Unit 16: Stochastic Differential Equations

Instructor: Quan Zhou

Fall 2023

Stochastic differential equations

Let ξ be a random variable independent of the one-dimensional Brownian motion $(B_t)_{t \geq 0}$. Consider the stochastic differential equation (SDE):

$$\begin{aligned} X_0 &= \xi, \\ \mathrm{d}X_t &= b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t, \end{aligned} \tag{1}$$

where $\sigma \colon \mathbb{R} \times [0, \infty) \to (0, \infty)$ and $b \colon \mathbb{R} \times [0, \infty) \to \mathbb{R}$.

Stochastic differential equations

Let ξ be a random variable independent of the one-dimensional Brownian motion $(B_t)_{t \geq 0}$. Consider the stochastic differential equation (SDE):

$$\begin{aligned} X_0 &= \xi, \\ \mathrm{d}X_t &= b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t, \end{aligned} \tag{1}$$

where $\sigma \colon \mathbb{R} \times [0, \infty) \to (0, \infty)$ and $b \colon \mathbb{R} \times [0, \infty) \to \mathbb{R}$.

To solve this SDE means to seek an adapted process $(X_t)_{t \geq 0}$ s.t. a.s.,

$$X_t = \xi + \int_0^t b(X_s, s)\mathrm{d}s + \int_0^t \sigma(X_s, s)\mathrm{d}B_s, \quad \forall t \geq 0.$$

Existence and uniqueness of the solution? How to choose the filtration?

Strong solution

Given ξ and B_t, we use $\mathcal{F}_t^{\xi,B}$ to denote the completion of the σ-algebra generated by $\sigma(\xi)$ and $\sigma((B_s)_{0\le s\le t})$.

Definition 16.1

Let ξ, B_t be defined on $(\Omega, \mathcal{F}, \mathrm{P})$. A strong solution to (1) is a stochastic process $(X_t)_{t\ge 0}$ with continuous sample paths s.t.

1. X is adapted to $(\mathcal{F}_t^{\xi,B})_{t\ge 0}$;
2. $\mathrm{P}(X_0 = \xi) = 1$;
3. for any $0 \le t < \infty$,
$$\mathrm{P}\left(\int_0^t \left\{|b(X_s, s)| + \sigma^2(X_s, s)\right\} \mathrm{d}s < \infty\right) = 1;$$
4. almost surely, $X_t = \xi + \int_0^t b(X_s, s)\mathrm{d}s + \int_0^t \sigma(X_s, s)\mathrm{d}B_s, \quad \forall t \ge 0.$

Examples

Example 1 (Geometric Brownian motion)

Consider the geometric Brownian motion with $S_0 > 0$:

$$\mathrm{d}S_t = rS_t\mathrm{d}t + aS_t\mathrm{d}B_t.$$

By Itô formula, one can check that the solution is

$$S_t = S_0 \exp\left\{\left(r - \frac{1}{2}a^2\right)t + aB_t\right\}.$$

Law of iterated logarithm

If $r > a^2/2$, $S_t \to \infty$ a.s.; if $r < a^2/2$, $S_t \to 0$, a.s. This can be quickly proved by using the law of iterated logarithm.

Theorem 16.2

For a standard Brownian motion, a.s.

$$\liminf_{t\to\infty} \frac{B_t}{\sqrt{2t\log(\log t)}} = -1, \quad \limsup_{t\to\infty} \frac{B_t}{\sqrt{2t\log(\log t)}} = 1.$$

Examples

Example 2 (Ornstein-Uhlenbeck process)

The solution to the following SDE is known as Ornstein-Uhlenbeck process:

$$\mathrm{d}X_t = rX_t\mathrm{d}t + \sigma\mathrm{d}B_t.$$

It is a continuous-time version of the AR(1) process.

Examples

Example 2 (Ornstein-Uhlenbeck process)

The solution to the following SDE is known as Ornstein-Uhlenbeck process:

$$\mathrm{d}X_t = rX_t\mathrm{d}t + \sigma\mathrm{d}B_t.$$

It is a continuous-time version of the AR(1) process.

The strong solution is given by

$$X_t = e^{rt}X_0 + \sigma \int_0^t e^{r(t-s)}\mathrm{d}B_s.$$

Fubini theorem for Itô integrals

To verify that X_t indeed solves $\mathrm{d}X_t = rX_t\mathrm{d}t + \sigma\mathrm{d}B_t$, one can use the following version of Fubini's theorem for Itô integrals.

Theorem 16.3

Let $g(x,t)\colon \mathbb{R}\times[0,\infty)\to\mathbb{R}$ be continuous and twice continuously differentiable in x. Then

$$\int_0^s \left(\int_0^t g(u,v)\,\mathrm{d}u\right)\mathrm{d}B_v = \int_0^t \left(\int_0^s g(u,v)\,\mathrm{d}B_v\right)\mathrm{d}u.$$

Examples

Example 3 (Brownian bridge)

Let $b \in \mathbb{R}$. Consider the following SDE with $t \in [0, 1)$:

$$\mathrm{d}X_t = \frac{b - X_t}{1 - t}\mathrm{d}t + \mathrm{d}B_t.$$

Assume $X_0 = 0$. The strong solution is given by

$$X_t = bt + (1 - t) \int_0^t \frac{1}{1 - s}\mathrm{d}B_s.$$

Existence and uniqueness of strong solution

Theorem 16.4

Assume that $\mathrm{E}\xi^2 < \infty$ *and there exists constant* $K < \infty$ *s.t. for any* $x, y \in \mathbb{R}$ *and* $0 \le t < \infty$*,*

1. *(Lipschitz)*
$$|b(x,t) - b(y,t)| + |\sigma(x,t) - \sigma(y,t)| \le K|x - y|;$$
2. *(linear growth)*
$$|b(x,t)| + |\sigma(x,t)| \le K(1 + |x|).$$

Then the SDE given in (1) *has a unique strong solution* $(X_t)_{t \ge 0}$*.*

Grönwall lemma

Theorem 16.5

Let f, g be integrable functions and $t \in (0, \infty)$. Suppose there exists a constant $C \in (0, \infty)$ such that

$$f(s) \leq g(s) + C \int_0^s f(u) \mathrm{d}u, \quad \forall s \in [0, t].$$

Then,

$$f(s) \leq g(s) + C \int_0^s e^{C(s-u)} g(u) \mathrm{d}u, \quad \forall s \in [0, t].$$

In particular, if $g(t) \equiv a$ is constant, then $f(s) \leq ae^{Cs}$ for $s \in [0, t]$.

Existence and uniqueness of strong solution

Proof of Theorem 16.4: uniqueness.

Let X and $\tilde{X}$ be two strong solutions with initial r.v. ξ and $\tilde{\xi}$. Using Itô isometry, Lipschitz condition and Cauchy-Schwarz inequality, we find that

$$\mathrm{E}|X_t - \tilde{X}_t|^2 \leq 3\mathrm{E}|\xi - \tilde{\xi}|^2 + 3(1+t)K^2 \int_0^t \mathrm{E}|X_s - \tilde{X}_s|^2 \mathrm{d}s.$$

Letting $f(t) = \mathrm{E}|X_t - \tilde{X}_t|^2$ and applying Grönwall's lemma, we get

$$\mathrm{E}|X_s - \tilde{X}_s|^2 \leq 3e^{3(1+t)K^2 s}\mathrm{E}|\xi - \tilde{\xi}|^2, \quad \forall\, 0 \leq s \leq t.$$

Since we must have $\xi = \tilde{\xi}$ a.s., this shows that $X = \tilde{X}$ a.s. on the time interval $[0, t]$. Since t is arbitrary, we have the uniqueness on $[0, \infty)$. □

Existence and uniqueness of strong solution

Proof of Theorem 16.4: existence

Picard iteration: define $X_t^0 = \xi$, and for each $n \geq 1$, define

$$X_t^n = \xi + \int_0^t b(X_s^{n-1}, s)\mathrm{d}s + \int_0^t \sigma(X_s^{n-1}, s)\mathrm{d}B_s.$$

Our goal is to show that X^n converges to the SDE solution on time interval $[0, t]$ for every fixed $t \in [0, \infty)$.

Using the linear growth condition, one can show that $\int_0^t \mathsf{E}|X_s^n|^2 \mathrm{d}s < \infty$ for each n. This implies that $\int_0^t \sigma(X_s^n, s)\mathrm{d}B_s$ is defined for each n.

Existence and uniqueness of strong solution

Proof of Theorem 16.4: existence

Fix $t < \infty$. We will show that there exists $C(t) < \infty$ s.t.

$$\Delta_n(t) = \mathrm{E}\left[\sup_{0 \le s \le t} |X_s^n - X_s^{n-1}|^2\right] \le \frac{C(t)^n}{n!}. \tag{2}$$

Proof of Theorem 16.4: existence

Fix $t < \infty$. We will show that there exists $C(t) < \infty$ s.t.

$$\Delta_n(t) = \mathrm{E}\left[\sup_{0 \le s \le t} |X_s^n - X_s^{n-1}|^2\right] \le \frac{C(t)^n}{n!}. \tag{2}$$

Assuming that (2) holds, we have

$$\sum_{n=1}^{\infty} \mathrm{P}\left(\sup_{0 \le s \le t} |X_s^n - X_s^{n-1}|^2 > 2^{-n}\right) \le \sum_{n=1}^{\infty} 2^n \Delta_n(t) \le e^{2C(t)} < \infty.$$

Borel-Cantelli lemma thus shows that for almost every ω, $X^n(\omega)$ converges to some $X(\omega)$ in the space $\mathcal{C}([0, t])$ w.r.t. to the sup norm. Denote this limit by X. Since each X^n is continuous and adapted, so is X.

Existence and uniqueness of strong solution

Proof of Theorem 16.4: existence

Another consequence of (2) is that for any $m \geq n$ and $s \in [0, t]$,

$$\|X_s^m - X_s^n\|_2 \leq \sum_{k=n+1}^{m} \|X_s^k - X_s^{k-1}\|_2 \leq \sum_{k=n+1}^{\infty} \sqrt{\frac{C(t)^k}{k!}} =: B_n.$$

Note $B_n < \infty$ and $\lim_{n\to\infty} B_n = 0$. By Fatou's lemma,

$$\mathrm{E} \int_0^t |X_s - X_s^n|^2 \mathrm{d}s \leq \liminf_{m\to\infty} \mathrm{E} \int_0^t |X_s^m - X_s^n|^2 \mathrm{d}s \leq B_n^2 t.$$

Hence, $\lim_{n\to\infty} \mathrm{E} \int_0^t |X_t - X_t^n|^2 \mathrm{d}t = 0$. Using Itô isometry and assumptions on b and σ, we can then show that X satisfies (3) and (4) in Definition 16.1 on $[0, t]$. Since t is arbitrary, X is a strong solution.

Existence and uniqueness of strong solution

Sketch of proof for (2)

Define $D_t^n = \int_0^t [\sigma(X_s^n, s) - \sigma(X_s^{n-1}, s)] \mathrm{d}B_s$, which is a continuous martingale. Hence, Doob's inequality yields

$$\mathrm{E}\left(\sup_{s \le t} \|D_s^n\|_2^2\right) \le 4\mathrm{E}\|D_t^n\|_2^2 \le 4K^2 \int_0^t \mathrm{E}|X_s^n - X_s^{n-1}|^2 \mathrm{d}s.$$

Define $F_t^n = \int_0^t [b(X_s^n, s) - b(X_s^{n-1}, s)] \mathrm{d}s$. Cauchy-Schwarz yields that

$$\mathrm{E}\left(\sup_{s \le t} \|F_s^n\|_2^2\right) \le tK^2 \int_0^t \mathrm{E}|X_s^n - X_s^{n-1}|^2 \mathrm{d}s.$$

Hence, there exists some $C(t) < \infty$ s.t. $\Delta_{n+1}(t) \le C(t) \int_0^t \Delta_n(s) \mathrm{d}s$. A routine induction argument completes the proof.

Itô diffusion

The solution to an SDE is often called (Itô) diffusion; b is called the drift coefficient, and σ the diffusion coefficient.

Itô diffusion

The solution to an SDE is often called (Itô) diffusion; b is called the drift coefficient, and σ the diffusion coefficient.

In the time-homogeneous case, we have an SDE

$$\mathrm{d}X_t = b(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}B_t. \tag{3}$$

The assumptions of Theorem 16.4 can be simplified to

$$\exists K < \infty, \text{ s.t. } |b(x) - b(y)| + |\sigma(x) - \sigma(y)| \leq K|x - y|, \ \forall x, y. \tag{4}$$

Strong Markovian property

Theorem 16.6

Assume (4) holds. The solution to (3) is a strong Markov process; that is, for any bounded f, any finite stopping time T w.r.t. the filtration defined by $\mathcal{F}_t^0 = \sigma((B_s)_{0 \le s \le t})$, and any $s > 0$, we have a.s.,

$$\mathrm{E}_x \left[f(X_{T+s}) \,|\, \mathcal{F}_T^0 \right] = \mathrm{E}_{X_T}[f(X_s)],$$

where E_x denotes the expectation corresponding to the probability measure P_x under which $\mathrm{P}_x(X_0 = x) = 1$.

Infinitesimal generator

Let X be the solution to (3). The infinitesimal generator of X, denoted by $\mathcal{A}$, is defined by

$$(\mathcal{A}f)(x) = \lim_{s \downarrow 0} \frac{\mathsf{E}_x[f(X_s)] - f(x)}{s}.$$

Let $\mathcal{D}(\mathcal{A}) = \{f : (\mathcal{A}f)(x) \text{ exists for every } x \in \mathbb{R}\}$.

Infinitesimal generator

Let X be the solution to (3). The infinitesimal generator of X, denoted by $\mathcal{A}$, is defined by

$$(\mathcal{A}f)(x) = \lim_{s \downarrow 0} \frac{\mathrm{E}_x[f(X_s)] - f(x)}{s}.$$

Let $\mathcal{D}(\mathcal{A}) = \{f : (\mathcal{A}f)(x) \text{ exists for every } x \in \mathbb{R}\}$.

Theorem 16.7

Assume (4) holds. Let $f : \mathbb{R} \to \mathbb{R}$ be twice continuously differentiable and have a bounded support. Then $f \in \mathcal{D}(\mathcal{A})$, and

$$(\mathcal{A}f)(x) = b(x)\frac{\partial f}{\partial x} + \frac{1}{2}\sigma^2(x)\frac{\partial^2 f}{\partial x^2}.$$

Dynkin's formula

The infinitesimal generator is often used to calculate the expectation of $f(X_T)$ for some stopping time T. The proof of the following theorem, known as Dynkin's formula, is similar to that of Theorem 16.8. One applies Itô's lemma and verifies that the stochastic integral involving $\mathrm{d}B_t$ has expectation zero.

Dynkin's formula

The infinitesimal generator is often used to calculate the expectation of $f(X_T)$ for some stopping time T. The proof of the following theorem, known as Dynkin's formula, is similar to that of Theorem 16.8. One applies Itô's lemma and verifies that the stochastic integral involving $\mathrm{d}B_t$ has expectation zero.

Theorem 16.8

Under the setting of Theorem 16.7, for any stopping time T such that $\mathsf{E}_x[T] < \infty$, we have

$$\mathsf{E}_x[f(X_T)] = f(x) + \mathsf{E}_x\left[\int_0^T (\mathcal{A}f)(X_s)\mathrm{d}s\right].$$

Weak solutions

Definition 16.9

A weak solution to the SDE

$$\mathrm{d}X_t = b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t,$$

with initial distribution μ is a triple $(X, B), (\Omega, \mathcal{F}, \mathrm{P})$ and $(\mathcal{F}_t)_{t\geq 0}$ s.t.

1. $(\Omega, \mathcal{F}, \mathrm{P})$ is a probability space, and $(\mathcal{F}_t)_{t\geq 0}$ is a right-continuous and complete filtration;
2. X is adapted to $(\mathcal{F}_t)_{\geq 0}$ and has continuous paths, and B is a standard Brownian motion w.r.t. $(\mathcal{F}_t)_{\geq 0}$;
3. $\mathrm{P} \circ X_0^{-1} = \mu$;
4. for any $0 \leq t < \infty$, $\mathrm{P}\left(\int_0^t \left\{|b(X_s, s)| + \sigma^2(X_s, s)\right\} \mathrm{d}s < \infty\right) = 1$;
5. almost surely, $X_t = X_0 + \int_0^t b(X_s, s)\mathrm{d}s + \int_0^t \sigma(X_s, s)\mathrm{d}B_s$ for $t \geq 0$.

Definition 16.10

We say that the weak solution to a SDE is unique in law if, for any two weak solutions $\{(X, B), (\Omega, \mathcal{F}, \mathrm{P}), (\mathcal{F}_t)\}$ and $\{(\tilde{X}, \tilde{B}), (\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathrm{P}}), (\tilde{\mathcal{F}}_t)\}$, we have $\mathrm{Law}(X) = \mathrm{Law}(\tilde{X})$.

Examples

Example 4 (Tanaka's SDE)

Let $\mathrm{sgn}(x) = \mathbb{1}_{[0,\infty)}(x) - \mathbb{1}_{(-\infty,0)}(x)$. Consider the following SDE

$$\mathrm{d}X_t = \mathrm{sgn}(X_t)\,\mathrm{d}B_t, \qquad \text{with } X_0 = 0.$$

Examples

Example 4 (Tanaka's SDE)

Let $\operatorname{sgn}(x) = \mathbb{1}_{[0,\infty)}(x) - \mathbb{1}_{(-\infty,0)}(x)$. Consider the following SDE

$$\mathrm{d}X_t = \operatorname{sgn}(X_t)\,\mathrm{d}B_t, \quad \text{with } X_0 = 0.$$

Here is a weak solution unique in law. Let X be a standard Brownian motion on $(\Omega, \mathcal{F}, \mathrm{P})$, and let $\mathcal{F}_t$ be the completion of $\sigma((X_s)_{0 \le s \le t})$. By Tanaka's SDE, we can define B_t by

$$B_t = \int_0^t \operatorname{sgn}(X_t)\,\mathrm{d}X_t.$$

It can be shown that B_t is indeed a Brownian motion adapted to $(\mathcal{F}_t)_{t \ge 0}$. However, there is no strong solution.

Filtering problems

Suppose we observe the process $(Z_t)_{t \geq 0}$ with dynamics given by

$$\mathrm{d}Z_t = b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t.$$

How to estimate $(X_t)_{t \geq 0}$? The estimate $\hat{X}_t$ must be measurable w.r.t. $\mathcal{F}_t^Z$, the completion of $\sigma((Z_s)_{0 \leq s \leq t})$.

Filtering problems

Suppose we observe the process $(Z_t)_{t \geq 0}$ with dynamics given by

$$\mathrm{d}Z_t = b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t.$$

How to estimate $(X_t)_{t \geq 0}$? The estimate $\hat{X}_t$ must be measurable w.r.t. $\mathcal{F}_t^Z$, the completion of $\sigma((Z_s)_{0 \leq s \leq t})$.

By the projection property of conditional expectation, the best estimator that minimizes $\mathsf{E}|X_t - \hat{X}_t|^2$ is given by

$$\hat{X}_t = \mathsf{E}[X_t \mid \mathcal{F}_t^Z].$$

Linear filtering

Consider the linear case:

$$\mathrm{d}X_t = f_t X_t \mathrm{d}t + \sigma_t \mathrm{d}B_t, \quad \text{with } X_0 \sim N(\mu_0, v_0),$$
$$\mathrm{d}Z_t = g_t X_t \mathrm{d}t + \rho_t \mathrm{d}\tilde{B}_t, \quad \text{with } Z_0 = 0,$$

where $f_t, g_t, \sigma_t, \rho_t$ are deterministic functions, and $B, \tilde{B}$ are two independent Brownian motions. Assume that

1. $f_t, g_t, \sigma_t, \rho_t$ are all bounded on $[0, n]$ for every $n < \infty$;
2. $\sigma_t \geq 0$ for all t, and $\inf_{t \geq 0} \rho_t > 0$.

Linear filtering

Theorem 16.11 (Kalman-Bucy filter)

For the linear filtering problem, the solution is given by $\hat{X}_0 = \mu_0$,

$$\mathrm{d}\hat{X}_t = \left(f_t - \frac{g_t^2 s_t}{\rho_t^2} \right) \hat{X}_t \, \mathrm{d}t + \frac{g_t s_t}{\rho_t^2} \mathrm{d}Z_t,$$

where $s_t = \mathsf{E}|X_t - \hat{X}_t|^2$ *satisfies* $s_0 = v_0$ *and*

$$\frac{\mathrm{d}s}{\mathrm{d}t} = -\frac{g_t^2}{\rho_t^2} s_t^2 + 2 f_t s_t + \sigma_t^2.$$

Optimal stopping

Let $g \colon \mathbb{R} \to [0, \infty)$ be given and X_t be given by

$$\mathrm{d}X_t = b(X_t, t)\mathrm{d}t + \sigma(X_t, t)\mathrm{d}B_t.$$

Optimal stopping means to find an optimal stopping time T^* that attains

$$\sup_T \mathsf{E}[g(X_T)],$$

where the supremum is taken over all stopping times w.r.t. the filtration generated by X. The function g is often known as the reward function.

Quickest detection

Let $\pi \in [0, 1)$ and $\theta \sim \pi\delta_0 + (1 - \pi)\mathrm{Exp}(\lambda)$ (where δ_0 denotes the Dirac measure at 0). Assume θ is unknown and independent of B. We observe the process X with $X_0 = 0$ and dynamics given by

$$\mathrm{d}X_t = \mu \mathbb{1}_{\{\theta \le t\}} \mathrm{d}t + \sigma \mathrm{d}B_t,$$

where μ, σ are known. For $\beta > 0$, the goal is to find T^* that attains

$$\inf_T \mathsf{P}(T < \theta) + \beta \mathsf{E}[(T - \theta)^+].$$

Let $\pi \in [0,1)$ and $\theta \sim \pi\delta_0 + (1-\pi)\mathrm{Exp}(\lambda)$ (where δ_0 denotes the Dirac measure at 0). Assume θ is unknown and independent of B. We observe the process X with $X_0 = 0$ and dynamics given by

$$\mathrm{d}X_t = \mu \mathbb{1}_{\{\theta \le t\}} \mathrm{d}t + \sigma \mathrm{d}B_t,$$

where μ, σ are known. For $\beta > 0$, the goal is to find T^* that attains

$$\inf_T \mathsf{P}(T < \theta) + \beta \mathsf{E}[(T-\theta)^+].$$

Why do we choose this objective function?

Stochastic control

In stochastic control problems, we can choose a stochastic process $(u_t)_{t \geq 0}$, known as the control, to modify the system dynamics. Assume that the controlled process, denoted by X^u, evolves by

$$\mathrm{d}X_t^u = b(X_t^u, t, u_t)\mathrm{d}t + \sigma(X_t^u, t, u_t)\mathrm{d}B_t. \tag{5}$$

We usually require that u_t be measurable w.r.t. $\mathcal{F}_t^B$ or w.r.t. $\mathcal{F}_t^X$.

Stochastic control

In stochastic control problems, we can choose a stochastic process $(u_t)_{t \geq 0}$, known as the control, to modify the system dynamics. Assume that the controlled process, denoted by X^u, evolves by

$$\mathrm{d}X_t^u = b(X_t^u, t, u_t)\mathrm{d}t + \sigma(X_t^u, t, u_t)\mathrm{d}B_t. \tag{5}$$

We usually require that u_t be measurable w.r.t. $\mathcal{F}_t^B$ or w.r.t. $\mathcal{F}_t^X$.

If we can write $u_t(\omega) = u_0(X_t^u(\omega), t)$ for some measurable function u_0, we say u is Markovian. Sometimes we only consider Markovian controls.

Stochastic control

Let T denote the time horizon of the problem. Some common choices are: $T \in (0, \infty)$, $T = \infty$, or $T = \inf\{t \geq 0 \colon (X_t^u, t) \notin \mathbb{C}\}$ for some bounded set $\mathbb{C} \subset \mathbb{R} \times [0, \infty)$.

Stochastic control

Let T denote the time horizon of the problem. Some common choices are: $T \in (0, \infty)$, $T = \infty$, or $T = \inf\{t \geq 0 \colon (X_t^u, t) \notin \mathbb{C}\}$ for some bounded set $\mathbb{C} \subset \mathbb{R} \times [0, \infty)$.

Given some measurable functions f, g, the goal is to find the control u^* that attains $\sup_u J(u)$, where

$$J(u) = \mathsf{E}\left[\int_0^T f(X_t^u, t, u_t)\mathrm{d}t + g(X_T^u)\mathbb{1}_{\{T<\infty\}}\right].$$

Stochastic control

Let T denote the time horizon of the problem. Some common choices are: $T \in (0, \infty)$, $T = \infty$, or $T = \inf\{t \geq 0 \colon (X_t^u, t) \notin \mathbb{C}\}$ for some bounded set $\mathbb{C} \subset \mathbb{R} \times [0, \infty)$.

Given some measurable functions f, g, the goal is to find the control u^* that attains $\sup_u J(u)$, where

$$J(u) = \mathsf{E}\left[\int_0^T f(X_t^u, t, u_t)\mathrm{d}t + g(X_T^u)\mathbb{1}_{\{T<\infty\}}\right].$$

A typical application of stochastic control is to find the optimal portfolio in a financial market.

Value function

It is often more convenient to find the optimal control for all possible initial states. Define the value function by

$$v(x,t) = \sup_u J_{x,t}(u), \quad \text{where}$$
$$J_{x,t}(u) = \mathrm{E}_{x,t}\left[\int_t^T f(X_t^u, t, u_t)\mathrm{d}s + g(X_T^u)\mathbb{1}_{\{T<\infty\}}\right].$$

The expectation $\mathrm{E}_{x,t}$ means that we consider the solution to the SDE (5) starting at $X_t^u = x$.

Hamilton-Jacobi-Bellman equation

Under certain regularity conditions, the optimal control is Markovian. To find the optimal control, one often begins with solving the so-called Hamilton-Jacobi-Bellman (HJB) equation,

$$\sup_u \left\{ f(x,t,u) + \frac{\partial v}{\partial t}(x,t) + b(x,t,u)\frac{\partial v}{\partial x} + \frac{1}{2}\sigma^2(x,t,u)\frac{\partial^2 v}{\partial x^2} \right\} = 0,$$

subject to the boundary condition $v(x,T) = g(x)$ (assuming T is fixed).

Under some conditions, one can prove that the solution v to the HJB equation is the value function we seek, and the control u that attains the supremum in the HJB equation is optimal. This technique is known as the verification theorem.

Exercises

Exercise 16.1

Consider the geometric Brownian motion S_t in Example 1. Prove that if $r < a^2/2$, $S_t \to 0$, a.s.

Exercise 16.2

Consider the Ornstein-Uhlenbeck process X_t in Example 2. Verify that X_t solves the SDE $\mathrm{d}X_t = rX_t\mathrm{d}t + \sigma\mathrm{d}B_t$.

Exercise 16.3

Show that the Brownian bridge X_t in Example 3 satisfies $\lim_{t\uparrow 1} X_t = b$ a.s.

Exercises

Exercise 16.4

Let θ be a parameter drawn from $N(\mu_0, v_0)$. Suppose we observe the process $(Z_t)_{t\geq 0}$ with dynamics

$$\mathrm{d}Z_t = \theta g_t \mathrm{d}t + \rho\, \mathrm{d}B_t.$$

where $\rho > 0$ is known and g_t is a known bounded function. Use Kalman-Bucy filter to find the estimate of θ at time t.

Exercise 16.5

Prove Theorem 16.7.

References

[1] Richard Durrett. *Stochastic calculus: a practical introduction*. CRC press, 2018.

[2] Ioannis Karatzas and Steven Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 2012.

[3] Achim Klenke. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.

[4] Bernt Øksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.