Derivative Estimation.

Model:

$$
\begin{aligned}
\theta \in \mathbb{R} &: \text{ Continuous parameter.} \\
(\Omega, \Sigma, P_\theta) &: \text{ Probability space.} \\
\omega \in \Omega &: \text{ Sample point (random element).} \\
h(\theta, \omega) &: \text{ Value of objective function (Σ-measurable).} \\
\alpha(\theta) &= \int_\Omega h(\theta, \omega) dP_\theta(\omega) \;=\; E_\theta[h(\theta, \omega)] \text{ (expectation).} \\
\alpha'(\theta) &: \text{ What we want to estimate.}
\end{aligned}
$$

WHY ?

— Sensitivity analysis;

— Stochastic optimization;

— Etc.

HOW ?

— Finite differences (FD), (FDC);

— Infinitesimal perturbation analysis (IPA);

— Likelihood ratios (LR or SF).

— Weak Derivatives (WD).

For $\theta \in \mathbb{R}^d$, just apply the method to each component of θ.

Example 1: A GI/G/1 Queue

Let θ be a scale parameter of service time distribution

$$B_\theta(s) = B(s/\theta).$$

$S_i = \theta Z_i =$ service time for customer i, where $Z_i : B$.

$A_i =$ interarrival times between customers i and $i+1$.

$W_i =$ waiting time for customer i.

$X_i = W_i + S_i =$ sojourn time for customer i.

Assuming system is initially empty, one has $W_1 = 0$,

$$X_i = W_i + \theta Z_i; \qquad W_{i+1} = (X_i - A_i)^+.$$

Let

$$h_t(\theta, \omega) = \sum_{i=1}^{t} X_i.$$

$\alpha_t(\theta)$ is the expected total sojourn time for first t customers.

Want to estimate $\alpha_t'(\theta)$.

Finite Differences (FD)

Independent random numbers (FDI). (Classical).

(Centered version):

Simulate at $\theta - c$ to obtain estimate $h(\theta - c, \omega^-)$;

Simulate at $\theta + c$ to obtain estimate $h(\theta + c, \omega^+)$;

$$Y = \frac{h(\theta + c, \omega^+) - h(\theta - c, \omega^-)}{2c}.$$

— Estimator is biased. If α''' continuous, then bias in $O(c^2)$.

— However, $\mathrm{var}[Y] = O(c^{-2}) \to \infty$ as $c \to 0$.

— Requires $2d$ simulations for d parameters.

— Also one-sided version: $d + 1$ simulations, but bias in $O(c)$ ...

Common random numbers (FDC).

Let $P_\theta \equiv P$ (independent of θ). Same ω for both runs.

That is, use same stream of random numbers, and synchronize.

Idea: compare two very similar systems under same conditions.

May reduce variance considerably.

If $P[h(\cdot, \omega)$ discont. in $(\theta - c, \theta + c)] = O(c)$, + some mild tech. conds.,

then $\mathrm{var}[Y] = O(c^{-1})$. CRN-I.

If $h(\cdot, \omega)$ is a.s. continuous, + some mild tech. conds.,

then $\mathrm{var}[Y] = O(1)$. CRN-II.

Can we save by simulating only once and tracing the few changes ?

...Leads to perturbation analysis.

Stochastic Derivative, or Infinitesimal Perturbation Analysis (IPA)

Let $P_\theta \equiv P$ be indep. of θ.

As a derivative estimator, consider the "process derivative":

$$h'(\theta,\omega) = \frac{\partial}{\partial\theta} h(\theta,\omega).$$

This is the basis of IPA derivative estimation.

(Ho, Cao, Glasserman, Suri, Heidelberger, Zazanis, etc.)

Under appropriate conditions, can interchange E and derivative:

$$E[h'(\theta,\omega)] = \alpha'(\theta).$$

Q: What about these appropriate conditions ?

A: Uniform integrability. See also next slide.

Continuity of $h(\cdot,\omega)$ is typically a key condition.

Q: How to compute $h'(\theta,\omega)$ for a given ω ?

A: A bunch of IPA techniques have been designed for that.

Basic idea of IPA:

— Generate ω and run simulation with $\theta = \theta_0$.

— In the process, for fixed ω, propagate infinitesimal perturbation on θ over the "history", assuming that sequence of events do not change, but that events only slide in time.

— Needs only one simulation to compute the derivative estimator.

Interchanging the derivative and expectation

ASSUMPTION A1(k), for $k \geq 1$:

$\exists \Upsilon$ neighborhood of θ_0, and $\Gamma : \Omega \to \mathbb{R}$, such that P-a.s.,

— $h(\cdot, \omega)$ exists and is continuous over Υ;

— $h(\cdot, \omega)$ is differentiable everywhere in $D(\omega) \subseteq \Upsilon$,

where $\Upsilon \setminus D(\omega)$ is at most denumerable and $\theta_0 \notin D(\omega)$;

— $\sup_{\theta \in D(\omega)} |h'(\theta, \omega)|^k \leq \Gamma(\omega)$ and $\int_\Omega \Gamma(\omega) dP(\omega) < \infty$.

THEOREM 1.

(a) Under A1(1), $\alpha(\theta)$ is differentiable at every $\theta \in \Upsilon$ and

$E[h'(\theta, \omega)] = \alpha'(\theta)$.

(b) Under A1(1), if $h'(\cdot, \omega)$ is continuous all over Υ for each $\omega \in \Xi$,

then α is continuously differentiable in Υ.

(c) Under A1(k), $E[(h'(\theta, \omega))^k] < \infty$.

PROOF. (see Glasserman 1988 and L'Ecuyer 1990).

From "generalized" mean value theorem (Dieudonné (1969)),

if θ_0 and $\theta_0 + c$ are in Υ and ω satisfies A1,

$$\left|\frac{h(\theta_0 + c, \omega) - h(\theta_0, \omega)}{c}\right| \;\le\; \sup_{\theta \in D(\omega)} |h'(\theta, \omega)| \;\le\; \Gamma(\omega).$$

For $k = 1$, from the dominated convergence theorem,

$$\begin{aligned}
\int_\Omega h'(\theta_0, \omega) dP(\omega) \;&=\; \int_\Omega \left(\lim_{c \to 0} \frac{h(\theta_0 + c, \omega) - h(\theta_0, \omega)}{c}\right) dP(\omega) \\
&= \lim_{c \to 0} \left(\frac{\alpha(\theta_0 + c) - \alpha(\theta_0)}{c}\right) \\
&= \nabla \alpha(\theta_0).
\end{aligned}$$

which proves (a). From DCT, (b–c) follow easily. ∎

Example 1: $GI/G/1$ queue (continuation).

Recall:

$$X_i = W_i + \theta Z_i; \qquad W_i = (X_{i-1} - A_{i-1})^+ \qquad h(\theta, \omega) = \sum_{i=1}^{t} X_i.$$

Can view ω as $(Z_1, A_1, \ldots, Z_t)$. Then

$$X_i' = W_i' + Z_i = X_{i-1}' I(W_i > 0) + Z_i.$$

In other words,

$$X_i' = \sum_{j=v_i}^{i} Z_j$$

where v_i is the first customer in same busy period as i.

The IPA estimator is then

$$h'(\theta, \omega) = \sum_{i=1}^{t} \sum_{j=v_i}^{i} Z_j.$$

Can be computed easily during a simulation.

Assumption A1 is easy to verify in this case.

Indeed, $h(\theta, \omega)$ is clearly continuous in θ.

Also differentiable in θ except when two events occur simultaneously, which happens at most for a finite number of values of θ. Also, $|h'(\theta, \omega)|^k \le \left(t \sum_{i=1}^{t} Z_i\right)^k$, which is integrable if $E[Z_i^k] < \infty$.

Example 2: $GI/G/1$ again.

Let $\alpha_i(\theta) = P[X_i > L]$ and $h_i(\theta, \omega) = I(X_i > L)$, for given i and L.

If ω is underlying sequence of $U(0, 1)$, then $h_j'(\theta, \omega) = 0$ a.s.,

so straightforward IPA fails in this case!

Smoothed Perturbation Analysis (SPA)

Idea: smooth out $h(\theta,\omega)$; replace it by a conditional expectation.

(Ho, Gong, Glasserman, Wardi, L'Ecuyer, etc.)

Equivalently, view ω as representing less information.

Example 2: (continuation).

Let ω_i "represent" $(Z_1, A_1, \ldots, Z_{i-1}, A_{i-1})$. Redefine

$$\begin{aligned} h_i(\theta,\omega_i) &= P[X_i > L \mid \theta,\omega_i] = P[X_i > L \mid W_i] \\ &= P[S_i > L - W_i \mid W_i] = 1 - B_\theta(L - W_i). \end{aligned}$$

The corresponding IPA estimator is (recall that $B_\theta(s) = B(s)/\theta$):

$$\begin{aligned} h_i'(\theta,\omega_i) &= -\left(\frac{\partial}{\partial\theta} B_\theta\right)(L - W_i) + b_\theta(L - W_i) W_i' \\ &= b_\theta(L - W_i)\left[(L - W_i)/\theta^2 + X_{i-1}' I(X_{i-1} \geq A_{i-1})\right]. \end{aligned}$$

If b_θ is bounded, then Theorem 1 applies (unbiased estimator).

Now, if we want to estimate

$$\alpha'(\theta) = \frac{1}{t}\sum_{i=1}^{t} \alpha_i'(\theta) = \frac{1}{t}\sum_{i=1}^{t} P'[X_i > L],$$

we can use $(1/t)\sum_{i=1}^{t} h_i'(\theta,\omega_i)$.

Alternative choice for the "meaning" of ω_i: $(Z_1, A_1, \ldots, A_{i-2}, Z_{i-1})$.

Will give less variance, since we condition on less.

For a given application, different SPA derivative estimators can often be obtained by conditioning on different things, i.e. with different choices of what ω represents.

Other variants of IPA: RPA, LTP, EPA, ... (Ho and Cao 1991).

Likelihood Ratio (LR) or Score Function (SF) Method

(Aleksandrov et al., Rubinstein, Glynn, Reiman and Weiss, ...)

Suppose that P_θ really depends on θ.

Then, how can we differentiate $\int_\Omega h(\theta,\omega)dP_\theta(\omega)$?

The following trick will bring us back to the IPA setting.

Example 1: $GI/GI/1$ queue (continuation).

Let ω represent $(S_1, A_1, \ldots, S_t)$. Now, $h(\theta,\omega) \equiv h(\omega)$.

Interarrival time density f.

Service time density b_θ, positive over $(0,\infty)$.

Let b be another service time density, indep. of θ. Then

$$\begin{aligned}
\alpha(\theta) &= E_\theta[h(\omega)] \;=\; \int_\Omega h(\omega)dP_\theta(\omega) \\
&= \int_0^\infty \cdots \int_0^\infty h(s_1, a_1, \ldots, s_t) b_\theta(s_1) f(a_1) \cdots b_\theta(s_t) ds_1 da_1 \cdots ds_t \\
&= \int_0^\infty \cdots \int_0^\infty h(s_1, a_1, \ldots, s_t) \left(\frac{b_\theta(s_1) f(a_1) \cdots b_\theta(s_t)}{b(s_1) f(a_1) \cdots b(s_t)}\right) \\
&\qquad\qquad \cdot\; b(s_1) f(a_1) \cdots b(s_t) ds_1 da_1 \cdots ds_t \\
&= \int_0^\infty \cdots \int_0^\infty H(s_1, a_1, \ldots, s_t) b(s_1) f(a_1) \cdots b(s_t) ds_1 da_1 \cdots ds_t \\
&= \int_\Omega H(\theta,\omega) dP(\omega) \;=\; E_b[H(\theta,\omega)].
\end{aligned}$$

where

$$H(\theta,\omega) \;=\; h(\omega) L(b,\theta,\omega) \;=\; h(\omega)\frac{b_\theta(S_1)\cdots b_\theta(S_t)}{b(S_1)\cdots b(S_t)}.$$

Now, consider the stochastic derivative of $H(\cdot, \omega)$.

Formally, same as IPA.

$$
\begin{aligned}
H'(\theta, \omega) &= h(\omega) L'(b, \theta, \omega) = h(\omega) L(b, \theta, \omega) S(\theta, \omega); \\
S(\theta, \omega) &= \frac{L'(b, \theta, \omega)}{L(b, \theta, \omega)} = \frac{\partial}{\partial \theta} \ln L(b, \theta, \omega) \\
&= \frac{\partial}{\partial \theta} \left[\sum_{i=1}^{t} \ln(b_\theta(S_i)) \right].
\end{aligned}
$$

Specific example: $GI/M/1$. Let S_i : exponential with mean θ.

Let $b = b_{\theta_0}$ for some fixed θ_0. Then,

$$
\begin{aligned}
L(b, \theta, \omega) &= \prod_{i=1}^{t} \frac{b_\theta(S_i)}{b_{\theta_0}(S_i)} = \left(\frac{\theta_0}{\theta}\right)^t \exp\left(\left(\frac{1}{\theta_0} - \frac{1}{\theta}\right) \sum_{i=1}^{t} S_i\right); \\
S(\theta, \omega) &= \sum_{i=1}^{t} \frac{\partial}{\partial \theta} \ln(b_\theta(S_i)) = -\frac{t}{\theta} + \frac{1}{\theta^2} \sum_{i=1}^{t} S_i.
\end{aligned}
$$

For any k, Assumption A1(k) holds for θ near θ_0.

Can estimate cost and deriv. around θ_0 by a single simul. at θ_0:

$$
\begin{aligned}
H(\theta, \omega) &= h(\omega) \left(\frac{\theta_0}{\theta}\right)^t \exp\left(\left(\frac{1}{\theta_0} - \frac{1}{\theta}\right) \sum_{i=1}^{t} S_i\right); \\
H'(\theta, \omega) &= h(\omega) \left(\frac{\theta_0}{\theta}\right)^t \exp\left(\left(\frac{1}{\theta_0} - \frac{1}{\theta}\right) \sum_{i=1}^{t} S_i\right) \left(-\frac{t}{\theta} + \frac{1}{\theta^2} \sum_{i=1}^{t} S_i\right).
\end{aligned}
$$

Note: this applies as well to example 2.

What about estimating cost by H, then differentiate ?

What about variance ?

Example 1 again, with

$$h(\omega) = \frac{1}{t} \sum_{i=1}^{t} X_i.$$

Simulate at θ_0 to estimate the derivative at all θ.

For this example, one can show that:

1. $E[H'(\theta, \omega)] = \alpha'(\theta)$ for all $\theta \in (0, \infty)$.
2. $E[(H'(\theta, \omega))^k] < \infty$ iff $(k-1)\theta < k\theta_0$.

$$\begin{aligned}
\mathrm{Var}[H(\theta,\omega)] &= \left(\frac{\theta_0\tilde{\theta}}{\theta^2}\right)^t E_{\tilde{\theta}}[h^2(\omega)] \;-\; \alpha^2(\theta) \\
&= O\left(\left(\theta_0\tilde{\theta}/\theta^2\right)^t t^{-1}\right) \qquad \text{for } \theta < \theta_0; \\
\mathrm{Var}[H'(\theta,\omega)] &= \left(\frac{\theta_0\tilde{\theta}}{\theta^2}\right)^t E_{\tilde{\theta}}\left[\frac{h^2(\omega)}{\theta^2}\left(\sum_{i=1}^{t}(S_i-\theta)\right)^2\right] \;-\; [\alpha'(\theta)]^2 \\
&= O\left(\left(\theta_0\tilde{\theta}/\theta^2\right)^t t\right) \qquad \text{for } \theta < \theta_0,
\end{aligned}$$

where $\tilde{\theta} = \theta\theta_0/(2\theta_0 - \theta)$.

Generalization.

Let P, indep. of θ, s.t. $B \in \Sigma$ and $P(B) = 0$ implies $P_\theta(B) = 0 \ \forall\theta$.

$$\alpha(\theta) = \int_\Omega h(\theta,\omega)dP_\theta(\omega) = \int_\Omega H(\theta,\omega)dP(\omega),$$

where

$$H(\theta,\omega) = h(\theta,\omega)\frac{dP_\theta}{dP}(\omega) = h(\theta,\omega)L(P,\theta,\omega).$$

This is related to importance sampling (IS). Now,

$$H'(\theta,\omega) = [h'(\theta,\omega) + h(\theta,\omega)S(\theta,\omega)]L(P,\theta,\omega)$$

where ω "follows" P and

$$S(\theta,\omega) = \frac{L'(P,\theta,\omega)}{L(P,\theta,\omega)} = \frac{\partial}{\partial\theta}\ln L(P,\theta,\omega).$$

If $h(\theta,\omega) \equiv h(\omega)$ and P_θ, P have densities f_θ, f, then

$$\begin{aligned}
H'(\theta,\omega) &= h(\omega)S(\theta,\omega)L(P,\theta,\omega) \\
L(P,\theta,\omega) &= f_\theta(\omega)/f(\omega) \\
S(\theta,\omega) &= \frac{\partial}{\partial\theta}\ln f_\theta(\omega).
\end{aligned}$$

From simul. based on f, can estimate α and α' "everywhere".

Note: Different ways of defining the probability space,

i.e., of deciding on what ω would represent.

Can we also estimate "everywhere" with IPA ?

Sometimes yes. Back to Example 1.

For LR/SF, we took the gradient AFTER changing the measure.

Can we do the opposite: differentiate first, then use the LR ?

This is applying IS (or LR) on top of IPA.

GI/M/1 example (continued).

First, interpret ω as the iid uniforms used to drive the simulation.

Let $S_j = \theta Z_j = -\theta \ln(1 - U_j)$. Then,

$$h'(\theta, \omega) \;=\; \frac{1}{t} \sum_{i=1}^{t} \sum_{j=\nu_i}^{i} Z_j \;=\; \frac{1}{t} \sum_{i=1}^{t} \sum_{j=\nu_i}^{i} S_j/\theta$$

is an unbiased estimator of $\alpha'(\theta) = E'\left[\frac{1}{t} \sum_{i=1}^{t} X_i\right]$,

where ν_i = first cust. in same busy period as cust. i.

To estimate the derivative everywhere, reinterpret ω as the sequence of interarrival and service times. Simulate at θ_0 and compute

$$H'(\theta, \omega) \;=\; \left(\frac{1}{t} \sum_{i=1}^{t} \sum_{j=\nu_i}^{i} S_j/\theta\right) \left(\frac{\theta_0}{\theta}\right)^{t} \exp\left(\left(\frac{1}{\theta_0} - \frac{1}{\theta}\right) \sum_{i=1}^{t} S_i\right).$$

Asymptotic variance of H':

$O((\theta_0^2/(2\theta_0 - \theta)\theta^2)^t\, t)$ for SF;

$O((\theta_0^2/(2\theta_0 - \theta)\theta^2)^t\, t^{-1})$ for IPA.

Functional estimation via LR/SF or IPA/SF.

Simulate using P to obtain $H'(\cdot, \omega)$ all over Θ.

Repeat N times, to get $H'_1(\cdot, \omega_1), \ldots, H'_N(\cdot, \omega_N)$. Let

$$\bar{H}'_N(\cdot) \;=\; \frac{1}{N} \sum_{i=1}^{N} H'_i(\cdot, \omega_i).$$

(Note: can do the same with $\bar{H}(\cdot, \omega)$).

Main practical difficulties:

1. Choice of P.
2. Storing and retrieving the information required for computing $\bar{H}'_N(\theta)$ at any θ.
3. If N not large enough, $\bar{H}'_N(\cdot)$ may be badly behaved.
4. Variance and bias ...
5. Getting aware that the variance is high for some θ's . . . .

Numerical illustration.

Take an $M/M/1$ queue with mean service time θ, $\lambda = 1$.

Let $t = 2,\ 10,\ 100$ and $N = 10^5$ replications.

Simulate at $\theta_0 = 0.8$ to estimate $\alpha_t(\cdot)$ and $\alpha_t'(\cdot)$ everywhere.

Estimate also the variance and compare with empirical variance.

Repeat with $\theta_0 = 0.5$.

Note:

Variance of $H'(\theta, \omega)$ is finite iff $\theta < 2\theta_0$;

Variance of variance estimator is finite iff $\theta < (4/3)\theta_0$.

When $\theta \to 0$, both the variance of H' and the variance of the sample variance tend to infinity.

SF-LR:

$$
\begin{aligned}
\theta = \theta_0 &\ :\ \mathrm{Var}[H'] = O(t) \\
\theta \neq \theta_0 &\ :\ \mathrm{Var}[H'] = O\left(\left[\theta_0^2/(\theta(2\theta_0 - \theta))\right]^t t\right)
\end{aligned}
$$

IPA-LR:

$$
\begin{aligned}
\theta = \theta_0 &\ :\ \mathrm{Var}[H'] = O(t^{-1}) \\
\theta \neq \theta_0 &\ :\ \mathrm{Var}[H'] = O\left(\left[\theta_0^2/(\theta(2\theta_0 - \theta))\right]^t t^{-1}\right)
\end{aligned}
$$

Example 3: (variant of example 1).

$0 < a < b$

$0 < \theta < 1$

$P[S_i = b] = \theta;$

$P[S_i = a] = 1 - \theta.$

a) $\omega = (U_1, \ldots, U_{2t-1})$ (IPA).

$U_{2i-1} \leq \theta \quad \Rightarrow \quad C_i = 1,\ S_i = b;$

$U_{2i-1} > \theta \quad \Rightarrow \quad C_i = 0,\ S_i = a;$

$h(\cdot, \omega)$ discontinuous at $\theta = U_{2i-1}$.

Direct IPA fails.

b) $\omega = (C_1, A_1 \ldots, C_t)$ (LR/SF).

Then, Theorem 1 applies.

$$\begin{aligned} h'(\theta, \omega) &= 0, \\ S(\theta, \omega) &= \sum_{i=1}^{t} \left(\frac{C_i}{\theta} + \frac{1 - C_i}{1 - \theta} \right). \end{aligned}$$

Example 4.

$0 < p < 1$

$0 < \theta$

$P[S_i = \theta] = p;$

$P[S_i = 2\theta] = 1 - p.$

a) $\omega = (S_1, A_1, \ldots, S_t), \qquad G = P_{\theta_0}.$

$\Rightarrow$ no neighborhood of θ_0 in which LR exists.

($H(\cdot, \omega)$ discontinuous at θ_0).

b) $\omega = (U_1, \ldots, U_{2t-1})$ (IPA).

$U_{2i-1} \leq p \quad \Rightarrow \quad S_i = \theta;$

$U_{2i-1} > p \quad \Rightarrow \quad S_i = 2\theta.$

$h'(\theta, \omega)$: same as in Example 1.

Example 5. $k\theta$ times a Bernouilli (θ).

$$
\begin{aligned}
U_{2i-1} \geq \theta \quad &\Rightarrow \quad C_i = 0,\ S_i = 0; \qquad P[S_i = 0] = 1 - \theta; \\
U_{2i-1} < \theta \quad &\Rightarrow \quad C_i = 1,\ S_i = k\theta; \qquad P[S_i = k\theta] = \theta.
\end{aligned}
$$

a) $\omega = (U_1, A_1, \ldots, U_{2t-1})$: Direct IPA fails.

b) $\omega = (S_1, A_1, \ldots, S_t)$: Direct LR/SF fails.

c) But $\omega = (C_1, A_1, \ldots, C_t)$ works.

$$
\begin{aligned}
h'(\theta, \omega) &\qquad \text{exactly as in Examples 1 and 4.} \\
S(\theta, \omega) &= \sum_{i=1}^{t} \frac{(C_i - \theta)}{\theta(1 - \theta)}.
\end{aligned}
$$

Experiment: $t = 10$, $k = 2$, 10^5 indep. replications. 95% c.i.

θ	True grad.	SF	FDC
0.2	1.038	$1.03 \pm .01$	$1.03 \pm .01$
0.5	4.529	$4.49 \pm .05$	$4.51 \pm .04$
0.9	16.671	$16.61 \pm .30$	$16.75 \pm .12$

Weak Derivatives (WD)

(Pflug 1989, 1991). P_θ still depends on θ.

Idea is to differentiate the probability measure P_θ directly.

Under appropriate conditions, one can write

$$\begin{aligned} \alpha'(\theta) &= \int_\Omega h'(\theta,\omega)dP_\theta(\omega) + \int_\Omega h(\theta,\omega)dP'_\theta(\omega), \\ P'_\theta &= c(\theta)(\dot{P}_\theta - \ddot{P}_\theta), \end{aligned}$$

where $\dot{P}_\theta$ and $\ddot{P}_\theta$ are probability measures on (Ω, Σ).

If ω, $\dot{\omega}$ and $\ddot{\omega}$ are generated from P_θ, $\dot{P}_\theta$ and $\ddot{P}_\theta$, respectively,

the WD derivative estimator is

$$h'(\theta,\omega) + c(\theta)(h(\theta,\dot{\omega}) - h(\theta,\ddot{\omega})).$$

For example, if $\Omega = \mathbb{R}$ and P_θ has density f_θ, then,

under appropriate regularity conditions,

$$\begin{aligned} c(\theta) &= c = \int_{-\infty}^{\infty} \left(\frac{\partial}{\partial\theta} f_\theta(\omega)\right)^+ d\omega = \int_{-\infty}^{\infty} \left(\frac{\partial}{\partial\theta} f_\theta(\omega)\right)^- d\omega, \\ \dot{f}_\theta &= \frac{1}{c}\left(\frac{\partial}{\partial\theta} f_\theta\right)^+, \\ \ddot{f}_\theta &= \frac{1}{c}\left(\frac{\partial}{\partial\theta} f_\theta\right)^-. \end{aligned}$$

Generate $\dot{\omega}$ and $\ddot{\omega}$ using densities $\dot{f}_\theta$ and $\ddot{f}_\theta$, respectively.

One disadvantage: needs more than one simulation run.

Convergence Rates (finite horizon).

Suppose we perform n independent replications for each method.

Yields iid sample $Y_1, \ldots, Y_n$. Gradient estimator:

$$\bar{Y}_n \;=\; \frac{1}{n}\sum_{i=1}^{n} Y_i.$$

Mean square error $= E[(\bar{Y}_n - \alpha'(\theta))^2] \;=\; \mathrm{Var}[\bar{Y}_n] + (\mathrm{Bias})^2$.

Asymptotic order of MSE:

IPA	:	$O(n^{-1})$	(canonical rate)
SF	:	$O(n^{-1})$	
WD	:	$O(n^{-1})$	
FDI (one-sided)	:	$O(n^{-1/2})$	with $c_n = O(n^{-1/4})$
FDI (centered)	:	$O(n^{-2/3})$	with $c_n = O(n^{-1/6})$
FDC-I (one-sided)	:	$O(n^{-2/3})$	with $c_n = O(n^{-1/3})$
		if $\mathrm{Var}[Y] = O(c^{-1})$;	
		Happens if $P[\text{discont.}] = O(c)$.	
FDC-I (centered)	:	$O(n^{-4/5})$	with $c_n = O(n^{-1/5})$
FDC-II (one-sided)	:	$O(n^{-1})$	with $c_n = O(n^{-1/2})$
		if $\mathrm{Var}[Y] = O(1)$.	
		Happens when IPA applies.	
FDC-II (centered)	:	$O(n^{-1})$	with $c_n = O(n^{-1/4})$.

Also: hidden constants...

Increasing the horizon

A Discrete-Time Markov Chain

Markov chain $\{X_i,\ i \geq 0\}$, $X_0 = s_0$, general state space S.

Chain evolves as follows:

At step i, generate ζ_i from density $q_\theta(\cdot \mid X_{i-1})$;

incur cost $C_i = c(\theta, X_{i-1}, \zeta_i)$;

next state is $X_i = \varphi(\theta, X_{i-1}, \zeta_i)$.

Deterministic Horizon t: $\omega = (\zeta_1, \ldots, \zeta_t)$,

$$h_t(\theta, \omega) = \sum_{i=1}^{t} C_i.$$

IPA:

If q_θ indep. of θ and if c and φ are "smooth enough":

$$h_t'(\theta, \omega) = \sum_{i=1}^{t} C_i'.$$

Typically (e.g., if the process is regenerative + finite second moments), then $\mathrm{Var}[h_t(\theta, \omega)/t] = O(1/t)$ and $\mathrm{Var}[h_t'(\theta, \omega)/t] = O(1/t)$.

LR/SF:

If q_θ depends on θ, choose $\{q(\cdot \mid s),\ s \in S\}$ and simulate using q instead of q_θ.

$$
\begin{aligned}
L_t(P,\theta,\omega) &= \prod_{i=1}^{t} \frac{q_\theta(\zeta_i \mid X_{i-1})}{q(\zeta_i \mid X_{i-1})}.\\
S_t(\theta,\omega) &= \sum_{i=1}^{t} \frac{\partial}{\partial\theta} \ln(q_\theta(\zeta_i \mid X_{i-1})).
\end{aligned}
$$

Major problem: variance of $S_t(\theta,\omega)$ increases linearly in t.

Suppose $h_t(\theta,\omega)/t$ converges to a constant K;

take $q = q_{\theta_0}$, and assume $h_t(\theta,\omega) \equiv h_t(\omega)$.

Then, under mild conditions,

$$\mathrm{Var}[H_t'(\theta_0,\omega)/t] = \mathrm{Var}[h_t(\omega)S_t(\theta_0,\omega)/t] = O(t).$$

Really bad ! But if $K = 0$, that variance is then $O(1)$.

To exploit that property, redefine

$$h_t(\omega) = \sum_{i=1}^{t} (C_i - K).$$

That does not change $\alpha_t'(\theta)$.

If K is unknown (usual case), replace it by an estimate.

Equivalent to using the score function as a control variate.

L'Ecuyer and Glynn (1994) give appropriate limit theorems.

Triangular LR/SF derivative estimator.

Estimate the derivative of $E_\theta[C_i]$ separately for each i.

Appropriate score function for C_i is $S_i(\theta, \omega)$.

Then, take the average. Triangular LR/SF estimator:

$$\sum_{i=1}^{t} C_i \sum_{j=1}^{i} \frac{\partial}{\partial\theta} \ln(q_\theta(\zeta_j \mid X_{j-1})).$$

Random Horizon.

Replace t by a (random) stopping time τ.

View ω as $\omega = (\zeta_1, \ldots, \zeta_\tau)$.

Now, for LR/SF, A1 is (generally) more difficult to verify, because the likelihood ratio is now the product of a random number of terms.

See Reiman and Weiss (1989) and L'Ecuyer (1994).

GSMP and Other Frameworks.

Generalization to semi-markovian or Markov renewal models.

Generalized Semi-Markov Process (GSMP), with denumerable state space.

See Glynn (1989), Glasserman (1991).

Steady-State Average Cost

Suppose Markov chain model is regenerative.

Let τ = cycle length; h_τ = total cost during cycle.

$$u(\theta) = E_\theta[h_\tau], \qquad \ell(\theta) = E_\theta[\tau], \qquad \alpha(\theta) = u(\theta)/\ell(\theta).$$

Assume $E_\theta[h_\tau^2 + \tau^2] < \infty$. From renewal-reward theorem:

$$\lim_{t\to\infty} \frac{1}{t} \sum_{i=1}^{t} C_i \stackrel{\text{a.s.}}{=} \alpha(\theta) = \lim_{t\to\infty} \frac{1}{t} \sum_{i=1}^{t} E_\theta[C_i].$$

Derivative of a Ratio of Expectations:

$$\alpha'(\theta) = \frac{u'(\theta)\ell(\theta) - \ell'(\theta)u(\theta)}{\ell^2(\theta)} = \frac{u'(\theta) - \alpha(\theta)\ell'(\theta)}{\ell(\theta)}.$$

Simulate n regenerative cycles of that process.

Let Y_j and T_j be values of h_τ and τ for cycle j;

Y_j' and T_j' estimators of $u'(\theta)$ and $\ell'(\theta)$ from cycle j.

(Can be obtained, for example, by LR/SF method.)

Consistent estimators for $\alpha(\theta)$ and $\alpha'(\theta)$ are

$$\begin{aligned}
\hat\alpha_n(\theta) &= \frac{\sum_{j=1}^n Y_j}{\sum_{j=1}^n T_j} \\
\hat\alpha_n'(\theta) &= \frac{\left(\sum_{j=1}^n Y_j'\right) - \hat\alpha_n(\theta)\sum_{j=1}^n T_j'}{\sum_{j=1}^n T_j}.
\end{aligned}$$

Note: with IPA, τ' is typically 0; not a valid estimator of $\ell'(\theta)$.

However, $\hat\alpha_n'(\theta)$ can nevertheless be a consistent estimator of $\alpha'(\theta)$.

Growing Horizon Estimators.

Simulate for long (fixed) horizon t, e.g.,

$$\hat{\alpha}'_t(\theta) = \begin{cases} \dfrac{1}{t} \sum\limits_{i=1}^{t} C'_i & \text{(IPA);} \\ \dfrac{1}{t} \sum\limits_{i=1}^{t} C_i S_i(\theta, \omega) & \text{(LR/SF);} \end{cases}$$

Bias is typically $O(1/t)$.

Repeat that n times for $t = t_n$. Computer budget $T = nt_n$.

Optimal MSE and rate of increase of t_n:

	c_n	t_n	MSE
FDI (one-sided)	$n^{-1/3}$	$n^{1/3}$	$T^{-1/2}$
FDI (centered)	$n^{-1/4}$	$n^{1/2}$	$T^{-2/3}$
FDC-I (one-sided)	$n^{-1/2}$	$n^{1/2}$	$T^{-2/3}$
FDC-I (centered)	$n^{-1/3}$	$n^{2/3}$	$T^{-4/5}$
FDC-II (one-sided)	$n^{-1/2}$	n	T^{-1}
FDC-II (centered)	$n^{-1/4}$	n	T^{-1}
IPA		n	T^{-1}
LR/SF		$n^{1/3}$	$T^{-1/2}$
LR/CSF		$n^{1/2}$	$T^{-2/3}$

$t_n = n^p$; $c_n = n^{-\gamma}$; T_n = computer budget (nb. customers).

	p	γ	$T_n = 10^1$	$T_n = 10^2$	$T_n = 10^3$	$T_n = 10^4$	$T_n = 10^5$
FDc	1/2	1/4	2.10	7.07E-1	2.42E-1	5.58E-2	2.25E-2
FDCc	2/3	1/3	9.36E-1	3.10E-1	7.08E-2	1.63E-2	4.05E-3
FDCc	1	1	7.62E-1	1.57E-1	2.25E-2	2.47E-3	2.44E-4
FDCc	2	2	7.60E-1	1.48E-1	1.95E-2	2.08E-3	1.92E-4
IPA	1		5.86E-1	1.39E-1	1.97E-2	2.12E-3	1.47E-4
IPA	2		6.86E-1	1.45E-1	1.93E-2	2.12E-3	1.32E-4
LR	1/3		3.08	8.43E-1	2.25E-1	7.26E-2	1.99E-2
LRD	1/3		2.64	6.88E-1	1.69E-1	5.34E-2	1.59E-2
CLR	1/2		1.90	5.54E-1	1.23E-1	2.99E-2	4.42E-3
CLRD	1/2		1.77	4.56E-1	9.48E-2	1.80E-2	2.85E-3
CLRD	1		2.71	7.95E-1	1.61E-1	2.78E-2	1.19E-2

	p	γ	$T_n = 10^1$	$T_n = 10^2$	$T_n = 10^3$	$T_n = 10^4$	$T_n = 10^5$
FDc	1/2	1/4	21.78	11.34	5.68	1.88	7.17E-1
FDCc	2/3	1/3	14.31	7.48	2.54	6.13E-1	1.23E-1
FDCc	1	1	9.80	3.26	6.48E-1	7.01E-2	7.38E-3
FDCc	2	2	7.27	1.99	4.08E-1	5.55E-2	4.84E-3
IPA	1		6.71	2.68	5.70E-1	7.20E-2	5.10E-3
IPA	2		5.58	1.85	4.15E-1	4.81E-2	6.88E-3
LR	1/3		10.58	7.05	4.15	2.46	1.13
LRD	1/3		10.10	6.86	4.09	2.44	1.11
CLR	1/2		8.84	5.19	2.53	9.34E-1	2.33E-1
CLRD	1/2		8.72	5.08	2.49	8.56E-1	2.14E-1
CLRD	1		9.02	4.41	1.62	3.28E-1	5.99E-2

Sample MSE, deriv. of mean sojourn time, M/M/1 queue.

	p	γ	$T_n = 10^1$	$T_n = 10^2$	$T_n = 10^3$	$T_n = 10^4$	$T_n = 10^5$
FDc	1/2	1/4	33.97	12.95	4.59	1.03	2.33E-1
FDCc	2/3	1/3	3.45	8.97E-1	1.64E-1	2.93E-2	3.85E-3
FDCc	1	1	8.68	3.54	1.18	4.14E-1	2.21E-1
FDCc	2	2	14.17	7.02	3.63	1.37	7.93E-1
LR	1/3		9.44E-1	2.51E-1	6.08E-2	1.87E-2	5.56E-3
LRD	1/3		9.09E-1	2.36E-1	5.61E-2	1.79E-2	5.10E-3
CLR	1/2		9.81E-1	2.39E-1	5.00E-2	1.05E-2	1.90E-3
CLRD	1/2		9.63E-1	2.15E-1	4.17E-2	6.99E-3	1.58E-3
CLRD	1		1.39	3.30E-1	5.83E-2	1.15E-2	2.75E-3

Table 1: Sample MSE, derivative of $P(\text{sojourn time} > 1)$, M/M/1 queue, $\rho = 1/4$.

Average Nb. customers in an open queueing network.

Open network of Q single-server queues (nodes).

Vector of parameters: $\theta = (\theta_1, \ldots, \theta_Q)$. Node q:

Service times $S_{q,j} : G_{\theta_q}$;

(External) interarrival times follow : F_q;

When leaving node i, go to node j with prob. p_{ij}

and leaves system with prob. $p_{i0} = 1 - \Sigma_{j=1}^{Q} p_{ij}$.

Assume system stable, regenerative.

$N_q(t) =$ nb of customers at node q at time t.

$t_{q,j} =$ time of the j-th event (arrival or departure) at node q.

$n_{q,j} =$ nb of customers at node q during $[t_{q,j-1},\ t_{q,j})$.

$M_q(t) = \max\{j \mid t_{q,j} \le t\} =$ nb events at node q during $[0, t]$.

Cost rate ν_q for each customer at node q.

Total cost during $[0, t]$:

$$\begin{aligned}
h_t(\theta, \omega) &= \int_0^t \left(\sum_{q=1}^{Q} \nu_q N_q(s) \right) ds \\
&= \sum_{q=1}^{Q} \nu_q \left(\sum_{j=1}^{M_q(t)} n_{q,j}(t_{q,j} - t_{q,j-1}) + N_q(t)(t - t_{q,M_q(t)}) \right) \\
&= \sum_{q=1}^{Q} \nu_q \left(\sum_{j=1}^{M_q(t)} t_{q,j}(n_{q,j} - n_{q,j+1}) + t N_q(t) \right).
\end{aligned}$$

Derivative:

$$h_t'(\theta, \omega) = \sum_{q=1}^{Q} \nu_q \sum_{j=1}^{M_q(t)} t_{q,j}'(n_{q,j} - n_{q,j+1}).$$

Now need derivative of $t_{q,j}$ with respect to each θ_j.

Event at $t_{q,j}$ may be

— external arrival: $t'_{q,j} = 0$;

— arrival from another queue (say r): $t'_{q,j} = t'_{r,l}$ if $t_{q,j} = t_{r,l}$;

— service completion at q: if service started at $t_{q,i}$ and lasts $S_{q,j}$, then $t'_{q,j} = t'_{q,i} + S'_{q,j}$.

Numerical illustration: Jackson network, $Q = 4$.

External arrival rates: $(0.3,\ 0.3,\ 0.3,\ 0.1)$.

Mean service times: $\theta_q = 1$ for all q.

Transition probabilities:

$$\begin{pmatrix} 0.7 & 0.1 & 0.1 & 0.1 & 0.0 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.6 & 0.1 & 0.1 & 0.1 & 0.1 \\ 0.4 & 0.2 & 0.1 & 0.2 & 0.1 \end{pmatrix}.$$

Total arrival rates and traffic intensities: $(0.590,\ 0.560,\ 0.590,\ 0.301)$.

Exact steady-state average cost: $\alpha(\theta) = 4.587$.

Gradient: $(3.518,\ 2.897,\ 3.518,\ 0.617)$.

Budget: $T_n = n^2$ cust.; n runs of length $t_n = n$. For FD: $c_n = n^{-1}$.

Each experiment was replicated r times to compute MSE.

FDC: for synchronization, one RNG for service times for each node, plus one RNG for interarrivals. Used CRN across table lines.

Typical results:

— Faster convergence when ρ and true deriv. are lower (see node 4).

— Convergence does not depend on number of nodes in network.

— Of course, for FDC, synchronization is crucial.