liv8.tex Week 8: 17.3.2014

Chapter V. STOCHASTIC PROCESSES IN CONTINUOUS TIME

§1. 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.

The regulatory authorities (Securities and Exchange Commission – SEC – in US; Financial Conduct Authority (FCA) and Prudential Regulation Authority (PRA, part of the Bank of England (BoE) in UK) monitor all trading electronically. Their software alerts them to patterns of suspicious trades. The software design (necessarily secret, in view of its value to criminals) involves all the necessary elements of Mathematical Finance in exaggerated form: economic and financial insight, plus: mathematics; statistics (especially pattern recognition, data mining and machine learning); numerics and computation.

§2. Classes of Processes.

1. *Martingales.*

The martingale property in continuous time is just that suggested by the discrete-time case:

$$E[X_t|\mathcal{F}_s] = X_s \qquad (s < t),$$

and similarly for submartingales and supermartingales. There are regularization results, under which one can take X_t right-continuous in t. Among the contrasts with the discrete case, we mention that the Doob-Meyer decomposition, easy in discrete time (III.8), is a deep result in continuous time. For background, see e.g.

MEYER, P.-A. (1966): *Probabilities and potentials.* Blaisdell

- and subsequent work by Meyer and the French school (Dellacherie & Meyer, *Probabilités et potentiel*, I-V, etc.

2. Gaussian Processes.

Recall the multivariate normal distribution $N(\mu, \mathbf{\Sigma})$ in n dimensions. If $\mu \in \mathbb{R}^n$, $\mathbf{\Sigma}$ is a non-negative definite $n \times n$ matrix, $\mathbf{X}$ has distribution $N(\mu, \mathbf{\Sigma})$

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

if it has characteristic function

$$\phi_{\mathbf{X}}(\mathbf{t}) := E\exp\{i\mathbf{t}^T.\mathbf{X}\} = \exp\{i\mathbf{t}^T.\mu - \frac{1}{2}\mathbf{t}^T\mathbf{\Sigma}\mathbf{t}\} \qquad (\mathbf{t} \in \mathbb{R}^n).$$

If further $\mathbf{\Sigma}$ is positive definite (so non-singular), $\mathbf{X}$ has density

$$f_{\mathbf{X}}(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{1}{2}n}|\Sigma|^{\frac{1}{2}}}\exp\{-\frac{1}{2}(\mathbf{x}-\mu)^T\mathbf{\Sigma}^{-1}(\mathbf{x}-\mu)\}$$

(*Edgeworth's Theorem* of 1893: F. Y. Edgeworth (1845-1926), English statistician).

A process $X = (X_t)_{t\geq 0}$ is *Gaussian* if all its finite-dimensional distributions are Gaussian. Such a process can be specified by:
(i) a measurable function $\mu = \mu(t)$ with $EX_t = \mu(t)$,
(ii) a non-negative definite function $\sigma(s,t)$ with

$$\sigma(s,t) = cov(X_s, X_t).$$

Gaussian processes have many interesting properties. Among these, we quote *Belayev's dichotomy*: with probability one, the paths of a Gaussian process are either continuous, or extremely pathological: for example, unbounded above and below on any time-interval, however short. Naturally, we shall confine attention in this course to continuous Gaussian processes.

3. Markov Processes.

X is *Markov* if for each t, each $A \in \sigma(X_s : s > t)$ (the 'future') and $B \in \sigma(X_s : s < t)$ (the 'past'),

$$P(A|X_t, B) = P(A|X_t).$$

That is, if you know where you are (at time t), how you got there doesn't matter so far as predicting the future is concerned – equivalently, past and future are conditionally independent given the present.

The same definition applied to Markov processes in discrete time.

X is said to be *strong Markov* if the above holds with the *fixed* time t replaced by a *stopping time* T (a random variable). This is a real restriction of the Markov property in the continuous-time case (though not in discrete time).

4. Diffusions.

A diffusion is a path-continuous strong-Markov process such that for each time t and state x the following limits exist:

$$\mu(t,x) := \lim_{h\downarrow 0}\frac{1}{h}E[(X_{t+h}-X_t)|X_t = x],$$

$$\sigma^2(t,x) := \lim_{h\downarrow 0}\frac{1}{h}E[(X_{t+h}-X_t)^2|X_t = x].$$

Then $\mu(t,x)$ is called the *drift*, $\sigma^2(t,x)$ the *diffusion coefficient*.

§3. Brownian Motion.

The Scottish botanist Robert Brown observed pollen particles in suspension under a microscope in 1828 and 1829 (though others had observed the phenomenon before him),[2] and observed that they were in constant irregular motion.

In 1900 L. Bachelier considered Brownian motion a possible model for stock-market prices:
BACHELIER, L. (1900): Théorie de la spéculation. *Ann. Sci. Ecole Normale Supérieure* **17**, 21-86
– the first time Brownian motion had been used to model financial or economic phenomena, and before a mathematical theory had been developed.

In 1905 Albert Einstein considered Brownian motion as a model of particles in suspension, and used it to estimate *Avogadro's number* ($N \sim 6\times 10^{23}$), based on the diffusion coefficient D in the *Einstein relation*

$$var X_t = Dt \qquad (t>0).$$

In 1923 Norbert WIENER defined and constructed Brownian motion rigorously for the first time. The resulting stochastic process is often called the *Wiener process* in his honour, and its probability measure (on path-space) is called *Wiener measure*.

We define *standard Brownian motion* on $\mathbb{R}$, BM or $BM(\mathbb{R})$, to be a stochastic process $X = (X_t)_{t\geq 0}$ such that
1. $X_0 = 0$,
2. X has *independent increments*: $X_{t+u} - X_t$ is independent of $\sigma(X_s : s \leq t)$

[2]The Roman author Lucretius observed this phenomenon in the gaseous phase – dust particles dancing in sunbeams – in antiquity: *De rerum naturae*, c. 50 BC.

for $u \geq 0$,
3. X has *stationary increments*: the law of $X_{t+u} - X_t$ depends only on u,
4. X has *Gaussian increments*: $X_{t+u} - X_t$ is normally distributed with mean 0 and variance u,

$$X_{t+u} - X_t \sim N(0, u),$$

5. X has *continuous paths*: X_t is a continuous function of t, i.e. $t \mapsto X_t$ is continuous in t.

For time t in a finite interval – $[0, 1]$, say – we can use the following filtered space:
$\Omega = C[0, 1]$, the space of all continuous functions on $[0, 1]$.
The points $\omega \in \Omega$ are thus random functions, and we use the coordinate mappings: X_t, or $X_t(\omega)$, $= \omega_t$.
The filtration is given by $\mathcal{F}_t := \sigma(X_s : 0 \leq s \leq t)$, $\mathcal{F} := \mathcal{F}_1$.
P is the measure on $(\Omega, \mathcal{F})$ with finite-dimensional distributions specified by the restriction that the increments $X_{t+u} - X_t$ are stationary independent Gaussian $N(0, u)$.

Theorem (WIENER, 1923). Brownian motion exists.

The best way to prove this is by construction, and one that reveals some properties. The proof that follows is originally due to Paley, Wiener and Zygmund (1933) and Lévy (1948), but is re-written in the modern language of *wavelet* expansions. We omit details; for these, see e.g. [BK] 5.3.1, or SP I20-22. The Haar system $(H_n) = (H_n(.))$ is a complete orthonormal system (cons) of functions in $L^2[0, 1]$. The Schauder System Δ_n) is obtained by integrating the Haar system. Consider the triangular function (or 'tent function')

$$\Delta(t) = \begin{cases} 2t & \text{on } [0, \frac{1}{2}), \\ 2(1-t) & \text{on } [\frac{1}{2}, 1], \\ 0 & \text{else.} \end{cases}$$

Write $\Delta_0(t) := t$, $\Delta_1(t) := \Delta(t)$, and define the nth *Schauder function* Δ_n by

$$\Delta_n(t) := \Delta(2^j t - k) \qquad (n = 2^j + k \geq 1).$$

Note that Δ_n has support $[k/2^j, (k+1)/2^j]$ (so is 'localized' on this dyadic interval, which is small for n, j large). We see that

$$\int_0^t H(u)du = \frac{1}{2}\Delta(t),$$

and similarly

$$\int_0^t H_n(u)du = \lambda_n \Delta_n(t),$$

where $\lambda_0 = 1$ and for $n \geq 1$,

$$\lambda_n = \frac{1}{2} \times 2^{-j/2} \qquad (n = 2^j + k \geq 1).$$

The Schauder system (Δ_n) is again a complete orthogonal system on $L^2[0,1]$. We can now formulate the next result; for proof, see the references above.

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. However, *existence* is really all we need, and this we shall take for granted. 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 shall henceforth 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 $c \in (0, \infty)$, $X(c^2 t)$ is $N(0, c^2 t)$, so $X_c(t) := c^{-1} X(c^2 t)$ 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 $c > 0$, $X_c(t) := c^{-1} X(c^2 t)$, 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.

§4. Quadratic Variation of Brownian Motion

Recall that for $\xi\ N(\mu, \sigma^2)$, ξ has moment-generating function (MGF)

$$M(t) := E \exp\{t\xi\} = \exp\{\mu t + \frac{1}{2}\sigma^2 t^2\}.$$

We take $\mu = 0$ below; we can recover the general case by adding μ back on. So, for $\xi\ N(0, \sigma^2)$,

$$\begin{aligned} M(t) := E \exp\{t\xi\} &= \exp\{\frac{1}{2}\sigma^2 t^2\} \\ &= 1 + \frac{1}{2}\sigma^2 t^2 + \frac{1}{2!}\left(\frac{1}{2}\sigma^2 t^2\right)^2 + O(t^6) \\ &= 1 + \frac{1}{2!}\sigma^2 t^2 + \frac{3}{4!}\sigma^4 t^4 + O(t^6). \end{aligned}$$

So as the Taylor coefficients of the MGF are the moments (hence the name MGF!),

$$E(\xi^2) = var\xi = \sigma^2, \qquad E(\xi^4) = 3\sigma^4, \qquad \text{so} \qquad var(\xi^2) = E(\xi^4) - [E(\xi^2)]^2 = 2\sigma^4.$$

For $B\ BM$, this gives in particular

$$EB_t = 0, \qquad varB_t = t, \qquad E[(B_t)^2] = t, \qquad var[(B_t)^2] = 2t^2.$$

In particular, for $t > 0$ *small*, this shows that the variance of B_t^2 is negligible compared with its expected value. Thus, the randomness in B_t^2 is negligible compared to its mean for t small.

This suggests that if we take a fine enough partition $\mathcal{P}$ of $[0, T]$ – a finite set of points

$$0 = t_0 < t_1 < \cdots < t_k = T$$

with $|\mathcal{P}| := \max |t_i - t_{i-1}|$ small enough – then writing

$$\Delta B(t_i) := B(t_i) - B(t_{i-1}), \qquad \Delta t_i := t_i - t_{i-1},$$

$\Sigma(\Delta B(t_i))^2$ will closely resemble $\Sigma E[(\Delta B(t_i)^2]$, which is $\Sigma \Delta t_i = \Sigma(t_i - t_{i-1}) = T$. This is in fact true a.s.:

$$\Sigma(\Delta B(t_i))^2 \to \Sigma \Delta t_i = T \qquad \text{as} \qquad \max |t_i - t_{i-1}| \to 0.$$

This limit is called the *quadratic variation* V_T^2 of B over $[0, T]$:

Theorem. The quadratic variation of a Brownian path over $[0, T]$ exists and equals T, a.s.

For details of the proof, see e.g. [BK], §5.3.2, SP L22, SA L7,8.

If we increase t by a small amount to $t + dt$, the increase in the quadratic variation can be written symbolically as $(dB_t)^2$, and the increase in t is dt. So, formally we may summarise the theorem as

$$(dB_t)^2 = dt.$$

Suppose now we look at the *ordinary* variation $\Sigma|\Delta B_t|$, rather than the *quadratic* variation $\Sigma(\Delta B_t)^2$. Then instead of $\Sigma(\Delta B_t)^2 \sim \Sigma \Delta t \sim t$, we get $\Sigma|\Delta B_t| \sim \Sigma\sqrt{\Delta t}$. Now for Δt small, $\sqrt{\Delta t}$ is of a larger order of magnitude that Δt. So if $\Sigma \Delta t = t$ converges, $\Sigma\sqrt{\Delta t}$ diverges to $+\infty$. This suggests – what is in fact true – the

Corollary. The paths of Brownian motion are of infinite variation - their variation is $+\infty$ on every interval, a.s.

The quadratic variation result above leads to Lévy's 1948 result, the Martingale Characterization of Brownian Motion. Recall that B_t is a continuous martingale with respect to its natural filtration $(\mathcal{F}_t)$ and with quadratic variation t. There is a remarkable converse:

Theorem (Lévy; Martingale Characterization of Brownian Motion). If M is any continuous local $(\mathcal{F}_t)$-martingale with $M_0 = 0$ and quadratic variation t, then M is an $(\mathcal{F}_t)$-Brownian motion.

This can be expressed differently:

Theorem (Lévy). If M is any continuous $(\mathcal{F}_t)$-martingale with $M_0 = 0$ and $M_t^2 - t$ a martingale, then M is an $(\mathcal{F}_t)$-Brownian motion.

For proof, see e.g. [RW1], I.2. Observe that for $s < t$,

$$B_t^2 = [B_s + (B_t - B_s)]^2 = B_s^2 + 2B_s(B_t - B_s) + (B_t - B_s)^2,$$

$$E[B_t^2|\mathcal{F}_s] = B_s^2 + 2B_s E[(B_t - B_s)|\mathcal{F}_s] + E[(B_t - B_s)^2|\mathcal{F}_s] = B_s^2 + 0 + (t - s) :$$

$$E[B_t^2 - t|\mathcal{F}_s] = B_s^2 - s :$$

$B_t^2 - t$ is a martingale.

Quadratic Variation.
The theory above extends to *continuous* martingales (bounded continuous martingales in general, but we work on a finite time-interval $[0, T]$, so continuity implies boundedness). We quote:

Theorem. A continuous martingale M is of finite quadratic variation $\langle M \rangle$, and $\langle M \rangle$ is the unique continuous increasing adapted process vanishing at zero with $M^2 - \langle M \rangle$ a martingale.

Corollary. A continuous martingale M has infinite variation.

For proof, see e.g. [RY], IV.1.

Quadratic Covariation. We write $\langle M, M \rangle$ for $\langle M \rangle$, and extend $\langle \ \rangle$ to a bilinear form $\langle ., . \rangle$ with two different arguments by the *polarization identity*:

$$\langle M, N \rangle := \frac{1}{4}(\langle M + N, M + N \rangle - \langle M - N, M - N \rangle).$$

If N is of *finite* variation, $M \pm N$ has the same quadratic variation as M, so $\langle M, N \rangle = 0$.

Itô's Lemma. We discuss Itô's Lemma in more detail in §6 below; we pause here to give the link with quadratic variation and covariation. We quote: if $f(t, x_1, \cdots, x_d)$ is C^1 in its zeroth (time) argument t and C^2 in its remaining d space arguments x_i, and $M = (M^1, \cdots, M^d)$ is a continuous vector martingale, then (writing f_i, f_{ij} for the first partial derivatives of f with respect to its ith argument and the second partial derivatives with respect to the ith and jth arguments) $f(M_t)$ has stochastic differential

$$df(M_t) = f_0(M)dt + \Sigma_{i=1}^d f_i(M_t)dM_t^i + \frac{1}{2}\Sigma_{i,j=1}^d f_{ij}(M_t)d\langle M^i, M^j\rangle_t.$$

Integration by Parts. If $f(t, x_1, x_2) = x_1 x_2$, we obtain

$$d(MN)_t = NdM_t + MdN_t + \frac{1}{2}\langle M, N\rangle_t.$$

Note. The integration-by-parts formula – a special case of Itô's Lemma, as above – is in fact *equivalent* to Itô's Lemma: either can be used to derive the other. Rogers & Williams [RW1, IV.32.4] describe the integration-by-parts formula as 'the cornerstone of stochastic calculus'; this description may also be applied to Itô's Lemma.

We shall need to extend quadratic variation and quadratic covariation to stochastic integrals (to be defined below). If

$$Z = \int HdM, \qquad dZ = HdM,$$

$$d\langle Z\rangle = (dZ)^2 = H^2(dM)^2 = H^2 d\langle M\rangle.$$

Similarly (or by polarization), if $Z_i = \int H_i dM_i$ $(i = 1, 2)$,

$$d\langle Z_1, Z_2\rangle = H_1 H_2 d\langle M_1, M_2\rangle.$$

Fractals Everywhere.

As we saw, a Brownian path is a *fractal* – a *self-similar* object. So too is its zero-set Z. Fractals were studied, named and popularised by the French mathematician Benôit B. Mandelbrot (1924-2010). See his books, and Michael F. Barnsley: *Fractals everywhere*. Academic Press, 1988.

Fractals *look the same at all scales*. This is diametrically opposite to the familiar functions of Calculus. In Differential Calculus, a differentiable function has a tangent; this means that locally, its graph *looks straight*; similarly in Integral Calculus.

While most continuous functions we encounter are differentiable, at least piecewise (i.e., except for 'kinks'), there is a sense in which the typical, or generic, continuous function is *nowhere differentiable*. Thus Brownian paths may look pathological at first sight – but in fact they are typical!