The Phantom SPA Method: an Inventory Example Revisited

Felisa J. Vázquez-Abad

Department of Computer Science and Operations Research
University of Montreal, Québec,
CANADA

7th INFORMS Computing Society Conference on OR and Computing Tools for the New Millennium.
Cancún, México, Enero 2000.

Outline of Presentation

1. Motivation: The Inventory Model
2. SPA Formulation
   - Mathematical Formulation
   - Difference Process
   - One-Sided Critical Rates
3. Improving Efficiency
   - Two-Sided Critical Rates
   - Difference Process via Parallel Phantoms
   - Strong Consistency of the Phantom SPA Estimator
4. Computer Simulations
5. Conclusions

The Inventory Model

Standard periodic review inventory model with backlog. The level of inventory as well as the demands are assumed *continuous* (Fu and Bashyam).

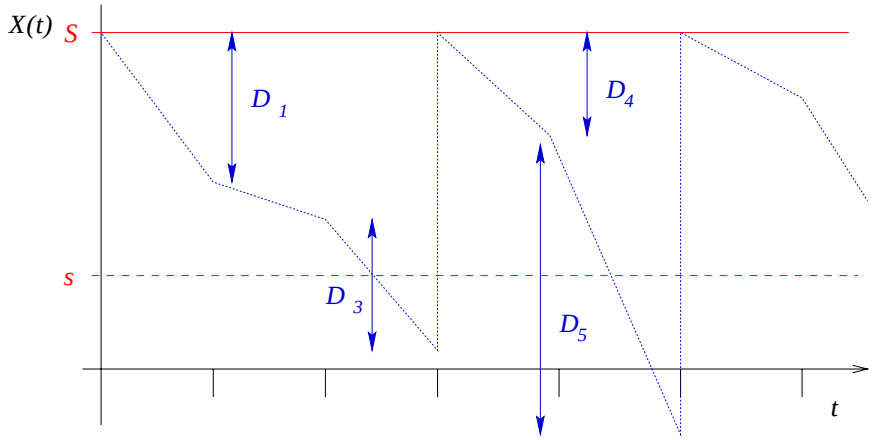

Costs at the end of period:

- Holding cost h per unit of inventory at end of period.
- Penalty cost p per unit of unsatisfied demand.
- Ordering cost K fixed if the decision is to order stock.

The inventory level $X(t)$ is a right-continous (cadlag) piecewise linear process, and the embedded Chain $\{X_n\}$ at the review epochs is Markovian:

$$X_{n+1} = \begin{cases} X_n - D_n & \text{if } X_n - D_n \geq s \\ S & \text{otherwise} \end{cases} \qquad \{D_n\} \text{ i.i.d.} \sim G$$

Control Problem: Find the optimal values of s and S to minimize the long term average cost per period $C(s, S)$. Suggested method for general distributions G and adaptive control is *gradient based stochastic approximation*.

$\Longrightarrow$ Estimation of $\dfrac{\partial}{\partial S} C(s, S)$

Mathematical Formulation: The Model

Probability space $(\Omega, \mathbb{P})$, $\{D_n\}$ i.i.d. random variables, $\mathfrak{F}_n = \sigma(D_1, \ldots, D_n)$.
$\Longrightarrow X_{n+1}$ is $\mathfrak{F}_n$-measurable.
Natural filtration of the process $\{X_n\}$ is given by $\mathbb{F} = \{\mathfrak{F}_{i-1}, i \geq 1\}$.

Assumption 1 *The demands $D_i, i = 1, 2, \ldots$ have a continuous, bounded density g, on $\mathbb{R}^+$.*

The cost associated with period n is:

$$C_n \equiv C(X_n, D_n) = \begin{cases} h(X_n - D_n) & \text{if } D_n < X_n - s \\ h(X_n - D_n) + K & \text{if } X_n - s < D_n < X_n \\ p(D_n - X_n) + K & \text{if } D_n > X_n \end{cases}$$

By construction, $X_n \in (s, S)$ is Harris recurrent and by Assumption 1 it is ergodic, therefore:

$$C(s, S) = \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} \mathrm{E}[C(X_n, D_n)].$$

Problem: How to estimate the stationary derivative

$$\frac{\partial}{\partial S} C(s, S) = \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} \frac{\partial}{\partial S} \mathrm{E}[C(X_n, D_n)]$$

Mathematical Formulation: Derivatives

For any demand d, the function $C(\cdot, d)$ is continuous, except at $x = d+s$. The cadlag process $C(X(t))$ is a Lipschitz continuous function of the process $X(t)$ in the sup-norm.

Stochastic derivative (IPA): for $\omega \in \Omega$,

$$\frac{\partial}{\partial S}C(X_n, D_n) = h\mathbf{1}_{\{X_n > D_n\}} + p\mathbf{1}_{\{X_n < D_n\}}$$

$$\frac{1}{N}\sum_{n=1}^{N}\frac{\partial}{\partial S}C(X_n, D_n) = \frac{hN^+ + pN^-}{N}.$$

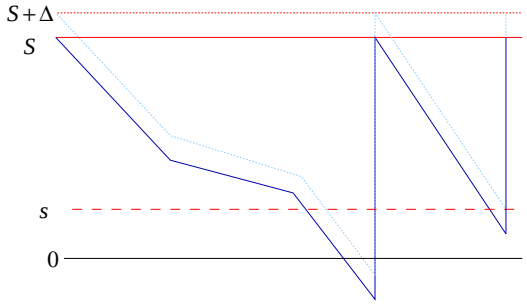

However, the stochastic derivative (IPA) is not consistent for the derivative, that is:

$$\frac{\partial}{\partial S}C(s, S) \neq \lim_{N\to\infty}\frac{h\,\mathrm{E}N^+ + p\,\mathrm{E}N^-}{N}.$$

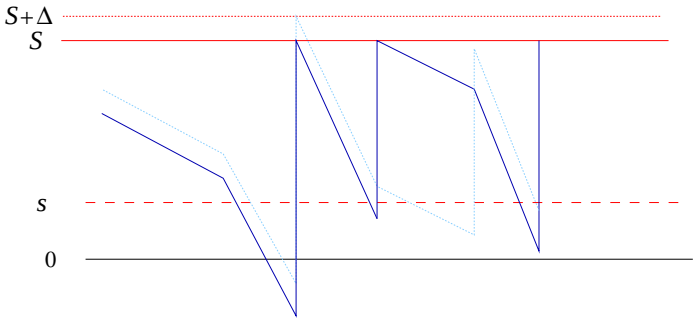

Reason: Because the process $X(t)$ *is not* Lipschitz continuous in S with probability one. To see this, let

$$\tau^+(\omega) = \min\{n > 1 : X_n - s < D_n \leq X_n + \Delta - s\}.$$

Then for all ω with $\tau^+(\omega) < N$, it holds that

$$\sup_{t\in[0,N)} \|X^{S+\Delta}(t,\omega) - X^S(t,\omega)\| \geq S - s - \Delta \neq \mathcal{O}(\Delta).$$

Filtered Conditioning: SPA

$$\tau^+(\omega) = \min\{n > 1 : X_n - s < D_n \leq X_n + \Delta - s\}.$$

Define the *critical events* as $\Omega_n = \{\omega : \tau^+(\omega) = n\}$, which help describe the pathwise discontinuities of $X(t)$. Express Ω as the union of disjoint sets:

$$\Omega = \Omega^* \cup \left(\cup_{n \leq N} \Omega_n\right),$$

and notice that Ω^* is set where the trajectories are Lipschitz continuous, while Ω_n is the set where there is a discontinuity at time n. Assumption 1 implies that critical events are *rare*: $\mathrm{P}(\Omega_n|\mathfrak{F}_{n-1}) = \mathrm{P}(\Omega_n|X_n) = \mathcal{O}(\Delta)$.

On the set Ω^* the IPA is valid for the pathwise derivative. But on the critical events sets Ω_n the processes are not Lipschitz continuous in the sup norm.

Filtered Monte Carlo conditions on the critical events.

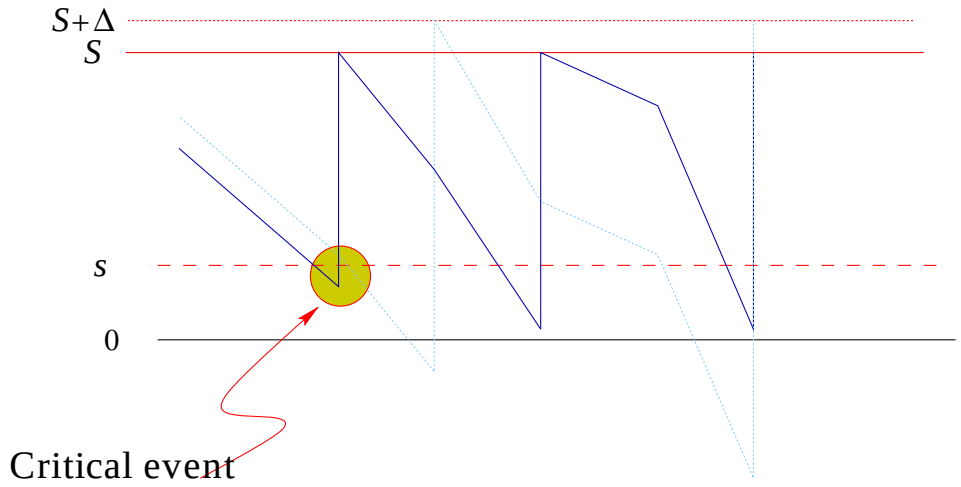

SPA Derivatives: Formulation

Finite Differences: Using the processes at $S+\Delta$ and S (right-sided w.l.o.g.), one has:

$$\frac{\partial}{\partial S}C(s,S) \;=\; \lim_{N\to\infty}\lim_{\Delta\to 0} D^+_\Delta(S),$$

$$\text{where}\;\; D^+_\Delta(S) \;=\; \frac{1}{N}\sum_{n=1}^{N} \mathrm{E}\left[\frac{C_i(S+\Delta)-C_i(S)}{\Delta}\right].$$

Using $\Omega = \Omega^* \cup (\cup_{n\le N}\Omega_n)$, the finite difference can be expressed by further conditioning:

$$D^+_\Delta(S) = \frac{h\,\mathrm{E}N^+ + p\,\mathrm{E}N^-}{N}\mathrm{P}[\Omega^*] + \sum_{n=1}^{N} \mathrm{E}\left(\frac{1}{N}\sum_{i=1}^{N}\mathrm{E}\left[\frac{C_i(S+\Delta)-C_i(S)}{\Delta}\mid \Omega_n\right]\mathrm{P}[\Omega_n|\mathfrak{F}_{n-1}]\right)$$

where $\mathrm{P}(\Omega^*) = 1-\mathcal{O}(\Delta)\to 1$ as $\Delta\to 0$.

Definition: The *SPA derivative estimator is a pathwise functional of the process that estimates the critical probability rates and in parallel, the difference processes started at the critical events.*

$$\lim_{\Delta\to 0}\sum_{n=1}^{N}\mathrm{E}\Big(\frac{1}{N}\sum_{i=1}^{N}\underbrace{\mathrm{E}[C_i(S+\Delta)-C_i(S)\mid\Omega_n]}_{\text{difference process}}\times\underbrace{\frac{\mathrm{P}[\Omega_n|\mathfrak{F}_{n-1}]}{\Delta}}_{\text{critcal rates}}\Big)$$

Difference Processes

Difference Processes: The following properties are essential to the simulation of SPA estimators and they hold for general conditional threshold problems.

Recall that $\Omega_n = \{\omega \in \Omega : \tau^+(\omega) = n\}$ is the set of trajectories where the first discontinuity in the two processes happens at time n. Conditioning on Ω_n,

- As $\Delta \to 0$, *both processes coincide* before n, that is, $C_i(S+\Delta) = C_i(S), i < n$.
- The limit processes for each $\omega \in \Omega_n$ start at the limit values, $X_n^+ = S$ (reorder and fill inventory), $X_n = s$ (no ordering). The rest of their evolution uses the same sequence of demands $D_i, i \geq n$ (CRN). From the Markovian properties, it follows that the difference process *is independent of* n *and of* $\mathfrak{F}_{n-1}$.
- After $\phi(n) = \min\{i > n : X_i^+ = X_i = S\}$, *both processes coincide* again and their difference is zero.

Therefore it is sufficient to estimate:

$$\mathrm{E}[\Delta C^+] = \lim_{N\to\infty} \sum_{i=n}^{N} \mathrm{E}\left[C_i(S+\Delta) - C_i(S) \mid \Omega_n\right] = \mathrm{E}\left\{\sum_{i=n}^{\phi(n)} [C(X_i^+, D_i) - C(X_i, D_i)]\right\}$$

Remark: Under negative perturbation the differences are $C_i(S) - C_i(S-\Delta)$, but in the limit $\Delta \to 0$, the difference process ΔC^- has the same distribution as ΔC^+, and similarly for a double sided perturbation $C_i(S+\Delta/2) - C_i(S-\Delta/2)$

Critical Rates: One-Sided

Definition: The *right-sided SPA* estimator is a conditional estimator of the rates:

$$\lim_{\Delta\to 0}\frac{1}{N}\sum_{n=1}^{N}\frac{\mathrm{P}[\Omega_n|\mathfrak{F}_{n-1}]}{\Delta},= \lim_{\Delta\to 0}\frac{1}{N}\sum_{n=1}^{N}\frac{\mathrm{P}[X_n-s<D_n\le X_n+\Delta-s|\mathfrak{F}_{n-1}]}{\Delta},$$

and the *left-sided SPA* estimator is a conditional estimator of the rates:

$$\lim_{\Delta\to 0}\frac{1}{N}\sum_{n=1}^{N}\frac{\mathrm{P}[\Omega_n|\mathfrak{F}_{n-1}]}{\Delta},= \lim_{\Delta\to 0}\frac{1}{N}\sum_{n=1}^{N}\frac{\mathrm{P}[X_n-\Delta-s<D_n\le X_n-s|\mathfrak{F}_{n-1}]}{\Delta},$$

Right-sided: The critical sets are defined as $\Omega_n^+=\{\omega: D_n>X_n-s\}\in\mathfrak{F}_n$. Only on Ω_n^+ there is a positive probability that $\tau^+(\omega)=n$.

$$\frac{\mathrm{P}[\Omega_i|\Omega_i^+,\mathfrak{F}_{i-1}]}{\Delta}=\frac{\mathrm{P}[D_i-Z_i\le\Delta|\Omega_i^+]}{\Delta}\to\frac{g(Z_i)}{1-G(Z_i)},\quad\text{with}\quad \boxed{Z_i=X_i-s}$$

The corresponding one-sided SPA estimators for the critical rates R are:

$$(R)^+_{SPA}=\frac{1}{N}\sum_{n=1}^{N}\frac{g(Z_n)}{1-G(Z_n)}\mathbf{1}_{\{D_n>Z_n\}},\qquad (R)^-_{SPA}=\frac{1}{N}\sum_{n=1}^{N}\frac{g(Z_n)}{G(Z_n)}\mathbf{1}_{\{D_n>Z_n\}}$$

and are typically estimated through a usual Filtered Monte Carlo, as the sequence of the consecutive demands $\{D_n\}$ is generated.

Efficiency of Estimation

In many Markov Chain simulations a CLT is applied to estimate confidence intervals. The *precision* of an estimator $\hat{Y}(N)$ is proportional to $\sqrt{\text{Var}[\hat{Y}(N)]/N}$, where N: sample size.

Variance reduction $\Rightarrow$ better precision, but CPU time may override the benefits.

Definition: The *efficiency* of a consistent estimator $\hat{Y}(N)$ is:

$$\mathcal{E}(\hat{Y}(N)) = \frac{1}{\text{CPU}[\hat{Y}(N)]\text{Var}[\hat{Y}(N)]}.$$

- Critical Rates: Uses further conditioning (on $\Omega_n^{\pm}$) which may reduce variance,

$$\text{P}[\Omega_i|\mathfrak{F}_{i-1}] = \text{E}\left(\text{P}[\Omega_i|\Omega_i^+, \mathfrak{F}_{i-1}]\right).$$

  Problem: For many distributions of interest the calculation of $G(Z_n)$ may be very costly and increases CPU.

  $\Longrightarrow$ Two-Sided SPA

- Difference Processes: Off-line simulations are used to estimate $\text{E}[\Delta C^+]$. Independent difference processes are simulated for N replications. Problem: May require long CPU time to generate all the demands, which may be costly for many distributions.

  $\Longrightarrow$ Parallel Phantoms

Two-Sided SPA Estimation

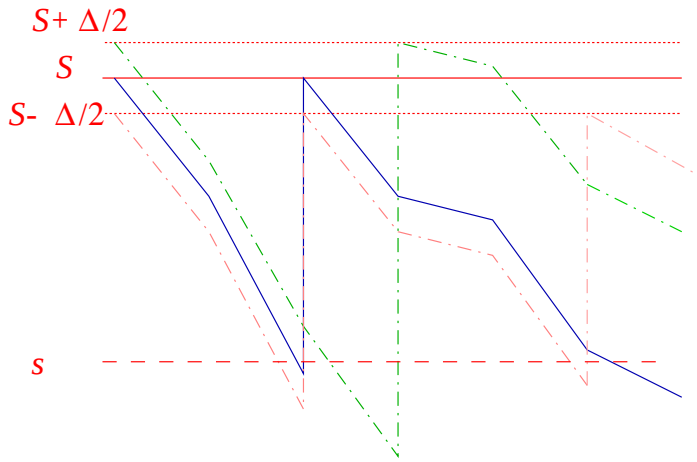

Nominal process $X_S(t)$, perturbation processes $X_{S+\Delta/2}(t)$ and $X_{S-\Delta/2}(t)$.
The critical events now are the first epochs where ordering decisions differ:

$$\tau(\omega) = \min\{n : -\Delta/2 < D_n - Z_n \leq \Delta/2\}.$$

The *two sided SPA* estimator is a conditional estimator of the rates:

$$\lim_{\Delta\to 0} \frac{1}{N}\sum_{n=1}^{N} \frac{\mathrm{P}[\Omega_n|\mathfrak{F}_{n-1}]}{\Delta}, = \lim_{\Delta\to 0} \frac{1}{N}\sum_{n=1}^{N} \frac{\mathrm{P}[X_n - \frac{\Delta}{2} - s < D_n \leq X_n + \frac{\Delta}{2} - s|\mathfrak{F}_{n-1}]}{\Delta}.$$

Using the density g of the demands D_n, the estimator for the critical rates R is:

$$R_{SPA} = \frac{1}{N}\sum_{n=1}^{N} g(Z_n)$$

Estimation through a usual Filtered Monte Carlo, as the sequence of the consecutive demands $\{D_n\}$ is generated. Only $g(\cdot)$ has to be evaluated and there is no random denominator.

Parallel Phantoms

Difference process $d_n = X_n^+ - X_n$ can be calculated recursively:

1. Start at $X_0^+ = S, X_0 = s$, set $d_0 = S - s$, $n = 1$.
2. If $d_n \neq 0$ set $n := n + 1$, else $\phi = n$ and END OF CYCLE.
3. Generate D_n, calculate X_{n+1} and update difference:

   If $d_n > 0$ then:

$$d_{n+1} = \begin{cases} d_n & \text{if } D_n < X_n^+ + d_n - s \text{ (both above reorder level } s) \\ X_n^+ - D_n - S < 0 & \text{if } X_n^+ - d_n - s < D_n < X_n^+ - s \text{ (one above and one below)} \\ 0 & \text{if } D_n > X_n^+ - s \text{ (both processes reach } S) \end{cases}$$

   If $d_n < 0$ then:

$$d_{n+1} = \begin{cases} d_n & \text{if } D_n < X_n^+ - s \text{ (both above reorder level } s) \\ S + D_n - X_n^+ - d_n > 0 & \text{if } X_n^+ - s < D_n < X_n^+ - d_n - s \text{ (one above and one below)} \\ 0 & \text{if } D_n > X_n^+ - d_n - s \text{ (both processes reach } S) \end{cases}$$

4. Go to 2.

Parallel Phantoms: Only the *nominal process* X_n^+ is simulated. Use X_n^+ and the consecutive demands to estimate in parallel the difference process by starting a *phantom process* with initial difference $d_1(j) = S - s$ *every time* that $X_j^+ = S$. Only *one simulation* $(D_1, D_2, \ldots, D_N)$ is used for estimating $\mathrm{E}[\Delta C^+]$. That is, a random number of phantom processes $d_n(j), j = 1, 2, \ldots$ is calculated along with the nominal simulation.

Parallel Phantom SPA

$$\frac{\partial}{\partial S}C(s,S) = \lim_{N\to\infty} \frac{h\,\mathrm{E}N^+ + p\,\mathrm{E}N^-}{N} + \mathrm{E}[\Delta C^+] \lim_{N\to\infty} \mathrm{E}\left(\frac{1}{N}\sum_{n=1}^{N} g(X_n - s)\right)$$

$M(N)$: number of steps $j = 1,\ldots,N$ such that $X_j = S$. The Phantom SPA is:

$$\widehat{(\Delta C^+)}_N = \frac{1}{M(N)}\sum_{j=1}^{N} \mathbf{1}_{\{X_j=S\}} \sum_{n=j}^{\phi(j)} d_n(j), \text{ and}$$

$$\widehat{C'(s,S)} = \frac{hN^+ + pN^-}{N} + \widehat{(\Delta C^+)}_N \sum_{n=1}^{N} \frac{g(X_n - s)}{N},$$

Proposition 1 *The Phantom SPA estimator* $\widehat{C'(s,S)}$ *is strongly consistent for* $\frac{\partial}{\partial S}C(s,S)$.

Proof : Given any index j and $n \geq j$, $\mathrm{P}[d_{n+1} = 0] \geq \mathrm{P}[D_n > S - s] \geq \pi > 0$. This implies that $\phi(j)$ is *stochastically dominated* by a geometric random variable with finite expectation $(1/\pi)$. Also, $\|d_n(j)\| < S - s$, w.p.1. Use Assumption 1, dominated convergence and uniqueness of the stationary measure $\mu(dx)$ for the process, to obtain the result:

$$\lim_{N\to\infty} \widehat{(\Delta C^+)}_N = \mathrm{E}[\Delta C^+], \quad \text{and} \lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N} g(X_n - s) = \int_s^S g(x-s)\mu(dx).$$

Computer Simulations

$K = 10, h = 10, p = 50, s = 10, S = 20$. $N = 5000$ nominal, $M = 5000$ cycles off-line.
Comparison of Two-Sided SPA with Fu's One-Sided Off-line estimators:

Results with Off-Line Simulations				
Distribution	Theoretical	SPA^-	SPA^+	SPA^*
$\text{Exp}(1/3)$	-4.531	-4.680 ± 0.26	-4.570 ± 0.10	-4.572 ± 0.08
$\text{Exp}(1/5)$	-2.862	-2.779 ± 0.145	-2.843 ± 0.088	-2.837 ± 0.082
$\text{Exp}(1/10)$	1.516	1.586 ± 0.114	1.569 ± 0.094	1.576 ± 0.101
$U(0, 6)$	–	-5.185 ± 0.171	-5.225 ± 0.171	-5.261 ± 0.122
$U(0, 12)$	–	-5.359 ± 0.211	-5.312 ± 0.088	-5.259 ± 0.044
$U(0, 20)$	–	-2.181 ± 0.100	-2.212 ± 0.084	-2.200 ± 0.079

Comparison of Two-Sided Off-line estimation with Parallel Phantom SPA:

	Off-Line		Parallel Phantoms	
Distribution	SPA^*	CPU (secs)	SPA^*	CPU (secs)
$\text{Exp}(1/3)$	-4.509 ± 0.081	25.54	-4.530 ± 0.027	11.91
$\text{Exp}(1/5)$	-2.816 ± 0.097	14.99	-2.829 ± 0.056	8.40
$\text{Exp}(1/10)$	-1.488 ± 0.088	10.5	-1.497 ± 0.078	6.75
$U(0, 6)$	-5.180 ± 0.085	45.91	-5.281 ± 0.025	21.69
$U(0, 12)$	-5.234 ± 0.036	11.31	-5.232 ± 0.018	7.30
$U(0, 20)$	2.242 ± 0.130	7.19	2.224 ± 0.084	5.32

Concluding Remarks

- Two-Sided vs One-Sided: amount of variance reduction seems to be problem dependent.
- Parallel Phantoms: The greatest contribution to efficiency improvement of the Phantom SPA comes from the parallel on-line systems.
- Difference Processes: Use of CRN via parallel phantom processes: exploit difference process recursions to decrease CPU time.
- Conjecture: high *negative* correlation between phantom systems.

  On going:

- General formulation for threshold problems with conditionally independent kernels:

$$\mathrm{P}[X_{n+1} \in ds | X_n] = H(D_n), \quad \{D_n\} \text{ i.i.d.}, \quad H(\cdot) \text{ has a mass at threshold values.}$$

  – Alternative formulation of two-sided critical rates via *weak differentiation* approach (in progress).

  – Extension of the properties: the distribution of the difference process is conditionally independent of (filtered) critical rates.

- Examples: barrier options in finance (estimation of the greeks), queueing systems under random polling policies, optimal replacement models, etc.