$  (despite the superior convergence rate of  $b_6$ ). This is because, as mentioned previously, the constant  $\psi_2$  in Theorem 6 is much larger than  $\psi_1$  in Theorem 5. In fact, the cross-over for which ultimately  $b_6(n) < b_5(n)$  happens for  $n > 10^{19}$  (not shown on Figure 3).

Second, from the right pane of Figure 3 we can see that the expectation of the Kolmogorov-Smirnov statistic is indeed the most stringent error criteria, because we need a very large  $n$  to guarantee an acceptably small error (at least  $n > 10^7$  to make  $b_5(n)$  smaller than  $10^{-2}$ ).

Third, we observe that since the transition kernel,  $\kappa_l$ , has a density (it is the transition pdf of a Gibbs Markov chain), Theorem 4 ensures the almost sure convergence of the empirical TV uniformly over the class  $\mathcal{A}$  of all convex subsets, that is,  $\sup_{A \in \mathcal{A}} |\hat{\mathbb{Q}}_n(A) - \mathbb{Q}(A)| \rightarrow 0$  with probability one.

Finally, we note that our convergence results do not theoretically quantify the speed of convergence of the Markov chains, induced by the kernels  $\kappa_l$ . This dynamics is captured by the moments of  $M$ , which we estimate empirically, but not theoretically. To analyze theoretically the growth of the moments of  $M$  will require an analysis of the speed of convergence of all Markov chains used in Algorithm 1.

### 5.3. Comparison with Sequential Monte Carlo for Rare Event Estimation

In the Bayesian context, the rare-event probability  $\ell(\gamma) = \mathbb{P}(\|\boldsymbol{\beta}\|_1 \leq \gamma)$  is the normalizing constant of the posterior (7), also called the *model evidence* or *marginal likelihood*, which is of importance in model selection and inference.

From equation (4) above, we can see that an estimator of  $\ell$  using  $n = 10^4$  independent runs of Algorithm 1 is  $\hat{\ell} := (\overset{\circ}{M}_1 + \dots + \overset{\circ}{M}_n) / (ns^{\tau-1})$  with relative error  $\sqrt{\text{Var}(\overset{\circ}{M})} / (\hat{\ell}\sqrt{n})$ . We obtained the estimate of  $\hat{\ell} = 2.4 \times 10^{-8}$  with estimated relative error of 3.6%.

For completeness, and as a benchmark to our results, we compared the performance of Algorithm 1 with the popular sequential Monte Carlo (SMC) method for rare-event estimation of Cérou et al. (2012), as described on top of page 798, column 1. For the SMC we used the same intermediate thresholds  $(\gamma_1, \gamma_2, \gamma_3, \gamma_4) = (1907, 1368, 1230, 1200)$  (in the notation on page 798, we have  $A_k := \{\|\boldsymbol{\beta}\|_1 < \gamma_k\}$ ) and a total simulation effort of  $6 \times 10^6$  particles to estimate  $\ell$ . This is roughly twice the average simulation effort for  $n$  runs of Algorithm 1, which is approximately  $n \times \sum_{k=1}^{\tau} \frac{1}{\rho_k} \approx 3.4 \times 10^6$ . Despite this, the relative error of the SMC estimator of  $\ell$  was estimated as 12%, or about three times larger than the relative error of  $\hat{\ell}$ .

The observation that the GS algorithm can, under certain conditions, perform better than sequential Monte Carlo methods is known and is already explained in Botev and Kroese (2012). Briefly, the GS sampler is expected to outperform standard SMC methods when the Markov chain induced by  $\kappa_l$  converges slowly to its stationary pdf (3). Conversely, when the Markov chain at each level  $l$  mixes fast (the particles follow the law of (3) almost exactly), then SMC methods are to be preferred. As previously explained (Botev and Kroese 2012), unlike standard SMC methods, the GS sampler does not have a bootstrap resampling step, which is advantageous when the transition kernel  $\kappa_l$  fails to create enough “diversity” in the samples (bootstrap resampling reduces the diversity). This advantage, however, disappears if the Markov chains at each level are mixing fast, and as a result using a fixed number of particles at each level (Cérou et al. 2012, Page 798) leads to superior accuracy compared to using a random number of particles (as in the GS Algorithm 1).

## 6. Summary and Conclusions

We presented two different implementations of the generalized splitting method that can be used to simulate approximately from a conditional density in high dimensions. In the

first implementation, we construct an empirical distribution  $\widehat{\mathbb{Q}}_n$  from  $n$  iid non-empty replications of the GS sampler (Algorithm 1). In the second implementation, we construct an empirical distribution  $\widehat{\mathbb{Q}}_{N(t)}$  by running Algorithm 2 until we have more than  $t$  states in total. In both implementations,  $\widehat{\mathbb{Q}}_n$  and  $\widehat{\mathbb{Q}}_{N(t)}$ , and their respective expectations  $\mathbb{Q}_n$  and  $\overline{\mathbb{Q}}_t$ , aim to approximate the true distribution  $\mathbb{Q}$ .

To assess the quality of the approximations we derived non-asymptotic bounds on three different error criteria: (1) the total variation errors of  $\mathbb{Q}_n$  and  $\overline{\mathbb{Q}}_t$ , widely used in MCMC convergence analysis; (2) the mean absolute errors of  $\widehat{\mathbb{Q}}_n$  and  $\widehat{\mathbb{Q}}_{N(t)}$ ; and (3) the expected total variation error of  $\widehat{\mathbb{Q}}_n$ .

The main take-away messages are as follows. First, the GS sampler in Algorithm 4, which samples until we have more than  $t$  states in total, converges faster than the GS sampler in Algorithm 3, which samples  $n$  iid non-empty replications.

Second, the proposed splitting samplers provide a simple qualitative method for assessing whether they are sampling accurately from the target distribution. Any unknown constants and terms in the theoretical error estimates depend only on moments of the number  $M$  of particles, which can be readily estimated from the simulation output. This allows us to make qualitative statements such as “choose  $n > 10^3$  to (approximately) obtain a total variation error of less than  $10^{-3}$ ”, or to rank the performance of different implementations of the algorithms.

Finally, we have confirmed that, under certain conditions, generalized splitting can be more efficient than sequential Monte Carlo in estimating rare-event probabilities. This observation extends not just to estimation, but approximate sampling as well, because if an algorithm is not the most efficient in estimating a rare-event probability, then it will also not be the most efficient algorithm to simulate conditional on the rare event.

## References

- Andrieu C, Doucet A, Holenstein R (2010) Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* 72:1–33.
- Botev ZI, Kroese DP (2012) Efficient Monte Carlo simulation via the generalized splitting method. *Statistics and Computing* 22(1):1–16.
- Botev ZI, L'Ecuyer P (2017) Simulation from the normal distribution truncated to an interval in the tail. *10th EAI International Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2016*, 23–29 (ACM), URL <http://dx.doi.org/DOI:10.4108/eai.25-10-2016.2266879>.

- Botev ZI, L'Ecuyer P, Tuffin B (2011) An importance sampling method based on a one-step look-ahead density from a markov chain. *Proceedings of the 2011 Winter Simulation Conference (WSC)*, 528–539 (IEEE).
- Botev ZI, L'Ecuyer P, Tuffin B (2012) Dependent failures in highly reliable static networks. *Proceedings of the 2012 Winter Simulation Conference (WSC)*, 1–12 (IEEE), URL <http://dx.doi.org/10.1109/WSC.2012.6465033>.
- Botev ZI, Ridder A (2014) Variance reduction. *Wiley StatsRef: Statistics Reference Online* 1–6.
- Botev ZI, Vaisman S, Rubinstein RY, L'Ecuyer P (2014) Reliability of stochastic flow networks with continuous link capacities. *Proceedings of the 2014 Winter Simulation Conference*, 543–552 (IEEE Press).
- Bréhier CE, Gazeau M, Goudenège L, Lelièvre T, Rousset M (2016) Unbiasedness of some generalized adaptive multilevel splitting algorithms. *The Annals of Applied Probability* 26(6):3559–3601.
- Cérou F, LeGland F, Del Moral P, Lezaud P (2005) Limit theorems for the multilevel splitting algorithm in the simulation of rare events. M E Kuhl FBA N M Steiger, Joines JA, eds., *Proceedings of the 2005 Winter Simulation Conference*, 682–691 (IEEE Press).
- Cérou F, Moral PD, Furon T, Guyader A (2012) Sequential Monte Carlo for rare event estimation. *Statistics and computing* 22(3):795–808.
- Dean T, Dupuis P (2009) Splitting for rare event simulation: A large deviation approach to design and analysis. *Stochastic Processes and their Applications* 119:562–587.
- Devroye L, Lugosi G (2001) *Combinatorial methods in density estimation* (Springer, New-York).
- Garvels MJ, Ommeren JKV, Kroese DP (2002) On the importance function in splitting simulation. *Transactions on Emerging Telecommunications Technologies* 13(4):363–371.
- Glasserman P, Heidelberger P, Shahabuddin P, Zajic T (1999) Multilevel splitting for estimating rare event probabilities. *Operations Research* 47(4):585–600.
- Jones GL, Hobert JP (2001) Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statistical Science* 16(4):312–334.
- Kahn H, Harris TE (1951) Estimation of particle transmission by random sampling. *National Bureau of Standards Applied Mathematics Series* 12:27–30.
- Kroese DP, Taimre T, Botev ZI (2011) *Handbook of Monte Carlo methods*, volume 706 (John Wiley & Sons).
- L'Ecuyer P, Botev ZI, Kroese DP (2018) On a generalized splitting method for sampling from a conditional distribution. *Proceedings of the 2018 Winter Simulation Conference*, 1694–1705 (IEEE Press).
- L'Ecuyer P, Demers V, Tuffin B (2007) Rare-events, splitting, and quasi-Monte Carlo. *ACM Transactions on Modeling and Computer Simulation* 17(2):Article 9, 45 pages.
- L'Ecuyer P, LeGland F, Lezaud P, Tuffin B (2009) Splitting techniques. Rubino G, Tuffin B, eds., *Rare Event Simulation Using Monte Carlo Methods*, 39–62 (Wiley), chapter 3.

- Park T, Casella G (2008) The Bayesian lasso. *Journal of the American Statistical Association* 103(482):681–686.
- Sauer N (1972) On the density of families of sets. *Journal of Combinatorial Theory, Series A* 13(1):145–147.
- Taimre T, Kroese DP, Botev ZI (2019) Monte Carlo methods. *Wiley StatsRef: Statistics Reference Online*  
DOI: 10.1002/9781118445112.stat03619.pub2.
- Tuffin B, Saggadi S, L'Ecuyer P (2014) An adaptive zero-variance importance sampling approximation for static network dependability evaluation. *Computers and Operations Research* 45:51–59.
- Vapnik V (2013) *The nature of statistical learning theory* (Springer-Verlag).