Approximating zero-variance importance sampling in a reliability setting

Pierre L'Ecuyer[1] and Bruno Tuffin[2]

1. DIRO, Université de Montréal, Canada
2. IRISA-INRIA, Rennes, France

Reuven's Birthday Conference, Sandbjerg, Denmark, July 2008

Discrete-time Markov chain (DTMC) model

Consider a simulation model represented as a DTMC $\{Y_j, j \geq 0\}$ with (large) state space $\mathcal{Y}$, and a set of absorbing states $\Delta \subset \mathcal{Y}$ so that the simulation stops when the chain hits Δ.

Transition kernel: $P(B \mid y) = \mathbb{P}[Y_j \in B \mid Y_{j-1} = y)$.

Stopping time: $\tau = \inf\{j : Y_j \in \Delta\}$.

Cost $c(y, y')$ for each transition $y \to y'$.

Total cost: $X = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j)$.

Discrete-time Markov chain (DTMC) model

Consider a simulation model represented as a DTMC $\{Y_j, j \geq 0\}$ with (large) state space $\mathcal{Y}$, and a set of absorbing states $\Delta \subset \mathcal{Y}$ so that the simulation stops when the chain hits Δ.

Transition kernel: $P(B \mid y) = \mathbb{P}[Y_j \in B \mid Y_{j-1} = y)$.

Stopping time: $\tau = \inf\{j : Y_j \in \Delta\}$.

Cost $c(y, y')$ for each transition $y \to y'$.

Total cost: $X = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j)$.

Expected cost-to-go from state y: $\mu(y) = \mathbb{E}[X \mid Y_0 = y]$.

We assume that $\mathbb{E}[\tau \mid Y_0 = y] < \infty$ and $\mu(y) < \infty$ for all $y \in \mathcal{Y}$.

We want to estimate $\mu(y_0)$ where y_0 is the initial state.

Importance sampling

We consider changing P to another transition kernel Q.

The estimator X is replaced by the IS estimator

$$X_{\mathrm{is}} = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j) \prod_{i=1}^{j} L(Y_{i-1}, Y_i),$$

where $L(Y_{i-1}, Y_i) = (dP/dQ)(Y_i \mid Y_{i-1})$.

Importance sampling

We consider changing P to another transition kernel Q.

The estimator X is replaced by the IS estimator

$$X_{\text{is}} = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j) \prod_{i=1}^{j} L(Y_{i-1}, Y_i),$$

where $L(Y_{i-1}, Y_i) = (dP/dQ)(Y_i \mid Y_{i-1})$.

Theorem. If we choose Q so that

$$dQ(y_1 \mid y) = \begin{cases} dP(y_1 \mid y) \dfrac{c(y, y_1) + \mu(y_1)}{\mu(y)} & \text{if } \mu(y) > 0, \\ dP(y_1 \mid y) & \text{if } \mu(y) = 0 \end{cases}$$

(this density integrates to 1), then X_{is} has zero variance.
Proof: By induction on j.

Simple special case: finite state space $\mathcal{Y}$

The DTMC has transition probabilities
$p(y_1 \mid y) = \mathbb{P}[Y_1 = y_1 \mid Y_0 = y]$, which are replaced by
$q(y_1 \mid y) = \mathbb{Q}[Y_1 = y_1 \mid Y_0 = y]$.
We have $L(y, y_1) = p(y_1 \mid y)/q(y_1 \mid y)$. For the zero variance:

$$q(y_1 \mid y) = \begin{cases} p(y_1 \mid y)\dfrac{c(y, y_1) + \mu(y_1)}{\mu(y)} & \text{if } \mu(y) > 0, \\ p(y_1 \mid y) & \text{if } \mu(y) = 0. \end{cases}$$

Simple special case: finite state space $\mathcal{Y}$

The DTMC has transition probabilities $p(y_1 \mid y) = \mathbb{P}[Y_1 = y_1 \mid Y_0 = y]$, which are replaced by $q(y_1 \mid y) = \mathbb{Q}[Y_1 = y_1 \mid Y_0 = y]$.
We have $L(y, y_1) = p(y_1 \mid y)/q(y_1 \mid y)$. For the zero variance:

$$q(y_1 \mid y) = \begin{cases} p(y_1 \mid y)\dfrac{c(y, y_1) + \mu(y_1)}{\mu(y)} & \text{if } \mu(y) > 0, \\ p(y_1 \mid y) & \text{if } \mu(y) = 0. \end{cases}$$

We approximate the unknown function μ by some v, either fixed or learned along the way. This gives:

$$q(y_1 \mid y) = \begin{cases} p(y_1 \mid y)\dfrac{c(y, y_1) + v(y_1)}{v(y)} & \text{if } v(y) > 0, \\ p(y_1 \mid y) & \text{if } v(y) = 0. \end{cases}$$

Model of Highly Reliable Markovian System (HRMS)

c component types, n_i components of type i.
Markov chain step: failure of repair of one component.
$Y_j = (Y_j^{(1)}, \ldots, Y_j^{(c)})$ = num. failed compon. of each type at step j.

$\{Y_j, j \geq 0\}$ is a DTMC with trans. probabilities
$p(y, y') = \mathbb{P}[Y_j = y' \mid Y_{j-1} = y]$.

Model of Highly Reliable Markovian System (HRMS)

c component types, n_i components of type i.
Markov chain step: failure of repair of one component.
$Y_j = (Y_j^{(1)}, \ldots, Y_j^{(c)})$ = num. failed compon. of each type at step j.

$\{Y_j, j \geq 0\}$ is a DTMC with trans. probabilities
$p(y, y') = \mathbb{P}[Y_j = y' \mid Y_{j-1} = y]$.
Suppose that failure probabilities are much smaller than repair probabilities. This is typical of highly reliable systems.

The state space $\mathcal{Y}$ is partitioned in: (1) a (decreasing) set of up states $\mathcal{U}$ and (2) the set of failure states $\mathcal{F}$.

For any set A, let τ_A = first hitting time of A.
For any state y, let

$$\mu(y) = \mathbb{P}[\tau_{\mathcal{F}} < \tau_{\mathbf{0}} \mid Y_0 = y],$$

the prob. of visiting $\mathcal{F}$ before returning to $\mathbf{0}$.

Goal: estimate $\mu(\mathbf{0})$. This can be difficult when $\mu(\mathbf{0})$ is very small.

Some proposed IS heuristics:

Balanced failure biasing (BFB) (Shahabuddin 1994) changes p to q as follows, for $x \notin B$:

$$q(x,y) = \begin{cases} \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) = 0; \\ \rho \, \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) > 0; \\ (1-\rho)\, \frac{p(x,y)}{p_{\mathrm{R}}(x)} & \text{if } y \in R(x); \\ 0 & \text{otherwise.} \end{cases}$$

Some proposed IS heuristics:

Balanced failure biasing (BFB) (Shahabuddin 1994) changes p to q as follows, for $x \notin B$:

$$q(x,y) = \begin{cases} \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) = 0; \\ \rho \, \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) > 0; \\ (1-\rho) \, \frac{p(x,y)}{p_{\mathrm{R}}(x)} & \text{if } y \in R(x); \\ 0 & \text{otherwise.} \end{cases}$$

Simple failure biasing (SFB) (Shahabuddin 1988): Replace $1/|F(x)|$ above by $p(x,y)/\sum_{y \in F(x)} p(x,y)$.

Some proposed IS heuristics:

Balanced failure biasing (BFB) (Shahabuddin 1994) changes p to q as follows, for $x \notin B$:

$$q(x,y) = \begin{cases} \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) = 0; \\ \rho \, \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) > 0; \\ (1-\rho) \, \frac{p(x,y)}{p_{\mathrm{R}}(x)} & \text{if } y \in R(x); \\ 0 & \text{otherwise.} \end{cases}$$

Simple failure biasing (SFB) (Shahabuddin 1988): Replace $1/|F(x)|$ above by $p(x,y)/\sum_{y \in F(x)} p(x,y)$.

SBLR (Alexopoulos and Shultes 2001) changes the probabilities in a way that over any cycle in the visited states during the simulation, the cumulated likelihood ratio remains bounded

Some proposed IS heuristics:

Balanced failure biasing (BFB) (Shahabuddin 1994) changes p to q as follows, for $x \notin B$:

$$q(x,y) = \begin{cases} \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) = 0; \\ \rho \, \frac{1}{|F(x)|} & \text{if } y \in F(x) \text{ and } p_{\mathrm{R}}(x) > 0; \\ (1-\rho) \, \frac{p(x,y)}{p_{\mathrm{R}}(x)} & \text{if } y \in R(x); \\ 0 & \text{otherwise.} \end{cases}$$

Simple failure biasing (SFB) (Shahabuddin 1988): Replace $1/|F(x)|$ above by $p(x,y)/\sum_{y \in F(x)} p(x,y)$.

SBLR (Alexopoulos and Shultes 2001) changes the probabilities in a way that over any cycle in the visited states during the simulation, the cumulated likelihood ratio remains bounded

These methods do not attempt to mimic zero-variance sampling.

Proposed approximation (ZVA)

Approximate μ by some easily computable function v, and plug into zero-variance formula.

For any state $y \in \mathcal{U}$, let $\Gamma(y)$ be the set of all paths $\pi = (y = y_0 \to y_1 \to \cdots \to y_k)$ where $y_1, \ldots, y_{k-1} \notin \mathcal{F} \cup \{\mathbf{0}\}$, $y_k \in \mathcal{F}$, and having positive probability

$$p(\pi) = \prod_{j=1}^{k} p(y_{j-1}, y_j) > 0.$$

Proposed approximation (ZVA)

Approximate μ by some easily computable function v, and plug into zero-variance formula.

For any state $y \in \mathcal{U}$, let $\Gamma(y)$ be the set of all paths $\pi = (y = y_0 \to y_1 \to \cdots \to y_k)$ where $y_1, \ldots, y_{k-1} \notin \mathcal{F} \cup \{\mathbf{0}\}$, $y_k \in \mathcal{F}$, and having positive probability

$$p(\pi) = \prod_{j=1}^{k} p(y_{j-1}, y_j) > 0.$$

Because these paths represent disjoint events, we have

$$\mu(y) = \sum_{\pi \in \Gamma(y)} p(\pi).$$

Proposed approximation (ZVA)

Approximate μ by some easily computable function v, and plug into zero-variance formula.

For any state $y \in \mathcal{U}$, let $\Gamma(y)$ be the set of all paths $\pi = (y = y_0 \to y_1 \to \cdots \to y_k)$ where $y_1, \ldots, y_{k-1} \notin \mathcal{F} \cup \{\mathbf{0}\}$, $y_k \in \mathcal{F}$, and having positive probability

$$p(\pi) = \prod_{j=1}^{k} p(y_{j-1}, y_j) > 0.$$

Because these paths represent disjoint events, we have

$$\mu(y) = \sum_{\pi \in \Gamma(y)} p(\pi).$$

This last sum may contain a huge (perhaps ∞) number of terms.

A very crude approximation is to just take the path with largest probability, i.e., approximate

$$\mu(y) = \sum_{\pi \in \Gamma(y)} p(\pi)$$

by its lower bound

$$v_0(y) = \max_{\pi \in \Gamma(y)} p(\pi).$$

Computing $v_0(y)$ amounts to computing a shortest path from y to $\mathcal{F}$, where the length of a link $y' \to y''$ is $-\log p(y', y'')$. Easy.

A very crude approximation is to just take the path with largest probability, i.e., approximate

$$\mu(y) = \sum_{\pi \in \Gamma(y)} p(\pi)$$

by its lower bound

$$v_0(y) = \max_{\pi \in \Gamma(y)} p(\pi).$$

Computing $v_0(y)$ amounts to computing a shortest path from y to $\mathcal{F}$, where the length of a link $y' \to y''$ is $-\log p(y', y'')$. Easy.

This would work fine if a single path dominates the sum (this may happen when failure transitions have very small probabilities), but this v_0 will often underestimate the bound significantly.

Refinements

Typically, the farther we are from $\mathcal{F}$, the more v_0 underestimates μ. Close to $\mathcal{F}$, things are fine, but not close to $\mathbf{0}$.

First simple correction:

1. Estimate $\mu(\mathbf{0})$ in preliminary runs with crude IS strategy;
2. Find constant $\alpha \le 1$ such that $(v_0(\mathbf{0}))^\alpha$ equals this estimate;
3. Use $v_1(y) = (v_0(y))^\alpha$ for all $y \in \mathcal{U}$ as approx. of $\mu(y)$.

This v_1 matches μ for $y \in \mathcal{F}$ and matches its estimate at $y = \mathbf{0}$.

Refinements

Typically, the farther we are from $\mathcal{F}$, the more v_0 underestimates μ. Close to $\mathcal{F}$, things are fine, but not close to $\mathbf{0}$.

First simple correction:

1. Estimate $\mu(\mathbf{0})$ in preliminary runs with crude IS strategy;
2. Find constant $\alpha \le 1$ such that $(v_0(\mathbf{0}))^{\alpha}$ equals this estimate;
3. Use $v_1(y) = (v_0(y))^{\alpha}$ for all $y \in \mathcal{U}$ as approx. of $\mu(y)$.

This v_1 matches μ for $y \in \mathcal{F}$ and matches its estimate at $y = \mathbf{0}$.

Second refinement: Replace α by a state-dependent exponent

$$\alpha(y) = 1 + [\alpha(\mathbf{0}) - 1]\frac{\log v_0(y)}{\log v_0(\mathbf{0})},$$

where $\alpha(\mathbf{0}) = \alpha$ as above. This $\alpha(y)$ changes progressively from 1 near $\mathcal{F}$ to $\alpha(\mathbf{0}) < 1$ in state $\mathbf{0}$. The correction here is milder than in the previous case when we are close to $\mathcal{F}$.
Let $v_2(y) = (v_0(y))^{\alpha(y)}$ be the resulting approximation.

Example: Three types of components

$c = 3$ and $n_1 = n_2 = n_3$.
Expon. repair times with mean 1.
Failure rate λ_i for component type i,
with $\lambda_1 = \varepsilon$, $\lambda_2 = 1.5\varepsilon$, and $\lambda_3 = 2\varepsilon^2$, for some small real number ε.

We will try different values of (n_i, ε).

$\mathcal{F}$ = states where at least one component type has fewer than 2 operational units.

Example: Three types of components

$c = 3$ and $n_1 = n_2 = n_3$.
Expon. repair times with mean 1.
Failure rate λ_i for component type i,
with $\lambda_1 = \varepsilon$, $\lambda_2 = 1.5\varepsilon$, and $\lambda_3 = 2\varepsilon^2$, for some small real number ε.

We will try different values of (n_i, ε).

$\mathcal{F}$ = states where at least one component type has fewer than 2 operational units.

To define $v_0(y)$, we consider all three paths to $\mathcal{F}$ that result from failures of a single component type, and sum their probabilities.

The table contains results with $n = 2^{20}$ runs.

Best estimate of $\mu(\mathbf{0})$: obtained from a large number of runs with our best IS strategies.

Mean

n_i	ε	$\mu(\mathbf{0})$	$v_0(\mathbf{0})$	BFB	SBLR
3	0.001	2.6×10^{-3}	1.3×10^{-3}	2.7×10^{-3}	2.6×10^{-3}
6	0.01	1.8×10^{-7}	3.4×10^{-8}	1.9×10^{-7}	$[9.9 \times 10^{-7}]$
6	0.001	1.7×10^{-11}	3.4×10^{-12}	1.8×10^{-11}	(1.8×10^{-16})
12	0.1	6.0×10^{-8}	3.2×10^{-9}	4.8×10^{-8}	1.3×10^{-8}
12	0.001	3.9×10^{-28}	3.5×10^{-29}	(1.8×10^{-40})	(2.9×10^{-45})

Variance

n_i	ε	BFB	SBLR
3	0.001	1.8×10^{-2}	8.0×10^{-3}
6	0.01	6.3×10^{-11}	(4.5×10^{-16})
6	0.001	8.8×10^{-19}	(2.0×10^{-26})
12	0.1	8.1×10^{-10}	1.7×10^{-10}
12	0.001	(3.2×10^{-74})	(3.5×10^{-84})

Mean

n_i	ε	$\mu(\mathbf{0})$	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)
3	0.001	2.6×10^{-3}	2.6×10^{-3}	2.6×10^{-3}	2.6×10^{-3}
6	0.01	1.8×10^{-7}	1.8×10^{-7}	1.8×10^{-7}	1.8×10^{-7}
6	0.001	1.7×10^{-11}	1.7×10^{-11}	1.7×10^{-11}	1.7×10^{-11}
12	0.1	6.0×10^{-8}	6.0×10^{-8}	6.2×10^{-8}	6.7×10^{-8}
12	0.001	3.9×10^{-28}	3.9×10^{-28}	3.9×10^{-28}	3.9×10^{-28}

Variance

n_i	ε	α	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)	RE(v_2)
3	0.001	0.906	6.5×10^{-4}	2.7×10^{-3}	9.3×10^{-9}	0.04
6	0.01	0.903	2.0×10^{-14}	1.2×10^{-14}	7.7×10^{-15}	0.48
6	0.001	0.939	1.2×10^{-23}	1.1×10^{-23}	7.6×10^{-24}	0.16
12	0.1	0.851	1.6×10^{-10}	2.9×10^{-10}	1.5×10^{-11}	64.50
12	0.001	0.963	1.4×10^{-55}	9.3×10^{-56}	9.4×10^{-56}	0.78

We have $\alpha \to 1$ when $\varepsilon \to 0$ or when $n_i \nearrow$

Example 2

Taken from Shahabuddin (1994).
Was chosen to illustrate the performance of BFB.

System:
Two sets of processors, with two units per set, $\lambda_i = 5 \times 10^{-5}$;
two sets of disk controllers, two units per set; $\lambda_i = 2 \times 10^{-5}$
six clusters of disks, four units per cluster, $\lambda_i = 2 \times 10^{-5}$.

Thus, $c = 10$ and each $n_i = 2$ or 4.
All repair rates are 1.

Example 2

Taken from Shahabuddin (1994).
Was chosen to illustrate the performance of BFB.

System:
Two sets of processors, with two units per set, $\lambda_i = 5 \times 10^{-5}$;
two sets of disk controllers, two units per set; $\lambda_i = 2 \times 10^{-5}$
six clusters of disks, four units per cluster, $\lambda_i = 2 \times 10^{-5}$.

Thus, $c = 10$ and each $n_i = 2$ or 4.
All repair rates are 1.

System is operational if we have at least of processor of each type, on controller of each set, and three disks from each cluster.

We have $\mu(\mathbf{0}) \approx 5.6 \times 10^{-5}$.

Empirical variance

α	BFB	SBLR	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)
0.949	5.8×10^{-8}	1.3×10^{-4}	2.3×10^{-12}	1.0×10^{-12}	1.2×10^{-12}

Asymptotic analysis

We want to characterize the asymptotic behavior when the failure rates converge to 0 in certain ways, while the rest remains fixed.

Following Shahabuddin, Nakayama, and collaborators, suppose

$$\lambda(y, y') = a(y, y')\varepsilon^{b(y,y')}$$

for some state-dependent parameters $a(y, y') \geq 0$ and $b(y, y') > 0$.
Repair rates are $\Theta(1)$.
We will look at what happens when $\varepsilon \to 0$.

Some asymptotic properties of estimators for $\varepsilon \to 0$

(Studied by L., Blanchet, Glynn, Tuffin (2008))

Definitions. An estimator $X(\varepsilon)$ with mean $\mu(\varepsilon)$ has bounded relative moment of order k (BRM-k) if

$$\limsup_{\varepsilon \to 0} \mathbb{E}[X^k(\varepsilon)]/\mu^k(\varepsilon) < \infty.$$

Note: BRM-2 means bounded relative error.

Some asymptotic properties of estimators for $\varepsilon \to 0$

(Studied by L., Blanchet, Glynn, Tuffin (2008))

Definitions. An estimator $X(\varepsilon)$ with mean $\mu(\varepsilon)$ has bounded relative moment of order k (BRM-k) if

$$\limsup_{\varepsilon \to 0} \mathbb{E}[X^k(\varepsilon)]/\mu^k(\varepsilon) < \infty.$$

Note: BRM-2 means bounded relative error.

It has logarithmic efficiency of order k (LE-k) if

$$\lim_{\varepsilon \to 0} \ln \mathbb{E}[X^k(\varepsilon)]/k \ln \mu(\varepsilon) = 1.$$

Note: LE-2 is often called asymptotically efficient.

Some asymptotic properties of estimators for $\varepsilon \longrightarrow 0$

(Studied by L., Blanchet, Glynn, Tuffin (2008))

Definitions. An estimator $X(\varepsilon)$ with mean $\mu(\varepsilon)$ has bounded relative moment of order k (BRM-k) if

$$\limsup_{\varepsilon \to 0} \mathbb{E}[X^k(\varepsilon)]/\mu^k(\varepsilon) < \infty.$$

Note: BRM-2 means bounded relative error.

It has logarithmic efficiency of order k (LE-k) if

$$\lim_{\varepsilon \to 0} \ln \mathbb{E}[X^k(\varepsilon)]/k \ln \mu(\varepsilon) = 1.$$

Note: LE-2 is often called asymptotically efficient.

It has vanishing relative centered moment of order k (VRCM-k) if

$$\limsup_{\varepsilon \to 0} \mathbb{E}[X^k(\varepsilon)]/\mu^k(\varepsilon) < \infty.$$

SFB and BFB

Proposition. In this HRMS framework, with SFB, BRM-k and LE-k are equivalent.
They are also equivalent for the gth empirical moment.

SFB and BFB

Proposition. In this HRMS framework, with SFB, BRM-k and LE-k are equivalent. They are also equivalent for the gth empirical moment.

Proposition. For an IS scheme such that $p(x, y, \varepsilon) = \Theta(\varepsilon^d)$ implies $q(x, y, \varepsilon) = \Theta(\varepsilon^\ell)$ for $\ell \leq d$, we have BRM-k of the g-th empirical moment if and only if for all integers m such that $r \leq m < ks_g$ and all sample paths $(x_0, \cdots, x_n)$ leading to B and having probability $\Theta(\varepsilon^m)$,

$$\mathbb{P}^*\{(X_0, \cdots, X_\tau) = (x_0, \cdots, x_n)\} = \Theta(\varepsilon^\ell)$$

for some $\ell \leq k(mg - s_g)/(kg - 1)$, where $\mathbb{E}[Y^g(\varepsilon)] = \Theta(\varepsilon^{s_g})$.

SFB and BFB

Proposition. In this HRMS framework, with SFB, BRM-k and LE-k are equivalent.
They are also equivalent for the gth empirical moment.

Proposition. For an IS scheme such that $p(x,y,\varepsilon)=\Theta(\varepsilon^d)$ implies $q(x,y,\varepsilon)=\Theta(\varepsilon^\ell)$ for $\ell \le d$, we have BRM-k of the g-th empirical moment if and only if for all integers m such that $r \le m < ks_g$ and all sample paths $(x_0,\cdots,x_n)$ leading to B and having probability $\Theta(\varepsilon^m)$,

$$\mathbb{P}^*\{(X_0,\cdots,X_\tau)=(x_0,\cdots,x_n)\}=\Theta(\varepsilon^\ell)$$

for some $\ell \le k(mg-s_g)/(kg-1)$, where $\mathbb{E}[Y^g(\varepsilon)]=\Theta(\varepsilon^{s_g})$].

Proposition. With SFB or BFB, one cannot achieve VRCM-k.

ZVA

Proposition. With our ZVA scheme, if we just take $v(y)$ as the probability of the most probable path to failure from y, then $v(y) = \Theta(\mu(y))$ for all y and we have BRM-2.

ZVA

Proposition. With our ZVA scheme, if we just take $v(y)$ as the probability of the most probable path to failure from y, then $v(y) = \Theta(\mu(y))$ for all y and we have BRM-2.

Proposition. If $v(y)/\mu(y) \to 1$ for each y when $\varepsilon \to 0$, then we have VRCM-2. This holds if $v(y)$ is the sum of probabilities of all the dominant paths from y (those with the smallest power of ε).

Conclusion

Approximating zero-variance IS via a crude approximation of the cost-to-go (or Bellman) function is viable in this setting. It turns out to beat all previously proposed IS strategy for the examples we have examined.