

18.1: Standard Brownian Motion

Basic Theory

History

In 1827, the botanist Robert Brown noticed that tiny particles from pollen, when suspended in water, exhibited continuous but very jittery and erratic motion. In his “miracle year” in 1905, Albert Einstein explained the behavior physically, showing that the particles were constantly being bombarded by the molecules of the water, and thus helping to firmly establish the atomic theory of matter. Brownian motion as a mathematical random process was first constructed in rigorous way by Norbert Wiener in a series of papers starting in 1918. For this reason, the Brownian motion process is also known as the *Wiener process*.

Run the two-dimensional Brownian motion simulation several times in single-step mode to get an idea of what Mr. Brown may have observed under his microscope.

Along with the Bernoulli trials process and the Poisson process, the Brownian motion process is of central importance in probability. Each of these processes is based on a set of idealized assumptions that lead to a rich mathematical theory. In each case also, the process is used as a building block for a number of related random processes that are of great importance in a variety of applications. In particular, Brownian motion and related processes are used in applications ranging from physics to statistics to economics.

Definition

A *standard Brownian motion* is a random process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with state space \mathbb{R} that satisfies the following properties:

1. $X_0 = 0$ (with probability 1).
2. \mathbf{X} has stationary increments. That is, for $s, t \in [0, \infty)$ with $s < t$, the distribution of $X_t - X_s$ is the same as the distribution of X_{t-s} .
3. \mathbf{X} has independent increments. That is, for $t_1, t_2, \dots, t_n \in [0, \infty)$ with $t_1 < t_2 < \dots < t_n$, the random variables $X_{t_1}, X_{t_2} - X_{t_1}, \dots, X_{t_n} - X_{t_{n-1}}$ are independent.
4. X_t is normally distributed with mean 0 and variance t for each $t \in (0, \infty)$.
5. With probability 1, $t \mapsto X_t$ is continuous on $[0, \infty)$.

To understand the assumptions physically, let's take them one at a time.

1. Suppose that we measure the position of a Brownian particle in one dimension, starting at an arbitrary time which we designate as $t = 0$, with the initial position designated as $x = 0$. Then this assumption is satisfied by convention. Indeed, occasionally, it's convenient to relax this assumption and allow X_0 to have other values.
2. This is a statement of *time homogeneity*: the underlying dynamics (namely the jostling of the particle by the molecules of water) do not change over time, so the distribution of the displacement of the particle in a time interval $[s, t]$ depends only on the length of the time interval.
3. This is an idealized assumption that would hold approximately if the time intervals are large compared to the tiny times between collisions of the particle with the molecules.
4. This is another idealized assumption based on the central limit theorem: the position of the particle at time t is the result of a very large number of collisions, each making a very small contribution. The fact that the mean is 0 is a statement of *spatial homogeneity*: the particle is no more or less likely to be jostled to the right than to the left. Next, recall that the assumptions of stationary, independent increments means that $\text{var}(X_t) = \sigma^2 t$ for some positive constant σ^2 . By a change in time scale, we can assume $\sigma^2 = 1$, although we will consider more general Brownian motions in the next section.
5. Finally, the continuity of the sample paths is an essential assumption, since we are modeling the position of a physical particle as a function of time.

Of course, the first question we should ask is whether there *exists* a stochastic process satisfying the definition. Fortunately, the answer is yes, although the proof is complicated.

There exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a stochastic process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on this probability space satisfying the assumptions in the [definition](#).

Proof sketch

The assumptions in the definition lead to a consistent set of finite dimensional distributions (which are given [below](#)). Thus by Kolmogorov existence theorem, there exists a stochastic process $\mathbf{U} = \{U_t : t \in [0, \infty)\}$ that has these finite dimensional distributions. However, \mathbf{U} does not have continuous sample paths, but we can construct from \mathbf{U} an equivalent process that does have continuous sample paths.

First recall that a *binary rational* (or *dyadic rational*) in $[0, \infty)$ is a number of the form $k/2^n$ where $k, n \in \mathbb{N}$. Let \mathbb{D}_+ denote the set of all binary rationals in $[0, \infty)$, and recall that \mathbb{D}_+ is countable but also *dense* in $[0, \infty)$ (that is, if $t \in [0, \infty) \setminus \mathbb{D}_+$ then there exists $t_n \in \mathbb{D}_+$ for $n \in \mathbb{N}_+$ such that $t_n \rightarrow t$ as $n \rightarrow \infty$).

Now, for $n \in \mathbb{N}_+$, let $X_n(t) = U_t$ if t is a binary rational of the form $k/2^n$ for some $k \in \mathbb{N}$. If t is not such a binary rational, define $X_n(t)$ by linear interpolation between the the closest binary rationals of this form on either side of t . Then $X_n(t) \rightarrow U(t)$ as $n \rightarrow \infty$ for every $t \in \mathbb{D}_+$, and with probability 1, the convergence is uniform on $\mathbb{D}_+ \cap [0, T]$ for each $T > 0$. It then follows that \mathbf{U} is continuous on \mathbb{D}_+ with probability 1.

For the last step, let $X_t = \lim_{s \rightarrow t, s \in \mathbb{D}_+} U_s$ for $t \in [0, \infty)$. The limit exists since \mathbf{U} is continuous on \mathbb{D}_+ with probability 1. The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is continuous on $[0, \infty)$ with probability 1, and has the same finite dimensional distributions as \mathbf{U} .

Run the simulation of the standard Brownian motion process a few times in single-step mode. Note the qualitative behavior of the sample paths. Run the simulation 1000 times and compare the empirical density function and moments of X_t to the true probability density function and moments.

Brownian Motion as a Limit of Random Walks

Clearly the underlying dynamics of the Brownian particle being knocked about by molecules suggests a random walk as a possible model, but with tiny time steps and tiny spatial jumps. Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be the symmetric simple random walk. Thus, $X_n = \sum_{i=1}^n U_i$ where $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent variables with $\mathbb{P}(U_i = 1) = \mathbb{P}(U_i = -1) = \frac{1}{2}$ for each $i \in \mathbb{N}_+$. Recall that $\mathbb{E}(X_n) = 0$ and $\text{var}(X_n) = n$ for $n \in \mathbb{N}$. Also, since \mathbf{X} is the partial sum process associated with an IID sequence, \mathbf{X} has stationary, independent increments (but of course in discrete time). Finally, recall that by the central limit theorem, X_n/\sqrt{n} converges to the standard normal distribution as $n \rightarrow \infty$. Now, for $h, d \in (0, \infty)$ the continuous time process

$$\mathbf{X}_{h,d} = \{dX_{\lfloor t/h \rfloor} : t \in [0, \infty)\} \quad (18.1.1)$$

is a jump process with jumps at $\{0, h, 2h, \dots\}$ and with jumps of size $\pm d$. Basically we would like to let $h \downarrow 0$ and $d \downarrow 0$, but this cannot be done arbitrarily. Note that $\mathbb{E}[X_{h,d}(t)] = 0$ but $\text{var}[X_{h,d}(t)] = d^2 \lfloor t/h \rfloor$. Thus, by the central limit theorem, if we take $d = \sqrt{h}$ then the distribution of $X_{h,d}(t)$ will converge to the normal distribution with mean 0 and variance t as $h \downarrow 0$. More generally, we might hope that all of requirements in the [definition](#) are satisfied by the limiting process, and if so, we have a standard Brownian motion.

Run the simulation of the random walk process for increasing values of n . In particular, run the simulation several times with $n = 100$. Compare the qualitative behavior with the standard Brownian motion process. Note that the scaling of the random walk in time and space is effectivelly accomplished by scaling the horizontal and vertical axes in the graph window.

Finite Dimensional Distributions

Let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be a standard Brownian motion. It follows from part (d) of the [definition](#) that X_t has probability density function f_t given by

$$f_t(x) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{x^2}{2t}\right), \quad x \in \mathbb{R} \quad (18.1.2)$$

This family of density functions determines the finite dimensional distributions of \mathbf{X} .

If $t_1, t_2, \dots, t_n \in (0, \infty)$ with $0 < t_1 < t_2 < \dots < t_n$ then $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ has probability density function f_{t_1, t_2, \dots, t_n} given by

$$f_{t_1, t_2, \dots, t_n}(x_1, x_2, \dots, x_n) = f_{t_1}(x_1) f_{t_2 - t_1}(x_2 - x_1) \cdots f_{t_n - t_{n-1}}(x_n - x_{n-1}), \quad (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (18.1.3)$$

Proof

This follows because \mathbf{X} has stationary, independent increments.

\mathbf{X} is a Gaussian process with mean function mean function $m(t) = 0$ for $t \in [0, \infty)$ and covariance function $c(s, t) = \min\{s, t\}$ for $s, t \in [0, \infty)$.

Proof

The fact that \mathbf{X} is a Gaussian process follows because X_t is normally distributed for each $t \in T$ and \mathbf{X} has stationary, independent increments. The mean function is 0 by assumption. For the covariance function, suppose $s, t \in [0, \infty)$ with $s \leq t$. Since X_s and $X_t - X_s$ are independent, we have

$$\text{cov}(X_s, X_t) = \text{cov}[X_s, X_s + (X_t - X_s)] = \text{var}(X_s) + 0 = s \quad (18.1.4)$$

Recall that for a Gaussian process, the finite dimensional (multivariate normal) distributions are completely determined by the mean function m and the covariance function c . Thus, it follows that a standard Brownian motion is characterized as a continuous Gaussian process with the mean and covariance functions in the last theorem. Note also that

$$\text{cor}(X_s, X_t) = \frac{\min\{s, t\}}{\sqrt{st}} = \sqrt{\frac{\min\{s, t\}}{\max\{s, t\}}}, \quad (s, t) \in [0, \infty)^2 \quad (18.1.5)$$

We can also give the higher moments and the moment generating function for X_t .

For $n \in \mathbb{N}$ and $t \in [0, \infty)$,

1. $\mathbb{E}(X_t^{2n}) = 1 \cdot 3 \cdots (2n-1)t^n = (2n)!t^n / (n!2^n)$
2. $\mathbb{E}(X_t^{2n+1}) = 0$

Proof

These moments follow from standard results, since X_t is normally distributed with mean 0 and variance t .

For $t \in [0, \infty)$, X_t has moment generating function given by

$$\mathbb{E}(e^{uX_t}) = e^{tu/2}, \quad u \in \mathbb{R} \quad (18.1.6)$$

Proof

Again, this is a standard result for the normal distribution.

Simple Transformations

There are several simple transformations that preserve standard Brownian motion and will give us insight into some of its properties. As usual, our starting place is a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Our first result is that reflecting the paths of \mathbf{X} in the line $x = 0$ gives another standard Brownian motion

Let $Y_t = -X_t$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \geq 0\}$ is also a standard Brownian motion.

Proof

Clearly the new process is still a Gaussian process, with mean function $\mathbb{E}(-X_t) = -\mathbb{E}(X_t) = 0$ for $t \in [0, \infty)$ and covariance function $\text{cov}(-X_s, -X_t) = \text{cov}(X_s, X_t) = \min\{s, t\}$ for $(s, t) \in [0, \infty)^2$. Finally, since \mathbf{X} is continuous, so is \mathbf{Y} .

Our next result is related to the Markov property, which we explore in more detail below. If we “restart” Brownian motion at a fixed time s , and shift the origin to X_s , then we have another standard Brownian motion. This means that Brownian motion is both *temporally* and *spatially* homogeneous.

Fix $s \in [0, \infty)$ and define $Y_t = X_{s+t} - X_s$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Since \mathbf{X} has stationary, independent increments, the process \mathbf{Y} is equivalent in distribution to \mathbf{X} . Clearly also \mathbf{Y} is continuous since \mathbf{X} is.

Our next result is a simple time reversal, but to state this result, we need to restrict the time parameter to a bounded interval of the form $[0, T]$ where $T > 0$. The upper endpoint T is sometimes referred to as a *finite time horizon*. Note that $\{X_t : t \in [0, T]\}$ still satisfies the [definition](#), but with the time parameters restricted to $[0, T]$.

Define $Y_t = X_{T-t} - X_T$ for $0 \leq t \leq T$. Then $\mathbf{Y} = \{Y_t : t \in [0, T]\}$ is also a standard Brownian motion on $[0, T]$.

Proof

\mathbf{Y} is a Gaussian process, since a finite, linear combination of variables from this process reduces to a finite, linear combination of variables from \mathbf{X} . Next, $\mathbb{E}(Y_t) = \mathbb{E}(X_{T-t}) - \mathbb{E}(X_T) = 0$. Next, if $s, t \in [0, T]$ with $s \leq t$ then

$$\text{cov}(Y_s, Y_t) = \text{cov}(X_{T-s} - X_T, X_{T-t} - X_T) = \text{cov}(X_{T-s}, X_{T-t}) - \text{cov}(X_{T-s}, X_T) - \text{cov}(X_T, X_{T-t}) + \text{cov}(X_T, X_T) \quad (18.1.7)$$

$$= (T-t) - (T-s) - (T-t) + T = s \quad (18.1.8)$$

Finally, $t \mapsto Y_t$ is continuous on $[0, T]$ with probability 1, since $t \mapsto X_t$ is continuous on $[0, T]$ with probability 1.

Our next transformation involves scaling \mathbf{X} both temporally and spatially, and is known as *self-similarity*.

Let $a > 0$ and define $Y_t = \frac{1}{a} X_{a^2 t}$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Once again, \mathbf{Y} is a Gaussian process, since finite, linear combinations of variables in \mathbf{Y} reduce to finite, linear combinations of variables in \mathbf{X} . Next, $\mathbb{E}(Y_t) = a\mathbb{E}(X_{a^2 t}) = 0$ for $t > 0$, and for $s, t > 0$ with $s < t$,

$$\text{cov}(Y_s, Y_t) = \text{cov}\left(\frac{1}{a}X_{a^2s}, \frac{1}{a}X_{a^2t}\right) = \frac{1}{a^2}\text{cov}(X_{a^2s}, X_{a^2t}) = \frac{1}{a^2}a^2s = s \quad (18.1.9)$$

Finally \mathbf{Y} is a continuous process since \mathbf{X} is continuous.

Note that the graph of \mathbf{Y} can be obtained from the graph of \mathbf{X} by scaling the time axis t by a factor of a^2 and scaling the spatial axis x by a factor of a . The fact that the temporal scale factor must be the square of the spatial scale factor is clearly related to Brownian motion as the limit of random walks. Note also that this transformation amounts to “zooming in or out” of the graph of \mathbf{X} and hence Brownian motion has a self-similar, fractal quality, since the graph is unchanged by this transformation. This also suggests that, although continuous, $t \mapsto X_t$ is highly irregular. We consider this in the next subsection.

Our final transformation is referred to as *time inversion*.

Let $Y_0 = 0$ and $Y_t = tX_{1/t}$ for $t > 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Clearly \mathbf{Y} is a Gaussian process, since finite, linear combinations of variables in \mathbf{Y} reduce to finite, linear combinations of variables in \mathbf{X} . Next, $\mathbb{E}(Y_t) = t\mathbb{E}(X_{1/t}) = 0$ for $t > 0$, and for $s, t > 0$ with $s < t$,

$$\text{cov}(Y_s, Y_t) = \text{cov}(sX_{1/s}, tX_{1/t}) = st \text{cov}(X_{1/s}, X_{1/t}) = st \frac{1}{t} = s \quad (18.1.10)$$

Since $t \mapsto X_t$ is continuous on $[0, \infty)$ with probability 1, $t \mapsto Y_t$ is continuous on $(0, \infty)$ with probability 1. Thus, all that remains is to show continuity at $t = 0$. Thus we need to show that with probability 1, $tX_{1/t} \rightarrow 0$ as $t \downarrow 0$, or equivalently, $X_s/s \rightarrow 0$ as $s \uparrow \infty$. But this last statement holds by the [law of the iterated logarithm](#), given below.

Irregularity

The [defining properties](#) suggest that standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ cannot be a smooth, differentiable function. Consider the usual difference quotient at t ,

$$\frac{X_{t+h} - X_t}{h} \quad (18.1.11)$$

By the stationary increments property, if $h > 0$, the numerator has the same distribution as X_h , while if $h < 0$, the numerator has the same distribution as $-X_{-h}$, which in turn has the same distribution as X_{-h} . So, in both cases, the difference quotient has the same distribution as $X_{|h|}/h$, and this variable has the normal distribution with mean 0 and variance $|h|/h^2 = 1/|h|$. So the variance of the difference quotient diverges to ∞ as $h \rightarrow 0$, and hence the difference quotient does not even converge in distribution, the weakest form of convergence.

The [temporal-spatial transformation](#) above also suggests that Brownian motion cannot be differentiable. The intuitive meaning of *differentiable* at t is that the function is *locally linear* at t —as we zoom in, the graph near t begins to look like a line (whose slope, of course, is the derivative). But as we zoom in on Brownian motion, (in the sense of the transformation), it always looks the same, and in particular, just as jagged. More formally, if \mathbf{X} is differentiable at t , then so is the transformed process \mathbf{Y} , and the chain rule gives $Y'(t) = aX'(a^2t)$. But \mathbf{Y} is also a standard Brownian motion for every $a > 0$, so something is clearly wrong. While not rigorous, these examples are motivation for the following theorem:

With probability 1, \mathbf{X} is nowhere differentiable on $[0, \infty)$.

Run the simulation of the standard Brownian motion process. Note the continuity but very jagged quality of the sample paths. Of course, the simulation cannot really capture Brownian motion with complete fidelity.

The following theorems give a more precise measure of the irregularity of standard Brownian motion.

Standard Brownian motion \mathbf{X} has Hölder exponent $\frac{1}{2}$. That is, \mathbf{X} is Hölder continuous with exponent α for every $\alpha < \frac{1}{2}$, but is not Hölder continuous with exponent α for any $\alpha > \frac{1}{2}$.

In particular, \mathbf{X} is not Lipschitz continuous, and this shows again that it is not differentiable. The following result states that in terms of Hausdorff dimension, the graph of standard Brownian motion lies midway between a simple curve (dimension 1) and the plane (dimension 2).

The graph of standard Brownian motion has Hausdorff dimension $\frac{3}{2}$.

Yet another indication of the irregularity of Brownian motion is that it has infinite total variation on any interval of positive length.

Suppose that $a, b \in \mathbb{R}$ with $a < b$. Then the total variation of \mathbf{X} on $[a, b]$ is ∞ .

The Markov Property and Stopping Times

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Recall that a Markov process has the property that the future is independent of the past, given the present state. Because of the stationary, independent increments property, Brownian motion has the property. As a minor note, to view \mathbf{X} as a Markov process, we sometimes need to relax Assumption 1 and let X_0 have an arbitrary value in \mathbb{R} . Let $\mathcal{F}_t = \sigma\{X_s : 0 \leq s \leq t\}$, the sigma-algebra generated by the process up to time $t \in [0, \infty)$. The family of σ -algebras $\mathcal{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is known as a filtration.

Standard Brownian motion is a time-homogeneous Markov process with transition probability density p given by

$$p_t(x, y) = f_t(y - x) = \frac{1}{\sqrt{2\pi t}} \exp\left[-\frac{(y - x)^2}{2t}\right], \quad t \in (0, \infty); x, y \in \mathbb{R} \quad (18.1.12)$$

Proof

Fix $s \in [0, \infty)$. The theorem follows from the fact that the process $\{X_{s+t} - X_s : t \in [0, \infty)\}$ is another standard Brownian motion, as shown [above](#), and is independent of \mathcal{F}_s .

The transition density p satisfies the following *diffusion equations*. The first is known as the *forward equation* and the second as the *backward equation*.

$$\frac{\partial}{\partial t} p_t(x, y) = \frac{1}{2} \frac{\partial^2}{\partial y^2} p_t(x, y) \quad (18.1.13)$$

$$\frac{\partial}{\partial t} p_t(x, y) = \frac{1}{2} \frac{\partial^2}{\partial x^2} p_t(x, y) \quad (18.1.14)$$

Proof

These results follow from standard calculus.

The diffusion equations are so named, because the spatial derivative in the first equation is with respect to y , the state *forward* at time t , while the spatial derivative in the second equation is with respect to x , the state *backward* at time 0.

Recall that a random time τ taking values in $[0, \infty]$ is a *stopping time* with respect to the process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ if $\{\tau \leq t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$. Informally, we can determine whether or not $\tau \leq t$ by observing the process up to time t . An important special case is the first time that our Brownian motion hits a specified state. Thus, for $x \in \mathbb{R}$ let $\tau_x = \inf\{t \in [0, \infty) : X_t = x\}$. The random time τ_x is a stopping time.

For a stopping time τ , we need the σ -algebra of events that can be defined in terms of the process up to the random time τ , analogous to \mathcal{F}_t , the σ -algebra of events that can be defined in terms of the process up to a fixed time t . The appropriate definition is

$$\mathcal{F}_\tau = \{B \in \mathcal{F} : B \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \geq 0\} \quad (18.1.15)$$

See the section on Filtrations and Stopping Times for more information on filtrations, stopping times, and the σ -algebra associated with a stopping time.

The *strong Markov property* is the Markov property generalized to stopping times. Standard Brownian motion \mathbf{X} is also a strong Markov process. The best way to say this is by a generalization of the [temporal and spatial homogeneity result](#) above.

Suppose that τ is a stopping time and define $Y_t = X_{\tau+t} - X_\tau$ for $t \in [0, \infty)$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is a standard Brownian motion and is independent of \mathcal{F}_τ .

The Reflection Principle

Many interesting properties of Brownian motion can be obtained from a clever idea known as the *reflection principle*. As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Let τ be a stopping time for \mathbf{X} . Define

$$W_t = \begin{cases} X_t, & 0 \leq t < \tau \\ 2X_\tau - X_t, & \tau \leq t < \infty \end{cases} \quad (18.1.16)$$

Thus, the graph of $\mathbf{W} = \{W_t : t \in [0, \infty)\}$ can be obtained from the graph of \mathbf{X} by reflecting in the line $x = X_\tau$ after time τ . In particular, if the stopping time τ is τ_a , the first time that the process hits a specified state $a > 0$, then the graph of \mathbf{W} is obtained from the graph of \mathbf{X} by reflecting in the line $x = a$ after time τ_a .

Open the simulation of reflecting Brownian motion. This app shows the process \mathbf{W} corresponding to the stopping time τ_a , the time of first visit to a positive state a . Run the simulation in single step mode until you see the reflected process several times. Make sure that you understand how the process \mathbf{W} works.

The reflected process $\mathbf{W} = \{W_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Run the simulation of the reflected Brownian motion process 1000 times. Compare the empirical density function and moments of W_t to the true probability density function and moments.

Martingales

As usual, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be a standard Brownian motion, and let $\mathcal{F}_t = \sigma\{X_s : 0 \leq s \leq t\}$ for $t \in [0, \infty)$, so that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is the natural filtration for \mathbf{X} . There are several important martingales associated with \mathbf{X} . We will study a couple of them in this section, and others in subsequent sections. Our first result is that \mathbf{X} itself is a martingale, simply by virtue of having stationary, independent increments and 0 mean.

\mathbf{X} is a martingale with respect to \mathfrak{F} .

Proof

Again, this is true of any process with stationary, independent increments and 0 mean, but we give the proof anyway, for completeness. Let $s, t \in [0, \infty)$ with $s < t$. Since X_s is measurable with respect to \mathcal{F}_s and $X_t - X_s$ is independent of \mathcal{F}_s we have

$$\mathbb{E}(X_t | \mathcal{F}_s) = \mathbb{E}[X_s + (X_t - X_s) | \mathcal{F}_s] = X_s + \mathbb{E}(X_t - X_s) = X_s \quad (18.1.17)$$

The next martingale is a little more interesting.

Let $Y_t = X_t^2 - t$ for $t \in [0, \infty)$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is a martingale with respect to \mathfrak{F} .

Proof

Let $s, t \in [0, \infty)$ with $s < t$. Then

$$Y_t = X_t^2 - t = [X_s + (X_t - X_s)]^2 - t = X_s^2 + 2X_s(X_t - X_s) + (X_t - X_s)^2 - t \quad (18.1.18)$$

Since X_s is measurable with respect to \mathcal{F}_s and $X_t - X_s$ is independent of \mathcal{F}_s we have

$$\mathbb{E}(Y_t | \mathcal{F}_s) = X_s^2 + 2X_s \mathbb{E}(X_t - X_s) + \mathbb{E}[(X_t - X_s)^2] - t \quad (18.1.19)$$

But $\mathbb{E}(X_t - X_s) = 0$ and $\mathbb{E}[(X_t - X_s)^2] = \text{var}(X_t - X_s) = t - s$ so $\mathbb{E}(Y_t | \mathcal{F}_s) = X_s^2 - s = Y_s$.

Maximums and Hitting Times

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. For $y \in [0, \infty)$ recall that $\tau_y = \min\{t \geq 0 : X_t = y\}$ is the first time that the process hits state y . Of course, $\tau_0 = 0$. For $t \in [0, \infty)$, let $Y_t = \max\{X_s : 0 \leq s \leq t\}$, the maximum value of \mathbf{X} on the interval $[0, t]$. Note that Y_t is well defined by the continuity of \mathbf{X} , and of course $Y_0 = 0$. Thus we have two new stochastic processes: $\{\tau_y : y \in [0, \infty)\}$ and $\{Y_t : t \in [0, \infty)\}$. Both have index set $[0, \infty)$ and (as we will see) state space $[0, \infty)$. Moreover, the processes are inverses of each other in a sense:

For $t, y \in (0, \infty)$, $\tau_y \leq t$ if and only if $Y_t \geq y$.

Proof

Since standard Brownian motion starts at 0 and is continuous, both events mean that the process hits state y in the interval $[0, t]$.

Thus, if we can compute the distribution of Y_t for each $t \in (0, \infty)$ then we can compute the distribution of τ_y for each $y \in (0, \infty)$, and conversely.

For $y > 0$, τ_y has the same distribution as y^2/Z^2 , where Z is a standard normal variable. The probability density function g_y is given by

$$g_y(t) = \frac{y}{\sqrt{2\pi t^3}} \exp\left(-\frac{y^2}{2t}\right), \quad t \in (0, \infty) \quad (18.1.20)$$

Proof

Let $t > 0$. From the [previous result](#), note that $X_t \geq y \implies Y_t \geq y \implies \tau_y \leq t$. Hence

$$\mathbb{P}(X_t \geq y) = \mathbb{P}(X_t \geq y, \tau_y \leq t) = \mathbb{P}(X_t \geq y | \tau_y \leq t) \mathbb{P}(\tau_y \leq t) \quad (18.1.21)$$

But from the [strong Markov property](#) above, $s \mapsto X(\tau_y + s) - y$ is another standard Brownian motion. Hence $\mathbb{P}(X_t \geq y | \tau_y \leq t) = \frac{1}{2}$. Therefore

$$\mathbb{P}(\tau_y \leq t) = 2\mathbb{P}(X_t \geq y) = \frac{2}{\sqrt{2\pi t}} \int_y^\infty e^{-x^2/2t} dx = \frac{2}{\sqrt{2\pi}} \int_{y/\sqrt{t}}^\infty e^{-z^2/2} dz \quad (18.1.22)$$

The second integral follows from the first by the change of variables $z = x/\sqrt{t}$. We can recognize this integral as $\mathbb{P}(y^2/Z^2 \leq t)$ where Z has a standard normal distribution. Taking the derivative of the integral with respect to t gives the PDF.

The distribution of τ_y is the *Lévy distribution* with scale parameter y^2 , and is named for the French mathematician Paul Lévy. The Lévy distribution is studied in more detail in the chapter on special distributions.

Open the hitting time experiment. Vary y and note the shape and location of the probability density function of τ_y . For selected values of the parameter, run the simulation in single step mode a few times. Then run the experiment 1000 times and compare the empirical density function to the probability density function.

Open the special distribution simulator and select the Lévy distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Standard Brownian motion is *recurrent*. That is, $\mathbb{P}(\tau_y < \infty) = 1$ for every $y \in \mathbb{R}$.

Proof

Suppose first that $y > 0$. From the proof of the [last theorem](#),

$$\mathbb{P}(\tau_y < \infty) = \lim_{t \rightarrow \infty} \mathbb{P}(\tau_y \leq t) = \frac{2}{\sqrt{2\pi}} \int_0^\infty e^{-z^2/2} dz = 1 \quad (18.1.23)$$

Note that the integral above is equivalent to the integral of the standard normal PDF over \mathbb{R} . In particular, the function g_y given above really is a valid PDF. If $y < 0$ then by symmetry, τ_y has the same distribution as τ_{-y} , so $\mathbb{P}(\tau_y < \infty) = 1$. Trivially, $\tau_0 = 0$.

Thus, for each $y \in \mathbb{R}$, \mathbf{X} eventually hits y with probability 1. Actually we can say more:

With probability 1, \mathbf{X} visits every point in \mathbb{R} .

Proof

By continuity, if \mathbf{X} reaches $y > 0$ then \mathbf{X} visits every point in $[0, y]$. By symmetry, a similar statement holds for $y < 0$. Thus the event that \mathbf{X} visits every point in \mathbb{R} is $\bigcap_{n=1}^\infty (\{\tau_n < \infty\} \cap \{\tau_{-n} < \infty\})$. The probability of a countable intersection of events with probability 1 still has probability 1.

On the other hand,

Standard Brownian motion is *null recurrent*. That is, $\mathbb{E}(\tau_y) = \infty$ for every $y \in \mathbb{R} \setminus \{0\}$.

Proof

By symmetry, it suffices to consider $y > 0$. From the result above on the [distribution of \$\tau_y\$](#) ,

$$\mathbb{E}(\tau_y) = \int_0^\infty \mathbb{P}(\tau_y > t) dt = \frac{2}{\sqrt{2\pi}} \int_0^\infty \int_0^{y/\sqrt{t}} e^{-z^2/2} dz dt \quad (18.1.24)$$

Changing the order of integration gives

$$\mathbb{E}(\tau_y) = \frac{2}{\sqrt{2\pi}} \int_0^\infty \int_0^{y^2/z^2} e^{-z^2/2} dt dz = \frac{2y^2}{\sqrt{2\pi}} \int_0^\infty \frac{1}{z^2} e^{-z^2/2} dz \quad (18.1.25)$$

Next we get a lower bound on the last integral by integrating over the interval $[0, 1]$ and noting that $e^{-z^2/2} \geq e^{-1/2}$ on this interval. Thus,

$$\mathbb{E}(\tau_y) \geq \frac{2y^2 e^{-1/2}}{\sqrt{2\pi}} \int_0^1 \frac{1}{z^2} dz = \infty \quad (18.1.26)$$

The process $\{\tau_x : x \in [0, \infty)\}$ has stationary, independent increments.

Proof

The proof relies on the temporal and spatial homogeneity of Brownian motion and the strong Markov property. Suppose that $x, y \in [0, \infty)$ with $x < y$. By continuity, \mathbf{X} must reach x before reaching y . Thus, $\tau_y = \tau_x + (\tau_y - \tau_x)$. But $\tau_y - \tau_x$ is the hitting time to $y - x$ for the process $t \mapsto X(\tau_x + t) - x$, and as [shown above](#), this process is also a standard Brownian motion, independent of $\mathcal{F}(\tau_x)$. Hence $\tau_y - \tau_x$ is independent of $\mathcal{F}(\tau_x)$ and has the same distribution as τ_{y-x} .

The family of probability density functions $\{g_x : x \in (0, \infty)\}$ is closed under convolution. That is, $g_x * g_y = g_{x+y}$ for $x, y \in (0, \infty)$.

Proof

This follows immediately from the previous theorem. A direct proof is an interesting exercise.

Now we turn our attention to the maximum process $\{Y_t : t \in [0, \infty)\}$, the “inverse” of the hitting process $\{\tau_y : y \in [0, \infty)\}$.

For $t > 0$, Y_t has the same distribution as $|X_t|$, known as the *half-normal distribution* with scale parameter t . The probability density function is

$$h_t(y) = \sqrt{\frac{2}{\pi t}} \exp\left(-\frac{y^2}{2t}\right), \quad y \in [0, \infty) \quad (18.1.27)$$

Proof

From the [inverse relation](#) and the [distribution](#) of τ_y , $\mathbb{P}(Y_t \geq y) = \mathbb{P}(\tau_y \leq t) = 2\mathbb{P}(X_t \geq y) = \mathbb{P}(|X_t| \geq y)$ for $y \geq 0$. By definition, $|X_t|$ has the half-normal distribution with parameter t . In particular,

$$\mathbb{P}(Y_t \geq y) = \frac{2}{\sqrt{2\pi t}} \int_y^\infty e^{-x^2/2t} dx \quad (18.1.28)$$

Taking the negative derivative of the integral above, with respect to y , gives the PDF.

The half-normal distribution is a special case of the folded normal distribution, which is studied in more detail in the chapter on special distributions.

For $t \geq 0$, the mean and variance of Y_t are

1. $\mathbb{E}(Y_t) = \sqrt{\frac{2t}{\pi}}$
2. $\text{var}(Y_t) = t \left(1 - \frac{2}{\pi}\right)$

Proof

These follow from standard results for the half-normal distribution.

In the standard Brownian motion simulation, select the maximum value. Vary the parameter t and note the shape of the probability density function and the location and size of the mean-standard deviation bar. Run the simulation 1000 times and compare the empirical density and moments to the true probability density function and moments.

Open the special distribution simulator and select the folded-normal distribution. Vary the parameters and note the shape and location of the probability density function and the size and location of the mean-standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

Zeros and Arcsine Laws

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Study of the zeros of \mathbf{X} lead to a number of probability laws referred to as *arcsine laws*, because as we might guess, the probabilities and distributions involve the arcsine function.

For $s, t \in [0, \infty)$ with $s < t$, let $E(s, t)$ be the event that \mathbf{X} has a zero in the time interval (s, t) . That is, $\mathbb{E}(s, t) = \{X_u = 0 \text{ for some } u \in (s, t)\}$. Then

$$\mathbb{P}[E(s, t)] = 1 - \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right) \quad (18.1.29)$$

Proof

Conditioning on X_s and using symmetry gives

$$\mathbb{P}[E(s, t)] = \int_{-\infty}^{\infty} \mathbb{P}[E(s, t) \mid X_s = x] f_s(x) dx = 2 \int_{-\infty}^0 \mathbb{P}[E(s, t) \mid X_s = x] f_s(x) dx \quad (18.1.30)$$

But by the homogeneity of \mathbf{X} in time and space, note that for $x > 0$, $\mathbb{P}[E(s, t) \mid X_s = -x] = \mathbb{P}(\tau_x < t - s)$. That is, a process in state $-x$ at time s that hits 0 before time t is the same as a process in state 0 at time 0 reaching state x before time $t - s$. Hence

$$\mathbb{P}[E(s, t)] = \int_0^\infty \int_0^{t-s} g_x(u) f_s(-x) du dx \quad (18.1.31)$$

where g_x is the PDF of τ_x given [above](#). Substituting gives

$$\mathbb{P}[E(s, t)] = \frac{1}{\pi\sqrt{s}} \int_0^{t-s} u^{-3/2} \int_0^\infty x \exp\left[-\frac{1}{2}x^2 \left(\frac{u+s}{us}\right)\right] dx du = \frac{\sqrt{s}}{\pi} \int_0^{t-s} \frac{1}{(u+s)\sqrt{u}} du \quad (18.1.32)$$

Finally substituting $v = \sqrt{u/s}$ in the last integral give

$$\mathbb{P}[E(s, t)] = \frac{2}{\pi} \int_0^{\sqrt{t/s-1}} \frac{1}{v^2 + 1} dv = \frac{2}{\pi} \arctan\left(\sqrt{\frac{t}{s}-1}\right) = 1 - \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right) \quad (18.1.33)$$

In particular, $\mathbb{P}[E(0, t)] = 1$ for every $t > 0$, so with probability 1, \mathbf{X} has a zero in $(0, t)$. Actually, we can say a bit more:

For $t > 0$, \mathbf{X} has infinitely many zeros in $(0, t)$ with probability 1.

Proof

The event that \mathbf{X} has infinitely many zeros in $(0, t)$ is $\bigcap_{n=1}^{\infty} E(0, t/n)$. The intersection of a countable collection of events with probability 1 still has probability 1.

The last result is further evidence of the very strange and irregular behavior of Brownian motion. Note also that $\mathbb{P}[E(s, t)]$ depends only on the ratio s/t . Thus, $\mathbb{P}[E(s, t)] = \mathbb{P}[E(1/t, 1/s)]$ and $\mathbb{P}[E(s, t)] = \mathbb{P}[E(cs, ct)]$ for every $c > 0$. So, for example the probability of at least one zero in the interval $(2, 5)$ is the same as the probability of at least one zero in $(1/5, 1/2)$ the same as the probability of at least one zero in $(6, 15)$, and the same as the probability of at least one zero in $(200, 500)$.

For $t > 0$, let Z_t denote the time of the last zero of \mathbf{X} before time t . That is, $Z_t = \max\{s \in [0, t] : X_s = 0\}$. Then Z_t has the *arcsine distribution* with parameter t . The distribution function H_t and the probability density function h_t are given by

$$H_t(s) = \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right), \quad 0 \leq s \leq t \quad (18.1.34)$$

$$h_t(s) = \frac{1}{\pi \sqrt{s(t-s)}}, \quad 0 < s < t \quad (18.1.35)$$

Proof

For $0 \leq s < t$, the event $Z_t \leq s$ is the same as $\neg [E(s, t)]^c$, that there are no zeros in the interval (s, t) . Hence the formula for H_t follows from the [result above](#). Taking the derivative of H_t and simplifying gives the formula for h_t .

The density function of Z_t is *u-shaped* and symmetric about the midpoint $t/2$, so the points with the largest density are those near the endpoints 0 and t , a surprising result at first. The arcsine distribution is studied in more detail in the chapter on special distributions.

The mean and variance of Z_t are

1. $\mathbb{E}(Z_t) = t/2$
2. $\mathbb{E}(Z_t) = t^2/8$

Proof

These are standard results for the arcsine distribution. That the mean is the midpoint $t/2$ also follows from symmetry, of course.

In the simulation of standard Brownian motion, select the last zero variable. Vary the parameter t and note the shape of the probability density function and the size and location of the mean-standard deviation bar. For selected values of t run the simulation in single step mode a few times and note the position of the last zero. Finally, run the simulation 1000 times and compare the empirical density function and moments to the true probability density function and moments.

Open the special distribution simulator and select the arcsine distribution. Vary the parameters and note the shape and location of the probability density function and the size and location of the mean-standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

Now let $Z = \{t \in [0, \infty) : X_t = 0\}$ denote the set of zeros of \mathbf{X} , so that Z is a random subset of $[0, \infty)$. The theorem below gives some of the strange properties of the random set Z , but to understand these, we need to review some definitions. A *nowhere dense* set is a set whose closure has empty interior. A *perfect set* is a set with no isolated points. As usual, we let λ denote Lebesgue measure on \mathbb{R} .

With probability 1,

1. Z is closed.
2. $\lambda(Z) = 0$
3. Z is nowhere dense.
4. Z is perfect.

Proof

1. Note that Z is the inverse image of the closed set $\{0\}$ under the function $t \mapsto X_t$. Since this function is continuous with probability 1, Z is closed with probability 1.

2. For each $t \in (0, \infty)$ note that $\mathbb{P}(t \in Z) = \mathbb{P}(X_t = 0) = 0$ since X_t has a continuous distribution. Using Fubini's theorem

$$\mathbb{E}[\lambda(Z)] = \mathbb{E}\left[\int_0^\infty \mathbf{1}_Z(t) d\lambda(t)\right] = \int_0^\infty \mathbb{E}[\mathbf{1}_Z(t)] d\lambda(t) = 0 \quad (18.1.36)$$

and hence $\mathbb{P}[\lambda(Z) = 0] = 1$,

3. Since Z is closed and has Lebesgue measure 0, its interior is empty (all of these statements with probability 1).

4. Suppose that $s \in Z$. Then by the [temporal and spatial homogeneity](#) properties, $t \mapsto X_{s+t}$ is also a standard Brownian motion. But then by the result above on [zeros](#), with probability 1, X has a zero in the interval $(s, s + 1/n)$ for every $n \in \mathbb{N}_+$. Hence s is not an isolated point of Z .

The following theorem gives a deeper property of Z . The Hausdorff dimension of Z is midway between that of a point (dimension 0) and a line (dimension 1).

Z has Hausdorff dimension $\frac{1}{2}$.

The Law of the Iterated Logarithm

As usual, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be standard Brownian motion. By definition, we know that X_t has the normal distribution with mean 0 and standard deviation \sqrt{t} , so the function $x = \sqrt{t}$ gives some idea of how the process grows in time. The precise growth rate is given by the famous *law of the iterated logarithm*

With probability 1,

$$\limsup_{t \rightarrow \infty} \frac{X_t}{\sqrt{2t \ln \ln t}} = 1 \quad (18.1.37)$$

Computational Exercises

In the following exercises, $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a standard Brownian motion process.

Explicitly find the probability density function, covariance matrix, and correlation matrix of $(X_{0.5}, X_1, X_{2.3})$.

This page titled [18.1: Standard Brownian Motion](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.