Motivation

Much of probability theory is devoted to describing the *macroscopic picture* emerging in random systems defined by a host of *microscopic random effects*. Brownian motion is the macroscopic picture emerging from a particle moving randomly in d-dimensional space without making very big jumps. On the microscopic level, at any time step, the particle receives a random displacement, caused for example by other particles hitting it or by an external force, so that, if its position at time zero is S_0, its position at time n is given as $S_n = S_0 + \sum_{i=1}^{n} X_i$, where the displacements $X_1, X_2, X_3, \ldots$ are assumed to be independent, identically distributed random variables with values in $\mathbb{R}^d$. The process $\{S_n : n \geqslant 0\}$ is a random walk, the displacements represent the microscopic inputs. When we think about the macroscopic picture, what we mean is questions such as:

- Does S_n drift to infinity?
- Does S_n return to the neighbourhood of the origin infinitely often?
- What is the speed of growth of $\max\{|S_1|, \ldots, |S_n|\}$ as $n \to \infty$?
- What is the asymptotic number of windings of $\{S_n : n \geqslant 0\}$ around the origin?

It turns out that not all the features of the microscopic inputs contribute to the macroscopic picture. Indeed, if they exist, only the *mean* and *covariance* of the displacements are shaping the picture. In other words, all random walks whose displacements have the same mean and covariance matrix give rise to the same macroscopic process, and even the assumption that the displacements have to be independent and identically distributed can be substantially relaxed. This effect is called *universality*, and the macroscopic process is often called a *universal object*. It is a common approach in probability to study various phenomena through the associated universal objects.

If the jumps of a random walk are sufficiently tame to become negligible in the macroscopic picture, in particular if it has finite mean and variance, any continuous time stochastic process $\{B(t) : t \geqslant 0\}$ describing the macroscopic features of this random walk should have the following properties:

(1) for all times $0 \leqslant t_1 \leqslant t_2 \leqslant \ldots \leqslant t_n$ the random variables

$$B(t_n) - B(t_{n-1}),\ B(t_{n-1}) - B(t_{n-2}),\ \ldots,\ B(t_2) - B(t_1)$$

are independent; we say that the process has *independent increments*,

(2) the distribution of the increment $B(t+h) - B(t)$ does not depend on t; we say that the process has *stationary increments*,
(3) the process $\{B(t) \colon t \geqslant 0\}$ has almost surely continuous paths.

It follows (with some work) from the central limit theorem that these features imply that there exists a vector $\mu \in \mathbb{R}^d$ and a matrix $\Sigma \in \mathbb{R}^{d \times d}$ such that

(4) for every $t \geqslant 0$ and $h \geqslant 0$ the increment $B(t+h) - B(t)$ is multivariate normally distributed with mean $h\mu$ and covariance matrix $h\Sigma\Sigma^{\mathrm{T}}$.

Hence any process with the features (1)-(3) above is characterised by just three parameters,

- the *initial distribution*, i.e. the law of $B(0)$,
- the *drift vector* μ,
- the *diffusion matrix* Σ.

The process $\{B(t) \colon t \geqslant 0\}$ is called a *Brownian motion with drift* μ *and diffusion matrix* Σ. If the drift vector is zero, and the diffusion matrix is the identity we simply say the process is a *Brownian motion*. If $B(0) = 0$, i.e. the motion is started at the origin, we use the term *standard Brownian motion*.

Suppose we have a standard Brownian motion $\{B(t) \colon t \geqslant 0\}$. If X is a random variable with values in $\mathbb{R}^d$, μ a vector in $\mathbb{R}^d$ and Σ a $d \times d$ matrix, then it is easy to check that $\{\tilde{B}(t) \colon t \geqslant 0\}$ given by

$$\tilde{B}(t) = \tilde{B}(0) + \mu t + \Sigma B(t), \text{ for } t \geqslant 0,$$

is a process with the properties (1)-(4) with initial distribution X, drift vector μ and diffusion matrix Σ. Hence the macroscopic picture emerging from a random walk with finite variance can be fully described by a standard Brownian motion.

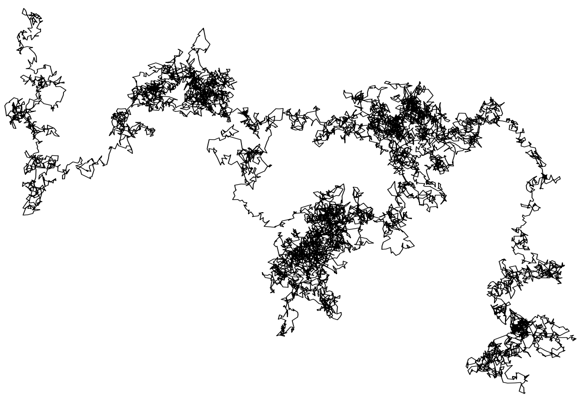

Fig. 0.1. The range of a planar Brownian motion $\{B(t) \colon 0 \leqslant t \leqslant 1\}$.

In *Chapter 1* we start exploring Brownian motion by looking at dimension $d = 1$. Here Brownian motion is a random continuous function and we ask about its *regularity*, for example: For which parameters α is the random function $B\colon [0,1] \to \mathbb{R}$ α-Hölder continuous? Is the random function $B\colon [0,1] \to \mathbb{R}$ differentiable? The surprising answer to the second question was given by Paley, Wiener and Zygmund in 1933: Almost surely, the random function $B\colon [0,1] \to \mathbb{R}$ is *nowhere* differentiable! This is particularly interesting, as it is not easy to construct a continuous, nowhere differentiable function without the help of randomness. We give a modern proof of the Paley, Wiener and Zygmund theorem, see Theorem 1.30.

In *Chapter 2* we move to general dimension d. We prove and explore the strong Markov property, which roughly says that at suitable random times Brownian motion starts afresh, see Theorem 2.16. Among the facts we derive from this property are that the set of all points visited by a Brownian motion in $d \geqslant 2$ has area zero, but the set of times when Brownian motion in $d = 1$ revisits the origin is uncountable. Besides these sample path properties, the strong Markov property is also the key to some fascinating distributional identities. It enables us to understand, for example, the process $\{M(t)\colon t \geqslant 0\}$ of the running maxima $M(t) = \max_{0 \leqslant s \leqslant t} B(s)$ of Brownian motion in $d = 1$, the process $\{T_a\colon a \geqslant 0\}$ of the first hitting times $T_a = \inf\{t \geqslant 0\colon B(t) = a\}$ of level a of a Brownian motion in $d = 1$, and the process of the vertical first hitting positions of the lines $\{(x,y) \in \mathbb{R}^2 : x = a\}$ by a Brownian motion in $d = 2$, as a function of a.

In *Chapter 3* we explore the rich relations of Brownian motion to harmonic analysis. In particular we learn how Brownian motion helps solving the classical *Dirichlet problem*.

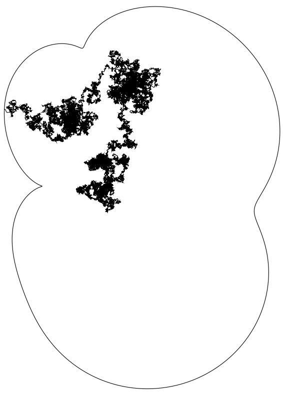

Fig. 0.2. Brownian motion and the Dirichlet problem

For its formulation in the planar case, fix a connected open set $U \subset \mathbb{R}^2$ with nice boundary, and let $\varphi\colon \partial U \to \mathbb{R}$ be continuous. The harmonic functions $f\colon U \to \mathbb{R}$ on the domain U are characterised by the differential equation

$$\frac{\partial^2 f}{\partial x_1^2}(x) + \frac{\partial^2 f}{\partial x_2^2}(x) = 0 \qquad \text{for all } x \in U.$$

The Dirichlet problem is to find, for a given domain U and boundary data φ, a continuous function $f\colon U \cup \partial U \to \mathbb{R}$, which is harmonic on U and agrees with φ on ∂U. In Theorem 3.12 we show that the unique solution of this problem is given as

$$f(x) = \mathbb{E}\big[\varphi(B(T)) \,\big|\, B(0) = x\big], \text{ for } x \in \overline{U},$$

where $\{B(t)\colon t \geqslant 0\}$ is a Brownian motion and $T = \inf\{t \geqslant 0\colon B(t) \notin U\}$ is the first exit time from U. We exploit this result, for example, to show exactly in which dimensions a particle following a Brownian motion drifts to infinity, see Theorem 3.20.

In *Chapter 4* we provide one of the major tools in our study of Brownian motion, the concept of Hausdorff dimension, and show how it can be applied in the context of Brownian motion. Indeed, when describing the sample paths of a Brownian motion one frequently encounters questions of the size of a given set: How big is the set of all points visited by a Brownian motion in the plane? How big is the set of double-points of a planar Brownian motion? How big is the set of times where Brownian motion visits a given set, say a point? For an example, let $\{B(t)\colon t \geqslant 0\}$ be Brownian motion on the real line and look at $\mathsf{Zeros} = \{t \geqslant 0\colon B(t) = 0\}$, the set of its zeros. Although $t \mapsto B(t)$ is a continuous function, Zeros is an infinite set. This set is *big*, as it is an uncountable set without isolated points. However, it is also *small* in the sense that its Lebesgue measure is zero. Indeed, Zeros is a fractal set and we show in Theorem 4.24 that its Hausdorff dimension is $1/2$.

In *Chapter 5* we explore the relationship of random walk and Brownian motion. We prove a theorem which justifies our initial point of view that Brownian motion is the macroscopic picture emerging from a large class of random walks: By *Donsker's invariance principle* one can obtain Brownian motion by taking scaled copies of a random walk and taking a limit in distribution. This result is called an invariance principle because all random walks whose increments have mean zero and finite variance essentially produce the same limit, a Brownian motion. Donsker's invariance principle is also a major tool in deriving results for random walks from those of Brownian motion, and vice versa. Both directions can be useful: In some cases the fact that Brownian motion is a continuous time process is an advantage over discrete time random walks. For example, as we discuss below, Brownian motion has scaling invariance properties, which can be a powerful tool in the study of its path properties. In other cases it is a major advantage that (simple) random walk is a discrete object and combinatorial arguments can be the right tool to derive important features. Chapter 5 offers a number of case studies for the mutually beneficial relationship between Brownian motion and random walks. Beyond Donsker's invariance principle, there is a second fascinating aspect of the relationship between random walk and Brownian motion: Given a Brownian motion in $d = 1$, we can sample from its path at certain carefully chosen times, and thus construct every random walk with mean zero and finite variance. Finding these times is called the *Skorokhod embedding problem* and we shall give two different solutions to it. The embedding problem is also the main tool in our proof of Donsker's invariance principle.

In *Chapter 6* we look again at Brownian motion in dimension $d = 1$. For a random walk on the integers running for a finite amount of time, we can define a 'local time' at a point $z \in \mathbb{Z}$ by simply counting how many times the walk visits z. Can we define an analogous quantity for Brownian motion? In Chapter 6 we show that this is possible, and offer an elegant construction of Brownian local time based on a random walk approximation. A first highlight of this chapter arises when we aim to describe the local times: If a Brownian path is started at some positive level $a > 0$ and stopped upon hitting zero, we can describe the process of local times in x as a function of x, for $0 \leqslant x \leqslant a$. The resulting process is distributed like the square of the modulus of a planar Brownian motion. This is the famous *Ray–Knight theorem*. The second highlight of this chapter is related to the nature of local time at a fixed point. The Brownian local time in x is no longer the number of visits to the point x by a Brownian motion – if x is visited at all, this number would be infinite – but we shall see that it can be described as the Hausdorff measure of the set of times at which the motion visits x.

Because Brownian motion arises as the scaling limit of a great variety of different random walks, it naturally has a number of invariance properties. One of the most important invariance properties of Brownian motion is *conformal invariance*, which we discuss in *Chapter 7*. To make this plausible think of an angle-preserving linear mapping $L\colon \mathbb{R}^d \to \mathbb{R}^d$, like a rotation followed by multiplication by a. Take a random walk started in zero with increments of mean zero and covariance matrix the identity, and look at its image under L. This image is again a random walk and its increments are distributed like LX. Appropriately rescaled as in Donsker's invariance principle, both random walks converge to a Brownian motion, the second one with a slightly different covariance matrix. This process can be identified as a time-changed Brownian motion $\{B(a^2t)\colon t \geqslant 0\}$. This easy observation has a deeper, local counterpart for planar Brownian motion: Suppose that $\phi\colon U \to V$ is a conformal mapping of a simply connected domain $U \subset \mathbb{R}^2$ onto a domain $V \subset \mathbb{R}^2$. Conformal mappings are locally angle-preserving and the Riemann mapping theorem of complex analysis tells us that a lot of such domains and mappings exist.

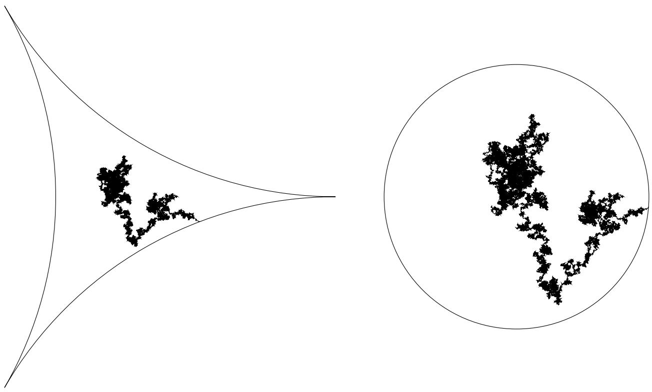

Fig. 0.3. A conformal mapping of Brownian paths

Suppose that $\{B(t)\colon t \geqslant 0\}$ is a standard Brownian motion started in some point $x \in U$ and $\tau = \inf\{t > 0\colon B(t) \notin U\}$ is the first exit time of the path from the domain U. Then it turns out that the image process $\{\phi(B(t))\colon 0 \leqslant t \leqslant \tau\}$ is a *time-changed* Brownian

Cambridge University Press
978-0-521-76018-8 - Brownian Motion
Peter Morters and Yuval Peres
Excerpt
More information

motion in the domain V, stopped when it leaves V, see Theorem 7.20. In order to prove this we have to develop a little bit of the theory of stochastic integration with respect to a Brownian motion, and we give a lot of further applications of this tool in Chapter 7.

In *Chapter 8* we develop the potential theory of Brownian motion. The problem which is the motivation behind this is, given a compact set $A \subset \mathbb{R}^d$, to find the probability that a Brownian motion $\{B(t): t \geqslant 0\}$ hits the set A, i.e. that there exists $t > 0$ with $B(t) \in A$. This problem is answered in the best possible way by Theorem 8.24, which is a modern extension of a classical result of Kakutani: The hitting probability can be approximated by the capacity of A with respect to the Martin kernel up to a factor of two.

With a wide range of tools at our hand, in *Chapter 9* we study the self-intersections of Brownian motion: For example, a point $x \in \mathbb{R}^d$ is called a double point of $\{B(t): t \geqslant 0\}$ if there exist times $0 < t_1 < t_2$ such that $B(t_1) = B(t_2) = x$. In which dimensions does Brownian motion have double points? How big is the set of double points? We show that in dimensions $d \geqslant 4$ no double points exist, in dimension $d = 3$ double points exist and the set of double points has Hausdorff dimension one, and in dimension $d = 2$ double points exist and the set of double points has Hausdorff dimension two. In dimension $d = 2$ we find a surprisingly complex situation: While every point $x \in \mathbb{R}^2$ is almost surely not visited by a Brownian motion, there exist (random) points in the plane, which are visited infinitely often, even uncountably often. This result, Theorem 9.24, is one of the highlights of this book.

Chapter 10 deals with exceptional points for Brownian motion and Hausdorff dimension spectra of families of exceptional points. To explain an example, we look at a Brownian motion in the plane run for one time unit, which is a continuous curve $\{B(t): t \in [0, 1]\}$. In Chapter 7 we see that, for any point on the curve, almost surely, the Brownian motion performs an infinite number of full windings in both directions around this point. Still, there exist random points on the curve, which are exceptional in the sense that Brownian motion performs no windings around them at all. This follows from an easy geometric argument: Take a point in $\mathbb{R}^2$ with coordinates (x_1, x_2) such that $x_1 = \min\{x: (x, x_2) \in B[0, 1]\}$, i.e. a point which is the leftmost on the intersection of the Brownian curve and the line $\{(z, y): z \in \mathbb{R}\}$, for some $x_2 \in \mathbb{R}$. Then Brownian motion does not perform any full windings around (x_1, x_2), as this would necessarily imply that it crosses the half-line $\{(x, x_2): x < x_2\}$, contradicting the minimality of x_1. One can ask for a more extreme deviation from typical behaviour: A point $x = B(t)$ is an α-cone point if the Brownian curve is contained in an open cone with tip in $x = (x_1, x_2)$, central axis $\{(x_1, x): x > x_2\}$ and opening angle α. Note that the points described in the previous paragraph are 2π-cone points in this sense. In Theorem 10.38 we show that α-cone points exist exactly if $\alpha \in [\pi, 2\pi]$, and prove that for every such α, almost surely,

$$\dim \big\{ x \in \mathbb{R}^2 : x \text{ is an } \alpha\text{-cone point} \big\} = 2 - \frac{2\pi}{\alpha}.$$

This is an example of a Hausdorff dimension spectrum, a topic which has been at the centre of some research activity at the beginning of the current millennium.

Cambridge University Press
978-0-521-76018-8 - Brownian Motion
Peter Morters and Yuval Peres
Excerpt
More information

1

Brownian motion as a random function

In this chapter we focus on one-dimensional, or linear, Brownian motion. We start with Paul Lévy's construction of Brownian motion and discuss two fundamental sample path properties, continuity and differentiability. We then discuss the Cameron–Martin theorem, which shows that sample path properties for Brownian motion with drift can be obtained from the corresponding results for driftless Brownian motion.

1.1 Paul Lévy's construction of Brownian motion

1.1.1 Definition of Brownian motion

Brownian motion is closely linked to the normal distribution. Recall that a random variable X is normally distributed with mean μ and variance σ^2 if

$$\mathbb{P}\{X > x\} = \frac{1}{\sqrt{2\pi\sigma^2}} \int_x^\infty e^{-\frac{(u-\mu)^2}{2\sigma^2}}\, du, \qquad \text{for all } x \in \mathbb{R}.$$

Definition 1.1. A real-valued stochastic process $\{B(t)\colon t \geqslant 0\}$ is called a **(linear) Brownian motion** with start in $x \in \mathbb{R}$ if the following holds:

- $B(0) = x$,
- the process has **independent increments**, i.e. for all times $0 \leqslant t_1 \leqslant t_2 \leqslant \ldots \leqslant t_n$ the increments $B(t_n) - B(t_{n-1}),\ B(t_{n-1}) - B(t_{n-2}),\ \ldots,\ B(t_2) - B(t_1)$ are independent random variables,
- for all $t \geqslant 0$ and $h > 0$, the increments $B(t+h) - B(t)$ are normally distributed with expectation zero and variance h,
- almost surely, the function $t \mapsto B(t)$ is continuous.

We say that $\{B(t)\colon t \geqslant 0\}$ is a **standard Brownian motion** if $x = 0$. $\diamond$

We will address the nontrivial question of the *existence* of a Brownian motion in Section 1.1.2. For the moment let us step back and look at some technical points. We have defined Brownian motion as a *stochastic process* $\{B(t)\colon t \geqslant 0\}$ which is just a family of (uncountably many) random variables $\omega \mapsto B(t,\omega)$ defined on a single probability space $(\Omega, \mathcal{A}, \mathbb{P})$. At the same time, a stochastic process can also be interpreted as a *random function* with the sample functions defined by $t \mapsto B(t,\omega)$. The *sample path properties* of a stochastic process are the properties of these random functions, and it is these properties we will be most interested in in this book.

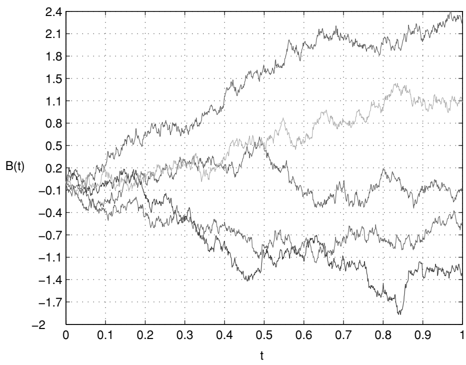

Fig. 1.1. Graphs of five sampled Brownian motions

By the **finite-dimensional distributions** of a stochastic process $\{B(t)\colon t \geqslant 0\}$ we mean the laws of all the finite dimensional random vectors

$$\big(B(t_1), B(t_2), \dots, B(t_n)\big), \text{ for all } 0 \leqslant t_1 \leqslant t_2 \leqslant \dots \leqslant t_n .$$

To describe these joint laws it suffices to describe the joint law of $B(0)$ and the increments

$$\big(B(t_1) - B(0), B(t_2) - B(t_1), \dots, B(t_n) - B(t_{n-1})\big), \text{ for all } 0 \leqslant t_1 \leqslant t_2 \leqslant \dots \leqslant t_n .$$

This is what we have done in the first three items of the definition, which specify the finite-dimensional distributions of Brownian motion. However, the last item, almost sure continuity, is also crucial, and this is information which goes beyond the finite-dimensional distributions of the process in the sense above, technically because the set $\{\omega \in \Omega\colon t \mapsto B(t,\omega) \text{ continuous}\}$ is in general not in the σ-algebra generated by the random vectors $(B(t_1), B(t_2), \dots, B(t_n))$, $n \in \mathbb{N}$.

Example 1.2 Suppose that $\{B(t)\colon t \geqslant 0\}$ is a Brownian motion and U is an independent random variable, which is uniformly distributed on $[0,1]$. Then the process $\{\tilde{B}(t)\colon t \geqslant 0\}$ defined by

$$\tilde{B}(t) = \begin{cases} B(t) & \text{if } t \neq U, \\ 0 & \text{if } t = U, \end{cases}$$

has the same finite-dimensional distributions as a Brownian motion, but is discontinuous if $B(U) \neq 0$, i.e. with probability one, and hence this process is not a Brownian motion. ⋄

We see that, if we are interested in the sample path properties of a stochastic process, we may need to specify more than just its finite-dimensional distributions. Suppose $\mathfrak{X}$ is a property a function might or might not have, like continuity, differentiability, etc. We say that a process $\{X(t) : t \geqslant 0\}$ **has property $\mathfrak{X}$ almost surely** if there exists $A \in \mathcal{A}$ such that $\mathbb{P}(A) = 1$ and $A \subset \big\{\omega \in \Omega\colon t \mapsto X(t,\omega) \text{ has property } \mathfrak{X}\big\}$. Note that the set on the right need not lie in $\mathcal{A}$.

1.1.2 Paul Lévy's construction of Brownian motion

It is a substantial issue whether the conditions imposed on the finite-dimensional distributions in the definition of Brownian motion allow the process to have continuous sample paths, or whether there is a contradiction. In this section we show that there is no contradiction and, fortunately, Brownian motion exists.

Theorem 1.3 (Wiener 1923) *Standard Brownian motion exists.*

We construct Brownian motion as a uniform limit of continuous functions, to ensure that it automatically has continuous paths. Recall that we need only construct a *standard* Brownian motion $\{B(t) \colon t \geqslant 0\}$, as $X(t) = x + B(t)$ is a Brownian motion with starting point x. The proof exploits properties of Gaussian random vectors, which are the higher-dimensional analogue of the normal distribution.

Definition 1.4. A random vector $X = (X_1, \dots, X_n)$ is called a **Gaussian random vector** if there exists an $n \times m$ matrix A, and an n-dimensional vector b such that $X^{\mathrm{T}} = AY + b$, where Y is an m-dimensional vector with independent standard normal entries. ◇

Basic facts about Gaussian random variables are collected in Appendix 12.2.

Proof of Wiener's theorem. We first construct Brownian motion on the interval $[0,1]$ as a random element on the space $\mathbf{C}[0,1]$ of continuous functions on $[0,1]$. The idea is to construct the right joint distribution of Brownian motion step by step on the finite sets

$$\mathcal{D}_n = \left\{ \tfrac{k}{2^n} : 0 \leqslant k \leqslant 2^n \right\}$$

of dyadic points. We then interpolate the values on $\mathcal{D}_n$ linearly and check that the uniform limit of these continuous functions exists and is a Brownian motion.

To do this let $\mathcal{D} = \bigcup_{n=0}^{\infty} \mathcal{D}_n$ and let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space on which a collection $\{Z_t \colon t \in \mathcal{D}\}$ of independent, standard normally distributed random variables can be defined. Let $B(0) := 0$ and $B(1) := Z_1$. For each $n \in \mathbb{N}$ we define the random variables $B(d)$, $d \in \mathcal{D}_n$ such that

(1) for all $r < s < t$ in $\mathcal{D}_n$ the random variable $B(t) - B(s)$ is normally distributed with mean zero and variance $t - s$, and is independent of $B(s) - B(r)$,
(2) the vectors $(B(d) \colon d \in \mathcal{D}_n)$ and $(Z_t : t \in \mathcal{D} \setminus \mathcal{D}_n)$ are independent.

Note that we have already done this for $\mathcal{D}_0 = \{0, 1\}$. Proceeding inductively we may assume that we have succeeded in doing it for some $n - 1$. We then define $B(d)$ for $d \in \mathcal{D}_n \setminus \mathcal{D}_{n-1}$ by

$$B(d) = \frac{B(d - 2^{-n}) + B(d + 2^{-n})}{2} + \frac{Z_d}{2^{(n+1)/2}}.$$

Note that the first summand is the linear interpolation of the values of B at the neighbouring points of d in $\mathcal{D}_{n-1}$. Therefore $B(d)$ is independent of $(Z_t : t \in \mathcal{D} \setminus \mathcal{D}_n)$ and the second property is fulfilled.

Moreover, as $\frac{1}{2}[B(d+2^{-n})-B(d-2^{-n})]$ depends only on $(Z_t : t \in \mathcal{D}_{n-1})$, it is independent of $Z_d/2^{(n+1)/2}$. By our induction assumptions both terms are normally distributed with mean zero and variance $2^{-(n+1)}$. Hence their sum $B(d) - B(d-2^{-n})$ and their difference $B(d+2^{-n}) - B(d)$ are independent and normally distributed with mean zero and variance 2^{-n} by Corollary 12.12.

Indeed, all increments $B(d) - B(d-2^{-n})$, for $d \in \mathcal{D}_n \setminus \{0\}$, are independent. To see this it suffices to show that they are pairwise independent, as the vector of these increments is Gaussian. We have seen in the previous paragraph that pairs $B(d) - B(d-2^{-n})$, $B(d+2^{-n}) - B(d)$ with $d \in \mathcal{D}_n \setminus \mathcal{D}_{n-1}$ are independent. The other possibility is that the increments are over intervals separated by some $d \in \mathcal{D}_{n-1}$. Choose $d \in \mathcal{D}_j$ with this property and minimal j, so that the two intervals are contained in $[d-2^{-j}, d]$, respectively $[d, d+2^{-j}]$. By induction the increments over these two intervals of length 2^{-j} are independent, and the increments over the intervals of length 2^{-n} are constructed from the independent increments $B(d) - B(d-2^{-j})$, respectively $B(d+2^{-j}) - B(d)$, using a disjoint set of variables $(Z_t : t \in \mathcal{D}_n)$. Hence they are independent and this implies the first property, and completes the induction step.

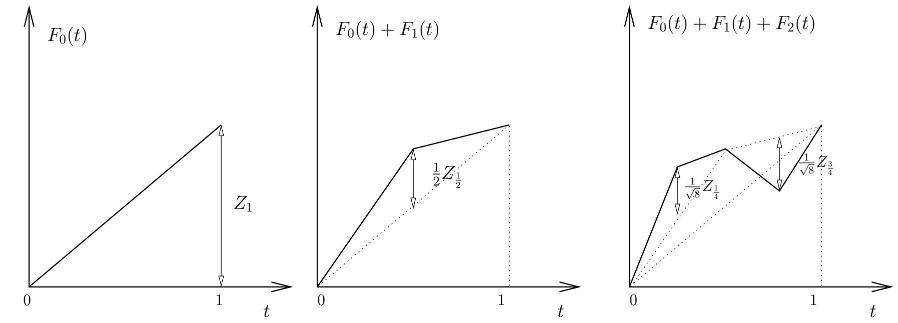

Fig. 1.2. The first three steps in the construction of Brownian motion

Having thus chosen the values of the process on all dyadic points, we interpolate between them. Formally, define

$$F_0(t) = \begin{cases} Z_1 & \text{for } t = 1, \\ 0 & \text{for } t = 0, \\ \text{linear} & \text{in between,} \end{cases}$$

and, for each $n \geqslant 1$,

$$F_n(t) = \begin{cases} 2^{-(n+1)/2} Z_t & \text{for } t \in \mathcal{D}_n \setminus \mathcal{D}_{n-1} \\ 0 & \text{for } t \in \mathcal{D}_{n-1} \\ \text{linear between consecutive points in } \mathcal{D}_n . \end{cases}$$

These functions are continuous on $[0, 1]$ and, for all n and $d \in \mathcal{D}_n$,

$$B(d) = \sum_{i=0}^{n} F_i(d) = \sum_{i=0}^{\infty} F_i(d), \tag{1.1}$$