Stochastic Approximation for Finance

Noufel Frikha

Université Paris 1 Panthéon-Sorbonne

5th January 2026

Contents

Introduction

Introduction

• **Abstract:** The purpose of this course is to introduce the fundamental concepts for for stochastic approximation schemes for finance. The course will cover essential theoretical and computational aspects of stochastic approximation algorithm theory.

• **Prerequisites:** The course requires advanced knowledge in stochastic calculus and numerical probability but many concepts are recalled.

• **References:**

- Algorithmes Stochastiques, Marie Duflo, 1996.
- Numerical Probability, Gilles Pagès, 2018.
- L.C.G. Rogers et D. Talay, editors : Numerical Methods in Finance. Publications of the Newton Institute. Cambridge University Press, 1997.
- Stochastic simulation and Monte Carlo mehods, C. Graham, D. Talay, 2013.

Presentation of the course

- Who am I?
  - Noufel FRIKHA, Full-time professor (professeur des universités), Mathematical Finance Department, Paris 1 Panthéon-Sorbonne.
  - Main research topics: Mathematical finance, Numerical Probability, Machine Learning for finance, Stochastic Analysis and PDEs, ...

Presentation of the course

- Who am I?
  - Noufel FRIKHA, Full-time professor (professeur des universités), Mathematical Finance Department, Paris 1 Panthéon-Sorbonne.
  - Main research topics: Mathematical finance, Numerical Probability, Machine Learning for finance, Stochastic Analysis and PDEs, ...
- The course is divided into two parts:
  - 7 lectures: 05/01; 09/01; 23/01; 30/01; 06/02; 13/02; 23/02.

Presentation of the course

- Who am I?
  - Noufel FRIKHA, Full-time professor (professeur des universités), Mathematical Finance Department, Paris 1 Panthéon-Sorbonne.
  - Main research topics: Mathematical finance, Numerical Probability, Machine Learning for finance, Stochastic Analysis and PDEs, ...

- The course is divided into two parts:
  - 7 lectures: 05/01; 09/01; 23/01; 30/01; 06/02; 13/02; 23/02.

- Mark consists in two practical works and a final exam of 1 or 2 hours on Friday 20th March:
  - Two practical sessions on computers.
  - Questions on lectures, test general/basic knowledge of the two parts of the lecture course + 1 or 2 exercices/problems.

- Caution: Sorry but...no retake!

Stochastic approximation: why?

○ Stochastic approximation schemes appear in many fields of applied mathematics, especially in mathematical finance for the computation of option prices and sensitivities and also in stochastic optimization/control.

$\rightsquigarrow$ axiomatics and stochastic calculus theory and numerical probability.

Stochastic approximation: why?

◦ Stochastic approximation schemes appear in many fields of applied mathematics, especially in mathematical finance for the computation of option prices and sensivities and also in stochastic optimization/control.

$\rightsquigarrow$ axiomatics and stochastic calculus theory and numerical probability.

◦ Pricing and hedging of financial options: In mathematical finance, the price at time t of a financial contract with vanilla payoff $h(X_T)$ at maturity T is given by the risk-neutral rule

$$p_t = e^{-r(T-t)}\mathbb{E}[h(X_T)|\mathcal{F}_t].$$

When the underlying process $(X_t)_{t\geqslant 0}$ is Markov then

$$p_t = p_t(X_t) \quad \text{with} \quad p_t(x) = e^{-r(T-t)}\mathbb{E}[h(X_T)|X_t = x].$$

In this case, the delta hedge is based on the computation of the process

$$\delta_t = \nabla_x p_t(X_t), \quad t \geqslant 0.$$

◦ Many financial derivatives are path-dependent (lookback, Asian, barrier options): their payoff depends on the whole trajectory $(X_s)_{0\leqslant s\leqslant T}$.

◦ For instance, financial products such as lookback or barrier options depend on the running maximum, minimum, or occupation time of the price process X:

$$h(X_{0:T}) = h\left(\max_{0\leqslant s\leqslant T} X_s,\ \min_{0\leqslant s\leqslant T} X_s,\ X_T\right).$$

Example 1: Down-and-Out Call.

The option becomes worthless if the underlying crosses the barrier L before maturity T:

$$h_{\mathrm{DO}}(X_{0:T}) = (X_T - K)_+ \mathbf{1}_{\{\min_{0\leqslant s\leqslant T} X_s > L\}}$$

Example 2: Lookback Call.

$$h_{\mathrm{LB}}(X_{0:T}) = (\max_{0\leqslant s\leqslant T} X_s - K)_+.$$

Example 3: Parisian Down-and-Out Call.

The option is knocked out only if X stays below a barrier L for a cumulative time exceeding a window τ:

$$h_{\mathrm{Parisian}}(X_{0:T}) = (X_T - K)_+ \mathbf{1}_{\left\{\int_0^T \mathbf{1}_{\{X_s < L\}}\, ds \leqslant \tau\right\}}$$

Bachelier model

In the Bachelier model, the stock price is stochastic and follows a Brownian motion

$$S_t = S_0 + \sigma W_t, \quad 0 \leqslant t \leqslant T,$$

where $S_0 > 0$ is the initial price of the asset, T is the investment horizon, and $\sigma > 0$.

- The price process is a martingale (no arbitrage/fair game in Bachelier terms)
- S can become negative but T is assumed to be short, so that the probability of the event $\{S_T < 0\}$ is very small.
- For the same reason, short rate $r = 0$.

Bachelier Model Simulation

The following code generates 100 trajectories of the Bachelier model with parameters $S_0 = 100$ and $\sigma = 30\%$.

```
import numpy as np
import matplotlib.pyplot as plt

# Parameters
S0 = 100     # Initial price
volatility = 0.3 # Volatility (30%)
T = 1        # Time horizon (1 year)
n_steps = 252 # Number of time steps (daily steps)
n_simulations = 100 # Number of trajectories
dt = T / n_steps # Time step size

# Generate 100 trajectories of the Bachelier model
np.random.seed(42) # For reproducibility
dW = np.random.normal(scale=np.sqrt(dt), size=(n_steps, n_simulations))
paths = np.cumsum(volatility * dW, axis=0) # Brownian increments
paths = S0 + paths # Add to initial price

# Time grid
time = np.linspace(0, T, n_steps)

# Plot the trajectories
plt.figure(figsize=(12, 6))
plt.plot(time, paths, alpha=0.7)
plt.title("100 Trajectories of the Bachelier Model (S0=100, Volatility=30%)")
plt.xlabel("Time (Years)")
plt.ylabel("Price")
plt.grid()
plt.show()
```

Trajectories of Bachelier's model

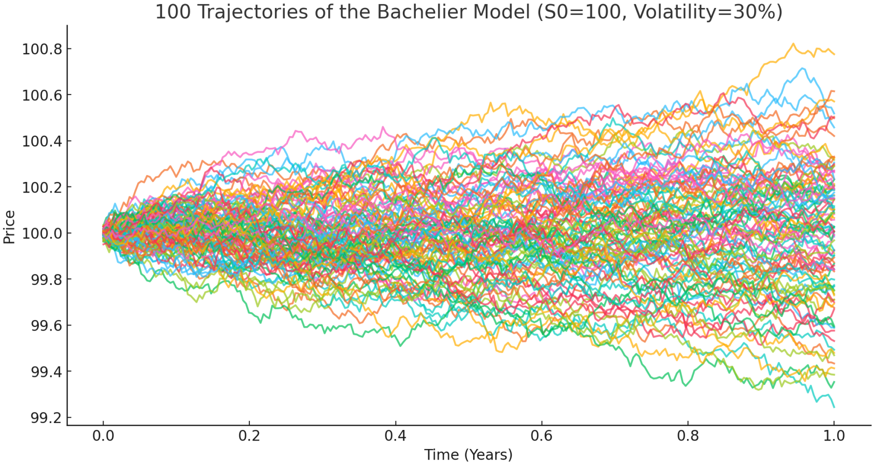

Bachelier model

Bachelier was interested in computing European call option prices with payoff $(S_T - K)^+$. Call option as protection:

- For Short Sellers: exposed to losses if the stock price rises.
- For Buyers of the Stock (Stock Replacement): limits their downside risk (the most they lose is the option premium)

$$C_t = \mathbb{E}\left[(S_T - K)^+ | \mathcal{F}_t\right]$$

Using properties of the Gaussian distribution, one can show that

$$C_t = (S_t - K)\Phi\left(\frac{S_t - K}{\sigma\sqrt{T-t}}\right) + \sigma\sqrt{T-t}\varphi\left(\frac{S_t - K}{\sigma\sqrt{T-t}}\right)$$

where

$$\varphi(x) = \frac{e^{-x^2/2}}{\sqrt{2\pi}} \quad \text{and} \quad \Phi(x) = \int_{-\infty}^{x} \varphi(y)\, dy.$$

◦ Exercise: Prove the above formula. Similar formula for the Put.

Black-Scholes Model

In the Black-Scholes model, the asset price follows the geometric SDE

$$dS_t = rS_t\,dt + \sigma S_t\,dW_t, \qquad 0 \leqslant t \leqslant T,$$

with initial price $S_0 > 0$, constant volatility $\sigma > 0$ and short rate r.

- The solution is the lognormal process

$$S_t = S_0 \exp\left(\left(r - \frac{\sigma^2}{2}\right) t + \sigma W_t\right).$$

- $S_t > 0$ almost surely: unlike in the Bachelier model, the price never becomes negative.
- Under the risk-neutral measure, the discounted process $(e^{-rt}S_t)_{t\geqslant 0}$ is a martingale (no-arbitrage).

Black–Scholes Model Simulation

100 trajectories of the B-S model with parameters $S_0 = 100$, $\sigma = 20\%$, $r = 2\%$.

```
import numpy as np
import matplotlib.pyplot as plt

# Parameters
S0 = 100      # Initial price
sigma = 0.2   # Volatility (20%)
r = 0.02      # Interest rate
T = 1         # Time horizon
n_steps = 252 # Daily time steps
n_sim = 100   # Number of paths
dt = T / n_steps

# Random increments
np.random.seed(42)
dW = np.random.normal(scale=np.sqrt(dt), size=(n_steps, n_sim))

# Geometric Brownian motion
S = np.zeros((n_steps+1, n_sim))
S[0] = S0

for t in range(n_steps):
   S[t+1] = S[t] * np.exp((r - 0.5*sigma**2)*dt + sigma * dW[t])

# Time grid
time = np.linspace(0, T, n_steps+1)

# Plot
plt.figure(figsize=(12,6))
plt.plot(time, S, alpha=0.7)
plt.title("100 Trajectories of the Black-Scholes Model (S0=100, sigma=20%)")
plt.xlabel("Time (Years)")
plt.ylabel("Price")
plt.grid()
```

Trajectories of the Black-Scholes Model

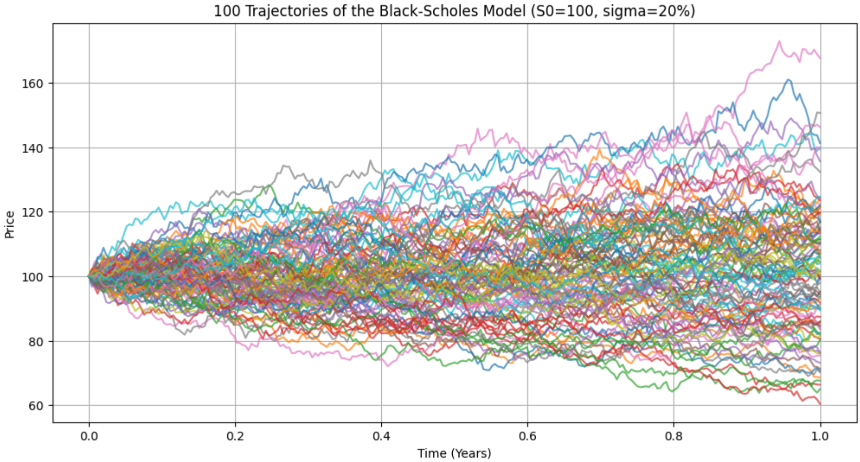

Black-Scholes Model

Black-Scholes computed the price of a European call option with payoff $(S_T - K)^+$:

$$C_t = e^{-r(T-t)}\,\mathbb{E}\big[(S_T - K)^+ \mid \mathcal{F}_t\big].$$

Since S_T is lognormally distributed,

$$C_t = S_t \Phi(d_+) - K e^{-r(T-t)} \Phi(d_-),$$

where

$$d_{+/-} = \frac{\ln(S_t/K) + (r \pm \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}.$$

- Exercise: Prove the Black-Scholes formula using the lognormal distribution.
- Similar expression for Put options (Put-Call parity).

◦ However, in general, the process X is given by the unique solution to a stochastic differential equation (SDE) with dynamics

$$X_t = x + \int_0^t b(s, X_s)\, ds + \int_0^t \sigma(s, X_s) dW_s.$$

◦ Except in some specific cases, like in the Black-Scholes or Bachelier models, no closed form solution exists.

◦ Their pricing requires simulating the full path of the SDE. Since closed-form solutions are rare, stochastic approximation schemes like the Euler-Maruyama time discretization scheme provide discrete-time approximations of the trajectory.

◦ Sensitivities (delta, vega) for such products are harder to estimate and rely heavily on stochastic approximation and variance-reduced algorithms.

Stochastic approximation algorithm: why?

In machine learning and mathematical finance, one often faces some optimization or zero search problems such as

- Estimate unknown relationship between X (observed variable) and Y (variable that should be explained) from a set of data $(X_1, Y_1), \cdots, (X_n, Y_n)$.
- Portfolio optimization: $\min_{u \in \Delta_d} \{\rho(-\langle u, X\rangle) + \lambda_d(u)\}$ where $\Delta_d := \left\{u \in (\mathbb{R}_+)^d : \sum_{i=1}^d u_i = 1\right\}$.
- Computation of VaR-ES: $\mathbb{P}(L \leqslant \xi) = \alpha$, $\text{ES}_\alpha(L) = \mathbb{E}[L | L \geqslant \text{VaR}_\alpha(L)]$.

Motivation from statistics

$\circ$ We are interested in the conditional distribution of Y given X. $\mathbb{P}_{\theta_0}(.|X)$ will represent the distribution of what we want to predict (the variable Y) given the value θ_0 and the value of the observation X.
$\circ$ We consider a sequence of i.i.d. observations $X_1, \cdots, X_n$ with a distribution function $f(x, \theta_0)$. We want to estimate the true value θ_0 of the parameter in a set of possible values Θ. The joint density of a n-sample $(x_1, \cdots, x_n)$ is

$$f(x_1, \theta) \times \cdots \times f(x_n, \theta).$$

$\circ$ The MLE estimator is defined as the value $\hat{\theta}_n \in \Theta$ that is the most likely to produce the set of observations $(x_1, \cdots, x_n)$

$$\hat{\theta}_n = \arg\max \left\{ l_n(\theta) := \sum_{i=1}^{n} \log(f(x_i, \theta)) \right\}.$$

Linear Gaussian model

As a standard example, we assume that X and Y are linked with the following Gaussian linear model

$$\mathcal{L}(Y|X) = \mathcal{N}(\langle \theta_0, X \rangle, \sigma^2),$$

where $X \in \mathbb{R}^p$, $\theta_0 \in \mathbb{R}^p$ is the unknown parameter and σ^2 is supposed to be known.

$$\hat{\theta}_n = \arg\min \left\{ -l_n(\theta) := \frac{1}{2\sigma^2} \sum_{i=1}^{n} (Y_i - \langle \theta, X \rangle)^2 + n \log(\sqrt{2\pi}\sigma) \right\}.$$

Denoting $X = [X_1, \cdots, X_n] \in \mathbb{R}^{n \times p}$ and $Y \in \mathbb{R}^{n \times 1}$, the MLE estimator is obtained by

$$\hat{\theta}_n = (X^t X)^{-1} X^t Y$$

$\circ$ In general models (e.g. inhomogeneous logistic model), no closed-form solution for the maximizer so that one must devise the gradient descent algorithm

$$\theta_{p+1} = \theta_p + \gamma_{p+1} \nabla l_n(\theta_p), \quad p \geqslant 0.$$

where $(\gamma_p)_{p\geqslant 1}$ is a positive deterministic sequence.

Mathematical Finance

Stochastic optimization & inverse problems

In mathematical finance, one often faces some optimization or zero search problems such as:

- Implied volatility: find σ, $\mathbb{E}[(S_T(\sigma) - K)_+] = P_{market}$.
- Computation of Value-at-Risk: For $\alpha \in (0,1)$, find ξ, $P(L \leqslant \xi) = \alpha$.
- Risk minimization: $\min_{u \in \Delta_d} \rho(-\langle u, X \rangle)$.

◦ General problem: for $h : \mathbb{R}^d \to \mathbb{R}^d$, find $\theta \in \mathbb{R}^d$ such that $h(\theta) = 0$. One can device the following Newton-Raphson like recursive schemes:

$$\theta_{n+1} = \theta_n - \gamma_{n+1} h(\theta_n), \quad \theta_0 \in \mathbb{R}^d,$$

where $(\gamma_n)_{n \geqslant 1}$ is a deterministic positive step sequence satisfying:

$$\sum_{n \geqslant 1} \gamma_n = +\infty \text{ and } \sum_{n \geqslant 1} \gamma_n^2 < +\infty. \tag{1}$$

If $\{h = 0\} = \{\theta^*\}$, h has a sub-linear growth and satisfies the mean-reverting assumption

$$\forall \theta \neq \theta^*, \quad \langle \theta - \theta^*, h(\theta) \rangle > 0$$

then

$$\theta_n \to \theta^*, \quad \text{as } n \to \infty.$$

However, this works only when one has straightforward access to the value of $h(\theta)$ which is not the case here since h can be written

$$h(\theta) = \mathbb{E}[H(\theta, U)], \quad H : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}^d.$$

The idea of Robbins and Monro (1951) consists in using the following randomized scheme

$$\theta_{n+1} = \theta_n - \gamma_{n+1} H(\theta_n, U^{(n+1)}), \quad n \geqslant 0$$

where $(U^{(n)})_{n \geqslant 1}$ is an i.i.d. sequence of r.v. independent of θ_0 with $\mathbb{E}[|\theta_0|^2] < \infty$.

A first version of the Robbins-Monro theorem

Assume that h satisfies the mean reverting assumption:

$$\forall \theta^* \in \{h = 0\}, \forall \theta \in \mathbb{R}^d \setminus \{h = 0\}, \quad \langle \theta - \theta^*, h(\theta) \rangle > 0,$$

and quadratic growth

$$\forall \theta \in \mathbb{R}^d, \quad \mathbb{E}[|H(\theta, U)|^2] \leqslant C(1 + |\theta|^2)$$

then, under the assumption (1) on the step $(\gamma_n)_{n \geqslant 1}$, as $n \uparrow \infty$, one has

$$\theta_n \to \theta_\infty, \quad a.s. \text{ and in } L^p(\mathbb{P}), \text{ for any } p < 2, \quad \text{as} n \to \infty.$$

where θ_∞ is an $\{h = 0\}$-valued random variable

We will study in more details stochastic approximation schemes in a dedicated chapter, see also the specific exercises and practical works on computers.

Conclusion

Take-home message

◦ In simple model (Bachelier or Black& Scholes), no stochastic approximation schemes are needed since one can rely on the explicit law of the marginal combined (if necessary) with Monte Carlo method.

◦ However, in general, dynamics must be discretized.

◦ Stochastic approximation schemes: Robbins-Monro, stochastic gradient, mirror descent, ... are now commonly used in stochastic optimization problems.

Thank you!