Cambridge University Press
978-0-521-77594-6 - Diffusions, Markov Processes, and Martingales Volume 1: Foundations, 2nd Edition
L. C. G. Rogers and David Williams
Excerpt
More information

CHAPTER I

Brownian Motion

1. INTRODUCTION

1. What is Brownian motion, and why study it? The first thing is to define Brownian motion. We assume given some probability triple $(\Omega, \mathscr{F}, \mathbf{P})$.

(*1.1*) DEFINITION. *A real-valued stochastic process* $\{B_t : t \in \mathbf{R}^+\}$ *is a* Brownian motion *if it has the properties*

(1.2) (i) $B_0(\omega) = 0, \ \forall \omega$;
(1.2) (ii) *the map* $t \mapsto B_t(\omega)$ *is a continuous function of* $t \in \mathbf{R}^+$ *for all* ω;
(1.2)(iii) *for every* $t, h \geqslant 0$, $B_{t+h} - B_t$ *is independent of* $\{B_u : 0 \leqslant u \leqslant t\}$, *and has a Gaussian distribution with mean* 0 *and variance* h.

The conditions (1.2)(ii) and (1.2)(iii) are the really essential ones; if $B = \{B_t : t \in \mathbf{R}^+\}$ is a Brownian motion, we frequently speak of $\{\xi + B_t : t \in \mathbf{R}^+\}$ as a Brownian motion (started at ξ); the starting point ξ can be a fixed real, or a random variable independent of B.

Now that we know what a Brownian motion is, questions of existence and uniqueness (answered in Section 6) are less important than an answer to the second question of the title, 'Why study it?' There are many answers to this question, but to us there seem to be four main ones:

(i) Virtually every interesting class of processes contains Brownian motion—Brownian motion is a martingale, a Gaussian process, a Markov process, a diffusion, a Lévy process, . . . ;
(ii) Brownian motion is sufficiently concrete that one can do explicit calculations, which are impossible for more general objects;
(iii) Brownian motion can be used as a building block for other processes (indeed, a number of the most important results on Brownian motion state that the most general process in a certain class can be obtained from Brownian motion by some sequence of transformations);
(iv) last but not least, Brownian motion is a rich and beautiful mathematical object in its own right.

The aim of this chapter is to expand on these reasons, and convince you that

Brownian motion is indeed worthy of study; and the rest of this introduction gives a brief outline of some of the main points of the chapter.

2. Brownian motion as a martingale. Let $\{B_t : t \geqslant 0\}$ be a Brownian motion, and define $\mathscr{B}_t = \sigma(\{B_s : s \leqslant t\})$. *Then* $(B_t, \mathscr{B}_t)_{t \geqslant 0}$ *is a martingale.* We shall have a lot more to say about martingales in Chapter II, but for now we need little of the theory developed there. Let us just check that $(B_t, \mathscr{B}_t)_{t \geqslant 0}$ *is* a martingale (cf. Section II.63); first, $B_t \in L^1$ for all t, because, from (1.2)(i) and (1.2)(iii), $B_t \sim N(0, t)$, and, secondly, for $0 \leqslant s \leqslant t$,

$$\mathbf{E}[B_t - B_s \mid \mathscr{B}_s] = 0, \qquad \text{equivalently, } \mathbf{E}[B_t \mid \mathscr{B}_s] = B_s,$$

since $B_t - B_s$ is independent of $\mathscr{B}_s$ by (1.2)(iii). Likewise, since $B_t - B_s \sim N(0, t-s)$ independently of $\mathscr{B}_s$, we have

$$\mathbf{E}[(B_t - B_s)^2 \mid \mathscr{B}_s] = t - s.$$

But

$$\mathbf{E}[(B_t - B_s)^2 \mid \mathscr{B}_s] = \mathbf{E}[B_t^2 - 2B_tB_s + B_s^2 \mid \mathscr{B}_s] = \mathbf{E}[B_t^2 \mid \mathscr{B}_s] - B_s^2,$$

using properties of conditional expectation (Section II.41), so since we have (almost surely) that $\mathbf{E}[B_t^2 - t \mid \mathscr{B}_s] = B_s^2 - s$, we conclude that

(2.1) $$B_t^2 - t \quad \textit{is a martingale.}$$

This simple fact is a pointer to the development of stochastic integrals; once that theory is developed, we shall be in a position to prove the following startling converse to (2.1).

(2.2) THEOREM. (Lévy) *Let* $(X_t)_{t \geqslant 0}$ *be a continuous martingale,* $X_0 = 0$, *and suppose that*

$$X_t^2 - t \quad \textit{is a martingale.}$$

Then X *is a Brownian motion.*

By a continuous martingale, we mean of course one such that $t \mapsto X_t(\omega)$ is a continuous map for all ω. We have not been too specific about the filtration $(\mathscr{F}_t)_{t \geqslant 0}$ with respect to which X is a martingale, but this is not necessary; if X is a martingale with respect to $(\mathscr{F}_t)_{t \geqslant 0}$, and satisfies the hypotheses of Theorem 2.2 then X is an $(\mathscr{F}_t)$ Brownian motion—that is, X satisfies (1.2)(i), (1.2)(ii) and the stronger condition

(1.2)(iii)′ *for any* $t, h \geqslant 0$, $X_{t+h} - X_t$ *is independent of* $\mathscr{F}_t$ *and has a Gaussian distribution with mean zero and variance* h.

The Kunita–Watanabe proof of Theorem 2.2 is given in Section IV.33; a more elementary proof without using stochastic calculus appears in Doob [1].

A remarkable consequence of Theorem 2.2 is that

(2.3) *every continuous martingale is a time-change of Brownian motion.*

For a statement and proof of this, see Section IV.34. One extremely useful consequence is that, since

$$\mathbf{P}\left[\limsup_{t\to\infty} B_t = +\infty,\ \liminf_{t\to\infty} B_t = -\infty\right] = 1$$

(as we shall see in Lemma 3.6), if X is a continuous martingale for which $\mathbf{P}(\liminf X_t = -\infty) > 0$ then we must have $\mathbf{P}(\limsup X_t = +\infty) > 0$. See Section IV.34 for a full discussion.

The elementary arguments that gave (2.1) also show that for any $\theta \in \mathbb{R}$ (or indeed, for $\theta \in \mathbb{C}$)

(2.4) $$\exp(\theta B_t - \tfrac{1}{2}\theta^2 t) \quad \textit{is a martingale};$$

all one needs is that $\mathbf{E}(\exp[\theta(B_t - B_s)]) = \exp[\frac{1}{2}\theta^2(t-s)]$ for $0 \leqslant s \leqslant t$, which is just the moment-generating function of a Gaussian distribution. These exponential martingales are extremely useful in many ways; in Section 9 we use them to compute the Brownian first-passage distribution to a level, and in Section 16 we derive the Law of the Iterated Logarithm using them.

One small point to note here in connection with the exponential martingales (2.4) is that if we define the Hermite polynomials $H_n(t, x)$ by

$$\exp(\theta x - \tfrac{1}{2}\theta^2 t) := \sum_{n \geqslant 0} \frac{\theta^n}{n!} H_n(t, x),$$

then, for $0 \leqslant s \leqslant t$,

$$\begin{aligned}\mathbf{E}(\exp(\theta B_t - \tfrac{1}{2}\theta^2 t) \mid \mathscr{B}_s) &= \sum_{n \geqslant 0} \frac{\theta^n}{n!} \mathbf{E}(H_n(t, B_t) \mid \mathscr{B}_s)\\ &= \exp(\theta B_s - \tfrac{1}{2}\theta^2 s)\\ &= \sum_{n \geqslant 0} \frac{\theta^n}{n!} H_n(s, B_s),\end{aligned}$$

so, by comparing coefficients of θ^n, we deduce that

$$H_n(t, B_t) \quad \textit{is a martingale for each } n.$$

It is easy to check that $H_1(t, x) = x$ and $H_2(t, x) = x^2 - t$, so, in particular, (2.4)$\Rightarrow$(2.1); Lévy's Theorem 2.2 is essentially the converse to this.

(2.5) *Remark.* If $(N_t)_{t \geqslant 0}$ is a standard Poisson process then $X_t := N_t - t$ satisfies all of the hypotheses of Theorem 2.2 except for continuity of the paths.

3. Brownian motion as a Gaussian process. In complete generality, a (real-valued) process $(X_t)_{t \in T}$ indexed by some set T is said to be a *Gaussian process* if, for any $t_1, \ldots, t_n \in T$, the law of $(X(t_1), \ldots, X(t_n))$ is multivariate Gaussian. Thus the law of the process X is specified by the functions

$$\mu(t) := \mathbf{E}X_t, \quad \rho(s,t) := \operatorname{cov}(X_s, X_t).$$

(By this, we mean no more than that if we were told μ and ρ, we could work out the law of $(X(t_1), \ldots, X(t_n))$ for any $t_1, \ldots, t_n \in T$.) In the study of Gaussian processes, one usually assumes that $\mu \equiv 0$, to which the general case can be reduced by considering the Gaussian process $X_t - \mu(t)$.

It is obvious that $(B_t)_{t \geqslant 0}$ is a Gaussian process, with mean zero, and covariance

(3.1) $$\rho(s,t) = s \wedge t \quad (s, t \geqslant 0).$$

Any continuous real-valued process $(X_t)_{t \geqslant 0}$ that is a zero-mean Gaussian process with covariance (3.1) *is* a Brownian motion—just check the definition! This simple fact turns out to be an extremely efficient means of checking when a process is a Brownian motion, and the following four simple but extremely important examples serve to illustrate this:

(3.2) the process $(-B_t)_{t \geqslant 0}$ is a Brownian motion;
(3.3) for any $a \geqslant 0$, the process $(B_{t+a} - B_a)_{t \geqslant 0}$ is a Brownian motion;
(3.4) for any $c \neq 0$, $(cB_{t/c^2})_{t \geqslant 0}$ is a Brownian motion (*Brownian scaling*);
(3.5) the process $(\tilde{B}_t)_{t \geqslant 0}$ defined by

$$\tilde{B}_0 = 0,$$
$$\tilde{B}_t = tB_{1/t} \quad \text{for } t > 0,$$

is a Brownian motion.

The proofs of these properties are trivial exercises, with the sole exception of the proof of continuity at 0 of $\tilde{B}$. But this is not difficult, because the event that $\tilde{B} \to 0$ at 0 is

$$\tilde{F} = \bigcap_n \bigcup_m \bigcap_{q \in \mathbb{Q} \cap (0, 1/m]} \left\{ |\tilde{B}_q| \leqslant \frac{1}{n} \right\},$$

since $\tilde{B}$ is certainly continuous in $(0, \infty)$. But the processes $(\tilde{B}_t)_{t>0}$ and $(B_t)_{t>0}$ are continuous, and have the same distribution (they are Gaussian processes with the same covariance!), so

$$\mathbf{P}(\tilde{F}) = \mathbf{P}(F) = 1,$$

where F is the event $\bigcap_n \bigcup_m \bigcap_{q \in \mathbb{Q} \cap (0,1/m]} \{|B_q| \leqslant 1/n\}$ that $B \to 0$ at 0, which, by the definition of B, is certain.

The most important by far of the properties (3.2)–(3.5) is the *Brownian scaling* property (3.4). We shall give here an easy but striking consequence.

(3.6) LEMMA. *We have*

$$\mathbf{P}\left(\sup_t B_t = +\infty, \quad \inf_t B_t = -\infty \right) = 1.$$

Proof. Let $Z := \sup_t B_t$. By Brownian scaling, for any $c > 0$, we have

$$cZ \overset{\mathcal{D}}{=} Z,$$

so the law of Z is concentrated on $\{0, +\infty\}$. Let $p = \mathbf{P}(Z = 0)$. Then

$$\mathbf{P}(Z = 0) \leqslant \mathbf{P}[B_1 \leqslant 0 \quad \text{and} \quad B_u \leqslant 0 \quad \text{for all} \quad u \geqslant 1]$$

$$= \mathbf{P}\left[B_1 \leqslant 0 \quad \text{and} \quad \sup_{t\geqslant 0}\{B_{1+t} - B_1\} = 0\right],$$

because $(B_{1+t} - B_1)_{t\geqslant 0}$ is a Brownian motion, whose supremum is therefore 0 or $+\infty$. But $(B_{1+t} - B_1)_{t\geqslant 0}$ is independent of $(B_u)_{u\leqslant 1}$, so we deduce that

$$p = \mathbf{P}(Z = 0) \leqslant \mathbf{P}(B_1 \leqslant 0)\mathbf{P}(Z = 0) = \tfrac{1}{2}p,$$

whence $p = 0$. Combining with (3.2) gives the stated result. □

Lemma 3.6 implies straight away that, almost surely, *for each* $a \in \mathbf{R}$, $\{t: B_t = a\}$ *is not bounded above*. Thus Brownian motion is recurrent—it keeps returning to its starting point.

We shall have more to say about Gaussian processes in Part 4 of this chapter, but point out now that the discussion there is by way of an interesting digression from our main theme; the general setting for Gaussian processes is too general to permit full exploitation of the special features of Brownian motion (notably a completely ordered index set).

4. Brownian motion as a Markov process. Brownian motion is a (time-homogeneous) Markov process; for any bounded Borel $f: \mathbf{R} \to \mathbf{R}$, and $s, t \geqslant 0$,

(4.1) $$\mathbf{E}[f(B_{t+s})|\mathscr{B}_s] = P_t f(B_s)$$

where the *transition semigroup* $(P_t)_{t\geqslant 0}$ is defined by

$$P_t f(x) := \begin{cases} \int_{-\infty}^{\infty} p_t(x, y) f(y)\, dy & (t > 0), \\ f(x) & (t = 0) \end{cases}$$

where

(4.2) $$p_t(x, y) := (2\pi t)^{-1/2} \exp\left[-\frac{(x - y)^2}{2t}\right]$$

is the Brownian *transition density*. The Markov property (4.1) is immediate from the definition of Brownian motion. It is easy to confirm that $(P_t)_{t\geqslant 0}$ is a semigroup:

(4.3) $$P_{t+s} = P_t P_s = P_s P_t \quad (s, t \geqslant 0),$$

the so-called *Chapman–Kolmogorov equations*. The semigroup property (4.3) suggests that we ought in some sense to have

(4.4) $$\frac{d}{dt}P_t = \lim_{s\downarrow 0}\frac{1}{s}(P_{t+s} - P_t) = P_t\mathscr{G} = \mathscr{G}P_t,$$

where

(4.5) $$\mathscr{G} := \lim_{s\downarrow 0}\frac{1}{s}(P_s - I)$$

is the *(infinitesimal) generator of* $(P_t)_{t\geq 0}$. This is indeed true in complete generality, when suitably interpreted; the suitable interpretation involves us in some fairly careful analysis, because in general $\mathscr{G}$ is not defined for all functions, and much of the classical early work on Markov processes struggled with these technicalities. This functional-analytic viewpoint has many merits, not least that it can suggest quickly what things are likely to be true, but we shall not stress it too much because it is not a very convenient framework in which to prove the conjectures to which it leads. But, for now, let us illustrate the notion by working out the generator of Brownian motion. From (4.5), we should define $\mathscr{G}f$ for suitable f by

$$\mathscr{G}f := \lim_{t\downarrow 0}\frac{1}{t}(P_t f - f),$$

and, indeed, if $f \in C_b^2(\mathbf{R})$ then

$$\lim_{t\downarrow 0}\frac{1}{t}(P_t f - f)(x) = \lim_{t\downarrow 0}\int_{-\infty}^{\infty}\frac{f(x+y\sqrt{t}) - f(x)}{t}\exp(-\tfrac{1}{2}y^2)\frac{dy}{\sqrt{2\pi}}$$

$$= \lim_{t\downarrow 0}\int_{-\infty}^{\infty}\frac{1}{t}\{y\sqrt{t}f'(x) + \tfrac{1}{2}y^2 t f''(x+\theta y\sqrt{t})\}\exp(-\tfrac{1}{2}y^2)\frac{dy}{\sqrt{2\pi}}$$

(where $\theta \in (0, 1)$ depends on $y\sqrt{t}$)

$$= \tfrac{1}{2}f''(x).$$

Thus the infinitesimal generator of Brownian motion is

$$\mathscr{G} = \frac{1}{2}\frac{d^2}{dx^2},$$

at least when applied to $C_b^2(\mathbf{R})$. From (4.4), we find that, for $f \in C_b^2(\mathbf{R})$,

$$\frac{\partial}{\partial t}P_t f(x) = \mathscr{G}P_t f(x) = \tfrac{1}{2}(P_t f)''(x),$$

which leads to *Kolmogorov's backward equation* for the Brownian transition density:

(4.6) $$\frac{\partial}{\partial t}p_t(x, y) = \frac{1}{2}\frac{\partial^2}{\partial x^2}p_t(x, y),$$

since f is arbitrary. Using the other part of (4.4) gives us

$$\frac{\partial}{\partial t} P_t f(x) = P_t \mathscr{G} f(x) = \tfrac{1}{2} P_t f''(x),$$

and an integration by parts now yields *Kolmogorov's forward equation* for the Brownian transition density:

$$\text{(4.7)} \qquad \frac{\partial}{\partial t} p_t(x, y) = \frac{1}{2} \frac{\partial^2}{\partial y^2} p_t(x, y).$$

This equation is familiar in physics, where it is known as the *heat equation*, or the *diffusion equation*, so called because it determines the physical flow of heat, or the physical diffusion of particles in solution, in a homogeneous medium. Many of the notions of diffusion that probabilists use everyday were known to physicists long ago, and amount to the same things in different language (see, for example, the classic book by Crank [1] for a physicists' exposition—and a broad selection of fascinating and challenging questions). It is however important to stress that we are not simply going to be rederiving results well known in physics; probability provides techniques for the study of *individual* diffusing particles, which are far more flexible and powerful than the classical analysis of the heat equation, which is only a statement about the *average* behaviour of a large number of diffusing particles.

5. Brownian motion as a diffusion (and martingale). Without trying to be too precise, a diffusion (on the real line for now) is a continuous time-homogeneous Markov process X that is 'characterised' in some sense by its local infinitesimal *drift* b and *variance* a: for small h,

$$\text{(5.1) (i)} \qquad \mathbf{E}[X_{t+h} - X_t | \mathscr{F}_t] \doteqdot hb(X_t),$$

$$\text{(5.1)(ii)} \qquad \mathbf{E}[\{X_{t+h} - X_t - hb(X_t)\}^2 | \mathscr{F}_t] \doteqdot ha(X_t).$$

If a and b were constant functions then

$$X_t = \sigma B_t + bt, \qquad \sigma := a^{1/2},$$

would satisfy the description (5.1); the more general diffusion is rather similar except that the drift and variance may now depend on position. It is unnecessary to impose conditions on moments of the increment $X_{t+h} - X_t$ beyond the second, which you will certainly accept as plausible if you recall Lévy's Theorem (2.2), which said that Brownian motion is characterised by $b \equiv 0$, $a \equiv 1$.

Broadly speaking, there are three approaches to diffusions: the stochastic differential equation (SDE) approach, the martingale-problem approach, and the partial differential equation (PDE) approach. Each has its merits and peculiar techniques.

The SDE approach constructs the diffusion X with given infinitesimal characteristics a and b by solving

(5.2) $$X_t = X_0 + \int_0^t \sigma(X_s)\,dB_s + \int_0^t b(X_s)\,ds,$$

where $\sigma := a^{1/2}$, an equation that is commonly written in 'differential' form

(5.3) $$dX_t = \sigma(X_t)\,dB_t + b(X_t)\,dt.$$

Thus X has infinitesimal drift b and infinitesimal variance, a, since the increment $X_{t+h} - X_t$ is (approximately)

$$\sigma(X_t)(B_{t+h} - B_t) + hb(X_t).$$

There is a lot of work involved in defining what the second term on the right-hand side of (5.2) means, and in verifying existence and uniqueness of a solution under suitable conditions on σ and b; we shall have almost nothing to say on this until Volume 2.

The martingale problem approach and the PDE approach both begin from the same trivial calculation based on (5.1). For any $f \in C_b^2$,

(5.4) $$\begin{aligned} &\mathbf{E}[f(X_{t+h}) - f(X_t)|\mathscr{F}_t] \\ &= \mathbf{E}[f'(X_t)(X_{t+h} - X_t) + \tfrac{1}{2}f''(\theta X_{t+h} + (1-\theta)X_t)(X_{t+h} - X_t)^2|\mathscr{F}_t] \\ &\qquad\qquad\qquad\qquad \text{(where } \theta \in (0,1) \text{ is random)} \\ &\doteqdot f'(X_t)hb(X_t) + \tfrac{1}{2}f''(X_t)[ha(X_t) + h^2 b(X_t)^2] \\ &= h\mathscr{L}f(X_t) + O(h^2), \end{aligned}$$

where $\mathscr{L}$ is the second-order elliptic operator,

(5.5) $$\mathscr{L}f(x) := \tfrac{1}{2}a(x)\frac{d^2 f}{dx^2}(x) + b(x)\frac{df}{dx}(x).$$

The martingale-problem approach takes (5.4) and re-expresses it as

$$\mathbf{E}\left[f(X_{t+h}) - f(X_t) - \int_t^{t+h} \mathscr{L}f(X_s)\,ds \,\middle|\, \mathscr{F}_t \right] = o(h),$$

so that the martingale-problem 'definition' of a diffusion X with drift b and variance a is that X is a continuous process such that, for all $f \in C_b^2$,

(5.6) $$f(X_t) - \int_0^t \mathscr{L}f(X_s)\,ds \quad \textit{is a martingale.}$$

The PDE approach takes expectations on both sides of (5.4) to get

(5.7) $$P_h f(x) - f(x) = h\mathscr{L}f(x) + O(h^2),$$

so that, dividing by h and letting $h \downarrow 0$,

(5.8) *the infinitesimal generator* $\mathscr{G}$ *of* X *is* $\mathscr{L}$.

The PDE approach is now ready to go, with all of the arsenal of PDE techniques at its disposal; for example, one may begin by looking for the fundamental solution $p_t(x, y)$ to

$$\frac{\partial}{\partial t} p_t(x, y) = \mathscr{L}_x p_t(x, y), \qquad p_0(x, y) = \delta_y(x), \qquad p_t \geqslant 0,$$

where $\mathscr{L}_x$ is the operator $\mathscr{L}$ acting on the x-variable, and δ_y is the Dirac delta function at y. This fundamental solution is the transition density of the diffusion, from which one can obtain much information; see Chapter 3 of Stroock and Varadhan [1], which is also the definitive account of the martingale-problem method applied to multidimensional diffusions. There are still real problems of definition, existence and uniqueness for each of the three approaches—least severe for the PDE approach. But the additional price to be paid for using stochastic methods *is* worth it; the conditions imposed on a and b to get a PDE result to work are generally of a *global* nature, whereas the diffusion, being continuous, should only care about *local* behaviour. The stochastic methods are just right for this—once a diffusion leaves a region where everything is nice, we can stop it, and solve in the nice region, thereby giving results under only local conditions. We have great admiration and respect for the PDE approach—the analysts' fine results are not just valid for second-order elliptic operators, which is the case with the probabilistic results. The last word for the moment on the comparison between the three methods must be with Sid Port: 'The one thing probabilists can do which analysts can't is *stop*—and they never forgive us for it.'

You will realise by now that one can perfectly well have diffusions in dimension greater than one, but the one-dimensional diffusion theory is essentially complete, thanks to *Brownian local time*. The existence and properties of Brownian local time form the first non-trivial result in the theory (after the existence of Brownian motion itself).

(5.9) THEOREM (Trotter). *There exists a process* $\{l(t, x): t \geqslant 0, x \in \mathbb{R}\}$ *such that*

(5.10) (i) $(t, x) \mapsto l(t, x)$ *is jointly continuous*;
(5.10)(ii) *for any bounded measurable* f, *and* $t \geqslant 0$,

$$\int_0^t f(B_s)ds = \int_{-\infty}^{\infty} f(x) l(t, x) dx.$$

This is a deep result, whose proof using stochastic calculus we shall finally give in Section IV.44. The key property (5.10)(ii) is the *occupation density formula*; we shall discuss some of the implications for the Brownian sample path in Section 10, but for now we describe the most general regular diffusion (see Sections V.44–54 for the whole story).

(5.11) THEOREM. *A (regular) one-dimensional diffusion X on an interval I can* *be obtained from Browian motion B as*

$$X_t = s^{-1}(B_{\tau_t})$$

where the scale function $s: I \to \mathbb{R}$ *is continuous and strictly increasing and* $\tau_t = \inf\{u : A_u > t\}$, *where*

$$A_u = \int m(dx)\, l(u, x)$$

for some measure m (*the* speed measure) *that puts positive finite mass on bounded non-empty open subintervals of I which exclude the endpoints of I.*

Any pair (s, m) gives rise to a regular diffusion, which is uniquely characterised by (s, m).

The theory of diffusions in dimension greater than one is still much less complete, and doubtless will remain so.

2. BASICS ABOUT BROWNIAN MOTION

6. Existence and uniqueness of Brownian motion. The existence proof for Brownian motion that we now give (due to Ciesielski [1]) is the ultimate refinement of Wiener's original idea of representing Brownian motion as a random Fourier series.

(*6.1*) THEOREM. *There exists a probability space on which it is possible to define a process* $(B_t)_{0 \leqslant t \leqslant 1}$ *with the properties*

(i) $B_0(\omega) = 0$ *for all* ω;
(ii) *the map* $t \mapsto B_t(\omega)$ *is a continuous function of* $t \in [0, 1]$ *for all* ω;
(iii) *for every* $0 \leqslant s \leqslant t \leqslant 1$, $B_t - B_s$ *is independent of* $\{B_u : u \leqslant s\}$ *and has a* $N(0, t-s)$ *distribution.*

Proof. Take some probability space on which there is defined an infinite sequence of independent $N(0, 1)$ random variables. For reasons that will soon be apparent, we assume that they are indexed as $\{Z_{k,n} : n \in \mathbf{Z}^+,\ k \text{ odd},\ k \leqslant 2^n\}$. Now define

$$g_{1,0}(t) = 1$$

$$g_{k,n}(t) = \begin{cases} 2^{(n-1)/2} & ((k-1)2^{-n} < t \leqslant k2^{-n}), \\ -2^{(n-1)/2} & (k2^{-n} < t \leqslant (k+1)2^{-n}), \\ 0 & \text{otherwise}, \end{cases}$$

for $n \geqslant 1$, $k \leqslant 2^n$, k odd. For notational convenience, let $S_n = \{(k, n) : k \text{ odd},\ k \leqslant 2^n\}$, $S = \bigcup_{n \geqslant 0} S_n$. The first thing to notice is that $\{g_{k,n} : (k, n) \in S\}$ *is a complete*