Examples comparing Importance Sampling and the Metropolis algorithm

Federico Bassetti

Università degli Studi di Pavia, Dipartimento di Matematica,

via Ferrata 1, 27100 Pavia, Italy

Persi Diaconis

Department of Mathematics and Statistics, Stanford University,

Stanford, California 94305, USA

Abstract

Importance sampling, particularly sequential and adaptive importance sampling, have emerged as competitive simulation techniques to Markov–chain Monte–Carlo techniques. We compare importance sampling and the Metropolis algorithm as two ways of changing the output of a Markov chain to get a different stationary distribution.

1   Introduction.

Let $\mathcal{X}$ be a finite set and $\pi(x)$ be a probability on $\mathcal{X}$. For $f : \mathcal{X} \to \mathbb{R}$, we want to approximate

$$\mu = \sum_x f(x)\pi(x). \tag{1}$$

Suppose we have available a reversible Markov chain $K(x,y)$ on $\mathcal{X}$ with stationary distribution $\sigma(x) > 0$ for all x in $\mathcal{X}$. Two classical procedures are available.

Metropolis: Change the output of the K chain to have stationary distribution π by constructing

$$M(x,y) = \begin{cases} K(x,y)A(x,y) & x \neq y \quad A(x,y) := \min\left(\frac{\pi(y)K(y,x)}{\pi(x)K(x,y)}, 1\right) \\ K(x,x) + \sum_{z \neq x} K(x,z)(1 - A(x,z)) & x = y \end{cases}$$

Generate Y_1 from π and then $Y_2, \ldots, Y_N$ from $M(x,y)$. It follows that

$$\hat{\mu}_M = \frac{1}{N}\sum_{i=1}^N f(Y_i) \tag{2}$$

is an unbiased estimator of μ, the *Metropolis estimate.*

Importance sampling: Generate X_1 from σ and then $X_2, \ldots, X_N$ from $K(x,y)$. Then

$$\hat{\mu}_I = \frac{1}{N}\sum_{i=1}^N \frac{\pi(X_i)}{\sigma(X_i)}f(X_i) \tag{3}$$

1

is an unbiased estimate of μ, the *importance sampling estimate*. Often one uses

$$\tilde{\mu}_I = \frac{1}{\sum_{i=1}^{N} \frac{\pi(X_i)}{\sigma(X_i)}} \sum_{i=1}^{N} \frac{\pi(X_i)}{\sigma(X_i)} f(X_i)$$

instead of $\hat{\mu}_I$. The advantage for choosing $\tilde{\mu}_I$ instead of $\hat{\mu}_I$ is that the importance sampling ratios only need to be evaluated up to an unknown constant.

The Metropolis algorithm was introduced by Metropolis, Rosenbluth, Roseenbluth, Teller and Teller in [12] and later generalized by Hastings [5] and Peskun [13, 14]. There exists a large body of literature on Metropolis algorithm, the interested reader is referred to [19], [10], [15] and references therein. For this introduction, we have started both Markov chains out in their stationary distribution. For the study of the rate of convergence of Metropolis algorithm see the survey [2]. The systematic development of importance sampling began in the 1950's with works by Khan [6, 7]. See also [17], [4], [11]. More recent references can be found in [10].

Both $\hat{\mu}_M$ and $\hat{\mu}_I$ take the output of the Markov chain K and reweight to get an unbiased estimate of μ. The work involved is comparable and it is natural to ask which estimate is better.

In this note we address this question through examples; a random walk on binary d–tuples with K based on changing a random coordinate (Section 3) and independence proposal chain (Section 4). Moreover in Section 5 we discuss a problem of Knuth on non self intersecting paths and develop the theory for monotone paths in fairly complete detail. In most of our examples the Metropolis algorithm is either comparable or else dominates, sometimes by an exponential amount. The proofs are based on explicit spectral decompositions which give exact expression for variances as determined in the following section.

2 Variance computation.

Let $P(x, y)$ be a reversible Markov chain on the finite set $\mathcal{X}$ with stationary distribution $p(x)$. Thus, $p(x)P(x, y) = p(y)P(y, x)$. Throughout we assume all Markov chains are ergodic so p is the unique stationary distribution for P. Let $L^2(p) = \{f : \mathcal{X} \to \mathbb{R}\}$ with $< f, g >_p = E_p(fg) = \sum_x f(x)g(x)p(x)$. Reversibility is equivalent to $P : L^2 \to L^2$ being self–adjoint. Here $Pf(x) = \sum_y f(y)P(x, y)$. The spectral theorem implies that P has real eigenvalues $1 = \beta_0 > \beta_1 \geq \beta_2 \geq \cdots \geq \beta_{|\mathcal{X}|-1} > -1$ with an orthonormal basis of eigen–functions

$\psi_i : \mathcal{X} \to \mathbb{R}$ $(P\psi_i(x) = \beta_i \psi_i(x),\ < \psi_i, \psi_j >_p = \delta_{ij})$.

Proposition 2.1 *Let $f \in L^2(p)$ have $\sum_x f(x)p(x) = 0$, expand $f(x) = \sum_{i \geq 1} a_i \psi_i(x)$ (with $a_i = <f, \psi_i>_p$). Let Z be chosen from p and $Z_1, \ldots, Z_N$ be a realization of the $P(x,y)$ chain. Then*

$$\hat{\mu}_P = \frac{1}{N} \sum_{i=1}^{N} f(Z_i)$$

has variance

$$Var_p(\hat{\mu}_P) = \frac{1}{N^2} \sum_{k \geq 1} |a_k|^2 W_N(k) \tag{4}$$

with

$$W_N(k) = \frac{N + 2\beta_k - N\beta_k^2 + 2\beta_k^{N+1}}{(1 - \beta_k)^2}. \tag{5}$$

Proof. Because $\hat{\mu}$ has mean zero,

$$Var_p = E_p(\hat{\mu}_p^2) = \frac{1}{N^2} \sum_{i,j} Ef(Y_i)f(Y_j).$$

For $i \leq j$,

$$Ef(Y_i)f(Y_j) = E_p\{(\sum_k a_k \psi_k(Y_i))(\sum_l a_l \psi_l(Y_l))\} = E_p\{(\sum_k a_k \psi_k(Y_i))E_p(\sum_l a_l \psi_l(Y_l)|Y_i)\}$$

$$= E_p\{(\sum_k a_k \psi_k(Y_i))(\sum_l a_l \beta_l^{j-i} \psi_l(Y_i))\} = \sum_k a_k^2 \beta_k^{j-i}.$$

The last equality uses orthonormality of ψ_j. The next to last equality uses $E_p(\psi_l(Y_j)|Y_i) = \beta_l^{j-i}\psi_l(Y_i)$. Summing over i, j, using the identity $\sum_{1 \leq i < j \leq N} x^{j-i} = \{(N-1)x - Nx^2 + x^{N+1}\}/(1-x)^2$, one gets

$$E_p(\hat{\mu}_P^2) = \frac{1}{N^2} \sum_k a_k^2 \{N + 2 \sum_{1 \leq i < j \leq N} \beta_k^{j-i}\} = \frac{1}{N^2} \sum_k a_k^2 W_N(k).$$

Remark 1 (a) If $\beta_k = 1 - h_k$ for N large and h_k small, $W_N(k) \sim 2N/h_k$ and $Var_p(\hat{\mu}_P) \sim \frac{2}{N} \sum_{k \geq 1} a_k^2/h_k$. Of course this last is just heuristic. We will see that it is accurate in examples.

More formally,

$$\sigma_\infty^2(\hat{\mu}_P) := \lim_{N \to +\infty} NVar_p(\hat{\mu}_P) = \sum_{k \geq 1} |a_k|^2 \frac{1 + \beta_k}{1 - \beta_k}$$

$$\leq \frac{2}{1 - \beta_1} \|f\|_{2,p}^2. \tag{6}$$

3

This last inequality is classical and it is the usual way of relating spectral gap to asymptotic variance. It is used to compare proposal chains [13] and as a standard bound or rough estimate of the actual variance of the estimator. For small state spaces and long runs, this is reasonable. However, for large state spaces and runs a few multiples of the relaxation time, it can be badly off.

(b) Laurent Saloff-Coste suggests that the asymptotic variance can also be bounded by

$$\sigma_\infty^2(\hat{\mu}_P) \le 2|a_*|^2 \sum_{k \ge 1} \frac{1}{1 - \beta_k} \tag{7}$$

with $a_* := \max_{i \ge 1} |a_i|$.

The following examples show that both bounds (6) and (7) are useful.

Example 1 Let $\mathcal{X} = \mathbb{Z}_n$, the integers modulo n, with $n = 2m - 1$ an odd number. Let

$$P(x, y) = \begin{cases} 1/2 & \text{if } |x - y| = 1 \\ 0 & \text{otherwise} \end{cases}$$

be the transition matrix for simple random walk. This has stationary distribution $p(x) = 1/n$, and the eigenvalues and orthonormal eigen functions are well known:

$$\beta_0 = 1, \psi_0 = 1, \quad \beta_j = \cos(2\pi j/n), \quad \psi_j^c(h) = \sqrt{2}\cos(2\pi jh/n), \ \psi_j^s(h) = \sqrt{2}\sin(2\pi jh/n),$$

$1 \le j \le (n-1)/2$ (the non–unit eigenvalues have multiplicity two). If $f(h) = \delta_0(h) - \frac{1}{n}$, then $a_j^c = \sqrt{2}/n, \ a_j^s = 0$. The asymptotic variance is

$$\sigma_\infty^2 = \sum_{j=1}^{\frac{n-1}{2}} \frac{2}{n^2} \frac{1 + \cos(2\pi j/n)}{1 - \cos(2\pi j/n)} \sim \frac{\pi^2}{3}.$$

The bound (6) is of order $n^2(1-1/n)/n \sim n$, thus here the easiest bound is off while the bound (7) is bounded in n. Similar results hold for $f(h) = \delta_{[-a,a]}(h)$. Saloff-Coste has suggested that the a_* bound (7) will be better for one and two dimensional random walk problems but will not be an improvement more generally.

Example 2 Let $\mathcal{X} = \mathbb{Z}_2^d$ be the hypercube. Let

$$P(x, y) = \begin{cases} 1/d & \text{if } |x - y| = 1 \\ 0 & \text{otherwise} \end{cases}$$

4

be the transition matrix for the nearest neighbor random walk. This has stationary distribution $p(x) = 1/2^d$. The eigenvalues and othonormal eigenfunctions are well known. It is convenient to index them by $x \in \mathbb{Z}_2^d$:

$$\beta_x = 1 - \frac{2|x|}{d}, \quad \psi_x(y) = (-1)^{x \cdot y}.$$

Here $|x| = \sum_{i=1}^d x_i$ is the Hamming–weight. Note that $1 - 2j/d$ has multiplicity $\binom{d}{j}$, $0 \le j \le d$. If $f(x) = \delta_0(x) - 1/2^d$, the fourier coefficients at $x \ne 0$ are $a_x = 1/2^d$. The asymptotic variance is

$$\sigma_\infty^2 = \frac{1}{2^{2d}} \sum_{j=1}^d \binom{d}{j} \frac{2 + 2j/d}{2j/d} = \frac{1}{2^{2d}} \sum_{j=1}^d [\frac{d}{j} \binom{d}{j} + 2^d] \le \frac{4}{2^d}.$$

On the other hand $\|f\|_2^2 = \frac{1}{2^d}(1 - \frac{1}{2^d})$. The crude upper bound from (6) is $2d \frac{1}{2^d}(1 - \frac{1}{2^d})$ which is off by a factor of d. The a_* bound from (7) for σ_∞^2 is

$$\frac{2}{2^{2d}} \sum_{j=1}^{d-1} \binom{d}{j} \frac{d}{2j} \le \frac{C}{2^d}.$$

This has the right order.

If $f(x) = \sum_{j=1}^d (-1)^{x_j} = d - 2|x|$, $a_x = 0$ if $|x| > 1$ and $a_x = 1$ if $|x| = 1$. Thus $\sigma_\infty^2 = d \frac{2 - 2/d}{2/d} = d(d-1)$. Here, $\|f\|_2^2 = d$. The crude bound from Remark 1 is $2d/(2/d) = d^2$. The a_* bound is

$$\sum_{j=1}^d \binom{d}{j} \frac{d}{2j} \sim c 2^d$$

for some c. This is wildly off.

3   The hypercube.

Let $\mathcal{X} = \mathbb{Z}_2^d$ be the set of binary d–tuples. Let

$$\pi(x) = \theta^{|x|}(1 - \theta)^{d - |x|}$$

with $1/2 \le \theta \le 1$ and $|x|$ the number of ones in the d–tuple x. Let the base chain $K(x, y)$ be given by "from x, pick a coordinate at random and change it to one or zero with probability p or $1 - p$" ($1/2 < p \le 1$). The K–chain has stationary distribution

$$\sigma(x) = p^{|x|}(1 - p)^{d - |x|}.$$

This example models a high–dimensional problem where the desired distribution $\pi(x)$ is concentrated in a small part of the space. We have available a sampling procedure – run the chain $K(x, y)$– where the stationary distribution is roughly right (if p is close to θ) but not spot on.

Let us begin by diagonalizing the Metropolis chain. This may be presented as

$$M = \frac{1}{d} \sum_{i=1}^{d} M_i \tag{8}$$

with M_i the Metropolis construction operating on the i–th coordinate. The transition matrix restricted to the i^{th} coordinate is

$$
\begin{array}{c c}
 & \begin{array}{cc} 0 & \quad 1 \end{array} \\
\begin{array}{c} 0 \\ 1 \end{array} &
\left(\begin{array}{cc} \bar{p} & p \\ p\bar{\theta}/\theta & 1 - p\bar{\theta}/\theta \end{array} \right)
\end{array}
\qquad \text{with } \bar{p} = 1 - p, \ \bar{\theta} = 1 - \theta.
$$

This matrix has stationary distribution $(\bar{\theta}, \theta)$ on $\{0, 1\}$. The eigen–values are $\beta_0 = 1$, $\beta_1 = 1 - p/\theta$, with normalized eigen–vectors

$$\psi_0(0) = \psi_0(1) = 1, \qquad \psi_1(0) = \sqrt{\theta/\bar{\theta}}, \ \psi_1(1) = -\sqrt{\bar{\theta}/\theta}. \tag{9}$$

Proposition 3.1 *The Metropolis chain (8) on $\mathbb{Z}_2^d$ has 2^d eigenvalues and eigenvectors which will be indexed by $\zeta \in \mathbb{Z}_2^d$. These are*

$$\beta_\zeta = 1 - \frac{|\zeta| p}{d\theta} \tag{10}$$

$$\psi_\zeta(x) = \prod_{i=1}^{d} \psi_{\zeta_i}(x_i) = \prod_{i=1}^{d} \left(\sqrt{\frac{\theta}{\bar{\theta}}} \right)^{\zeta_i(1-x_i)} \left(-\sqrt{\frac{\bar{\theta}}{\theta}} \right)^{\zeta_i x_i} \tag{11}$$

with ψ_i defined in (9). The eigenvectors are orthonormal in $L^2(\pi)$.

Proof. This is a straight–forward verification from (9) and the basic structure of the product chains. For more details, see [1] Sec. 5. ◇

Using these tools we may compute the variance of the Metropolis algorithm for a variety of functions f. We take f to be the number of ones normalized to have mean zero.

Proposition 3.2 *On $\mathbb{Z}_2^d$, let*

$$f(x) = \sum_{i=1}^{d} (x_i - \theta). \tag{12}$$

Under the Metropolis chain (8), with $\hat{\mu}_M$ defined by (2) we have $\mu = 0$ and

$$Var_\pi(\hat{\mu}_M) = \frac{2d^2 \bar{\theta}\theta^2}{Np} - \frac{d\theta\bar{\theta}}{N} + \frac{2d^3\theta^3\bar{\theta}}{N^2 p^2} \left(1 - \frac{p}{d\theta} \right)^{N+1} + \frac{2d^3\bar{\theta}\theta^3}{N^2 p^2} \left(1 - \frac{p}{d\theta} \right).$$

Here

$$\sigma_\infty^2(\hat{\mu}_M) \sim \frac{2d^2\bar{\theta}\theta^2}{p} \qquad (d \to +\infty).$$

Proof. To use Proposition 2.1, we must compute the expansion of f in (12) with respect to the eigenbasis ψ_ζ of (11). For $\zeta \neq 0$ by orthogonality of ψ_ζ with $\psi_0 \equiv 1$, $\sum_x \psi_\zeta(x)f(x)\pi(x) = \sum_x(\sum_i x_i)\psi_\zeta(x)\pi(x)$. For fixed i, write x^i for $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_d)$. If $\zeta_i = 0$, $\sum_x x_i\psi_\zeta(x)\pi(x) = 0$. If $\zeta_i = 1$,

$$\sum_x x_i\psi_\zeta(x)\pi(x) = \sum_{x^i} \pi(x^i)\psi_{\zeta^i}(x^i)(0\bar{\theta}\psi_{\zeta_i}(0) - \theta\psi_{z_i}(1))$$

$$= \begin{cases} -\sqrt{\theta\bar{\theta}} & if \; \zeta^i = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Hence, $a_\zeta = -\sqrt{\theta\bar{\theta}}$ if $|\zeta| = 1$ and $a_z = 0$ otherwise. Now, Proposition 2.1 and (10) give

$$Var_\pi(\hat{\mu}_M) = \frac{d\bar{\theta}\theta}{N^2}\left(\frac{2Nh - Nh^2 + 2(1-h) + 2(1-h)^{N+1}}{h^2}\right)$$

with $h = (\theta d)^{-1}p$. $\diamond$

Consider next the importance sampling chain with the same K and σ considered above. Represent K as

$$K = \frac{1}{d}\sum_{i=1}^d K_i \tag{13}$$

with K_i having matrix (restricted to the i^{th} coordinate)

$$\begin{array}{cc} & \begin{array}{cc} 0 & 1 \end{array} \\ \begin{array}{c} 0 \\ 1 \end{array} & \left(\begin{array}{cc} \bar{p} & p \\ \bar{p} & p \end{array}\right) \end{array} \qquad \text{with } \bar{p} = 1 - p.$$

This matrix has stationary distribution $(\bar{p}, p)$ on $\{0, 1\}$. The eigenvalues are $\beta_0^* = 1$, $\beta_1^* = 0$ with normalized eigenvectors $\psi_0^*(0) = \psi_1^*(1) = 1$, $\psi_1^*(0) = \sqrt{p/\bar{p}}$, $\psi_1^*(1) = -\sqrt{\bar{p}/p}$. Arguing as for Proposition 3.2 we have the following

Proposition 3.3 *The Markov chain (13) on $\mathbb{Z}_2^d$ has 2^d eigenvalues and eigenvectors indexed by ζ in $\mathbb{Z}_2^d$. These are*

$$\beta_\zeta^* = 1 - \frac{|\zeta|}{d} \tag{14}$$

$$\psi_z^*(x) = \prod_{i=1}^d \psi_{\zeta_i}^*(x_i) = \prod_{i=1}^d \left(\sqrt{\frac{p}{\bar{p}}}\right)^{\zeta_i(1-x_i)}\left(-\sqrt{\frac{\bar{p}}{p}}\right)^{\zeta_i x_i}. \tag{15}$$

The eigenvectors are orthonormal in $L^2(\sigma)$ where $\sigma(x) = p^{|x|}(1-p)^{d-|x|}$.

7

The importance sampling estimate is

$$\hat{\mu}_I = \frac{1}{N}\sum_{i=1}^{N}\frac{\pi(X_i)}{\sigma(X_i)}f(X_i) \tag{16}$$

with

$$\frac{\pi(x)}{\sigma(x)} = ab^{|x|}$$

for $a = ((1-\theta)/(1-p))^d$, $b = (\theta(1-p)/(p(1-\theta)))^d$.

To use the machinery above, we need to compute the spectral coefficients.

Proposition 3.4 *Let ψ_ζ^* be defined as in (15) and*

$$g(x) = \Big(\sum_{i=1}^{d}(x_i - \theta)\Big)\frac{\pi(x)}{\sigma(x)}.$$

Then $< g, \psi_\zeta^* >_\pi = -\alpha|\zeta|\beta^{|\zeta|}$ with $\alpha := (\theta\bar{\theta}\sqrt{p/\bar{p}} + \bar{\theta}\theta\sqrt{\bar{p}/p})/(\bar{\theta}\sqrt{p/\bar{p}} - \theta\sqrt{\bar{p}/p})$ and $\beta := \bar{\theta}\sqrt{p/\bar{p}} - \theta\sqrt{\bar{p}/p}$.

Proof. Write $< g, \psi_\zeta^* >_\pi = \sum_{i=1}^{d} E_\sigma \frac{\pi(x)}{\sigma(x)}\psi_\zeta^*(x)(x_i - \theta)$. Under σ, the coordinates are independent taking values one or zero with probability p and $1-p$. The integrand is a product and we may compute it componentwise. Consider the i^{th} term in the sum. For $j \neq i$, the expectation of the j^{th} term in the product is 1 if $\zeta_j = 0$ and $\bar{\theta}\sqrt{p/\bar{p}} - \theta\sqrt{\bar{p}/p}$ if $\zeta_j = 1$. For $j = i$, the expectation of the i^{th} term in the product is 0 if $\zeta_i = 0$ and $-\theta\bar{\theta}\sqrt{p/\bar{p}} - \theta\bar{\theta}\sqrt{\bar{p}/p}$ if $\zeta_i = 1$. Computing the product and summing in i gives the stated result. $\Diamond$

Combining these results gives a formula for the variance.

Proposition 3.5 *For the Markov chain K of (13) and $f(x) = \sum_{i=1}^{d}(x_i - \theta)$, the importance sampling estimate $\hat{\mu}_I$ of (16) has mean zero and variance*

$$Var_\sigma(\hat{\mu}_I) = \frac{\alpha^2}{N^2}\sum_{i=1}^{d}\binom{d}{i}i^2\beta^{2i}W_N^*(i),$$

with α and β from Proposition 3.4 *and*

$$W_n^*(i) = \frac{2Nd}{i} - N + \frac{2d^2}{i^2}(1 - i/d) + \frac{2d^2}{i^2}(1 - i/d)^{N+1}.$$

Remark 2 (a) The lead term in $Var_\sigma(\hat{\mu}_I)$ is

$$\frac{\sigma^2}{N^2}\sum_{i=1}^{d}\binom{d}{i}i^2\beta^{2i}\frac{2Nd}{i} = \frac{2\alpha^2d}{N}\sum_{i=1}^{d}\binom{d}{i}i\beta^{2i} = 2\beta^2\frac{\alpha^2d^2}{N}(1 + \beta^2)^{d-1}.$$

8

The last equality used $d(1+x)^{d-1}x = \sum_{i=0}^{d} \binom{d}{i} i x^i$. For our running example, $\theta = 7/8$, $p = 3/4$, $\beta = -0.2887$, $1 + \beta^2 = 1.0833$; for large d and N, $Var_\sigma(\hat{\mu}_I)$ is exponentially worse (in d) than the Metropolis variance.

(b) The next term is

$$-\frac{\alpha^2}{N^2} \sum_{i=1}^{d} \binom{d}{i} i^2 \beta^{2i} N = -\frac{\alpha^2}{N}(1 + d\beta^2) s \beta^2 (1 + \beta^2)^{d-2}$$

the sum of these two lead terms is

$$\frac{\alpha^2 d \beta^2}{N}(1 + \beta^2)^{d-2} \{2(1 + \beta^2)d - (1 + d\beta^2)\}.$$

(c) For the next term we need

$$\sum_{i=1}^{d} \binom{d}{i} i^2 \beta^{2i} \frac{d^2}{i^2}(1 - \frac{i}{d}) = d^2((1 + \beta^2)^d - 1) - d^2(1 + \beta^2)^{d-1}\beta^2 = d^2((1 + \beta^2)^{d-1} - 1).$$

From this, the third term is $2d^2\alpha^2((1 + \beta^2)^{d-1} - 1)/N^2$.

(d) For the last term we need

$$2d^2 \sum_{i=1}^{d} \binom{d}{i} \beta^{2i} (1 - \frac{i}{d})^{N+1} \le 2d^2 \sum_{i=1}^{d} \binom{d}{i} \beta^{2i} e^{-i(N+1)/d}$$

$$\le 2d^2((1 + \beta^2 e^{-(N+1)/d})^d - 1).$$

From this the last term is positive and bounded above by $2d^2\alpha^2((1 + \beta^2 e^{-(N+1)/d})^d - 1)/N^2$.

(e) The bottom line;

$$Var_\sigma(\hat{\mu}_I) \sim \frac{2\alpha^2 d^2 \beta^2}{N}(1 + \beta^2)^{d-1}.$$

This is exponentially worse (in d) than $Var_\pi(\hat{\mu}_M) \sim \frac{2d^2\theta^2\bar{\theta}}{Np}$ from Proposition 3.2.

(f) From Remark 2 (a), the variance of the importance sampling estimator blows up as $d^2(1 + \beta^2)^d$. With $\beta^2 = \bar{\theta}^2 p/\bar{p} + \bar{p}/p\theta^2 - 2\theta\bar{\theta}$. It is natural to ask how close p must be to θ so that this doesn't blow up. If $p = \theta + \epsilon$, a straightforward calculation gives

$$\beta^2 = \frac{\epsilon^2}{\theta\bar{\theta}} + O(\epsilon^2).$$

It follows that ϵ of order $1/\sqrt{d}$ is required to keep the variance from exponential explosion.

For another example we take $f(x) = \delta_d(|x|) - \theta^d$. Again we compute the spectral coefficients.

Proposition 3.6 *Let ψ_ζ and ψ_ζ^* be defined as in (10) and (15). Let $f(x) = \delta_d(|x|) - \theta^d$ and $g(x) = f(x)\pi(x)/\sigma(x)$, then $\mu = E_\pi(f) = 0$ and $Var_\pi(f) = \theta^d(1 - \theta^d)$. Moreover*

$$a_\zeta := < f, \psi_\zeta >_\pi = (-1)^d \left(\sqrt{\frac{\bar\theta}{\theta}}\right)^{|\zeta|} \theta^d$$

$$a_\zeta^* := < f, \psi_\zeta^* >_\pi = (-1)^d \left(\sqrt{\frac{\bar\theta}{\theta}}\right)^{|\zeta|} \theta^d \left\{ 1 - \left(\frac{\theta - p}{1 - p}\right)^{|\zeta|} \right\}$$

holds true for every ζ with $|\zeta| \neq 0$, and $a_0 = a_0^* = 0$.

Proof. This is a straight–forward verification. Indeed, by orthogonality, $< f, \psi_\zeta >_\pi = < \delta_d(|x|), \psi_\zeta >_\pi$, and then $< f, \psi_\zeta >_\pi = \theta^d \psi_\zeta(\mathbf{1})$ where $\mathbf{1} = (1, 1, \ldots, 1)$. Moreover, arguing as in the proof of Proposition 3.4, we get

$$a_\zeta^* = < \delta_d(|x|), \psi_\zeta^* >_\pi - \theta^d < 1, \psi_z^* >_\pi$$

$$= (-1)^{|\zeta|} \left(\sqrt{\frac{\bar p}{p}}\right)^{|\zeta|} \theta^d - \theta^d (\bar\theta \sqrt{\frac{p}{\bar p}} - \theta \sqrt{\frac{\bar p}{p}})^{|\zeta|}$$

$$= (-1)^{|\zeta|} \left(\sqrt{\frac{\bar p}{p}}\right)^{|\zeta|} \theta^d \left(1 - \left(\frac{\theta - p}{1 - p}\right)^{|\zeta|}\right).$$

$\Diamond$

Combining the previous results we get

Proposition 3.7 *On $\mathbb{Z}_2^d$, let $f(x) = \delta_d(|x|) - \theta^d$. Under the Metropolis chain (8), with $\hat\mu_M$ defined by (2) we have $\mu = 0$ and*

$$Var_\pi(\hat\mu_M) = \frac{\theta^{2d}}{N^2} \sum_{i=1}^d \binom{d}{i} \left(\frac{\bar\theta}{\theta}\right)^i W_N^*(ip/\theta),$$

with

$$W_n^*(i) = \frac{2Nd}{i} - N + \frac{2d^2}{i^2}(1 - i/d) + \frac{2d^2}{i^2}(1 - i/d)^{N+1}.$$

Moreover for the Markov chain K of (13) the importance sampling estimate $\hat\mu_I$ of (16) has mean zero and variance

$$Var_\sigma(\hat\mu_I) = \frac{\theta^{2d}}{N^2} \sum_{i=1}^d \binom{d}{i} \left(\frac{\bar p}{p}\right)^i \left(1 - \left(\frac{\theta - p}{1 - p}\right)^i\right)^2 W_N^*(i).$$

Remark 3 (a) The lead term in $Var_\pi(\hat\mu_M)$ is

$$\frac{\theta^{2d}}{N} \sum_{i=1}^d \binom{d}{i} \left(\frac{\bar\theta}{\theta}\right)^i \left(\frac{2d\theta}{p}\frac{1}{i} - 1\right) = \frac{2\theta^{2d}\theta}{Np} dA_d(\bar\theta/\theta) - \theta^d(1 - \theta^d)/N,$$

10

with $A_d(x) := \sum_{i=1}^{d} \binom{d}{i} x^i / i$. Since $A_d(x) = (x(d+1))^{-1} \sum_{i=2}^{d+1} \binom{d+1}{i} x^i i/(i-1)$, for every positive x, we can write

$$\frac{[(1+x)^{d+1} - 1 - x(d+1)]}{x(d+1)} \leq A_d(x) \leq 2\frac{[(1+x)^{d+1} - 1 - x(d+1)]}{x(d+1)}. \tag{17}$$

Hence, the lead term in $Var_\pi(\hat{\mu}_M)$ can be written as $N^{-1}(2\theta^{d+1} d B_d(\bar{\theta}/\theta)/(p\bar{\theta}(d+1)) - Var_\pi(f))$ with $B_d(\bar{\theta}/\theta)$ bounded in d. More exactly (17) gives $[1 - \theta^{d+1} - \bar{\theta}\theta^d(d+1)] \leq B_d(\bar{\theta}/\theta) \leq 2[1 - \theta^{d+1} - \bar{\theta}\theta^d(d+1)]$. This suggests that $NVar_\pi(\hat{\mu}_M) = O(Var_\pi(f))$ for $d \to +\infty$. This will be formalized in the next proposition.

(b) The first term in $Var_\sigma(\hat{\mu}_I)$ can be written as

$$\frac{2\theta^{2d} d}{N} \sum_{i=1}^{d} \binom{d}{i} \left(\frac{\bar{p}}{p}\right)^i \frac{1}{i} \left[1 + \left(\frac{\theta - p}{1-p}\right)^{2i} - 2\left(\frac{\theta-p}{1-p}\right)^i\right]$$

$$= \frac{2\theta^{2d} d}{N} \left[A_d(\bar{p}/p) + A_d\left(\frac{(\theta-p)^2}{p(1-p)}\right) - 2A_d\left(\frac{\theta-p}{p}\right)\right]$$

$$= \frac{2\theta^d d}{N(d+1)} \left(\frac{\theta}{p}\right)^d B_d^*(p, \theta)$$

with $B_d^*(p, \theta) = (d+1)p^d [A_d(\bar{p}/p) + A_d\left(\frac{(\theta-p)^2}{p(1-p)}\right) - 2A_d\left(\frac{\theta-p}{p}\right)]$. Again using (17), it is easy to see that $d \mapsto B_d^*(p, \theta)$ is a bounded function. The second term in $Var_\sigma(\hat{\mu}_I)$ is

$$-\frac{\theta^d}{N} \left(\frac{\theta}{p}\right)^d \left[1 - 2\theta^d + \left(\frac{p\bar{p} + (\theta - p)^2}{\bar{p}}\right)^d\right]$$

$$= -\frac{\theta^d}{N} \left(\frac{\theta}{p}\right)^d \left[1 - 2\theta^d + \left(\frac{p + \theta^2 - 2\theta p}{\bar{p}}\right)^d\right]$$

$$= -\frac{\theta^d}{N} \left(\frac{\theta}{p}\right)^d C_d^*(p, \theta)$$

with $C_d^*(p, \theta) = [1 - 2\theta^d + \left(\frac{p+\theta^2-2\theta p}{\bar{p}}\right)^d]$. If $1/2 < p < \theta < 1$, then $(p + \theta^2 - 2\theta p)/\bar{p} < 1$. Hence, $C_d^*(p, \theta) = 1 + o(1)$ for $d \to +\infty$. While some simple computations show that $\theta(p + \theta^2 - 2\theta p)/(p(1-p)) < 1$ if $1/2 < \theta < p < 1$.

Combining the previous remarks we get

Proposition 3.8 *For f as in Proposition 3.7,*

$$\sigma_\infty^2(\hat{\mu}_M)/Var_\pi(f) = \frac{2\theta}{p\bar{\theta}(1 - \theta^d)} \frac{d}{d+1} B_d(\bar{\theta}/\theta) - 1 \sim K_1(\theta, p)$$

with $d \mapsto B_d(\bar{\theta}/\theta)$ bounded, K_1 being a suitable constant. Moreover,

$$\sigma_\infty^2(\hat{\mu}_I)/Var_\pi(f) = \frac{1}{1 - \theta^d} \left(\frac{\theta}{p}\right)^d \left[\frac{2d}{d+1} B_d^*(p, \theta) - C_d^*(p, \theta)\right] \sim K_2(\theta, p) \left(\frac{\theta}{p}\right)^d$$

11

with $d \mapsto [\frac{2d}{d+1} B_d^*(p, \theta) - C^*(p, \theta)]$ *bounded and strictly positive, for every p and θ such that* $1/2 < p < \theta < 1$, *while*

$$\sigma_\infty^2(\hat{\mu}_I)/Var_\pi(f) \sim K_3(\theta, p) \left[\left(\frac{\theta p + \theta^3 - 2\theta^2 p}{p(1-p)} \right)^d + \left(\frac{\theta}{p} \right)^d \right]$$

with $\frac{\theta p + \theta^3 - 2\theta^2 p}{p(1-p)} < 1$ *if* $1/2 < \theta < p < 1$. K_2, K_3 *being suitable constants.*

Remark 4 The last proposition shows that, for $1/2 < p < \theta < 1$, the normalized asymptotic variance of $\hat{\mu}_I$ is exponentially worse (in d) than the normalized asymptotic variance of $\hat{\mu}_M$, while it is exponentially better if $1/2 < \theta < p < 1$. On the one hand, this fact agrees with the heuristic that the importance sampling estimator of $\pi(A)$ performs better than the Monte Carlo estimator, whenever the importance distribution σ puts more weight on A than π. On the other hand, the last example shows that even a small "loss of weight" in A can cause an exponentially worse behavior.

4 Independence sampling.

In this section the proposal chain is a sequence of independent and identically distributed variables with common probability density σ. Because of this, the structure of the state space does not matter. Throughout we take $\mathcal{X} = \{0, 1, \ldots, m-1\}$ with $\sigma(i) > 0$ fixed and $\pi(i)$ the desired distribution. Without loss, suppose the states are numbered so that the importance weights $\pi(i)/\sigma(i)$ are decreasing, i.e.

$$\frac{\pi(0)}{\sigma(0)} \geq \frac{\pi(1)}{\sigma(1)} \geq \cdots \geq \frac{\pi(m-1)}{\sigma(m-1)}. \tag{18}$$

This section makes use of an explicit diagonalization of the Metropolis chain due to Jun Liu [9].

Metropolis. For the proposal chain of independent picks from σ, the Metropolis chain starts with a pick from π. From state i it proceeds by choosing j from σ; if $j \leq i$ the chain moves to j. If $j > i$ the chain stays at j with probability $(\pi(j)\sigma(i)/\pi(i)\sigma(j))$ and remains at i with probability $(1 - \pi(j)\sigma(i))/\pi(i)\sigma(j))$. Liu [9] proves that this chain has eigenvalue 1 and, for $1 \leq k \leq m-1$

$$\beta_k = \sum_{i \geq k} \left(\sigma(i) - \pi(i) \frac{\sigma(k)}{\pi(k)} \right), \tag{19}$$

$$\psi_k = (\underbrace{0,\ldots,0}_{k-1}, S_\pi(k+1), -\pi(k), \ldots, -\pi(k)), \qquad S_\pi(k+1) := \sum_{j=k+1}^{m-1} \pi(j). \tag{20}$$

These eigen–vectors are orthogonal in $L^2(\pi)$. From these facts and Proposition 2.1 we get the following

Proposition 4.1 *For the Metropolis algorithm based on independent proposals with distribution $\sigma(i)$ and stationary distribution $\pi(i)$, let $f : \mathcal{X} \to \mathbb{R}$ have representation $f(i) = \sum_{k=1}^{m-1} a_k \psi_k(i)$, with $\mu = \sum_{i=0}^{m-1} f(i)\pi(i) = 0$. Then the Metropolis estimator $\hat{\mu}_M$ of (2) satisfies*

$$Var_\pi(\hat{\mu}_M) = \frac{1}{N^2} \sum_{k=1}^{m-1} b_k^2 W_N(k) \tag{21}$$

with $W_N(k)$ given by (5) using the eigenvalues of (19) and

$$b_k := \zeta_k \Big(f(k) S_\pi(k+1)\pi(k) - \pi(k) \sum_{j \geq k} f(j)\pi(j) \Big), \qquad \zeta_k := (\pi(k) S_\pi(k) S_\pi(k+1))^{-1/2}$$

The following lemma supplements J.Liu results

Lemma 4.2 *With notation as above, the eigenvalues β_k in (19) satisfy*

$$1 - \frac{\sigma(0)}{\pi(0)} \geq \beta_1 \geq \beta_2 \geq \cdots \geq \beta_{m-1} = 0$$

Proof. $\beta_k = \sigma(k) + \cdots + \sigma(m-1) - \frac{\sigma(k)}{\pi(k)}(\pi(k) + \ldots \pi(m-1)) \geq \sigma(k+1) + \cdots + \sigma(m-1) - \frac{\sigma(k+1)}{\pi(k+1)}(\pi(k+1) + \ldots \pi(m-1)) = \beta_k$. For β_1, note that $\beta_1 = 1 - \sigma(0) - \frac{\sigma(0)}{\pi(0)}(1 - \pi(0)) \leq 1 - \frac{\sigma(0)}{\pi(0)}$. For β_{m-1} note that $\beta_{m-1} = \sigma(m-1) - \frac{\sigma(m-1)}{\pi(m-1)}\pi(m-1) = 0$. $\diamond$

The variance of the importance sampling estimator is (for $\mu = 0$)

$$Var_\sigma(\hat{\mu}_I) = \frac{1}{N^2} \sum_l \left(\frac{\pi(l)}{\sigma(l)} f(l) \right)^2 \sigma(l). \tag{22}$$

We record a different expression for this, similar to J.Liu's above.

Lemma 4.3 *Let P be the Markov chain on $\{0, 1, \ldots, m-1\}$ with all rows equal to σ. Then, P has one eigenvalue $\beta_0 = 1$ with $\psi_0(i) \equiv 1$ and $m-1$ eigenvalues 0. An orthogonal basis for the zero eigenspace in $L^2(\sigma)$ is*

$$\psi_k = (0, \ldots, 0, S_\sigma(k+1), -\sigma(k), \ldots, -\sigma(k)), \qquad 1 \leq k \leq m-1. \tag{23}$$

Remark 5 If $\sigma = \pi$, this basis agrees with the Metropolis basis of (20). This must be because then P commutes with the Metropolis chain.

13

Proposition 4.4 *Let $g(i)$ satisfy $\sum_{i=0}^{m-1} g(i)\sigma(i) = 0$. Then*

$$Var_\sigma(g) = \sum_{i=0}^{m-1} g(i)^2 \sigma(i) = \sum_{i=1}^{m-1} \Big(g(i)S_\sigma(i+1) - \sigma(i)\sum_{j\geq i} g(j)\sigma(j) \Big)^2 / (\sigma(i)S_\sigma(i)S_\sigma(i+1)).$$

Example 3 Let us compare the Metropolis and the importance estimator in the following case. For $\mathcal{X} = \{0,1,\ldots,m-1\}$, let $\pi(i) = a^i c(a)$ for fixed a with $0 < a < 1$ with $c(a) = (1-a)/(1-a^m)$. For the proposal chain take $\sigma(i) = 1/m$, $0 \leq i \leq m-1$. Take $f(i) = (i - \mu)$ with $\mu = c(a)a(1 - a^{m-1}(a + (1-a)m))/(1-a)^2$, thus $\mu = E_\pi(f) = 0$. For a fixed and m large $\mu \sim a/(1-a)$.

Metropolis. We must compute the expansion of f in the basis ψ_k, $1 \leq k \leq m-1$. We have

$$a_k = \sum_j f(j)\psi_k(j)\pi(j) = (k - \mu)S_\pi(k+1)\pi(k) - \pi(k)\sum_{j=k+1}^{m-1}(j - \mu)\pi(j).$$

Here $S_\pi(k+1) = c(a)(1 - a^{m-k-1})a^{k+1}/(1-a)$, $\sum_{j=k+1}^{m-1}(j-\mu)\pi(j) = c(a)(a^k((k+1)(1-a) - a) - a^{m-1}(m(1-a) - a))/(1-a)^2$. It follows that

$$a_k = a^{2k}b(a,m,k)$$

with $b(a)$ bounded uniformly in m, k, $(1-a)^2$. It thus follow that b_k^2 in (21) equals $a^{2k}c(m,k)$ with $c(m,k)$ uniformly bounded. The eigenvalues β_k of (19) become

$$\beta_k = 1 - \frac{k}{m} - \frac{1}{m}\Big(\frac{1 - a^{m-h}}{1 - a}\Big).$$

Plugging into (21) yields the following proposition.

Proposition 4.5 *Fix a with $0 < a < 1$, for independent proposal Metropolis sampling with uniform proposals on $\{0,1,\ldots,m-1\}$ for $\pi(i) = a^i c(a)$ the metropolis algorithm has $Var_\pi(\hat{\mu}_M(f)) = (m/N)A(a)$ with A continuous and bounded uniformly in m and N.*

Importance sampling. From (22),

$$Var_\sigma(\hat{\mu}_I) = \frac{1}{N}\sum_{i=0}^{m-1}\Big(\frac{\pi(i)}{\sigma(i)}(i - \mu)\Big)^2 \sigma(i) = \frac{m}{N}\sum_{i=0}^{m-1}\Big(\pi(i)(i - \mu)\Big)^2.$$

This expression is of the form

$$\frac{m}{N}B(a)$$

with $B(a)$ continuous in a and bounded uniformly in m for fixed $a \in (0,1)$. It follows that importance sampling and the Metropolis algorithm are roughly comparable for this example.

14

Remark 6 The example above helps us to calibrate two obvious bounds. From (6) for independent Metropolis sampling applied to f of mean zero

$$\sigma_\infty^2(\hat{\mu}_M) \leq 2\frac{\|f\|_{2,\pi}^2}{1-\beta_1} \leq 2\frac{\pi(0)}{\sigma(0)}\|f\|_{2,\pi}^2,$$

$$Var_\sigma(\hat{\mu}_I) = \frac{1}{N}\sum_i \left(\frac{\pi(i)}{\sigma(i)}f(i)\right)^2 \sigma(i) = \frac{1}{N}\sum_i \frac{\pi(i)^2}{\sigma(i)}f(i)^2 \leq \frac{1}{N}\frac{\pi(0)}{\sigma(0)}\|f\|_{2,\pi}^2.$$

5   Non–self intersecting paths.

Our interest in this area started with Donald Knuth's [8] study of non–self intersecting paths in a grid. Knuth considered a 10×10 grid. He wanted to estimate the number of non–self intersecting lattice paths γ that start at $(0,0)$ and end at $(10,10)$. He used the following sequential importance sampling (SIS) estimate: build a path γ starting at $(0,0)$ sequentially, each time choosing one of the available nearest neighbors with equal probability. As the path grows, the past is recorded and only non–self intersecting choices are considered. Thus, the first step may go up or to the right with probability $1/2$. Suppose it goes up. The next step can go up or to the right. Suppose it goes to the right. The third step has three possibilities (up, right, down) chosen with probability $1/3$, and so on. If the algorithm gets stuck, it simply starts again at $(0,0)$. Let $\sigma(\gamma)$ be the probability of a successful path. Thus $\sigma(\gamma) = \frac{1}{2}\frac{1}{2}\frac{1}{3}\cdots$ in the example. This is easily computed as the path is created. Let $X_i = 0$ if the i^{th} trial fails and $X_i = 1/\sigma(\gamma)$ if the i^{th} trial produces a legal path γ. Observe that

$$E(X_i) = \sum_\gamma \frac{1}{\sigma(\gamma)}\sigma(\gamma) = \text{number of paths}.$$

Thus if $X_1, X_2, \ldots X_N$, is the result of N trials

$$\hat{\mu} = \frac{1}{N}\sum_{i=1}^N X_i$$

is an unbiased estimate of the number of paths. Knuth used this sequential importance sampling algorithm to give the following estimates:

- number of paths $= (1.6 \pm 0.3)10^{24}$

- av. path length $= 92 \pm 5$

- proportion of paths through $(5,5) = 81 \pm 10$ percent.

In a personal communication he noted that usually $1/\sigma(\gamma)$ was between 10^{11} and 10^{17}. A few of his sample values were much larger, accounting for the 10^{24}. It is hard to bound or asses the variance of the sequential importance sampling estimator.

There is something to understand: D.Bressanini has found the exact answer to Knuth problem in sequence A00764 in the on line version of Sloane's handbook of integer sequences. This contains further references [18]. The number of non–self intersecting paths in a 10×10 grid equals

$$1568758030464750013214106 = 1.5687 \times 10^{24}.$$

This is in good agreement with Knuth $(1.6 \pm 0.3)10^{24}$.

5.1   Monotone paths.

As a contribution to understanding Knuth's use of sequential importance sampling we consider an easier problem where all calculations can be carried out. Let $\mathcal{X}$ be the set of all monotone paths from $(0,0)$ to (n,n) in the usual lattice. Here, paths are only allowed to go up or to the right. Thus, if $n = 2$ there are 6 paths.
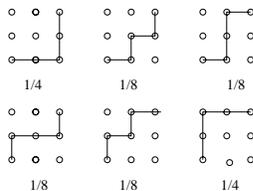

Figure 1:

In general, $|\mathcal{X}| = \binom{2n}{n}$. Shown underneath the example is $\sigma(\gamma)$ for the sequential importance sampling applied to this setting: choose one of the two available next steps with probability $1/2$ until the walk hits the top or right side of the $n \times n$ "box" when the remainder of the walk is forced. If $T(\gamma)$ is the first time the walk hits the top or right side, $\sigma(\gamma) = 2^{-T(\gamma)}$. Both the uniform distribution $\pi(\gamma) = 1/\binom{2n}{n}$ and the distribution $\sigma(\gamma) = 2^{-T(\gamma)}$ have the property that, conditional on $T(\gamma) = j$, they are uniformly distributed. Thus things are determined by the behavior of the distribution of $T(\gamma)$. The following Proposition determines this for π and σ.

Proposition 5.1 *For the monotone paths on an $n \times n$ grid:*

16

(a) *Under the uniform distribution π*

$$\pi\{T(\gamma) = j\} = 2\frac{\binom{j-1}{n-1}}{\binom{2n}{n}} \qquad n \le j \le 2n - 1.$$

(b) *For n large and any fixed k*

$$\pi\{T(\gamma) = 2n - 1 - k\} \to \frac{1}{2^{k+1}} \qquad 0 \le k < +\infty.$$

(c) *Under the importance sampling distribution σ*

$$\sigma\{T(\gamma) = j\} = 2^{1-j}\binom{j-1}{n-1} \qquad n \le j \le 2n - 1.$$

(d) *For n large and any fixed positive x*

$$\sigma\{\frac{2n - 1 - T(\gamma)}{\sqrt{n}} \le x\} \to \frac{1}{\pi}\int_0^x e^{-y^2/4}dy.$$

Proof. Paths with $T(\gamma) = j$ may be coded as sequences of zeros and ones with n zeros and n ones having all zeros (or all ones) before j. This is twice the number with all ones before j. To count these, put zero at j and the remaining $n - 1$ zeros in the remaining $j - 1$ places. This proves (a),(c). Part (b) is simple. For (d) we prove a local limit theorem. In (c), take $j = 2n - 1 - a$. Then,

$$\binom{j-1}{n-1} = \frac{(n-1)(n-2)\ldots(n-1-a+1)}{(2n-2)(2n-3)\ldots(2n-a+1)}\binom{2n-2}{n-1}.$$

Using Stirling's formula $\binom{2n-2}{n-1}2^{2n-2} \sim 1/\sqrt{\pi n}$. What is left is 2^{a-1} times

$$\frac{(n-1)(n-2)\ldots(n-1-a+1)}{(2n-2)(2n-3)\ldots(2n-a+1)} = \frac{(1 - \frac{1}{n-1})\ldots(1 - \frac{a-1}{n-1})}{(1 - \frac{1}{2(n-1)})\ldots(1 - \frac{a-1}{2(n-1)})}2^a.$$

Write the product in the numerator as

$$\prod_{i=1}^{a-1}(1 - \frac{i}{n-1}) = \exp\{\sum_{i=1}^{a-1}\log(1 - \frac{i}{n-1})\} \sim e^{-\frac{1}{n}\binom{a}{2}}$$

where the asymptotics are valid for $a \ll n^{2/3}$. This may be justified and refined as in [3] Chapter 5. For the denominator,

$$\prod_{i=1}^{a-1}(1 - \frac{i}{2(n-1)}) = \exp\{\sum_{i=1}^{a-1}\log(1 - \frac{i}{2(n-1)})\} \sim e^{-\frac{1}{2n}\binom{a}{2}}.$$

Putting things together, for $j = 2n - 1 - a$,

$$\sigma\{T(\gamma) = j\} \sim \frac{1}{\sqrt{\pi n}}e^{-\frac{1}{4n}a^2}.$$

For $a = x\sqrt{n}$, this last is $\frac{1}{\sqrt{\pi n}}e^{-x^2/4}$. Further details are omitted. $\diamond$

17

Remark 7 Thus, under the uniform distribution π, $T(\gamma)$ is close to its maximum $2n - 1$. Under the importance sampling distribution σ, $T(\gamma)$ is usually $\sqrt{n}$ away from $2n - 1$. We see below how our two algorithms deal with this. First, we treat the analog for estimating the size of the state space $|\mathcal{X}|$.

Let $\mu = |\mathcal{X}| = \binom{2n}{n}$. Generate paths γ_i, $1 \le i \le N$, independently from $\sigma(x)$ and set $\hat{\mu}_{SIS} = \frac{1}{N} \sum_{i=1}^{N} 1/\sigma(\gamma)$. As above, $E_\sigma(\hat{\mu}_{SIS}) = \mu$. The variance of the estimator $\hat{\mu}_{SIS}$ is given next.

Proposition 5.2 *For μ the number of monotone paths in an $n \times n$ grid*

$$Var_\sigma(\hat{\mu}_{SIS}) = \frac{1}{N} \frac{16^n}{4\sqrt{n}} (1 + O(\frac{1}{n})).$$

Proof. $Var_\sigma(\hat{\mu}_{SIS}) = \frac{1}{N} \{E_\sigma(1/\sigma(\gamma)^2 - \mu^2\}$. We have

$$E_\sigma(1/\sigma(\gamma)^2) = \sum_\gamma 1/\sigma(\gamma) = \sum_{j=n}^{2n-1} 2^{j+1} \binom{j-1}{n-1}.$$

This sum is dominated by its largest term. Thus

$$E(1/\sigma(\gamma)^2) = 2^{2n} \binom{2n-2}{n-1} (1 + O(\frac{1}{n})).$$

On the other hand, $\mu^2 = \binom{2n}{n} \sim 16^n/(\pi n)$ is of lower order. $\diamond$

Remark 8 While the variance is exponentially large in n, the *relative* variance is

$$Var_\sigma(\frac{\hat{\mu}_{SIS}}{\mu}) = \frac{1}{N} \sqrt{\pi n} (1 + O(\frac{1}{n})).$$

Thus a relatively small sample size N suffices to get a useful relative error.

We next compare importance sampling and the Metropolis algorithm for estimating the mean of some simple functions of monotone paths. In our examples, the Metropolis algorithm dominates.

Importance sampling. Again $\mathcal{X}$ is the set of monotone paths from $(0,0)$ to (n,n). Let $f : \mathcal{X} \to \mathbb{R}$ be any function. Let $\gamma_1, \ldots, \gamma_N$ be chosen independently using the sequential importance distribution $\sigma(\gamma) = 2^{-T(\gamma)}$. Then,

$$\hat{\mu}_{SIS}(f) := \frac{1}{N} \sum_{i=1}^{N} \frac{2^{T(\gamma_i)}}{\binom{2n}{n}} f(\gamma_i)$$

18

is an unbiased estimator of $\mu = \sum_\gamma \frac{1}{\binom{2n}{n}} f(\gamma)$. We have

$$Var_\sigma(\hat{\mu}_{SIS}(f)) = \frac{1}{N}\Big\{ E\Big(\frac{2^{T(\gamma)}}{\binom{2n}{n}} f(\gamma)\Big)^2 - \mu^2 \Big\}.$$

Using Proposition 5.2 above we may calculate this for simple functions f.

Example 4 Let $f(\gamma) = T(\gamma) =$ *the first hitting time of a uniformly chosen random path γ to the top or right side of an $n \times n$ grid. Then*

$$\mu = E_\pi(T) = (2 - \frac{2}{n+1})n$$

and

$$Var_\sigma(\hat{\mu}_{SIS}(T)) \sim \frac{\sqrt{\pi}n^{5/2}}{N}.$$

Indeed under the uniform distribution,

$$\mu = \sum_{j=n}^{2n-1} 2j \binom{j-1}{n-1}/\binom{2n}{n} = 2n \sum_{j=n}^{2n-1} \binom{j}{n}/\binom{2n}{n}$$

$$= (2 - \frac{2}{n+1})n.$$

For the variance,

$$E\Big(\frac{2^{T(\gamma)}}{\binom{2n}{n}} T(\gamma)\Big)^2 = \frac{1}{\binom{2n}{n}^2} \sum_\gamma 2^{2T(\gamma)} T^2(\gamma) 2^{-T(\gamma)} = \frac{2}{\binom{2n}{n}^2} \sum_{j=n}^{2n-1} 2^j j^2 \binom{j-1}{n-1}.$$

As before, the sum is dominated by its largest term. This is

$$2^{2n}(2n-1)^2 \binom{2n-2}{n-1}/\binom{2n}{n}^2 \sim \sqrt{\pi}n^{5/2}. \quad \Diamond$$

Remark 9 Here, the expected squared relative error is

$$E(\frac{\hat{\mu}}{\mu} - 1)^2 \sim \frac{\sqrt{\pi n}}{4N}.$$

Thus again, a sample size N of order only larger than $\sqrt{n}$ suffices to give acceptable relative error.

Metropolis sampling. Using independence Metropolis sampling as in Section 4 with $\pi(\gamma) = 1/\binom{2n}{n}$ and proposal distribution $\sigma(\gamma) = 2^{-T(\gamma)}$, we may use Liu's result. To proceed, we must order the state space of paths with decreasing importance weights $\pi(\gamma)/\sigma(\gamma)$ and thus by largest values of $T(\gamma)$. Using the binary coding of paths introduced in the proof

of Proposition 5.2 we may order via $\gamma \leq \gamma'$ if $T(\gamma) \geq T(\gamma')$. If $T(\gamma) = T(\gamma')$, use the lexiographical order. We break paths into groups by $T(\gamma)$:

      group one with $T(\gamma) = 2n - 1$ of size $A_1 := 2\binom{2n-2}{n-1}$

      group two with $T(\gamma) = 2n - 2$ of size $A_2 := 2\binom{2n-3}{n-1}$

      ...

      group i with $T(\gamma) = 2n - i$ of size $A_i := 2\binom{2n-1-i}{n-1}$

      group n with $T(\gamma) = n$ of size $A_n := 2$

Proposition 5.3 *The independence proposal Markov chain on monotone paths with proposal distribution $\sigma(\gamma) = 2^{-T(\gamma)}$ and stationary distribution $\pi(\gamma) = 1/\binom{2n}{n}$ has n distinct eigenvalues on each of the n groups above with multiplicity the size of the i^{th}-group. If $s(i) = 2^{-(2n-1-i)}$, the eigenvalues are*

$\beta_1 = 1 - s(0)\binom{2n}{n}, \qquad multiplicity\ A_1 - 1$

$\beta_2 = 1 - s(1)\binom{2n}{n} + A_0(s(1) - s(0)), \qquad multiplicity\ A_2$

$\ldots$

$\beta_i = 1 - s(i-1)\binom{2n}{n} + A_0(s(i-1) - s(0)) + \cdots + A_{i-2}(s(i-1) - s(i-2)), \qquad multiplicity\ A_i$

$\ldots$

Proof. This follows from equation (19) by elementary manipulations. $\diamond$

Remark 10 (a) Using Stirling's formula, for fixed j

$$\beta_j = 1 - \frac{2^j}{\sqrt{\pi n}} + O_j(\frac{1}{n})$$

From Lemma 4.2, β_1 is the second largest eigen–value.

(b) The eigenvectors of this chain are simple to write down from (20). We do not do this explicitly here. Using the eigenvalues, eigenvectors and convergence results from [16], we have proved that the relaxation time of the Metropolis algorithm is of order $n^{3/2}$ and there is a sharp cut off.

To conclude this section we bound the asymptotic variance of the Metropolis estimator of the function $T(\gamma)$ and show that it improves on the importance sampling estimator of Example 4. Recall from (6) $\sigma_\infty^2(\hat{\mu}) = \lim_N NVar(\hat{\mu})$.

20

Example 5 Let $f(\gamma) = T(\gamma) =$ *the first hitting time of a uniformly chosen random path to the top or right side of an $n \times n$ grid*. Then,

$$\sigma_\infty^2(\hat{\mu}_{Met}(T)) \leq \{8\sqrt{\pi}n^{3/2} + O(n).\}$$

To prove the last claim, use (6) to write

$$\sigma_\infty^2(\hat{\mu}_{Met}) \leq \frac{2}{1-\beta_1}Var_\pi(T).$$

We have shown above that $1/(1-\beta_0) \leq \sqrt{\pi n} + o(1)$. Now,

$$Var_\pi(T) = E_\pi(T^2) - (2 - \frac{2}{n+1})^2 n^2,$$

$$E_\pi(T^2) = \frac{2n}{\binom{2n}{n}}\sum_{i=0}^{n-1}(n+i)\binom{n+1}{n}.$$

By elementary calculations

$$4n^2(1 - \frac{1}{n})\Big(\frac{n + \frac{1}{2}}{n+1}\Big) \leq E_\pi(T^2) \leq \frac{4n^3}{n+1}.$$

Hence, $Var_\pi(T) = 4n + O(1)$. Combing bounds gives the result.

Remark 11 Roughly $Var_\pi(\hat{\mu}_{SIS}) \asymp \frac{n^{5/2}}{N}$ while $Var_\pi(\hat{\mu}_{MET}) \asymp \frac{n^{3/2}}{N}$.

References

[1] P. Diaconis and L. Saloff-Coste. Comparison techniques for random walk on finite groups. *Ann. Probab.*, 21(4):2131–2156, 1993.

[2] P. Diaconis and L. Saloff-Coste. What do we know about the Metropolis algorithm? *J. Comput. System Sci.*, 57(1):20–36, 1998. 27th Annual ACM Symposium on the Theory of Computing (STOC'95) (Las Vegas, NV).

[3] W. Feller. *An introduction to probability theory and its applications. Vol. I.* John Wiley and Sons, Inc., New York, 1968. 3nd ed.

[4] J. M. Hammersley and D. C. Handscomb. *Monte Carlo methods.* Methuen & Co. Ltd., London, 1965.

[5] W.K. Hastings. Monte carlo sampling methods using markov chains and their application. *Biometrika*, 57:97–109, 1970.

[6] H. Kahn. Modification of the Monte Carlo method. In *Proceedings, Seminar on Scientific Computation, November, 1949*, pages 20–27. International Business Machines Corp., New York, N. Y., 1950.

[7] H. Kahn. Use of different Monte Carlo sampling techniques. In *Symposium on Monte Carlo methods, University of Florida, 1954*, pages 146–190. John Wiley and Sons, Inc., New York, 1956.

[8] D. E. Knuth. Mathematics and computer science: coping with finiteness. *Science*, 194(4271):1235–1242, 1976.

[9] J. Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, 6:113–119, 1996.

[10] J. S. Liu. *Monte Carlo strategies in scientific computing.* Springer Series in Statistics. Springer-Verlag, New York, 2001.

[11] N. Madras and M. Piccioni. Importance sampling for families of distributions. *Ann. Appl. Probab.*, 9(4):1202–1225, 1999.

[12] N. Metrpolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state caluclations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.

[13] P. H. Peskun. Optimum Monte-Carlo sampling using Markov chains. *Biometrika*, 60:607–612, 1973.

[14] P. H. Peskun. Guidelines for choosing the transition matrix in Monte Carlo methods using Markov chains. *J. Comput. Phys.*, 40(2):327–344, 1981.

[15] C. P. Robert and G. Casella. *Monte Carlo statistical methods.* Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2004.

[16] L. Saloff-Coste. Lectures on finite Markov chains. In *Lectures on probability theory and statistics (Saint-Flour, 1996)*, volume 1665 of *Lecture Notes in Math.*, pages 301–413. Springer, Berlin, 1997.

[17] D. Siegmund. Importance sampling in the Monte Carlo study of sequential tests. *Ann. Statist.*, 4(4):673–684, 1976.

[18] N. J. A. Sloane. The on-line encyclopedia of integer sequences. Published electronically at http://www.research.att.com/∼njas/sequences.

[19] A. Sokal. Monte Carlo methods in statistical mechanics: foundations and new algorithms. In *Functional integration (Cargèse, 1996)*, volume 361 of *NATO Adv. Sci. Inst. Ser. B Phys.*, pages 131–192. Plenum, New York, 1997.