m3f22l21tex
Lecture 21 24.11.2016

Theorem (PWZ theorem: Paley-Wiener-Zygmund, 1933). For $(Z_n)_0^\infty$ independent $N(0,1)$ random variables, λ_n, Δ_n as above,

$$W_t := \sum_{n=0}^{\infty} \lambda_n Z_n \Delta_n(t)$$

converges uniformly on $[0,1]$, a.s. The process $W = (W_t : t \in [0,1])$ is Brownian motion.

Thus the above description does indeed define a stochastic process $X = (X_t)_{t\in[0,1]}$ on $(C[0,1], \mathcal{F}, (\mathcal{F}_t), P)$. The construction gives X on $C[0,n]$ for each $n = 1, 2, \cdots$, and combining these: X exists on $C[0,\infty)$. It is also unique (a stochastic process is uniquely determined by its finite-dimensional distributions and the restriction to path-continuity).

No construction of Brownian motion is easy: one needs both some work and some knowledge of measure theory. But *existence* is really all we need, and we assume this. For background, see any measure-theoretic text on stochastic processes. The classic is Doob's book, quoted above (see VIII.2 there). Excellent modern texts include Karatzas & Shreve [KS] (see particularly §2.2-4 for construction and §5.8 for applications to economics), Revuz & Yor [RY], Rogers & Williams [RW1] (Ch. 1), [RW2] Itô calculus – below).

We denote standard Brownian motion $BM(\mathbb{R})$ – or just BM for short – by $B = (B_t)$ (B for Brown), though $W = (W_t)$ (W for Wiener) is also common. Standard Brownian motion $BM(\mathbb{R}^d)$ in d dimensions is defined by $B(t) := (B_1(t), \cdots, B_d(t))$, where $B_1, \cdots, B_d$ are *independent* standard Brownian motions in one dimension (*independent copies* of $BM(\mathbb{R})$).

Zeros.

It can be shown that Brownian motion *oscillates*:

$$\limsup_{t\to\infty} X_t = +\infty, \qquad \liminf_{t\to\infty} X_t = -\infty \qquad a.s.$$

Hence, for every n there are zeros (times t with $X_t = 0$) of X with $t \geq n$ (indeed, infinitely many such zeros). So if

$$Z := \{t \geq 0 : X_t = 0\}$$

denotes the zero-set of $BM(\mathbb{R})$:
1. Z is an *infinite* set.
Next, if t_n are zeros and $t_n \to t$, then by path-continuity $B(t_n) \to B(t)$; but $B(t_n) = 0$, so $B(t) = 0$:
2. Z is a *closed* set (Z contains its limit points).
Less obvious are the next two properties:
3. Z is a *perfect* set: every point $t \in Z$ is a limit point of points in Z. So there are *infinitely many* zeros in *every* neighbourhood of *every* zero (so the paths must oscillate amazingly fast!).
4. Z is a (Lebesgue) *null* set: Z has Lebesgue measure zero.

In particular, the diagram above (or any other diagram!) grossly distorts Z: *it is impossible to draw a realistic picture of a Brownian path.*

Brownian Scaling.

For each $c \in (0, \infty)$, $X(c^2t)$ is $N(0, c^2t)$, so $X_c(t) := c^{-1}X(c^2t)$ is $N(0, t)$. Thus X_c has all the defining properties of a Brownian motion (check). So, X_c **IS** a Brownian motion:

Theorem. If X is BM and $c > 0$, $X_c(t) := c^{-1}X(c^2t)$, then X_c is again a BM.

Corollary. X is *self-similar* (reproduces itself under scaling), so a Brownian path $X(.)$ is a *fractal.* So too is the zero-set Z.

Brownian motion owes part of its importance to belonging to *all* the important classes of stochastic processes: it is (strong) Markov, a (continuous) martingale, Gaussian, a diffusion, a Lévy process (process with stationary independent increments), etc.

§2. Filtrations; Finite-Dimensional Distributions

The underlying set-up is as before, but now *time is continuous rather than discrete*; thus the time-variable will be $t \geq 0$ in place of $n = 0, 1, 2, \ldots$. The information available at time t is the σ-field $\mathcal{F}_t$; the collection of these as $t \geq 0$ varies is the filtration, modelling the information flow. The underlying probability space, endowed with this filtration, gives us the stochastic basis (filtered probability space) on which we work,

We assume that the filtration is *complete* (contains all subsets of null-sets as null-sets), and *right-continuous*: $\mathcal{F}_t = \mathcal{F}_{t+}$, i.e.

$$\mathcal{F}_t = \cap_{s>t}\mathcal{F}_s$$

(the 'usual conditions' – right-continuity and completeness – in Meyer's terminology).

A stochastic process $X = (X_t)_{t\geq 0}$ is a family of random variables defined on a filtered probability space with X_t $\mathcal{F}_t$-measurable for each t: thus X_t is known when $\mathcal{F}_t$ is known, at time t.

If $\{t_1, \cdots, t_n\}$ is a finite set of time-points in $[0, \infty)$, $(X_{t_1}, \cdots, X_{t_n})$, or $(X(t_1), \cdots, X(t_n))$ (for typographical convenience, we use both notations interchangeably, with or without ω: $X_t(\omega)$, or $X(t, \omega)$) is a random n-vector, with a distribution, $\mu(t_1, \cdots, t_n)$ say. The class of all such distributions as $\{t_1, \cdots, t_n\}$ ranges over all finite subsets of $[0, \infty)$ is called the class of all *finite-dimensional distributions* of X. These satisfy certain obvious consistency conditions:
(i) deletion of one point t_i can be obtained by 'integrating out the unwanted variable', as usual when passing from joint to marginal distributions,
(ii) permutation of the t_i permutes the arguments of the measure $\mu(t_1, \cdots, t_n)$ on $\mathbb{R}^n$.
Conversely, a collection of finite-dimensional distributions satisfying these two consistency conditions arises from a stochastic process in this way (this is the content of the DANIELL-KOLMOGOROV Theorem: P. J. Daniell in 1918, A. N. Kolmogorov in 1933).

Important though it is as a general existence result, however, the Daniell-Kolmogorov theorem does not take us very far. It gives a stochastic process X as a random function on $[0, \infty)$, i.e. a random variable on $\mathbb{R}^{[0,\infty)}$. This is a vast and unwieldy space; we shall usually be able to confine attention to much smaller and more manageable spaces, of functions satisfying regularity conditions. The most important of these is *continuity*: we want to be able to realise $X = (X_t(\omega))_{t\geq 0}$ as a random *continuous* function, i.e. a member of $C[0, \infty)$; such a process X is called *path-continuous* (since the map $t \to X_t(\omega)$ is called the sample path, or simply path, given by ω) – or more briefly, *continuous*. This is possible for the extremely important case of *Brownian motion* (below), for example, and its relatives. Sometimes we need to allow our random function $X_t(\omega)$ to have jumps. It is then customary, and convenient, to require X_t to be *right-continuous with left limits* (rcll), or càdlàg (continu à droite, limite à gauche) – i.e. to have X in the space

$D[0, \infty)$ of all such functions (the *Skorohod space*). This is the case, for instance, for the *Poisson process* and its relatives.

General results on realisability – whether or not it is possible to *realise*, or obtain, a process so as to have its paths in a particular function space – are known, but it is usually better to *construct* the processes we need directly on the function space on which they naturally live.

Given a stochastic process X, it is sometimes possible to improve the regularity of its paths without changing its distribution (that is, without changing its finite-dimensional distributions). For background on results of this type (separability, measurability, versions, regularization, ...) see e.g. Doob's classic book [D].

The continuous-time theory is technically much harder than the discrete-time theory, for two reasons:
(i) questions of path-regularity arise in continuous time but not in discrete time,
(ii) *uncountable* operations (like taking sup over an interval) arise in continuous time. But measure theory is constructed using *countable* operations: uncountable operations risk losing measurability.

Filtrations and Insider Trading

Recall that a filtration models an information flow. In our context, this is the information flow on the basis of which market participants – traders, investors etc. – make their decisions, and commit their funds and effort. All this is information in the *public* domain – necessarily, as stock exchange prices are publicly quoted. Again necessarily, many people are involved in major business projects and decisions (an important example: mergers and acquisitions, or M&A) involving publicly quoted companies. Frequently, this involves price-sensitive information. People in this position are – rightly – prohibited by law from profiting by it directly, by trading on their own account, in publicly quoted stocks but using private information. This is rightly regarded as theft at the expense of the investing public.[1] Instead, those involved in M&A etc. should seek to benefit legitimately (and indirectly) – enhanced career prospects, commission or fees, bonuses etc.

[1] The plot of the film *Wall Street* revolves round such a case, and is based on real life – recommended!