STOCHASTIC CALCULUS

A brief set of introductory notes on stochastic calculus and stochastic differential equations.

1. Processes.

We are concerned with continuous-time, real-valued stochastic processes $(X_t)_{0 \le t < \infty}$. These may be thought of as random functions – for each outcome of the random element, we have a real-valued function of a real variable t. These possible outcomes (functions) are called *realizations* or *sample paths*.

Examples: (i) Random walk (step up or down, probability 1/2 each, at each integer time), with linear interpolation. (ii) Polynomial with random coefficients. (iii) X_t = number of events in a Poisson process occurring by time t. (iv) Brownian motion.

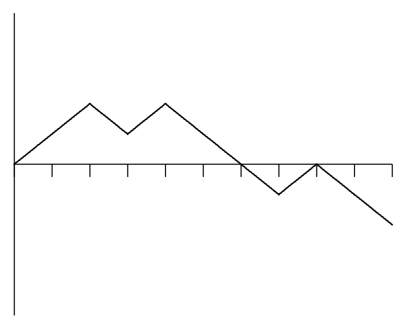

(i)

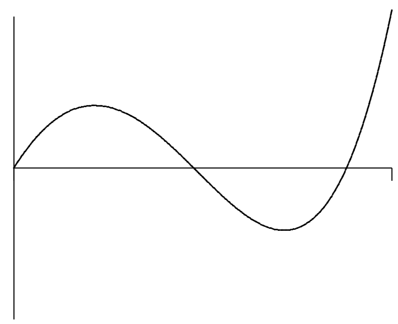

(ii)

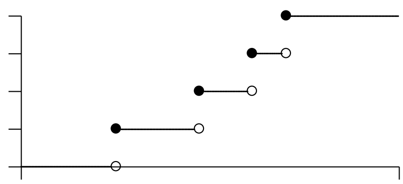

(iii)

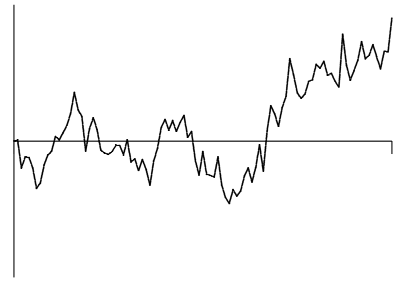

(iv)

Most of the interesting processes have continuous sample paths (e.g. (i),(ii), and (iv) above). If we want to allow jump discontinuities (as in (iii) above), we normally specify that the functions be *càdlàg* (continue à droite, limites à gauche) – they should be right-continuous, and have limits from the left, at every point.

2. Brownian motion.

The most interesting processes have something random happening at every time t. The archetype among these is *Brownian motion* $(W_t)_{0 \le t < \infty}$, a sort of random walk with infinitesimal steps. Its sample paths look like (iv) above.

- A Brownian motion has *independent*, *normal* increments:

$$W(t + \Delta t) - W(t) \sim N(0, \Delta t)$$

- Brownian sample paths are not differentiable. Note that the increment $W(t + \Delta t) - W(t)$ has magnitude of order $(\Delta t)^{1/2}$, rather than Δt as would be the case for a smooth function.

- Note that the normal is the only possible distribution for the increments, since each increment is the sum of many independent smaller ones. (Central Limit Theorem.)
- There are several ways to construct a Brownian motion. Here's one: let $\{Z_{ij}\}_{i,j=0}^{\infty}$ be i.i.d. standard normal random variables. Set $W(0) = 0$ and $W(1) = Z_{00}$. Then set $W(1/2) = (W(0) + W(1))/2 + 2^{-1/2}Z_{11}$. Then set $W(1/4) = (W(0)+W(1/2))/2+2^{-1}Z_{21}$ and $W(3/4) = (W(1/2)+W(1))/2 + 2^{-1}Z_{22}$. Then set $W(1/8) = (W(0) + W(1/4))/2 + 2^{-3/2}Z_{31}$, etc. After each step we can construct a piecewise linear approximation W_n of the Brownian motion on $[0, 1]$ by linearly interpolating between the values constructed so far. Then $W = \lim_n W_n$ gives the Brownian motion itself.

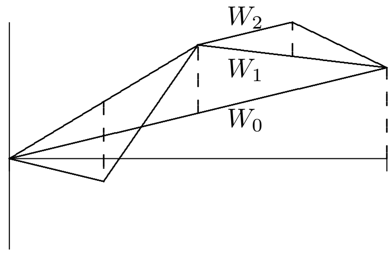

3. Itô diffusions.

Itô diffusions are the main objects of study in stochastic calculus. An Itô diffusion $(X_t)_{0\leq t<\infty}$ is a stochastic process described by a formal equation of the form:

$$dX_t = U_t dt \;+\; V_t dW_t$$

where (W_t) is a Brownian motion. Loosely speaking, this means that for small increments:

$$X(t + \Delta t) - X(t) \stackrel{\sim}{\approx} N(U_t\Delta t\;,\; V_t^2\Delta t),$$

where $\stackrel{\sim}{\approx}$ means "has approximately the distribution". This is a conditional distribution given the values of the processes (W, U, V, and X) on $[0, t]$.

- (U_t) is the *drift* process. It measures the expected rate of change of X_t in a similar way to the conventional derivative of a function.
- (V_t) is the *speed* process. It measures the amount of random diffusing around that is going on.
- Both (U_t) and (V_t) should be *adapted* or *non-anticipating* stochastic processes, meaning that U_t and V_t should not depend on what the driving process W does at times later than t (though they may depend on W_s for $0 \leq s \leq t$).
- Itô diffusions always have continuous sample paths.
- It is quite easy to construct a process (X_t) with $dX_t = U_t dt$: just set $X_t = \int_0^t U_s\, ds$. To construct (X_t) with $dX_t = V_t dW_t$ requires the stochastic integral (see below).

4. The stochastic integral.

For any adapted process (V_t), define its stochastic integral (X_t) with respect to a Brownian motion (W_t) by

$$\begin{aligned} X_t &= \int_0^t V_s\, dW_s \\ &= \lim_n \sum_{i=1}^{n} V(t_{i-1})\left(W(t_i) - W(t_{i-1})\right), \end{aligned}$$

the limit being taken as the partition $0 = t_0 < t_1 < \cdots < t_n = t$ is refined.

- This limit is slightly tricky – it may not exist for all (or even almost all) realizations of the sample paths of V and W. However, under suitable assumptions, the limit will always exist in quadratic mean and in probability. If V has sample paths of bounded variation, the limit does exist for almost all realizations of sample paths.

- It is significant that V is evaluated at the left-hand endpoint t_{i-1}, rather than any other point of $[t_{i-1}, t_i]$ (which would give different values for $\int_0^t V_s\, dW_s$). For example, replacing $V(t_{i-1})$ by $V(t_i)$ would be tantamount to adding in the term

$$\sum_{i=1}^{n} \left(V(t_i) - V(t_{i-1})\right)\left(W(t_i) - W(t_{i-1})\right),$$

which could be significant for the (very nonsmooth) functions we are dealing with here.
- Example: $\int_0^t W_s\, dW_s = (W_t^2 - t)/2$. Exercise for the reader: show this using the above limit definition.
- An Itô diffusion may be written in explicit integral form using a stochastic integral:

$$\begin{aligned} dX_t &= U_t dt \;+\; V_t dW_t \\ \Longleftrightarrow\; X_t &= X_0 + \int_0^t U_s\, ds + \int_0^t V_s\, dW_s \end{aligned}$$

5. Itô's lemma.

Itô's lemma is analogous to the chain rule for a change of variables in conventional calculus. For $X_t = h(t, W_t)$,

$$dX_t = h_t(t, W_t)dt + h_w(t, W_t)dW_t + \frac{1}{2}h_{ww}(t, W_t)dt,$$

or in integral form

$$X_t = X_0 + \int_0^t h_w(s, W_s)\, dW_s + \int_0^t \left(h_t(s, W_s) + \frac{1}{2}h_{ww}(s, W_s)\right) ds.$$

Here h_t, h_w, and h_{ww} are partial derivatives of h.

The most significant feature is the extra term involving h_{ww}, which arises because, loosely speaking, even a squared increment of Brownian motion is large enough to be significant: $(dW_t)^2 = dt$.

More generally, if (X_t) is an Itô diffusion with $dX_t = U_t dt + V_t dW_t$, and $Y_t = h(t, X_t)$, then

$$dY_t = h_t(t, X_t)dt + h_x(t, X_t)dX_t + \frac{1}{2}h_{xx}(t, X_t)\,(dX_t)^2,$$

where $(dX_t)^2 = V_t^2 dt$ is called the *quadratic variation* of X.

6. Stochastic differential equations.

A stochastic differential equation (SDE) for an Itô diffusion (X_t) takes the form

$$dX_t = f(t, X_t)dt \;+\; g(t, X_t)dW_t,$$

for some functions f and g.

- Note that $f(t, X_t)$ and $g(t, X_t)$ depend only on the current value of X_t, not past values. This means that the solution X_t will be a *Markov process*. The functions f and g are analogous to the transition probabilities of a Markov chain.
- Example. The main SDE of mathematical finance:

$$dX_t = rX_t dt \;+\; \sigma X_t dW_t.$$

This models the evolution of the price of a risky asset (e.g. a stock) with growth rate r (analogous to an interest rate) and volatility σ. We can solve this equation by trying for a solution of form $X_t = h(t, W_t)$: applying Itô's lemma shows that we need $h_w = \sigma h$ and $h_t + (1/2)h_{ww} = rh$, and this gives us

$$X_t = X_0 \exp\left(\sigma W_t + (r - (1/2)\sigma^2)t\right).$$

This process is called *geometric Brownian motion*.

- Example. The Ornstein-Uhlenbeck equation.

$$dX_t = -X_t dt \; + \; dW_t.$$

This is a mean-reverting process – the more different X_t is from 0, the more strongly it tends to move back towards 0. The equation has solution

$$X_t = e^{-t}\left(X_0 + \int_0^t e^s \, dW_s\right).$$

7. Martingales.

A stochastic process (X_t) is a *martingale* if

$$E\left[X_r \,|\, X_s, \; 0 \le s \le t\right] = X_t \qquad\qquad \text{for } t < r.$$

This means that the process has no average tendency to rise or fall. Martingales are often used as models for the fortunes of a player in a fair gambling game, who is as likely to lose as win. The main significance of being a martingale is that $E\left[X_t\right] = E\left[X_0\right]$ for all times t. This is even true of certain random times.

An Itô diffusion is a martingale iff it satisfies $dX_t = V_t dW_t$. (In particular, Brownian motion itself is a martingale.) This means

$$X_t = \int_0^t V_s \, dW_s \approx \sum_{i=1}^n V(t_{i-1})(W(t_i) - W(t_{i-1})),$$

so X_t can be thought of as the fortune of a gambler who wagers a variable amount $V(t_{i-1})$ on each increment $W(t_i) - W(t_{i-1})$, or an investor who owns V_t shares of a stock at each time t. For such a diffusion

$$E\left[X_t\right] = 0$$

$$E\left[X_t^2\right] = \int_0^t E\left[V_s^2\right] \, ds, \qquad \text{the } \textit{isometry property} \text{ of stochastic integration.}$$

8. Diffusions as limits of Markov chains.

In addition to being useful models themselves, Itô diffusions can also serve as approximations of the long-run behaviour of discrete processes.

Suppose that for each n, $(X_k^n)_{k=0}^\infty$ is a discrete-time Markov chain whose state space is (a subset of) the real numbers. If

$$E\left[X_{k+1}^n \,\middle|\, X_k^n = x\right] \approx x + \frac{1}{n} b(x)$$

$$E\left[(X_{k+1}^n - x)^2 \,\middle|\, X_k^n = x\right] \approx \frac{1}{n} \sigma(x)^2$$

then for large n, (X_k^n) behaves like the solution of the SDE

$$dX_t = b(X_t)dt \; + \; \sigma(X_t)dW_t,$$

with n steps per unit time.

- Example. The Ehrenfest urn. Two urns contain a total of $2n$ balls; at each step a ball is picked at random and moved to the other urn. Let N_k^n be the number of balls in the first urn after k steps, and $X_k^n = (N_k^n - n)/\sqrt{n}$. For large n, (X_k^n) behaves like the solution of $dX_t = -X_t dt + dW_t$, i.e. like the Ornstein-Uhlenbeck process.
- Example. The Wright-Fisher process. A population consists of n genes, which are of two types (A and B). At each step, a new population is formed by sampling with replacement from the old one. Let X_k^n be the proportion of type-A genes in the population after k steps. For large n, (X_k^n) behaves like the solution of $dX_t = \sqrt{X_t(1 - X_t)} dW_t$.

9. Multivariate stochastic calculus.

A vector-valued Itô diffusion $X_t = (X_1(t), X_2(t))$ can be driven by the usual Brownian motion W_t. This occurs, for example, in the following system of SDEs:

$$\begin{aligned} dX_1(t) &= X_2(t)dt \\ dX_2(t) &= -\alpha^2 X_1(t)dt \; + \; \sigma dW(t), \end{aligned}$$

which represents the motion of a "Brownian mass" on the end of a spring (X_1 =position, X_2 =velocity), or a noisy electronic oscillator.

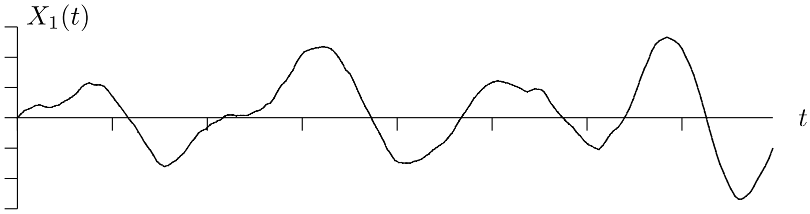

It is also possible to have vector-valued diffusions driven by more than one Brownian motion. As an example, consider the two-dimensional Brownian motion $W(t) = (W_1(t), W_2(t))$ itself, consisting of components W_1 and W_2 which are independent one-dimensional Brownian motions. It is a kind of continuous-time random walk in the plane.

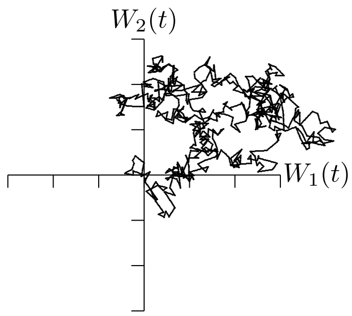

Similarly for Brownian motion in any number of dimensions.

10. References.

- K. L. Chung and R. Williams, *Introduction to stochastic integration*, 2nd ed. (1990). Birkhäuser.
- S. Ethier and T. Kurtz, *Markov processes: characterization and convergence* (1986). Wiley.
- I. Karatzas and S. Shreve, *Brownian motion and stochastic calculus* 2nd. ed. (1991). Springer.
- L. C. G. Rogers and D. Williams, *Diffusions, Markov processes, and martingales* (1987). Wiley.