Unit 6b: Stochastic Calculus

If X is a time-homogeneous diffusion we can write, for small h,

$$X_{t+h} = X_t + \mu(X_t)h + \sigma(X_t)[B(t+h) - B(t)] + o(h).$$

Check that this works: Markov property, continuity, moments of the increments $X_{t+h} - X_t$ conditional on $\mathcal{F}_t$.

This result is usually written as

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dB(t).$$

For many applications it is helpful to have an explicit representation of a diffusion process in terms of Brownian motion. We will formalise what it means for a process X to satisfy this equation.

The Itô Integral

Assume we have a filtration $\{\mathcal{F}_t\}$ and a BM $\{B(t)\}$ adapted to it. Let $Y(t)$ be a simple stochastic process:

$$Y(t) = \begin{cases} Y_0 & \text{for } 0 \le t < t_1 \\ Y_1 & \text{for } t_1 \le t < T \end{cases}$$

where Y_0 and Y_1 are random variables, $Y_0 \in \mathcal{F}_0$, $Y_1 \in \mathcal{F}_{t_1}$, $\mathbb{E}Y_i^2 < \infty$ $(i = 0, 1)$.

The *Itô integral* of Y is

$$I_t(Y) = \begin{cases} Y_0 B(t) & \text{if } 0 \le t < t_1 \\ Y_0 B(t_1) + Y_1(B(t) - B(t_1)) & \text{if } t_1 \le t \le T \end{cases}$$

This can clearly be extended to cover the case where Y takes finitely many values:

$$Y(t) = Y_i \text{ for } t \in [t_i, t_{i+1}),$$

where $0 = t_0 < t_1 < \cdots < t_n$, $Y_i \in \mathcal{F}_{t_i}$ and $\mathbb{E}Y_i^2 < \infty$. (This is a *step function.*)

Extension to general $Y(t)$

Consider a general $\mathcal{F}_t$-adapted square-integrable stochastic process $Y(t)$ with at most finitely many jumps.

We can find a sequence of step functions $Y^{(n)}(t)$ such that

$*$ $Y^{(n)}(t) \to Y(t)$ as $n \to \infty$ for each t.

$*$ $\mathbb{E}\left(Y(t) - Y^{(n)}(t)\right)^2 \to 0$ as $n \to \infty$ for each t.

Under these conditions it can be shown that
∗ $I_t(Y^{(n)})$ converges to a limit for each t
∗ If two different sequences $Y^{(n)}$ and $\overline{Y}^{(n)}$ have the same limit Y then $I_t(Y^{(n)})$ and $I_t(\overline{Y}^{(n)})$ have the same limit, which we can therefore denote $I_t(Y)$.

Since $I_t(c) = cB(t)$ for a constant c, it is reasonable to use the notation

$$I_t(Y) = \int_0^t Y(s)\, dB(s).$$

Properties of the Itô Integral

∗ $I_t(Y)$ has the Markov property
∗ $I_t(Y)$ is a martingale
∗ $\mathbb{E}[I_t(Y)] = 0$
∗ $\mathrm{Var}[I_t(Y)] = \mathbb{E} \int_0^t Y(s)^2\, ds$
∗ If Y is deterministic, then $I_t(Y)$ is Normally distributed.

These are easily proved in the case of step functions. The properties carry across to the general case.

Itô processes

An Itô process is any process of the form

$$X(t) = X_0 + \int_0^t b_s\, ds + \int_0^t \sigma_s\, dB(s)$$

where $X_0 \in \mathcal{F}_0$, $b_t \in \mathcal{F}_t$, $\sigma_t \in \mathcal{F}_t$ and $\int_0^t |b_s|\, ds$ and $\int_0^t \sigma_s^2\, ds$ are a.s. finite.

Stochastic differential

A shorthand notation for the above is

$$dX(t) = b_t\, dt + \sigma_t\, dB(t)$$

We will be encountering stochastic differentials where b_t and σ_t are functions of $X(t)$ as well as t:

$$dX(t) = b(t, X(t))\, dt + \sigma(t, X(t))\, dB(t)$$

These is a *stochastic differential equation* (SDE). To solve it we will need a special technique.

Itô's Lemma

If X_t is an Itô process satisfying

$$dX(t) = b(t, X(t))\, dt + \sigma(t, X(t))\, dB(t)$$

and if

$$Y(t) = f(t, X(t)),$$

where $f \in C^{12}(\mathbb{R}_+ \times \mathbb{R})$, then $Y(t)$ is an Itô process satisfying

$$dY(t) = \frac{\partial f}{\partial t}\, dt + \frac{\partial f}{\partial x}\, dX(t) + \frac{1}{2}\sigma^2(t, X(t)) \frac{\partial^2 f}{\partial x^2}\, dt$$

Proof. The Taylor expansion gives

$$dY(t) = \frac{\partial f}{\partial t}\, dt + \frac{\partial f}{\partial x}\, dX(t) + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}\, (dX(t))^2 + o(t).$$

Both $(dt)^2$ and $(dt)(dX(t))$ are $o(t)$, but

$$(dB(t))^2 = dt.$$

Examples of Itô processes

If $Y(t) = B(t)^2$ then

$$dY(t) = 2B(t)\, dB(t) + dt.$$

Therefore

$$\int_0^t 2B(s)\, dB(s) = B(t)^2 - t;$$

recall that the RHS is known to be a martingale.

Geometric BM

If $S(t) = S_0 e^{\mu t + \sigma B(t)}$, then

$$\begin{aligned} dS(t) &= \mu S(t)\, dt + \sigma S(t)\, dB(t) + \tfrac{1}{2}\sigma^2 S(t)\, dt \\ &= \left(\mu + \tfrac{1}{2}\sigma^2\right) S(t)\, dt + \sigma S(t)\, dB(t) \end{aligned}$$

Solving Stochastic DEs

Trial and error. The standard technique is to work out what the solution would be in the case of an ODE and see if this helps.

Example: Geometric BM

Suppose S satisfies

$$\frac{dS(t)}{S(t)} = b\, dt + \sigma\, dB(t).$$

We find $d(\log S(t))$:

$$\begin{aligned} d(\log S(t)) &= \frac{1}{S(t)}\, dS(t) - \frac{1}{2S(t)^2}\sigma^2 S(t)^2\, dt \\ &= \left(b - \tfrac{1}{2}\sigma^2\right)\, dt + \sigma\, dB(t). \end{aligned}$$

Therefore

$$\log S(t) - \log S(0) = \left(b - \tfrac{1}{2}\sigma^2\right) t + \sigma B(t).$$

Further solved SDEs

Ornstein-Uhlenbeck process
We have

$$dU(t) = -\gamma U(t)\,dt + \sigma\,dB(t)$$

Try $Y(t) = e^{\gamma t}U(t)$. Then

$$dY(t) = \gamma Y(t)\,dt + e^{\gamma t}\,dU(t),$$

with solution

$$e^{\gamma t}U(t) - U(0) = \sigma \int_0^t e^{\gamma s}\,dB(s).$$

This expression proves that $U(t)$ is Normally distributed.

It is not always possible to find a solution to an SDE: consider, for example, the *Cox-Ingersoll-Ross process*

$$dX(t) = \alpha(b - X(t))\,dt + \sigma\sqrt{X(t)}\,dB(t).$$

Moments of Itô processes

If we define $m_t = \mathbb{E}[X(t)|X_0 = a]$ then

$$\begin{aligned} m_{t+dt} - m_t &= \mathbb{E}[dX(t)|X_0 = a] \\ &= \mathbb{E}[b(t, X(t))|X_0 = a]\,dt \end{aligned}$$

In other words, $\frac{dm}{dt} = \mathbb{E}[b]$. If $b(t, X(t))$ is a linear function of $X(t)$, we have an ODE for m.

If we want $q_t = \mathbb{E}[X(t)^2|X_0 = a]$, then we must first define $Y_t = X(t)^2$, so that

$$dY_t = 2X(t)[b\,dt + \sigma\,dB(t)] + \sigma^2\,dt.$$

Hence

$$\frac{dq}{dt} = \mathbb{E}[2X_t b(t, X_t) + \sigma^2(t, X_t)|X_0 = a],$$

which again will be soluble for appropriate b and σ.

Lévy Processes

The BM and associated processes are continuous and fail to explain all market movements. A general process with stationary independent increments is a *Lévy process.*

Any Lévy process can be decomposed as a sum of:
∗ μt, a deterministic linear part;
∗ $\sigma B(t)$, a continuous random part;
∗ a purely discontinuous random part.

The final part is like a compound Poisson process. If we take any $c > 0$ and only look at jumps of size at least c then it is exactly a CPP.
The difference is that the discontinuous part of a Lévy process is permitted to have infinitely many jumps in a small time interval, whereas a Poisson process is not.

Stable laws

Stable laws ("law" = "distribution") are distributions which arise as the limiting distribution of a quantity of the form

$$\frac{X_1 + X_2 + \cdots + X_n - \alpha_n}{\beta_n}$$

where the X_i are i.i.d. and α_n and β_n are deterministic.

By the Central Limit Theorem, the Normal distribution is stable with $\beta_n = \sqrt{n}$.
Student's t distribution on 1 degree of freedom (Cauchy distribution) is also stable with $\beta_n = n$.

These form a subclass of the *infinitely divisible distributions*: a random variable Y has an i.d. distribution if, for any n, we can write Y as a sum $Y = \sum_{i=1}^{n} X_i$ of i.i.d. random variables.

Any Lévy process has an infinitely divisible distribution, in common with a Poisson process and a Brownian motion.

Modelling Brownian motion

A standard BM has no flexibility and is not suitable for modelling. The BM with drift μ and volatility σ (Wiener process) is better.

If observations are equally spaced, at times $k\Delta$, then we can define

$$Y_j = X(j\Delta) - X((j-1)\Delta)$$

so the Y_j are i.i.d. with mean $\mu\Delta$, variance $\sigma^2\Delta$ and the sample mean and variance of the Y_j are estimates of $\mu\Delta$ and $\sigma^2\Delta$.

Even if the observations are not equally spaced, maximum likelihood can still be applied.

Fitting a geometric BM to data is also easy: the log of a GBM is a Wiener process, so take the log of the data and proceed as before.

Testing goodness of fit has two parts:
∗ Are the increments Normally distributed?
∗ Are the increments independent of each other?
∗ Are the increments independent of the process?

Modelling a Lévy process

Again assume regular observations $X(k\Delta)$. The difficulty here lies in distinguishing between increments due to normal Brownian variation and increments due to jumps in the compound Poisson process.
It is usual to set a threshold u and assume that any time interval $((j-1)\Delta, j\Delta)$ for which $|X(j\Delta) - X((j-1)\Delta)| > u$ contains a jump of the CPP, whereas other time intervals do not.
The parameters μ and σ can be estimated using increments from the jump-free intervals; the distribution of jump heights is estimated from the intervals containing jumps.
The choice of the threshold u is subjective.

A common model is the geometric Lévy process, where $\log X(t)$ is a Lévy process. Simply log-transform the data and apply the procedure above.

Simulating Brownian motion

Fix a small time increment Δ, simulate a sequence Z_i of $N(0,1)$ variables and set

$$X(j\Delta) = X((j-1)\Delta) + \mu\Delta + \sigma\sqrt{\Delta}Z_j$$

to get a general BM (Wiener process).

Simulating a diffusion is the same, except that μ and σ are functions of $j\Delta$ (and possibly t).

To simulate a Lévy process we also need a simulation of a compound Poisson process.
Let V_i be a sequence of exponential pseudo-random variables, Y_i an independent sequence of jump heights.
Define $T_i = \sum_{j=1}^{i} V_j$, the time of the jth jump. Then

$$\begin{aligned} X(j\Delta) &= X((j-1)\Delta) + \mu\Delta + \sigma\sqrt{\Delta}Z_j \\ &\quad + \sum_i Y_i 1_{(j-1)\Delta < T_i \leq j\Delta} \end{aligned}$$