

Week 4

Brownian motion and the heat equation

Jonathan Goodman

September 30, 2013

1 Introduction to the material for the week

A *diffusion process* is a Markov process in continuous time with a continuous state space and continuous sample paths. This course is largely about diffusion processes. Partial differential equations (*PDEs*) of *diffusion type* are important tools for studying diffusion processes. Conversely, diffusion processes give insight into solutions of diffusion type partial differential equations. We have seen two diffusion processes so far, Brownian motion and the Ornstein Uhlenbeck process. This week, we discuss the partial differential equations associated with these two processes.

We start with the forward equation associated with Brownian motion. Let X_t be a standard Brownian motion with probability density $u(x, t)$. This probability density satisfies the *heat equation*, or *diffusion equation*, which is

$$\partial_t u = \frac{1}{2} \partial_x^2 u . \quad (1)$$

This PDE allows us to solve the *initial value problem*. Suppose s is a time and the probability density $X_s \sim u(x, s)$ is known, then (1) determines $u(x, t)$ for $t \geq s$. The initial value problem has a solution for more or less any *initial condition* $u(x, s)$. If $u(x, s)$ is a probability density, you can find $x(x, t)$ for $t > s$ by: first choosing $X_s \sim u(\cdot, s)$, then letting X_t for $t > s$ be a Brownian motion. The probability density of X_t , $u(x, t)$ satisfies the heat equation. By contrast, the heat generally cannot be run “backwards”. If you give a probability density $u(x, s)$, there probably is no function $u(x, t)$ defined for $t < s$ that satisfies the heat equation for $t < s$ and the specified values $u(x, s)$. Running the heat equation backwards is *ill posed*.¹

The Brownian motion interpretation provides a solution formula for the heat equation

$$u(x, t) = \frac{1}{\sqrt{2\pi(t-s)}} \int_{-\infty}^{\infty} e^{-(x-y)^2/2(t-s)} u(y, s) ds . \quad (2)$$

¹Stating a problem or task is *posing* the problem. If the task or mathematical problem has no solution that makes sense, the problem is poorly stated, or *ill posed*.

This formula may be expressed more abstractly as

$$u(x, t) = \int_{-\infty}^{\infty} G(x - y, t - s)u(y, s) ds , \quad (3)$$

where the function

$$G(x, t) = \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t}$$

is called the *fundamental solution*, or the *heat kernel*, or the *transition density*. You recognize it as the probability of a Gaussian with mean zero and variance t . This is the probability density of X_t if $X_0 = 0$ and X is standard Brownian motion.

We can think of the function $u(x, t)$ as an abstract vector and write it $u(t)$. We did this already in Week 2, where the occupation probabilities $u_{n,j} = P(X_n = j)$ were thought of as components of the row vector u_n . The solution formula (3) produces the function $u(t)$ from the data $u(s)$. We write this abstractly as

$$u(t) = G(t - s)u(s) . \quad (4)$$

The *operator*, G , is something like an infinite dimensional matrix. The abstract expression (4) is shorthand for the more concrete formula (3), just as matrix multiplication is shorthand for the actual sums involved. The particular operators $G(t)$ have the *semigroup property*

$$G(t) = G(t - s)G(s) , \quad (5)$$

as long as t , s , and $t - s$ are all positive.² This is because $u(t) = G(t)u(0)$, and $u(s) = G(s)u(0)$, so $u(t) = G(t - s)[G(s)u(0)] = [G(t - s)G(s)]u(0)$.

You can make a long list of ways the heat equation helps understand the behavior of Brownian motion. We can write formulas for *hitting probabilities* by writing solutions of (1) that satisfy the correct *boundary conditions*. This will allow us to explain the simulation results in question (4c) of assignment 3. You do not have to understand probability to check that a function $u(x, t)$ satisfies the heat equation, only calculus.

The *backward equation* for Brownian motion is

$$\partial_t f + \frac{1}{2}\partial_x^2 f = 0 . \quad (6)$$

This is the equation satisfied by expected values

$$f(x, t) = E[V(X_T)|X_t = x] , \quad (7)$$

if $T \geq t$. The *final condition* is the obvious statement that $f(x, T) = V(x)$. The PDE (6) allows you to move backwards in time to determine values of f

²A mathematical *group* is a collection of objects that you can multiply and invert, like the group of invertible matrices of a given size. A *semigroup* allows multiplication but not necessarily inversion. If the operators $G(t)$ were defined for $t < 0$ and the formula (5) still applied, then the operators would form a group. Our operators are only half a group because they are defined only for $t \geq 0$.

for $t < T$ from $f(T)$. The backward equation differs from the forward equation only by a sign, but this is a big difference. Moving forward with the backward equation is just as ill posed as moving backward with the forward. In particular, suppose you have a desired function $f(x, 0)$ and you want to know what function $V(x)$ gives rise to it using (7). Unless your function $f(0)$ is very special (details below), there is no V at all.

2 The heat equation

This section describes the heat equation and some of its solutions. This will help us understand Brownian motion, both qualitatively (general properties) and quantitatively (specific formulas).

The heat equation is used to model things other than probability. For example it can be the flow of heat in a metal rod. Here, $u(x, t)$ is the temperature at location x at time t . The temperature is modeled by $\partial_t u = D \partial_x^2 u$, where the *diffusion coefficient*, D , depends on the material (metal, stone, ..), and the units (seconds, days, centimeters, meters, degrees C , ..). The heat equation has the value $D = \frac{1}{2}$. Changing units, or *rescaling*, or *non-dimensionalizing* can replace D with $\frac{1}{2}$. For example, you can use $t' = Dt$, or $x' = \sqrt{D}x$.

The heat flow picture suggests that heat will flow from high temperature to low temperature regions. The fluctuations in $u(x, t)$ will smooth out and relax over time and the heat redistributes itself. The total amount of heat in an interval $[a, b]$ at time t is³

$$\int_a^b u(x, t) dx .$$

You understand the flow of heat by differentiating with respect to time and using the heat equation

$$\frac{d}{dt} \int_a^b u(x, t) dx = \int_a^b \partial_t u(x, t) dx = \frac{1}{2} \int_a^b \partial_x^2 u(x, t) dx = \frac{1}{2} (\partial_x u(b, t) - \partial_x u(a, t)) .$$

The *heat flux*,

$$F(x, t) = -\frac{1}{2} \partial_x u(x, t) , \tag{8}$$

puts this into the *conservation form*

$$\frac{d}{dt} \int_a^b u(x, t) dx = \int_a^b \partial_t u(x, t) dx = F(a, t) - F(b, t) . \tag{9}$$

The heat flux (8) is the rate at which heat is flowing across x at time t . If F is positive, heat flows from left to right. The specific formula (8) is *Fick's law*, which says that heat flows downhill toward lower temperature at a rate

³If you are one of those people who knows the technical distinction between heat and temperature, I say “choose units of temperature in which the specific heat is one”.

proportional to the temperature gradient. If $\partial_x u > 0$, then heat flows from right to left in the direction opposite the temperature gradient. The conservation equation (9) gives the rate of change of the amount of heat in $[a, b]$ as the rate of flow in, $F(a, t)$, minus the rate of flow out, $F(b, t)$. Of course, either of these numbers could be negative.

The heat equation has a family of solutions that are exponential in “space” (the x variable). These are

$$u(x, t) = A(t)e^{ikx} . \quad (10)$$

This is an *ansatz*, which is a hypothesized functional form for the solution. Calculating the time and space derivatives, this ansatz satisfies the heat equation if (here $\dot{A} = dA/dt$)

$$\dot{A}e^{ikx} = \frac{1}{2}(-k^2)Ae^{ikx} .$$

We cancel the common exponential factor and see that (10) is a solution if

$$\dot{A} = -\frac{1}{2}k^2 A .$$

This leads to $A(t) = A(0)e^{-k^2 t/2}$, and

$$u(x, t) = e^{ikx} e^{-k^2 t/2} .$$

The formula $e^{i\theta} = \cos(\theta) + i \sin(\theta)$ tells us that the real part of u is

$$v(x, t) = \cos(kx) e^{-k^2 t/2} .$$

You cannot v as a probability density because it has negative values. But it gives insight into that the heat equation does. A large k , which is a high *wave number*, or (less accurately) *frequency*, leads to rapid decay, $e^{-k^2 t/2}$. This is because positive and negative “heat” is close together and does not have to diffuse far to cancel out.

Another function that satisfies the heat equation is

$$u(x, t) = t^{-1/2} e^{-x^2/(2t)} . \quad (11)$$

The relevant calculations are

$$\begin{aligned} u &\xrightarrow{\partial_t} \left\{ \partial_t t^{-1/2} \right\} e^{-x^2/(2t)} + t^{-1/2} \left\{ e^{-x^2/(2t)} \right\} \\ &= -\frac{1}{2} t^{-1} u + \frac{x^2}{2t^2} u , \end{aligned}$$

and

$$\begin{aligned} u &\xrightarrow{\partial_x} -\frac{x}{t} t^{-1/2} e^{-x^2/(2t)} \\ &\xrightarrow{\partial_x} -\frac{1}{t} u + \frac{x^2}{t^2} u . \end{aligned}$$

This shows that $\partial_t u$ does equal $\frac{1}{2}\partial_x^2 u$. This solution illustrates the spreading of heat. The maximum of u is $1/\sqrt{t}$, which is at $x = 0$. This is large for small t and goes to zero as $t \rightarrow \infty$. We see the characteristic width by writing u as

$$u(x, t) = \frac{1}{\sqrt{t}} e^{-\frac{1}{2}\left(\frac{x}{\sqrt{t}}\right)^2} .$$

This gives $u(x, t)$ as a function of the *similarity variable* $\frac{x}{\sqrt{t}}$, except for the outside overall scale factor $\frac{1}{\sqrt{t}}$. Therefore, the characteristic width is on the order of \sqrt{t} . This is the order of the distance in x you have to go to get from the maximum value ($x = 0$) to, say, half the maximum value.

The heat equation preserves total heat in the sense that

$$\frac{d}{dt} \int_{-\infty}^{\infty} u(x, t) dx = 0 . \quad (12)$$

This follows from the conservation law (8) and (9) if $\partial_x u \rightarrow 0$ as $x \rightarrow \pm\infty$. You can check by direct integration that the Gaussian solution (11) satisfies this global conservation. But there is a dumber way. The total mass of the bump shaped Gaussian heat distribution (11) is roughly equal to the height multiplied by the width of the bump. The height is $t^{-1/2}$ and the width is $t^{1/2}$. The product is a constant.

There are methods for building general solutions of the heat equation from particular solutions such as the plane wave (10) or the Gaussian (11). The heat equation PDE is *linear*, which means that if $u_1(x, t)$ and $u_2(x, t)$ are solutions, then $u(x, t) = c_1 u_1(x, t) + c_2 u_2(x, t)$ is also a solution. This is the *superposition principle*. The graph of $c_1 u_1 + c_2 u_2$ is the *superposition* (one on top of the other) of the graphs of u_1 and u_2 . The equation is *translation invariant*, or *homogeneous in space and time*, which means that if $u(x, t)$ is a solution, then $v(x, t) = u(x - x_0, t - t_0)$ is a solution. The equation has the *scaling property* that if $u(x, t)$ is a solution, then $u_\lambda(x, t) = u(\lambda x, \lambda^2 t)$ is a solution. This scaling relation is “one power of t is two powers of x ”, or x^2 scales like t .

Here are some simple illustrations. You can put a Gaussian bump of any height and with any center:

$$\frac{c}{\sqrt{t}} e^{-(x-x_0)^2/2t} .$$

You can combine (superpose) them to make “multibump” solutions

$$u(x, t) = \frac{c_1}{\sqrt{t}} e^{-(x-x_1)^2/2t} + \frac{c_2}{\sqrt{t}} e^{-(x-x_2)^2/2t} .$$

The $k = 1$ plane wave $u(x, t) = \sin(x)e^{-t/2}$ may be rescaled to give the general plane wave: $u_\lambda(x, t) = \sin(\lambda x)e^{-\lambda^2 t/2}$, which is the same as (10). Changing the length scale by a factor of λ changes the time scale, which is the decay rate in this case, by a factor of λ^2 . The Gaussian solutions are *self similar* in the sense that $u_\lambda(x, t) = C_\lambda u(x, t)$. The exponent calculation is $x^2/t \rightarrow (\lambda^2 x^2)/(\lambda^2 t)$.

The solution formula (2) is an application of the superposition principle with integrals instead of sums. We explain it here, and take $s = 0$ for simplicity. The total heat (or total probability, or total mass, depending on the interpretation) of the Gaussian bump is $\sqrt{2\pi}$. You can see that simply by taking $t = 1$. It is simpler to work with a Gaussian solution with total mass equal to one. When you center the normalized bump at a point y , you get

$$\frac{1}{\sqrt{2\pi t}} e^{-(x-y)^2/2t} . \quad (13)$$

As $t \rightarrow 0$, this solution concentrates all its heat in a collapsing neighborhood of y . Therefore, it is the solution that results from an initial condition that concentrates a unit amount of heat at the point y . This is expressed using the *Dirac delta* function as $u(x, t) \rightarrow \delta(x - y)$ as $t \rightarrow 0$. It shows that the normalized, centered Gaussian (13) is the solution to the initial value problem for the heat equation with initial condition $u(x, 0) = \delta(x - y)$. More generally, the formula $(c/\sqrt{2\pi t}) e^{-(x-y)^2/2t}$ says what happens at later time to an amount c of heat at y at time zero. For general initial heat distribution $u(y, 0)$, the amount of heat in a dy neighborhood of y is $u(y, 0)dy$. This contributes $(u(y, 0)/\sqrt{2\pi t}) e^{-(x-y)^2/2t}$ to the solution $u(x, t)$. We get the total solution by adding all these contribution. The result is

$$u(x, t) = \int_{y=-\infty}^{\infty} u(y, 0) \frac{1}{\sqrt{2\pi t}} e^{-(x-y)^2/2t} .$$

This is the formula (2).

3 The forward equation and Brownian motion

We argue that the probability density of Brownian motion satisfies the heat equation (1). Suppose $u_0(x)$ is a probability density and we choose $X_0 \sim u_0$. Suppose we then start a Brownian motion path from X_0 . Then $X_t - X_0 \sim \mathcal{N}(0, t)$ and the joint density of X_0 and X_t is

$$u(x_0, x_t, t) = u_0(x_0) \frac{1}{\sqrt{2\pi t}} e^{-(x_t - x_0)^2/2t} .$$

The probability density of X_t is the integral of the joint density

$$u(x_t, t) = \int u(x_0, x_t, t) dx_0 = \int u_0(x_0) \frac{1}{\sqrt{2\pi t}} e^{-(x_t - x_0)^2/2t} dx_0 .$$

If you substitute y for x_0 and x for x_t , you get (2). This shows that the probability density of X_t is equal to the solution of the heat equation evaluated at time t .

4 Hitting probabilities and hitting times

If X_t is a stochastic process with continuous sample paths, the *hitting time* for a closed set A is $\tau_A = \min \{t \mid X_t \in A\}$. This is a random variable because

the hitting time depends on the path. For one dimensional Brownian motion starting at $X_0 = 0$, we define $\tau_a = \min \{t \mid X_t = a\}$. Let $f_a(t)$ be the probability density of τ_a . We will find formulas for $f_a(t)$ and the *survival probability* $S_a(t) = P(\tau_a \geq t)$. Clearly $f_a(t) = -\partial_t S_a(t)$, the survival is (up to a constant) the negative of the CDF of τ .

There are two related approaches to hitting times and survival probabilities for Brownian motion in one dimension. One uses the heat equation with a boundary condition at a . The other uses the *Kolmogorov reflection principle*. The reflection principle seems simpler, but it has two drawbacks. One is that I have no idea how Kolmogorov could have discovered it without first doing it the hard way, with the PDE. The other is that the PDE method is more general.

The PDE approach makes use of the PDF of surviving particles. This is defined by

$$P(X_t \in [x, x + dx] \mid \tau_a > t) = u_a(x, t) dx . \quad (14)$$

Stopped Brownian motion gives a different description of the same thing. This is

$$Y_t = \begin{cases} X_t & \text{if } t < \tau_a \\ a & \text{if } t \geq \tau_a \end{cases}$$

The process moves with X_t until X touches a , then it stops. The density (14) is the density of Y_t except at $x = a$, where the Y density has a δ component. Another notation for stopped Brownian motion uses the *wedge* notation $t \wedge s = \min(t, s)$. The formula is $Y_t = X_{t \wedge \tau_a}$. If $t < \tau_a$, this gives $Y_t = X_t$. If $t \geq \tau_a$, this gives $Y_t = X_{\tau_a}$, which is a because τ_a is the hitting time of a .

The conditional probability density $u_a(x, t)$ satisfies the heat equation except at a . We do not give a proof of this, only some plausibility arguments. First, (see this week's homework assignment), it is true for a stopped random walk approximation to stopped Brownian motion. Second, if X_t is not at a and has not been stopped, then it acts like ordinary Brownian motion, at least for a short time. In particular,

$$u_a(x, t + \Delta t) \approx \int G(x - y, \Delta t) u_a(y, t) dy ,$$

if Δt is small. The right side satisfies the heat equation, so the left should as well if Δt is small.

The conditional probability satisfies the *boundary condition* $u_a(x, t) \rightarrow 0$ as $x \rightarrow a$. This would be the same as $u(a, t) = 0$ if we knew that u was continuous (it is but we didn't show it). The boundary condition $u(a, t) = 0$ is called an *absorbing* boundary condition because it represents the physical fact that particles that touch a get stuck and do not re-enter the region $x \neq a$. We will not give a proof that the density for stopped Brownian motion satisfies the absorbing boundary condition, but we give two plausibility arguments. The first is that it is true in the approximating stopped random walk. The second involves the picture of Brownian motion as constantly moving back and forth. It (almost) never moves in the same direction for a positive amount of time. If $X_t = a$, then (almost surely) there are times $t_1 < t$ and $t_2 < t$ so that $X_{t_1} > a$

and $X_{t_2} < a$. The closer y is close to a , the less likely it is that $X_s \neq a$ for all $s < t$.

Accepting the two above claims, we can find hitting probabilities by finding solutions of the heat equation with absorbing boundary conditions. Let us assume that $X_0 = 0$ and the *absorbing boundary* is at $a > 0$. We want a function $u_a(x, t)$ that is defined for $x \leq a$ that satisfies the initial condition $u_a(x, t) \rightarrow \delta(x)$ as $t \rightarrow 0$, for $(x < a)$ and the absorbing boundary condition $u_a(a, t) = 0$. The trick that does this is the *method of images* from physics. A point $x < a$ has an *image point*, $x' > a$, that is the same distance from a . The image point is $x' = a + (a - x) = 2a - x$. If $x < a$, then $x' > a$ and $|x - a| = |x' - a|$, and $x' \rightarrow a$ as $x \rightarrow a$. The density function $u_a(x, t)$ starts out defined only for $x \leq a$. The trick is to *extend* the definition of u_a beyond a by *odd reflection*. That is,

$$u_a(x', t) = -u_a(x, t) . \quad (15)$$

The oddness of the extended function implies that $u_a(x, t) \rightarrow 0$ as $x \rightarrow a$ from either direction. The only direction we originally cared about was from $x < a$, but the other is true also.

We create an odd solution of the heat equation by taking odd initial data. We know $u_a(x, 0)$ needs a point mass at $x = 0$. To make the initial data odd, we add a negative point mass also at the image of 0, which is $x_* = 2a$. The resulting initial data is

$$u_a(x, 0) = \delta(x) - \delta(x - 2a) .$$

The initial data has changed, but the part for $x \leq a$ is the same. The solution is the superposition of the pieces from the two delta functions:

$$u_a(x, t) = \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t} - \frac{1}{\sqrt{2\pi t}} e^{-(x-2a)^2/2t} . \quad (16)$$

This function satisfies all three of our requirements. It has the right initial data, at least for $x \leq a$. It satisfies the heat equation for all $x \leq a$. It satisfies the heat equation also for $x > a$, which is interesting but irrelevant. It satisfies the absorbing boundary condition. It is a continuous function of x for $t > 0$ and has $u_a(a, t) = 0$.

The formula (16) answers many questions about stopped Brownian motion and absorbing boundaries. The survival probability at time t is

$$S_a(t) = \int_{-\infty}^a u_a(x, t) dx . \quad (17)$$

This is because u_a was the probability density of surviving Brownian motion paths. You can check that the method of images formula (16) has $u_a(x, t) > 0$

if $x < a$. The probability density of τ_a is

$$\begin{aligned}
 f_a(t) &= -\frac{d}{dt} S_a(t) \\
 &= -\int_{-\infty}^a \partial_t u_a(x, t) dx \\
 &= -\frac{1}{2} \int_{-\infty}^a \partial_x^2 u_a(x, t) dx \\
 f_a(t) &= -\frac{1}{2} \partial_x u_a(a, t) .
 \end{aligned} \tag{18}$$

Without using the specific formula (16) we know the right side of (18) is positive. That is because $u_a(x, t)$ is going from positive values for $x < a$ to zero when $x = a$. That makes $\partial_x u_a(a, t)$ negative (at least not positive) and $f(t)$ positive (at least not negative). The formula (18) reinforces the interpretation (see (8) of $-\frac{1}{2} \partial_x u$ as a probability flux. It is the rate at which probability leaves the *continuation region*, $x < 0$.

The formula for the hitting time probability density is found by differentiating (16) with respect to x and setting $x = a$. The two terms from the right turn out to be equal.

$$f_a(t) = \sqrt{\frac{2}{\pi}} \frac{a}{t^{3/2}} e^{-a^2/2t} . \tag{19}$$

The reader is invited to verify by explicit integration that

$$\sqrt{\frac{2}{\pi}} a \int_0^\infty \frac{1}{t^{3/2}} e^{-a^2/2t} dt = 1 .$$

This illustrates some features of Brownian motion.

Look at the formula as $t \rightarrow 0$. The exponent has t in the denominator, so $f_a(t) \rightarrow 0$ as $t \rightarrow 0$ exponentially. It is extremely, exponentially, unlikely for X_t to hit a in a short time. The probability starts being significantly different from zero when the exponent is not a large negative number, which is when t is on the order of a^2 . This is the (time) = (length)² aspect of Brownian motion.

Now look at the formula as $t \rightarrow \infty$. The exponent converges to zero, so $f_a(t) \approx Ct^{-3/2}$. (We know the constant, but it just gets in the way.) This integrates to a statement about the survival probability

$$S_a(t) = \int_t^\infty f_a(t') dt' \approx C \int_t^\infty t'^{3/2} dt' = Ct^{-1/2} .$$

The probability to survive a long time goes to zero as $t \rightarrow \infty$, but slowly as $1/\sqrt{t}$.

The maximum of a Brownian motion up to time t is

$$M_t = \max_{0 \leq s \leq t} X_s .$$

The hitting time formulas above also give formulas for the distribution of M_t . Let $G_t(a) = P(M_t \leq a)$ be the CDF of M_t . This is nearly the same as a survival

probability. Suppose $X_0 = 0$ and $a > 0$ as above. Then the event $M_t < a$ is the same as $X_s < a$ for all $s \in [0, t]$, which is the same as $\tau_a > t$. Let $g_t(a)$ be the PDF of M_t . Then $g(a) = \frac{d}{da}G(a)$. Since $G_t(a) = S_a(t)$, we can find g by differentiating the formula (17) with respect to a and using (16). This is not hard, but it is slightly involved because $S_a(t)$ depends on a in two ways – the limit of integration and the integrand $u_a(x, t)$.

There is another approach through the *Kolmogorov reflection principle*. This is a re-interpretation of the survival probability integral (17). Start with the observation that

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t} dx = 1 .$$

The integral (17) is less than 1 for two reasons. One reason is that the survival probability integral omits the part of the above integral from $x > a$. The other is the negative contribution from the image “charge”. It is obvious (draw a picture) that

$$\int_a^{\infty} \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t} dx = \int_{-\infty}^a \frac{1}{\sqrt{2\pi t}} e^{-(x-2a)^2/2t} dx$$

Also, the left side is $P(X_t > a)$. Therefore

$$S_a(t) = 1 - 2P(X_t > a) . \tag{20}$$

We derived this formula by calculating integrals. But once we see it we can look for a simple explanation.

The simple explanation given by Kolmogorov depends on two properties of Brownian motion: it is symmetric (as likely to go up by ΔX as down by ΔX), and it is Markov (after it hits level a , it continues as a Brownian motion starting at a). Let \mathcal{P}_a be the set of paths that reach the level a before time t . The *reflection principle* is the symmetry condition that

$$P(X_t > a \mid X_{[0,t]} \in \mathcal{P}_a) = P(X_t < a \mid X_{[0,t]} \in \mathcal{P}_a) .$$

This says that a path that touches level a at some time $\tau < t$ is equally likely to be outside at time t ($X_t > a$) as inside ($X_t < a$). If τ_a is the hitting time, then $X_{\tau_a} = a$. If $\tau_a < t$ then the probabilities for the path from time τ_a to time t are symmetric about a . In particular, the probabilities to be above a and below a are the same. A more precise version of this argument would say that if $s < t$, then

$$P(X_t > a \mid \tau_a = s) = P(X_t < a \mid \tau_a = s) ,$$

then integrate over s in the range $0 \leq s \leq t$. But it takes some mathematical work to define the conditional probabilities, since $P(\tau_a = s) = 0$ so you cannot use Bayes’ rule directly. Anyway, the reflection principle says that exactly half of the paths (half in the sense of probability) that ever touch the level a are above level a at time t . That is exactly (20).

5 Backward equation for Brownian motion

The backward equation is a PDE satisfied by conditional probabilities. Suppose there is a “reward” function $V(x)$ and you receive $V(X_t)$ depending on the value of a Brownian motion path. The *value function* is the conditional expectation of the reward given a location at time $t < T$:

$$f(x, t) = \mathbb{E}[V(X_T) \mid X_t = x] . \quad (21)$$

There other common notations for this. The expression $\mathbb{E}_*[\dots]$ means that the expectation is taken with respect to the probability distribution in the subscript. For example, if $Y \sim \mathcal{N}(\mu, \sigma^2)$, we might write

$$\mathbb{E}_{\mu, \sigma^2}[e^Y] = e^{\mu + \sigma^2/2} .$$

We write $\mathbb{E}_{x,t}[\dots]$ for expectation with respect to paths X with $X_t = x$. The value function in this notation is

$$f(x, t) = \mathbb{E}_{x,t}[V(X_T)] .$$

A third equivalent way uses the filtration associated with X , which is \mathcal{F}_t . The random variable $\mathbb{E}[\dots \mid \mathcal{F}_t]$ is a function of $X_{[0,t]}$. The Markov property simplifies $X_{[0,t]}$ to X_t if the random variable depends only on the future of t . Therefore, $\mathbb{E}[V(X_T) \mid \mathcal{F}_t]$ is a function of X_t , which we call $f(x, t)$. Therefore the following definition is equivalent to (21):

$$f(X_t, t) = \mathbb{E}[V(X_T) \mid \mathcal{F}_t] . \quad (22)$$

The backward equation satisfied by f may be derived using the tower property. This can be used to compare $f(\cdot, t)$ to $f(\cdot, t + \Delta t)$ for small Δt . The “physics” behind this is that ΔX will be small too, so $f(x, t)$ can be determined from $f(x + \Delta x, t + \Delta t)$, at least approximately, using Taylor series. These relations become exact in the limit $\Delta t \rightarrow 0$.

The σ -algebra $\mathcal{F}_{t+\Delta t}$ has a little more information than \mathcal{F}_t . Therefore, if Y is any random variable

$$\mathbb{E}[\mathbb{E}[Y \mid \mathcal{F}_{t+\Delta t}] \mid \mathcal{F}_t] = \mathbb{E}[Y \mid \mathcal{F}_t] .$$

We apply this general principle with $Y = V(X_T)$ and make use of (22), which leads to

$$\mathbb{E}[f(X_{t+\Delta t}, t + \Delta t) \mid \mathcal{F}_t] = f(X_t) .$$

We write $X_{t+\Delta t} = X_t + \Delta x$ and expand $f(X_{t+\Delta t}, t + \Delta t)$ in a Taylor series.

$$\begin{aligned} f(X_{t+\Delta t}, t + \Delta t) &= f(X_t, t) \\ &\quad + \partial_x f(X_t, t) \Delta X \\ &\quad + \partial_t f(X_t, t) \Delta t \\ &\quad + \frac{1}{2} \partial_x^2 f(X_t, t) \Delta X^2 \\ &\quad + O(|\Delta X|^3) + O(|\Delta X| \Delta t) + O(\Delta t^2) . \end{aligned}$$

The three remainder terms on the last line are the sizes of the three lowest order Taylor series terms left out. Now take the expectation of both sides conditioning on \mathcal{F}_t and pull out of the expectation anything that is known in \mathcal{F}_t :

$$\begin{aligned} \mathbb{E}[f(X_{t+\Delta t}, t + \Delta t | \mathcal{F}_t)] &= f(X_t, t) \\ &+ \partial_x f(X_t, t) \mathbb{E}[\Delta X | \mathcal{F}_t] \\ &+ \partial_t f(X_t, t) \Delta t \\ &+ \frac{1}{2} \partial_x^2 f(X_t, t) \mathbb{E}[\Delta X^2 | \mathcal{F}_t] \\ &+ O\left(\mathbb{E}\left[|\Delta X|^3 | \mathcal{F}_t\right]\right) + O(\mathbb{E}[|\Delta X| | \mathcal{F}_t] \Delta t) + O(\Delta t^2). \end{aligned}$$

The two terms on the top line are equal because of the tower property. The next line is zero because Brownian motion is symmetric and $\mathbb{E}[\Delta X | \mathcal{F}_t] = 0$. For the fourth line, use the independent increments property and the variance of Brownian motion increments to get $\mathbb{E}[\Delta X^2 | \mathcal{F}_t] = \Delta t$. We also know the scaling relations $\mathbb{E}[|\Delta X|] = C\Delta t^{1/2}$ and $\mathbb{E}[|\Delta X|^3] = C\Delta t^{3/2}$. Put all of these in and cancel the leading power of Δt :

$$0 = \partial_t f(X_t, t) + \frac{1}{2} \partial_x^2 f(X_t, t) + O\left(\Delta t^{1/2}\right).$$

Taking $\Delta t \rightarrow 0$ shows that f satisfies the backward equation (6).

We can find several explicit solutions to the backward equation that illustrate the properties of Brownian motion. One is $f(x, t) = x^2 + T - t$. This corresponds to final conditions $V(X_T) = X_T^2$. It tells us that if $X_0 = 0$, then $\mathbb{E}[V(X_T)] = \mathbb{E}[X_T^2] = f(0, 0) = T$. This is the variance of standard Brownian motion. Another well known calculation is the expected value of e^{aX_T} starting from $X_0 = 0$. For this, we want $f(x, t)$ that satisfies (6) and final condition $f(x, T) = e^{ax}$. We try the *ansatz* $f(x, t) = Ae^{ax-bt}$. Putting this into the equation gives

$$-bAe^{ax-bt} + \frac{1}{2}a^2Ae^{ax-bt} = 0.$$

Therefore, $f(x, t) = Ae^{ax-bt}$. Matching the final condition gives

$$e^{ax} = Ae^{ax-a^2T/2} \implies A = e^{a^2T/2}.$$

The final solution is

$$f(x, t) = e^{ax+a^2(T-t)/2}.$$

If $X_0 = 0$, we find $\mathbb{E}[e^{aX_T}] = e^{a^2T/2}$. We verify that this is the right answer by noting that $Y = aX_t \sim \mathcal{N}(0, a^2T)$.

You can add boundary conditions to the backward equation to take into account absorbing boundaries. Suppose you get a reward $W(t)$ if you first touch a barrier at time t , which is $\tau_a = t$. Consider the problem: run a Brownian motion starting at $X_0 = 0$ to time $\tau_a \wedge T$. For $a > 0$, the value function is defined for $t \leq T$ and $x \leq a$. The final condition at $t = T$ is $f(x, T) = V(x)$ as before. The *boundary condition* at $x = a$ is $f(a, t) = W(t)$. A more precise statement

of the boundary condition is $f(x, t) \rightarrow W(t)$ as $x \rightarrow a$. This is similar to the boundary condition $u(x, t) \rightarrow 0$ as $x \rightarrow a$. As you approach a , your probability of not hitting a in a short amount of time goes to zero. This implies that as $X_t \rightarrow a$, the conditional probability that $\tau < t + \epsilon$ goes to zero. You might think that the survival probability calculations above would prove this. But those were based on the boundary $u(a, t) = 0$ boundary condition, which we did not prove. It would be a circular argument.