

Gradient Estimation Methods for Discrete Event Systems Performance Measures

Felisa J. Vázquez-Abad

Département d'informatique et recherche opérationnelle
Université de Montréal, Québec, CANADA
courriel: vazquez@IRO.UMontreal.CA

Department of Electronic and Electrical Engineering
The University of Melbourne

5th WODES, Ghent August 21 and 22, 2000.

Course Objectives

Upon taking this course, you will:

- **understand** the problem of gradient estimation for stochastic dynamical systems driven by discrete events
- **study** the basic concepts required to read most books available in the subject (Cassandras and Lafortune '99, Fu and Hu '97, Glasserman '91, Ho and Cao '91, Pflug '96, Rubinstein '9, etc.)
- **learn** about the popular estimators for sensitivity analysis:
 - Infinitesimal Perturbation Analysis
 - Smoothed Perturbation Analysis, and beyond...
 - Score Function Method
 - Weak Derivatives, and beyond...
- **distinguish** between the **methods** and the **estimators**
- **acquire** the methodology to construct estimators for specific applications following our methodological approach

Course Outline

PART I

- Introduction and Motivation
- Review of Probability Concepts
- The Pathwise *vs* the Distributional Approaches: overview

PART II

- Main Results: Pathwise Analysis (Lipschitz continuity and conditional MC)
- Main Results: Distributional Analysis (weak differentiation of measures)
- Examples of Application

Gradient Estimation Problem

Model: Complex stochastic systems, **control variable** $\theta \in \mathbb{R}^s$.

- **Physical process:** Stochastic process $\{X_n(\theta)\}$ on $(\Omega, \mathfrak{F}, \mathbb{P})$.
- **Sample performance** : $L(X_1(\theta, \omega), \dots, X_n(\theta, \omega)) = h(\theta, \omega)$. A real-valued functional of the trajectories of the process.
- **Performance function** : $J(\theta) = \mathbb{E}[h(\theta)]$

Gradient Estimation Problem

Model: Complex stochastic systems, **control variable** $\theta \in \mathbb{R}^s$.

- **Physical process:** Stochastic process $\{X_n(\theta)\}$ on $(\Omega, \mathfrak{F}, \mathbb{P})$.
- **Sample performance** : $L(X_1(\theta, \omega), \dots, X_n(\theta, \omega)) = h(\theta, \omega)$. A real-valued functional of the trajectories of the process.
- **Performance function** : $J(\theta) = \mathbb{E}[h(\theta)]$

Problem: Performance function not available in closed analytical form, but can be **estimated** : $U_1, \dots, U_M \sim \mathbb{P}$ independent trajectories, use sample average.

Gradient Estimation Problem

Model: Complex stochastic systems, **control variable** $\theta \in \mathbb{R}^s$.

- **Physical process:** Stochastic process $\{X_n(\theta)\}$ on $(\Omega, \mathfrak{F}, \mathbb{P})$.
- **Sample performance** : $L(X_1(\theta, \omega), \dots, X_n(\theta, \omega)) = h(\theta, \omega)$. A real-valued functional of the trajectories of the process.
- **Performance function** :

$$J(\theta) = \mathbb{E}[h(\theta)]$$

Problem: Performance function not available in closed analytical form, but can be **estimated** : $U_1, \dots, U_M \sim \mathbb{P}$ independent trajectories, use sample average.

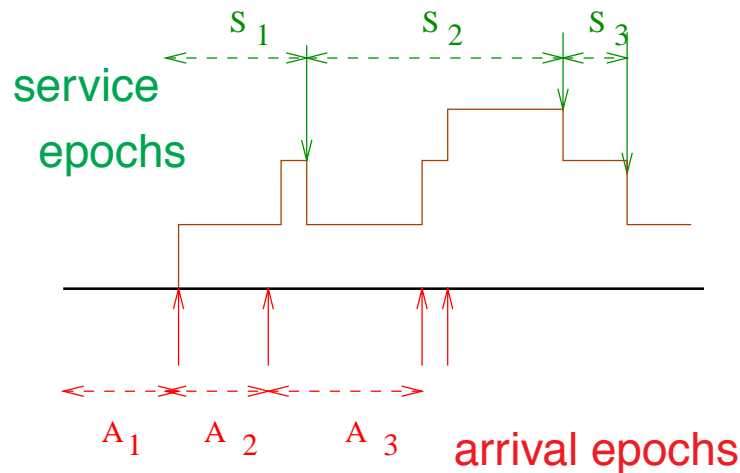
Goal: Estimate also the **sensitivity** of the performance function w.r.t θ : Construct a functional $G \in \mathbb{R}^s$ of a **derivative process** $\{Z_n(\theta)\}$, adapted to the natural filtration and such that

$$\nabla_{\theta} J(\theta) = \mathbb{E}[G(Z_1(\theta), \dots, Z_n(\theta))]$$

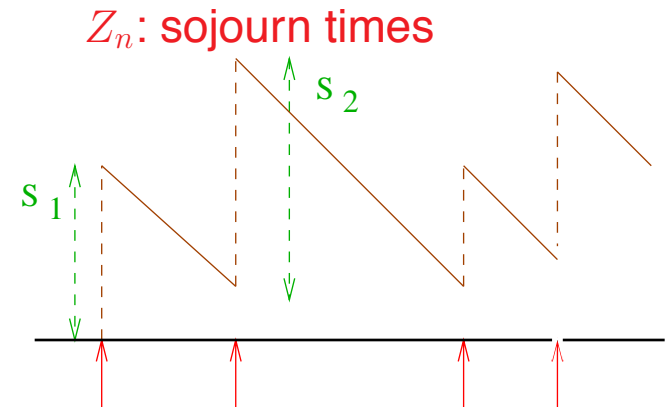
Applications: Queueing Systems

- **Physical state:** queue lengths $X_t \in \mathbb{R}^s$
- **Residual clocks:** each queue has residual clocks $Y_t \in \mathbb{R}^d$.
- **Events:** arrivals and services (polling, etc).

Occupancy process: X_t



Storage Process:

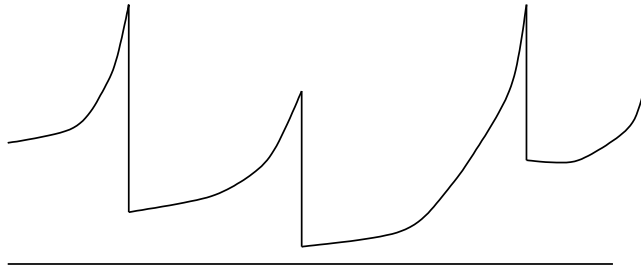


$$Z_{n+1} = (Z_n - A_n)^+ + S_n$$

- **Performance:** Average sojourn times, loss probabilities, etc.
- **Control variables of interest:** arrival rates, service parameters, buffer sizes, etc.

Applications: Risk Theory

U_t : surplus of insurance portfolio at time t

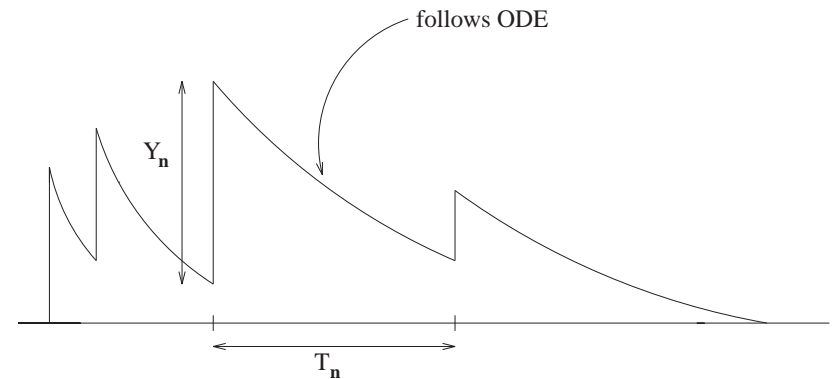


Surplus process:

$$dU(t) = [c + \delta U(t)]dt - dS(t)$$

δ : interest rate

- **Performance:** ruin probability, expected amount of ruin, etc.
- **Control variables of interest:** arrival rates, expected claim amount, premium c , etc.



Use a **dual process**: X_t a storage process (queue):

$$X_n = \left(X_{n-1} e^{-\delta A_n} - \frac{c}{\delta} (1 - e^{-\delta A_n}) \right)^+ + Y_n$$

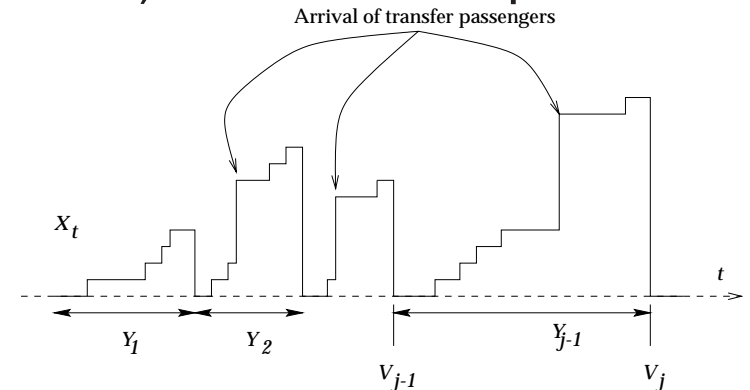
Applications: Transportation

Platform occupancy $\{X_t\}$

Residual clocks: train and passenger arrivals, train departure at platform.

Possible events: Arrivals (per source) and train departure.

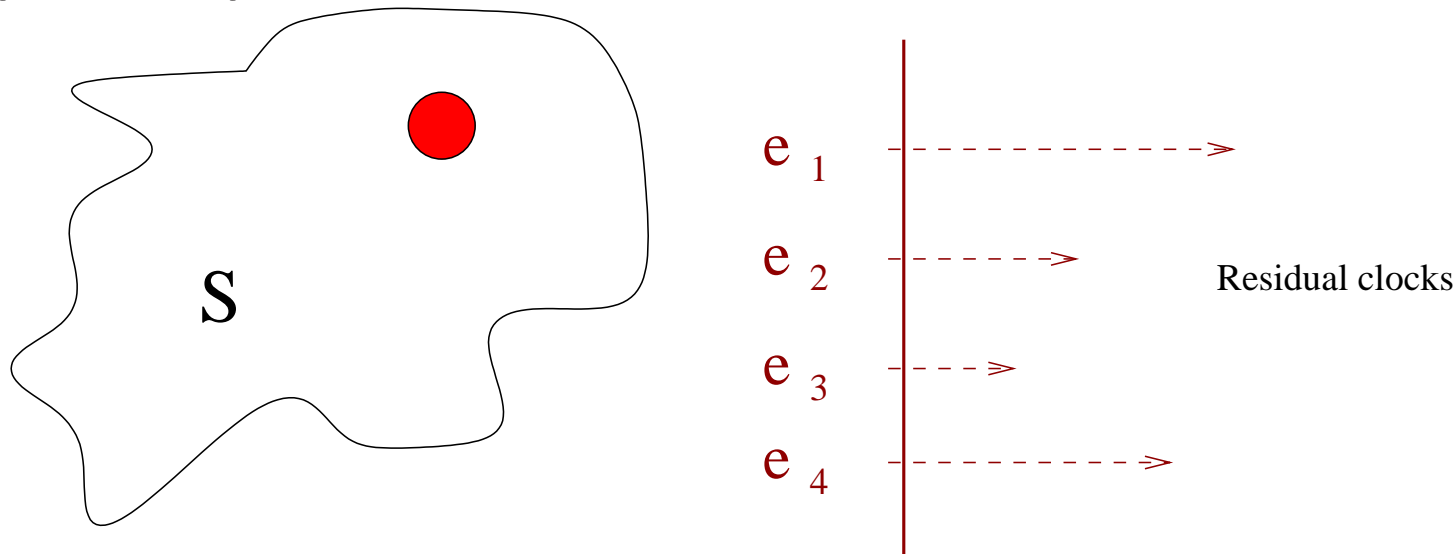
- Performance: waiting times, costs.
- Control variables of interest: line frequency, arrival rates, etc.



Process Description

- **Physical state space** $S \subset \mathbb{R}^N$, d **event-timing clocks** $\mathcal{T} = (\mathbb{R}^+)^d$
- **Event space** E , possible **events** at state $x \in S$ are $\Gamma(x) \subset E$.
- **Next event** $e(x, y) = \arg \min\{y_i : i \in \Gamma(x)\}$, $(x, y) \in S \times \mathcal{T}$.
- **Elapsed time:** $\tau(x, y) = Y_{e(x,y),t}$.

System's dynamics:

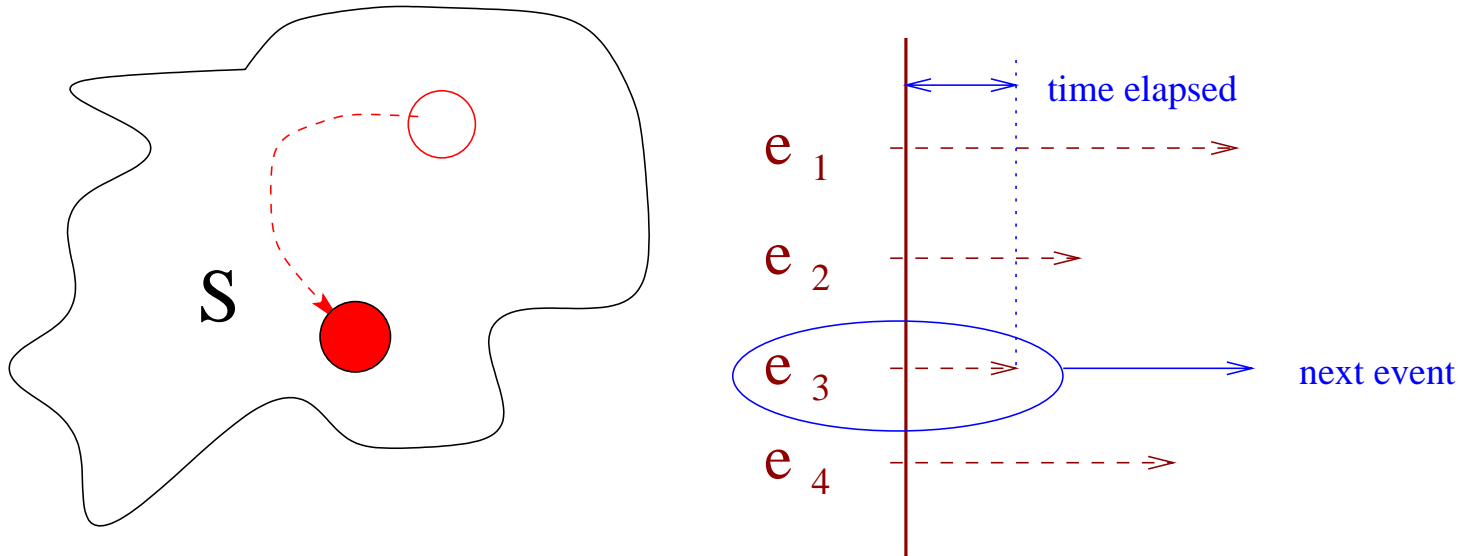


$\{X_{t+s}; 0 \leq s < \tau\}$ follows **deterministic dynamics** (ODE).

Process Description

- **Physical state space** $S \subset \mathbb{R}^N$, d **event-timing clocks** $\mathcal{T} = (\mathbb{R}^+)^d$
- **Event space** E , possible **events** at state $x \in S$ are $\Gamma(x) \subset E$.
- **Next event** $e(x, y) = \arg \min\{y_i : i \in \Gamma(x)\}$, $(x, y) \in S \times \mathcal{T}$.
- **Elapsed time:** $\tau(x, y) = Y_{e(x,y),t}$.

System's dynamics:



$\{X_{t+s}; 0 \leq s < \tau\}$ follows **deterministic dynamics** (ODE).

Process Description

- **Physical state space** $S \subset \mathbb{R}^N$, d **event-timing clocks** $\mathcal{T} = (\mathbb{R}^+)^d$
- **Event space** E , possible **events** at state $x \in S$ are $\Gamma(x) \subset E$.
- **Next event** $e(x, y) = \arg \min\{y_i : i \in \Gamma(x)\}$, $(x, y) \in S \times \mathcal{T}$.
- **Elapsed time:** $\tau(x, y) = Y_{e(x,y),t}$.

System's dynamics:

$$\text{Prob}[X_{t+\tau} \in A \mid X_t = x, Y_t = y] = p_\theta(A; x, e(x, y)),$$

$$\text{Prob}[Y_{t+\tau}(e) \in B \mid X_t = x, e(x, y) = e] = \mathbb{G}_{e,\theta}(B, x)$$

Result: Define $Z_t(\theta) = (X_t(\theta), Y_t(\theta))$ on $S \times \mathcal{T}$. The **embedded process** $\{Z_n(\theta)\}$ observed only at event epochs is a Markov chain in a general state space.

\Rightarrow
 $\{Z_n(\theta)\}$ is a **Generalized Semi-Markov Process**.

A Matter of Representation

- Process $\{Z_n\}$ in a state space $S \subset \mathbb{R}^d$.
- Sample performance $L(Z_1, Z_2, \dots)$
- Probabilistic description: $Z_n : \Omega \rightarrow \mathbb{R}^d$ is a mapping for each n .
- Sample performance is a function: $h(\theta, \omega) = L(Z_1(\omega), Z_2(\omega), \dots)$: a **random variable** .
- Performance measure is an expectation, $E[L(Z)] = \int L(z)P(dz)$.

$$J(\theta) = E_{\theta}[L(Z_1, Z_2, \dots)] = E[h(\theta)]$$

How does the performance measure depend on θ ?? Is it through the sample process, or through the distribution??

Mathematical formulation: **probability theory**

Probability Theory: a review

Information Structure

- A set Ω of **states of nature**
- A set of subsets of Ω , a σ -field \mathfrak{F} that **describes** all the possible **observables** of the model (events)

Probability Theory: a review

Information Structure

- A set Ω of **states of nature**
- A set of subsets of Ω , a σ -field \mathfrak{F} that **describes** all the possible **observables** of the model (events)

EXAMPLE 1.

Description of whether, uncertainty. Ω are all the states of nature: climatological conditions.

$\mathfrak{F} = \{\emptyset, \{\omega_1\}, \{\omega_2\}, \{\omega_1, \omega_2\}, \{\omega_1, \omega_2\}^c, \Omega\}$ is used to describe the two observables ω_1 : **rainy day** , and ω_2 : **sunny day** .

\mathfrak{F} describes the information structure that these two observables provide.

If we want to describe also observables as **sunny in the morning and rain in the afternoon** , \mathfrak{F} is not sufficient.

\implies

Information structure (Ω, \mathfrak{F}) depends on **model** .

Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$.

Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$.

Definition: A **random variable** on (Ω, \mathfrak{F}) is a function $X : \Omega \rightarrow \mathbb{R}$ that is **measurable** with respect to \mathfrak{F} , that is,

$$\{\omega \in \Omega : X(\omega) \leq x\} \in \mathfrak{F}, \quad \forall x \in \mathbb{R}$$

EXAMPLE 2. $X(\omega_i) = i$ enumerates the possible observables:

$$\begin{aligned} \omega_1 &= \text{Rainy day}, & \omega_3 &= \text{Cloudy in the morning}, \\ \omega_2 &= \text{Sunny day}, & \omega_4 &= \text{Rain in the afternoon}, \end{aligned}$$

then X is not a r.v. There is an information structure $\sigma(X)$ **generated by the mapping X** , it contains all the observables of Ω that are described by X .

Any mapping $X : \Omega \rightarrow \mathbb{R}$ is a r.v. on $(\Omega, \sigma(X))$



Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$. Information structures $\sigma(X) \rightarrow \mathcal{B}(\mathbb{R})$.
- Probability Measures

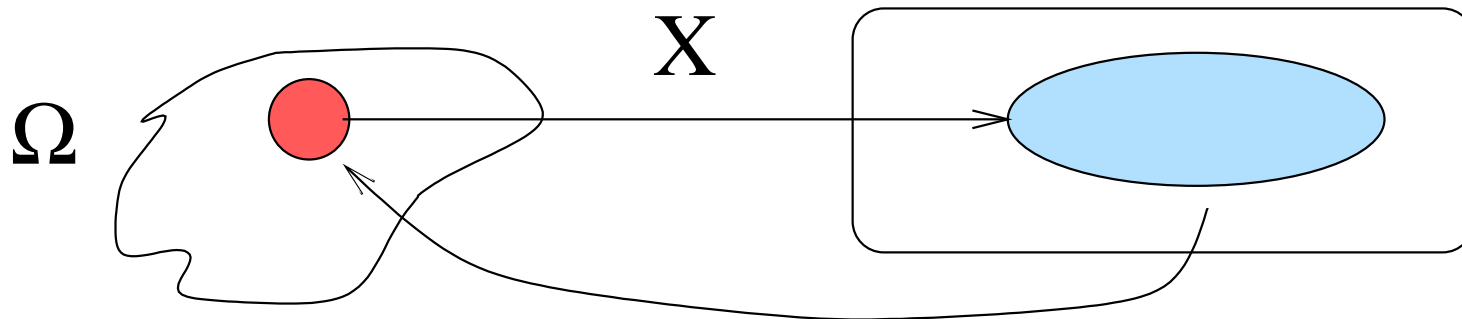
Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$. Information structures $\sigma(X) \rightarrow \mathcal{B}(\mathbb{R})$.
- Probability Measures

Definition: A **probability measure** \mathbb{P} on (Ω, \mathfrak{F}) is a function $\mathbb{P} : \mathfrak{F} \rightarrow [0, 1]$

$$\mathbb{P}(\emptyset) = 0, \quad \mathbb{P}(\Omega) = 1, \quad \mathbb{P} \left(\bigcap_{n=1}^{\infty} A_n \right) = \sum_{n=1}^{\infty} \mathbb{P}(A_n).$$



Probabilities : inverse mapping

State space $S \subset \mathbb{R}$. Induced probability measure \mathbb{P}_X on $(S, \mathcal{B}(S))$.

Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$. Information structures $\sigma(X) \rightarrow \mathcal{B}(\mathbb{R})$.
- Probability Measures: \mathbb{P} on (Ω, \mathfrak{F}) and P_X on $(S, \mathcal{B}(S))$
- Expectations

Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$. Information structures $\sigma(X) \rightarrow \mathcal{B}(\mathbb{R})$.
- Probability Measures: \mathbb{P} on (Ω, \mathfrak{F}) and P_X on $(S, \mathcal{B}(S))$
- Expectations

Definition: $F_X(x) = P_X[X \leq x] = \mathbb{P}\{\omega \in \Omega : X(\omega) \leq x\}$ is the **disribution function** of a random variable X . For any function F os X , the **expectation** of $F(X)$ is defined as

$$\mathbb{E}[F(X)] = \int_S F(x) P_X(dx) = \int_{\Omega} F[X(\omega)] \mathbb{P}(d\omega)$$

Probability Theory: a review

Information Structure and Random Processes

- A set Ω of **states of nature**
- A σ -field \mathfrak{F} that **describes** all the possible **observables** of the model.
- Random variables $X : \Omega \rightarrow \mathbb{R}$. Information structures $\sigma(X) \rightarrow \mathcal{B}(\mathbb{R})$.
- Probability Measures: \mathbb{P} on (Ω, \mathfrak{F}) and P_X on $(S, \mathcal{B}(S))$
- Expectations
- Random Processes $\{Z_n\}$ and performance measures

Definition: A **random process** on (Ω, \mathfrak{F}) is a sequence X_n such that for each $n \in \mathbb{N}$, $X_n : \Omega \rightarrow \mathbb{R}$, and $\mathfrak{F}_n \equiv \sigma(Z_i; i \leq n) \subset \mathfrak{F}$.

A **filtration** is a sequence of increasing σ -fields on Ω . In particular, $\mathbb{F} = (\mathfrak{F}_n; n \in \mathbb{N})$ is called the **natural filtration** of the process.

A **functional** $L(Z_1, \dots, Z_n)$ is also a **random variable** h on (Ω, \mathfrak{F}_n) . The performance measure is then: $J(\theta) = \mathbf{E}_\theta[L(Z_1, Z_2, \dots)] = \mathbf{E}[h(\theta)]$

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on** θ :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.

What is Ω ????

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on θ** :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.

What is Ω ????

EXAMPLE 3.

Model statement: “ X is a random variable with exponential distribution of intensity θ ”.

Sufficient to describe a random variable: $F(x) = 1 - e^{-\theta x}$.

... the actual set Ω becomes irrelevant.

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on** θ :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.

Theorem 1 (Skorohod Representation) Given a mapping $X : \Omega \rightarrow S$ on $(\Omega, \mathfrak{F}, \mathbb{P})$, there exists $\tilde{X} : [0, 1] \rightarrow S$ on $([0, 1], \mathcal{B}([0, 1]), \tilde{\mathbb{P}})$, where $\tilde{\mathbb{P}}$ is the Lebesgue measure (or uniform distribution on $[0, 1]$), such that $X \stackrel{\mathcal{L}}{=} \tilde{X}$ (equality in distribution).

EXAMPLE 4. Generation of a r.v. $X \sim F$: $X(U) = F^{-1}(U)$, $U \sim U[0, 1]$.
Exponential distribution: $X_1(U) = -\frac{1}{\theta} \ln(1 - U)$, $X_2(U) = -\frac{1}{\theta} \ln(U)$. Both on $([0, 1], \mathcal{B}([0, 1]), \tilde{\mathbb{P}})$, with distribution $F(x) = 1 - e^{-\theta x}$ but $X_1 \neq X_2$ a.s. ***

The representation is not unique.

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on** θ :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.
- Probabilistic description $n \in \mathbb{N}$, $Z_n(\theta) : \Omega \rightarrow \mathbb{R}^d$
- Sample performance $h(\theta, \omega) = L(Z_1(\theta, \omega), Z_2(\theta, \omega), \dots)$: a r.v.
- Performance measure
 $J(\theta) = \int h(\theta, \omega)\mathbb{P}(d\omega)$.

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on θ** :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.
- Probabilistic description $n \in \mathbb{N}$, $Z_n(\theta) : \Omega \rightarrow \mathbb{R}^d$
- Sample performance $h(\theta, \omega) = L(Z_1(\theta, \omega), Z_2(\theta, \omega), \dots)$: a r.v.
- Performance measure
 $J(\theta) = \int h(\theta, \omega)\mathbb{P}(d\omega)$.
- Use **sample path approach** via the pathwise analysis to obtain:

$$\frac{dJ(\theta)}{d\theta} = \int \frac{d}{d\theta} h(\theta, \omega)\mathbb{P}(d\omega)$$

- **Methods:** IPA + SPA, etc.

A Matter of Representation

- Markov process $\{Z_n\}$ in $S \subset \mathbb{R}^d$.
Transition kernels **depend on θ** :
 $P_\theta[Z_{n+1} \in A \mid Z_n = z]$
- Sample performance $L(Z)$
- Performance measure
 $J(\theta) = \int L(z)P_\theta(dz)$.
- Use **distributional approach** via the weak topology to obtain:
$$\frac{dJ(\theta)}{d\theta} = \int L(z) \frac{d}{d\theta} P_\theta(dz)$$
- **Methods:** score function, weak derivatives, etc.
- Probabilistic description $n \in \mathbb{N}$, $Z_n(\theta) : \Omega \rightarrow \mathbb{R}^d$
- Sample performance $h(\theta, \omega) = L(Z_1(\theta, \omega), Z_2(\theta, \omega), \dots)$: a r.v.
- Performance measure
 $J(\theta) = \int h(\theta, \omega) \mathbb{P}(d\omega)$.
- Use **sample path approach** via the pathwise analysis to obtain:
$$\frac{dJ(\theta)}{d\theta} = \int \frac{d}{d\theta} h(\theta, \omega) \mathbb{P}(d\omega)$$
- **Methods:** IPA + SPA, etc.

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate** θ

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \quad \{A_n\} \sim \exp(\theta), \{S_n\} \text{ iid} \sim G$$

What Ω ?

IPA: Canonical representation, inverse function for each interarrival

$$A_i = -\frac{1}{\theta} \ln(U_i), \{U_i\} \text{ iid} \sim U[0, 1]$$

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate θ**

•

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \text{ use telescopic sum: } X_{n+1} = \sum_{i=\alpha_k}^k S_i - \sum_{i=\alpha_k}^{k-1} A_i$$

$$L(X_1, \dots, X_n) = \frac{1}{n} \sum_{k=1}^n X_k \Rightarrow h(\theta) = \frac{1}{n} \sum_{k=1}^n \left[\sum_{i=\alpha_k}^k S_i + \frac{1}{\theta} \sum_{i=\alpha_k}^{k-1} \ln(U_i) \right]$$

$$h'(\theta) = -\frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} \frac{1}{\theta} \ln(U_i) = \frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} A_i$$

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate θ** .

.

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \text{ use telescopic sum: } X_{n+1} = \sum_{i=\alpha_k}^k S_i - \sum_{i=\alpha_k}^{k-1} A_i$$

$$L(X_1, \dots, X_n) = \frac{1}{n} \sum_{k=1}^n X_k \Rightarrow h(\theta) = \frac{1}{n} \sum_{k=1}^n \left[\sum_{i=\alpha_k}^k S_i + \frac{1}{\theta} \sum_{i=\alpha_k}^{k-1} \ln(U_i) \right]$$

$$h'(\theta) = -\frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} \frac{1}{\theta} \ln(U_i) = \frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} A_i$$

Estimation: The IPA formula for $h'(\theta)$ gives an unbiased estimator for $J'(\theta)$. Calculation in parallel as the system evolves (adapted to the filtration).

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate θ** .

.

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \text{ use telescopic sum: } X_{n+1} = \sum_{i=\alpha_k}^k S_i - \sum_{i=\alpha_k}^{k-1} A_i$$

$$L(X_1, \dots, X_n) = \frac{1}{n} \sum_{k=1}^n X_k \Rightarrow h(\theta) = \frac{1}{n} \sum_{k=1}^n \left[\sum_{i=\alpha_k}^k S_i + \frac{1}{\theta} \sum_{i=\alpha_k}^{k-1} \ln(U_i) \right]$$

$$h'(\theta) = -\frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} \frac{1}{\theta} \ln(U_i) = \frac{1}{n\theta} \sum_{k=1}^n \sum_{i=\alpha_k}^{k-1} A_i$$

Estimation: The IPA formula for $h'(\theta)$ gives an unbiased estimator for $J'(\theta)$. Calculation in parallel as the system evolves (adapted to the filtration).

How do we know it is unbiased?

Pathwise Analysis: Main Results

Unbiasedness of IPA estimator:

$$\lim_{\Delta\theta \rightarrow 0} \frac{\frac{d}{d\theta} \int_{\Omega} h(\theta, \omega) \mathbb{P}(d\omega) \mathbf{E}[h(\theta + \Delta\theta)] - \mathbf{E}[h(\theta)]}{\theta} \stackrel{?}{=} \int_{\Omega} h'(\theta, \omega) \mathbb{P}(d\omega) \stackrel{?}{=} \mathbf{E} \left[\lim_{\Delta\theta \rightarrow 0} \frac{h(\theta + \Delta\theta) - h(\theta)}{\theta} \right]$$

Pathwise Analysis: Main Results

Unbiasedness of IPA estimator:

$$\lim_{\Delta\theta \rightarrow 0} \frac{\frac{d}{d\theta} \int_{\Omega} h(\theta, \omega) \mathbb{P}(d\omega) \quad \stackrel{?}{=} \quad \int_{\Omega} h'(\theta, \omega) \mathbb{P}(d\omega)}{\mathbb{E}[h(\theta + \Delta\theta)] - \mathbb{E}[h(\theta)]} \quad \stackrel{?}{=} \quad \mathbb{E} \left[\lim_{\Delta\theta \rightarrow 0} \frac{h(\theta + \Delta\theta) - h(\theta)}{\Delta\theta} \right]$$

EXAMPLE 5. Take $\Omega = [0, 1]$, $\mathfrak{F} = \mathcal{B}(\Omega)$, \mathbb{P} the uniform measure.

$X(\omega) \equiv 0$, and

$$X_n(\omega) = \begin{cases} n & \text{if } \omega < 1/n \\ 0 & \text{otherwise} \end{cases}$$

Then the **pathwise limit** is $\lim_{n \rightarrow \infty} X_n(\omega) = X(\omega)$ for every $\omega > 0$, and $\mathbb{P}[\omega = 0] = 0$, so $X_n \rightarrow X$ **almost surely**.

Alas! $\mathbb{E}[X_n] = 1$ for all $n \geq 1$, so $\mathbb{E}[X_n] \not\rightarrow \mathbb{E}[X]$.

Pathwise Analysis: Main Results

Unbiasedness of IPA estimator:

$$\begin{aligned} \frac{d}{d\theta} \int_{\Omega} h(\theta, \omega) \mathbb{P}(d\omega) &\stackrel{?}{=} \int_{\Omega} h'(\theta, \omega) \mathbb{P}(d\omega) \\ \lim_{\Delta\theta \rightarrow 0} \frac{\mathbb{E}[h(\theta + \Delta\theta)] - \mathbb{E}[h(\theta)]}{\Delta\theta} &\stackrel{?}{=} \mathbb{E} \left[\lim_{\Delta\theta \rightarrow 0} \frac{h(\theta + \Delta\theta) - h(\theta)}{\Delta\theta} \right] \end{aligned}$$

EXAMPLE 6. Take $\Omega = [0, 1]$, $\mathfrak{F} = \mathcal{B}(\Omega)$, \mathbb{P} the uniform measure.

$X(\omega) \equiv 0$, and

$$X_n(\omega) = \begin{cases} n & \text{if } \omega < 1/n \\ 0 & \text{otherwise} \end{cases}$$

Then the **pathwise limit** is $\lim_{n \rightarrow \infty} X_n(\omega) = X(\omega)$ for every $\omega > 0$, and $\mathbb{P}[\omega = 0] = 0$, so $X_n \rightarrow X$ **almost surely**.

Alas! $\mathbb{E}[X_n] = 1$ for all $n \geq 1$, so $\mathbb{E}[X_n] \not\rightarrow \mathbb{E}[X]$.

IPA uses the **pathwise limit** $h'(\theta, \omega) = \lim_{\Delta\theta \rightarrow 0} \frac{h(\theta + \Delta\theta) - h(\theta)}{\Delta\theta}$. When can the limit and the expectation can be interchanged?

Mathematical formulation and theory

Pathwise Analysis: Main Results

Theorem 2 (*Dominated Convergence*) Let $(X_n, n \in \mathbb{N})$, X be random variables on $(\Omega, \mathfrak{F}, \mathbb{P})$. If $X_n \rightarrow X$ a.s. and there is a random variable K such that $\mathbf{E}(K) < \infty$ and $|X_n| \leq K$ a.s. for all $n \in \mathbb{N}$, then $\lim_{n \rightarrow \infty} \mathbf{E}[X_n] = \mathbf{E}[X]$.

Pathwise Analysis: Main Results

Theorem 3 (*Dominated Convergence*) Let $(X_n, n \in \mathbb{N})$, X be random variables on $(\Omega, \mathfrak{F}, \mathbb{P})$. If $X_n \rightarrow X$ a.s. and there is a random variable K such that $\mathbf{E}(K) < \infty$ and $|X_n| \leq K$ a.s. for all $n \in \mathbb{N}$, then $\lim_{n \rightarrow \infty} \mathbf{E}[X_n] = \mathbf{E}[X]$.

Definition: A random variable $h(\theta, \omega)$ on $(\Omega, \mathfrak{F}, \mathbb{P})$ is said to be a.s. Lipschitz continuous in θ if there is a random variable $K < \infty$ with $\mathbf{E}[K] < \infty$ such that for each ω ,

$$\sup_{\theta \in \Theta: \theta + \Delta\theta \in \Theta} \|h(\theta + \Delta\theta, \omega) - h(\theta, \omega)\| \leq K(\omega) \Delta\theta,$$

Pathwise Analysis: Main Results

Theorem 4 (Dominated Convergence) Let $(X_n, n \in \mathbb{N})$, X be random variables on $(\Omega, \mathfrak{F}, \mathbb{P})$. If $X_n \rightarrow X$ a.s. and there is a random variable K such that $E(K) < \infty$ and $|X_n| \leq K$ a.s. for all $n \in \mathbb{N}$, then $\lim_{n \rightarrow \infty} E[X_n] = E[X]$.

Definition: A random variable $h(\theta, \omega)$ on $(\Omega, \mathfrak{F}, \mathbb{P})$ is said to be **a.s. Lipschitz continuous in θ** if there is a random variable $K < \infty$ with $E[K] < \infty$ such that for each ω ,

$$\sup_{\theta \in \Theta: \theta + \Delta\theta \in \Theta} \|h(\theta + \Delta\theta, \omega) - h(\theta, \omega)\| \leq K(\omega) \Delta\theta,$$

- Suppose $h(\theta)$ is a.s. Lipschitz continuous
- then finite difference is uniformly bounded a.s.
- Use Dominated Convergence Theorem
- The expectation of the limit is the limit of the expectation

If $h(\theta)$ is a.s. Lipschitz continuous in θ , then $E[h'(\theta)] = J'(\theta)$.

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate θ**

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \quad f_\theta(a_k) = \theta e^{-\theta a}, \quad \{S_n\} \text{ iid } \sim G \perp \theta$$

Performance measure:

$$J(\theta) = \int \left(\int L(a_1, \dots, a_{n-1}; s_1, \dots, s_n) \prod_{k=1}^n g(s_k) d\bar{s}_n \right) \prod_{k=1}^{n-1} f_\theta(a_k) da_1, \dots, da_{n-1}$$

$$\bar{L}(a_1, \dots, a_{n-1}) = \int L(a_1, \dots, a_{n-1}; s_1, \dots, s_n) \prod_{k=1}^n g(s_k) d\bar{s}_n$$

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate θ**

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \quad f_\theta(a_k) = \theta e^{-\theta a}, \quad \{S_n\} \text{ iid} \sim G \perp \theta$$

Performance measure:

$$\begin{aligned} J(\theta) &= \int \bar{L}(a_1, \dots, a_{n-1}) \prod_{k=1}^{n-1} f_\theta(a_k) da_1, \dots, da_{n-1} \\ &= \int \bar{L}(a_1, \dots, a_{n-1}) \prod_{k=1}^{n-1} \frac{f_\theta(a_k)}{f_1(a_k)} f_1(a_k) da_1, \dots, da_{n-1} \\ J'(\theta) &= \sum_{k=1}^n \int \bar{L}(a_1, \dots, a_{n-1}) f'_\theta(a_k) \prod_{i=1, i \neq k}^{n-1} \frac{f_\theta(a_i)}{f_1(a_i)} f_1(a_i) da_1, \dots, da_{n-1} \\ &= \int \bar{L}(a_1, \dots, a_{n-1}) \frac{d}{d\theta} \sum_{k=1}^n \ln f_\theta(a_k) \prod_{i=1}^{n-1} f_\theta(a_i) da_1, \dots, da_{n-1} \end{aligned}$$

Overview: examples

Queueing system: $M/G/1$: The performance is the **average sojourn time** of the first n customers. The control parameter is the **arrival rate** θ

$$X_{n+1} = (X_n - A_n)^+ + S_{n+1}, \quad f_\theta(a_k) = \theta e^{-\theta a}, \quad \{S_n\} \text{ iid} \sim G \perp \theta$$

Result:

$$J'(\theta) = \mathbf{E}[L(X_1, \dots, X_n) S(\theta; A_1, \dots, A_{n-1})]$$

Score Function:

$$\begin{aligned} S(\theta, a_1, \dots, a_n) &= \frac{d}{d\theta} \ln \left(\prod_{i=1}^n f_\theta(a_i) \right) \\ &= \frac{d}{d\theta} \sum_{i=1}^n [\ln \theta - \theta a_i] = \left(\frac{n}{\theta} - \sum_{i=1}^n a_i \right) \end{aligned}$$

This is true for general performance functions but ... **may not work for some distributions.**

Overview: Examples

A Threshold example: $X \sim \exp(\theta)$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$.

- $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.
- Score function $S(\theta, x) = 1/\theta - x$, use the estimator $L(X)S(\theta, X)$

Overview: Examples

A Threshold example: $X \sim \exp(\theta)$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$.

- $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.
- Score function $S(\theta, x) = 1/\theta - x$, use the estimator $L(X)S(\theta, X)$
- **Representations:** simplify the problem with $Y \sim \text{Bernoulli}(p(\theta))$,
 $p(\theta) = e^{-\theta \bar{x}}, Y = \mathbf{1}_{\{x \geq \bar{x}\}}$

How do we calculate the Score Function here???

Overview: Examples

A Threshold example: $X \sim \exp(\theta)$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$.

- $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.
- Score function $S(\theta, x) = 1/\theta - x$, use the estimator $L(X)S(\theta, X)$
- **Representations:** simplify the problem with $Y \sim \text{Bernoulli}(p(\theta))$,
 $p(\theta) = e^{-\theta \bar{x}}, Y = \mathbf{1}_{\{x \geq \bar{x}\}}$

For any function $L(Y)$ one has:

$$\mathbb{E}_\theta[L(Y)] = L(0) \frac{(1 - p(\theta))}{1 - p(1)} (1 - p(1)) + L(1) \frac{p(\theta)}{p(1)} p(1)$$

The **density** of the $\text{Bernoulli}(p(\theta))$ w.r.t. the $\text{Bernoulli}(1)$ is:

$$f_\theta(y) = \begin{cases} \frac{1-p(\theta)}{1-p(1)} & y = 0 \\ \frac{p(\theta)}{p(1)} & y = 1 \end{cases}$$

The **score function** is

$$S(\theta, y) = \begin{cases} \frac{\bar{x} e^{-\theta \bar{x}}}{(1 - e^{-\theta \bar{x}})} & y = 0 \\ -\bar{x} & y = 1 \end{cases}$$

- Use $L(X)S[\theta, Y(X)]$ as an estimator.

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and
 $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$. $h(\theta)$ is not a.s. Lipschitz continuous,
 $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and
 $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$. $h(\theta)$ is not a.s. Lipschitz continuous,
 $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.

- **Finite Difference:** Express the derivative as:

$$J'(\theta) = \lim_{\Delta\theta \rightarrow 0} \frac{1}{2\Delta\theta} \int_0^1 [h(\theta + \Delta\theta, \omega) - h(\theta, \omega)] d\omega \quad \text{CRN}$$

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$. $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.

- **Finite Difference:** Express the derivative as:

$$J'(\theta) = \lim_{\Delta\theta \rightarrow 0} \frac{1}{2\Delta\theta} \int_0^1 [h(\theta + \Delta\theta, \omega) - h(\theta, \omega)] d\omega \quad \text{CRN}$$

- **Critical Sets:** Identify the critical observables in $\Omega = [0, 1]$ that cause the discontinuities:

$$\mathcal{C}(\Delta\theta) = \{\omega : e^{-(\theta+\Delta\theta)\bar{x}} < \omega \leq e^{-(\theta-\Delta\theta)\bar{x}}\}, \mathbb{P}[\mathcal{C}(\Delta\theta)] = \mathcal{O}(\Delta\theta) \text{ rare events}$$

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$. $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.

- **Finite Difference:** Express the derivative as:

$$J'(\theta) = \lim_{\Delta\theta \rightarrow 0} \frac{1}{2\Delta\theta} \int_0^1 [h(\theta + \Delta\theta, \omega) - h(\theta, \omega)] d\omega \quad \text{CRN}$$

- **Critical Sets:** Identify the critical observables in $\Omega = [0, 1]$ that cause the discontinuities:

$$\mathcal{C}(\Delta\theta) = \{\omega : e^{-(\theta+\Delta\theta)\bar{x}} < \omega \leq e^{-(\theta-\Delta\theta)\bar{x}}\}, \mathbb{P}[\mathcal{C}(\Delta\theta)] = \mathcal{O}(\Delta\theta) \text{ rare events}$$

- **Conditioning:** Use the law of total probability:

$$\mathbb{E} \left[\frac{h_{\Delta\theta}(\theta, \omega)}{2\Delta\theta} \mathbb{P}(\mathcal{C}(\Delta\theta)^c) \right] + \mathbb{E}[h_{\Delta\theta}(\theta, \omega) | \mathcal{C}(\Delta\theta)] \left(\frac{\mathbb{P}(\mathcal{C}(\Delta\theta))}{2\Delta\theta} \right)$$

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$. $h(\theta)$ is not a.s. Lipschitz continuous, $h'(\theta, \omega) = 0, \omega \neq e^{-\theta \bar{x}}$.

- **Finite Difference:** Express the derivative as:

$$J'(\theta) = \lim_{\Delta\theta \rightarrow 0} \frac{1}{2\Delta\theta} \int_0^1 [h(\theta + \Delta\theta, \omega) - h(\theta, \omega)] d\omega \quad \text{CRN}$$

- **Critical Sets:** Identify the critical observables in $\Omega = [0, 1]$ that cause the discontinuities:

$$\mathcal{C}(\Delta\theta) = \{\omega : e^{-(\theta+\Delta\theta)\bar{x}} < \omega \leq e^{-(\theta-\Delta\theta)\bar{x}}\}, \mathbb{P}[\mathcal{C}(\Delta\theta)] = \mathcal{O}(\Delta\theta) \text{ rare events}$$

- **Conditioning:** Use the law of total probability:

$$\mathbb{E} \left[\frac{h_{\Delta\theta}(\theta, \omega)}{2\Delta\theta} \mathbb{P}(\mathcal{C}(\Delta\theta)^c) \right] + \mathbb{E}[h_{\Delta\theta}(\theta, \omega) | \mathcal{C}(\Delta\theta)] \left(\frac{\mathbb{P}(\mathcal{C}(\Delta\theta))}{2\Delta\theta} \right)$$

- **Difference process and critical rates:** Evaluate the limit:

$$J'(\theta) = \text{“IPA”} + (c_1 - c_2)\bar{x}e^{-\theta\bar{x}}.$$

Zero variance!!!

Known as SPA estimation

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and

$$L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}.$$

- **Representations:** simplify the problem with $Y \sim \text{Bernoulli}(p(\theta))$,
 $p(\theta) = e^{-\theta \bar{x}}$, $Y = \mathbf{1}_{\{x \geq \bar{x}\}}$

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and

$$L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}.$$

- **Representations:** simplify the problem with $Y \sim \text{Bernoulli}(p(\theta))$,
 $p(\theta) = e^{-\theta \bar{x}}$, $Y = \mathbf{1}_{\{x \geq \bar{x}\}}$
- **Derivative of the measure:** A direct dsitributional approach

$$\begin{aligned} \frac{d}{d\theta} \mathbf{E}_\theta[L(Y)] &= \frac{d}{d\theta} [L(0) (1 - p(\theta)) + L(1) p(\theta)] \\ &= -L(0) p'(\theta) + L(1) p'(\theta) = \bar{x} e^{-\theta \bar{x}} (L(0) - L(1)) \end{aligned}$$

Overview: Examples

A Threshold example: $X = -1/\theta \ln(\omega)$, $\mathbb{P} \sim U[0, 1]$ and
 $L(x) = c_1 \mathbf{1}_{\{x < \bar{x}\}} + c_2 \mathbf{1}_{\{x \geq \bar{x}\}}$.

- **Representations:** simplify the problem with $Y \sim \text{Bernoulli}(p(\theta))$,
 $p(\theta) = e^{-\theta \bar{x}}$, $Y = \mathbf{1}_{\{x \geq \bar{x}\}}$
- **Derivative of the measure:** A direct dsitributional approach

$$\begin{aligned} \frac{d}{d\theta} \mathbf{E}_\theta[L(Y)] &= \frac{d}{d\theta} [L(0) (1 - p(\theta)) + L(1) p(\theta)] \\ &= -L(0) p'(\theta) + L(1) p'(\theta) = \bar{x} e^{-\theta \bar{x}} (L(0) - L(1)) \end{aligned}$$

- **Expectations:** Interpret in terms of expectations of **new** random variables

$$\frac{d}{d\theta} \mathbf{E}_\theta[L(Y)] = \bar{x} e^{-\theta \bar{x}} \mathbf{E}[L(Y^+) - L(Y^-)], \quad Y^+ \sim \delta_0, Y^- \sim \delta_1$$

Recovers SPA estimator : zero variance

- **Estimation:** Usually called **weak derivative** estimator.

Summary

- Pathwise analysis:

$$J'(\theta) = \int h'(\theta, \omega) \mathbb{P}(d\omega)$$

Lipschitz continuity or conditioning on discontinuities: **careful construction of perturbations of finite difference to establish unbiasedness**

- Distributional analysis:

$$J'(\theta) = \int L(z) \frac{d}{d\theta} \mathbb{P}_\theta(dz)$$

Re-interpretation of the derivatives of measures as expectations.

- **Score function:** re-writes in terms of the original distribution,
- **Weak derivatives:** use other distributions for the derivative.

Dual spaces and weak topology: true for all bounded and continuous sample performance. **Extensions**