

PDE for Finance Notes, Spring 2014 – Section 1.

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, MATH-GA2706. Prepared in 2003, minor updates made in 2011 and 2014.

Links between stochastic differential equations and PDE. A stochastic differential equation, together with its initial condition, determines a diffusion process. We can use it to define a deterministic function of space and time in two fundamentally different ways:

- (a) by considering the expected value of some “payoff,” as a function of the initial position and time; or
- (b) by considering the probability of being in a certain state at a given time, given knowledge of the initial state and time.

Students of finance will be familiar with the Black-Scholes PDE, which amounts to an example of (a). Thus in studying topic (a) we will be exploring among other things the origin of the Black-Scholes PDE. The basic mathematical ideas here are the *backward Kolmogorov equation* and the *Feynman-Kac formula*.

Viewpoint (b) is different from (a), but not unrelated. It is in fact *dual* to viewpoint (a), in a sense that we will make precise. The evolving probability density solves a different PDE, the *forward Kolmogorov equation* – which is actually the adjoint of the backward Kolmogorov equation.

It is of interest to consider how and when a diffusion process crosses a barrier. This arises in thinking subjectively about stock prices (e.g. what is the probability that IBM will reach 200 at least once in the coming year?). It is also crucial for pricing barrier options. Probabilistically, thinking about barriers means considering *exit times*. On the PDE side this will lead us to consider *boundary value problems* for the backward and forward Kolmogorov equations.

For a fairly accessible treatment of much of this material see Gardiner (the chapter on the Fokker-Planck Equation). Parts of my notes draw from Oksendal, however the treatment there is much more general and sophisticated so not easy to read.

Our main tool will be Ito’s formula, coupled with the fact that any Ito integral of the form $\int_a^b f dw$ has expected value zero. (Equivalently: $m(t) = \int_a^t f dw$ is a martingale.) Here w is Brownian motion and f is non-anticipating. The stochastic integral is defined as the limit of Ito sums $\sum_i f(t_i)(w(t_{i+1}) - w(t_i))$ as $\Delta t \rightarrow 0$. The sum has expected value zero because each of its terms does: $E[f(t_i)(w(t_{i+1}) - w(t_i))] = E[f(t_i)]E[w(t_{i+1}) - w(t_i)] = 0$.

Expected values and the backward Kolmogorov equation. Here’s the most basic version of the story. Suppose $y(t)$ solves the scalar stochastic differential equation

$$dy = f(y, s)ds + g(y, s)dw,$$

and let

$$u(x, t) = E_{y(t)=x} [\Phi(y(T))]$$

be the expected value of some payoff Φ at maturity time $T > t$, given that $y(t) = x$. Then u solves

$$u_t + f(x, t)u_x + \frac{1}{2}g^2(x, t)u_{xx} = 0 \text{ for } t < T, \text{ with } u(x, T) = \Phi(x). \quad (1)$$

The proof is easy: for any function $\phi(y, t)$, Ito's lemma gives

$$\begin{aligned} d(\phi(y(s), s)) &= \phi_y dy + \frac{1}{2}\phi_{yy} dy dy + \phi_s ds \\ &= (\phi_s + f\phi_y + \frac{1}{2}g^2\phi_{yy})dt + g\phi_y dw. \end{aligned}$$

Choosing $\phi = u$, the solution of (1), we get

$$u(y(T), T) - u(y(t), t) = \int_t^T (u_t + fu_y + \frac{1}{2}g^2u_{yy})ds + \int_t^T gu_y dw.$$

Taking the expected value and using the PDE gives

$$E_{y(t)=x} [\Phi(y(T))] - u(x, t) = 0$$

which is precisely our assertion.

That was the simplest case. It can be jazzed up in many ways. We discuss some of them:

Vector-valued diffusion. Suppose y solves a vector-valued stochastic differential equation

$$dy_i = f_i(y, s)ds + \sum_j g_{ij}(y, s)dw_j,$$

where each component of w is an independent Brownian motion. Then

$$u(x, t) = E_{y(t)=x} [\Phi(y(T))]$$

solves

$$u_t + \mathcal{L}u = 0 \text{ for } t < T, \text{ with } u(x, T) = \Phi(x),$$

where \mathcal{L} is the differential operator

$$\mathcal{L}u(x, t) = \sum_i f_i \frac{\partial u}{\partial x_i} + \frac{1}{2} \sum_{i,j,k} g_{ik}g_{jk} \frac{\partial^2 u}{\partial x_i \partial x_j}.$$

The justification is just as in the scalar case, using the multidimensional version of Ito's lemma. The operator \mathcal{L} is called the "infinitesimal generator" of the diffusion process $y(s)$.

The Feynman-Kac formula. We discuss the scalar case first, for clarity. Consider as above the solution of

$$dy = f(y, s)ds + g(y, s)dw$$

but suppose we are interested in a suitably “discounted” final-time payoff of the form:

$$u(x, t) = E_{y(t)=x} \left[e^{-\int_t^T b(y(s), s) ds} \Phi(y(T)) \right] \quad (2)$$

for some specified function $b(y)$. Then u solves

$$u_t + f(x, t)u_x + \frac{1}{2}g^2(x, t)u_{xx} - b(x, t)u = 0 \quad (3)$$

instead of (1). (Its final-time condition is unchanged: $u(x, T) = \Phi(x)$.) If you know some finance you’ll recognize that when y is log-normal and b is the interest rate, (3) is precisely the Black-Scholes partial differential equation. Also: if $b(y(s), s)$ is the spot interest rate, then (3) with $\Phi = 1$ gives the time- t value of a zero-coupon bond with maturity T , given that the spot interest rate at time t is $b(x, t)$.

To explain (3), we must calculate the stochastic differential $d[z_1(s)\phi(y(s), s)]$ where $z_1(s) = e^{-\int_t^s b(y(r), r) dr}$. The multidimensional version of Ito’s lemma gives

$$d[z_1(s)z_2(s)] = z_1 dz_2 + z_2 dz_1 + dz_1 dz_2.$$

We apply this with z_1 as defined above and $z_2(s) = \phi(y(s), s)$. Ito’s lemma (or ordinary differentiation) gives

$$dz_1(s) = -z_1 b(y(s), s) ds$$

and we’re already familiar with the fact that

$$\begin{aligned} dz_2(s) &= (\phi_s + f\phi_y + \frac{1}{2}g^2\phi_{yy})ds + g\phi_y dw \\ &= (\phi_s + \mathcal{L}\phi)ds + g\phi_y dw. \end{aligned}$$

Notice that $dz_1 dz_2 = 0$. Applying the above with $\phi = u$, the solution of the PDE (3), gives

$$\begin{aligned} d \left(e^{-\int_t^s b(y(r), r) dr} u(y(s), s) \right) &= z_1 dz_2 + z_2 dz_1 \\ &= z_1 [(u_s + \mathcal{L}u)ds + g u_y dw] - z_1 u b ds \\ &= z_1 g u_y dw. \end{aligned}$$

The right hand side has expected value 0, so

$$E_{y(t)=x} [z_1(T)z_2(T)] = z_1(t)z_2(t) = u(x, t)$$

as asserted.

A moment’s thought reveals that vector-valued case is no different. The discounted expected payoff (2) solves the PDE

$$u_t + \mathcal{L}u - bu = 0$$

where \mathcal{L} is the infinitesimal generator of the diffusion y .

Running payoff. Suppose we are interested in

$$u(x, t) = E_{y(t)=x} \left[\int_t^T \Psi(y(s), s) ds \right]$$

for some specified function Ψ . Then u solves

$$u_t + \mathcal{L}u + \Psi(x, t) = 0.$$

The final-time condition is $u(x, T) = 0$, since we have included no final-time term in the “payoff.” The proof is hardly different from before: by Ito’s lemma,

$$\begin{aligned} d[u(y(s), s)] &= (u_t + \mathcal{L}u)ds + \nabla u \cdot g \cdot dw \\ &= -\Psi(y(s), s)ds + \nabla u \cdot g \cdot dw. \end{aligned}$$

Integrating and taking the expectation gives

$$E_{y(t)=x} [u(y(T), T)] - u(x, t) = E_{y(t)=x} \left[- \int_t^T \Psi(y(s), s) ds \right].$$

This gives the desired assertion, since $u(y(T), T) = 0$.

In valuing options, “running payoffs” are relatively rare. However terms of this type will be very common later in the course, when discuss optimal control problems.

Boundary value problems and exit times. The preceding examples use stochastic integration from time t to a fixed time T , and they give PDE’s that must be solved for all $x \in R^n$. It is also interesting to consider integration from time t to the first time y exits from some specified region. The resulting PDE must be solved on this region, with suitable boundary data.

Let D be a region in R^n . Suppose y is an R^n -valued diffusion solving

$$dy = f(y, s)ds + g(y, s)dw \text{ for } s > t, \text{ with } y(t) = x$$

with $x \in D$. Let

$$\begin{aligned} \tau(x) &= \text{the first time } y(s) \text{ exits from } D, \text{ if} \\ &\text{prior to } T; \text{ otherwise } \tau(x) = T. \end{aligned}$$

This is an example of a *stopping time*. (Defining feature of a stopping time: the statement “ $\tau(x) < t$ ” is \mathcal{F}_t -measurable; in other words, the decision whether to stop or not at time t depends only on knowledge of the process up to time t . This is clearly true of the exit time defined above.)

Here is the basic result: the function

$$u(x, t) = E_{y(t)=x} \left[\int_t^{\tau(x)} \Psi(y(s), s) ds + \Phi(y(\tau(x)), \tau(x)) \right]$$

solves

$$u_t + \mathcal{L}u + \Psi = 0 \text{ for } x \in D$$

with boundary condition

$$u(x, t) = \Phi(x, t) \text{ for } x \in \partial D \quad (4)$$

and final-time condition

$$u(x, T) = \Phi(x, T) \text{ for all } x \in D. \quad (5)$$

The justification is entirely parallel to our earlier examples. The only change is that we integrate, in the final step, to the stopping time τ rather than the final time T . (This is permissible for any stopping time satisfying $E[\tau] < \infty$. The statement that $E[\int_t^\tau f dw] = 0$ when $E[\tau] < \infty$ is known as Dynkin's theorem.)

A subtlety is hiding here: the hypothesis that $E[\tau] < \infty$ is not a mere technicality. Rather, there are simple and interesting examples where it is false and $E[\int_t^\tau f dw] \neq 0$. One such example is related to the “gambler’s ruin” paradox. Consider the standard Brownian process $w(s)$, starting at $w(0) = 0$. Let τ_* be the first time $w(s)$ reaches 1. Then $w(\tau_*) - w(0) = 1 - 0 = 1$ certainly does not have mean 0, so $E[\int_0^{\tau_*} dw] \neq 0$ in this case. This doesn't contradict Dynkin's theorem; it just shows that $E[\tau_*] = \infty$. To understand the situation better, consider $\tau_n =$ the time of first exit from $[-n, 1]$. You'll show on HW1 that $E[\tau_n] < \infty$ for each n , but $E[\tau_n] \rightarrow \infty$ as $n \rightarrow \infty$. The Brownian motion process eventually reaches 1 with probability one, but it may make extremely large negative excursions before doing so. Here's the coin-flipping version of this situation: consider a gambler who decides to bet by flipping coins and never quitting till he's ahead by a fixed amount. If there is no limit on the amount he is permitted to lose along the way, then he'll eventually win with probability one. But if there is a threshold of losses beyond which he must stop then there is a nonzero probability of ruin and his expected outcome is 0.

There's something slightly misleading about our notation in (4)-(5). We use the same notation Φ for both the boundary condition (4) and the final-time condition (5) because they come from the same term in the payoff: $\Phi(y(\tau), \tau)$ where τ is the time the curve $(y(s), s)$ exits from the cylinder $D \times [0, T]$. But Φ should be thought of as representing two *distinct* functions – one at the spatial boundary $\partial D \times [0, T]$, the other at the final time boundary $D \times \{T\}$ (see the figure). These two functions need have nothing to do with one




Figure 1: *Distinguishing between the two different parts of Φ .*

another. Often one is chosen to be zero, while the other is nontrivial. [A financial example:

when one values a barrier option using the risk-neutral expectation of the payoff, Φ is zero at the knock-out price, and it equals the payoff at the maturity time.]

Elliptic boundary-value problems. Now suppose f and g in the stochastic differential equation don't depend on t , and for $x \in D$ let

$$\tau(x) = \text{the first time } y(s) \text{ exits from } D.$$

(Unlike the previous example, we do not impose a final time T . Clearly this amounts to taking $T = \infty$ in the previous definition.) Suppose furthermore the process does eventually exit from D , (and more: assume $E[\tau(x)] < \infty$, for all $x \in D$). Then

$$u(x) = E_{y(0)=x} \left[\int_0^{\tau(x)} \Psi(y(s)) ds + \Phi(y(\tau(x))) \right]$$

solves

$$\mathcal{L}u + \Psi = 0 \text{ for } x \in D,$$

with boundary condition

$$u = \Phi \text{ for } x \in \partial D.$$

The justification is again entirely parallel to our earlier examples.

Applications: some properties of the Brownian motion process. Let us use these results to deduce – by solving appropriate PDE's – some properties of the Brownian motion process. (This discussion is taken from Oksendal's example 7.4.2. Related material on exit times will be explored in HW1.)

QUESTION 1. Consider n -dimensional Brownian motion starting at x . What is the mean time it takes to exit from a ball of radius R , for $R > |x|$? Answer: apply the last example with $f = 0$, $g = \text{identity matrix}$, $\Psi = 1$, $\Phi = 0$. It tells us the mean exit time is the solution $u(x)$ of

$$\frac{1}{2} \Delta u + 1 = 0$$

in the ball $|x| < R$, with $u = 0$ at $|x| = R$. The (unique) solution is

$$u(x) = \frac{1}{n} (R^2 - |x|^2).$$

(To do this calculation we must know in advance that the expected exit time is finite. We'll justify this as Question 3 below.)

QUESTION 2. Consider the scalar lognormal process

$$dy = \mu y dt + \sigma y dw$$

with μ and σ constant. Starting from $y(0) = x$, what is the mean exit time from a specified interval (a, b) with $a < x < b$? Answer: the mean exit time $u(x)$ solves

$$\mu x u_x + \frac{1}{2} \sigma^2 x^2 u_{xx} + 1 = 0 \text{ for } a < x < b$$

with boundary conditions $u(a) = u(b) = 0$. The solution is

$$u(x) = \frac{1}{\frac{1}{2} \sigma^2 - \mu} \left(\log(x/a) - \frac{1 - (x/a)^{1-2\mu/\sigma^2}}{1 - (b/a)^{1-2\mu/\sigma^2}} \log(b/a) \right)$$

(readily verified by checking the equation and boundary conditions). This answer applies only if $\mu \neq \frac{1}{2} \sigma^2$. See HW1 for the case $\mu = \frac{1}{2} \sigma^2$.

QUESTION 3: Returning to the setting of Question 1, how do we know the mean exit time is finite? Answer: assume D is a bounded domain in R^n , and $y(s)$ is multidimensional Brownian motion starting at $x \in D$. Recall that by Ito's lemma, $t \rightarrow \phi(y(t))$ satisfies

$$d\phi = \nabla \phi dw + \frac{1}{2} \Delta \phi dt \tag{6}$$

for any function ϕ . Let's apply this with $\phi(y) = |y|^2$, integrating in time up to the stopping time

$$\tau_T(x) = \min\{\tau(x), T\} = \begin{cases} \text{first time } y(s) \text{ exits from } D & \text{if less than } T \\ T & \text{otherwise.} \end{cases}$$

We get

$$\begin{aligned} E \left[|y(\tau_T(x))|^2 \right] - |x|^2 &= \frac{1}{2} E \int_0^{\tau_T(x)} \Delta \phi(y(s)) ds \\ &= n E [\tau_T(x)] \end{aligned} \tag{7}$$

since $\Delta \phi = 2n$. Now let $T \rightarrow \infty$. The left hand side of (7) stays finite, since we're considering a *bounded* domain, and by definition $y(\tau_T(x))$ is either in D or on the boundary of D . Thus we conclude that

$$\lim_{T \rightarrow \infty} E [\tau_T(x)] < \infty.$$

It follows (using the monotone convergence theorem, from real variables) that the exit time $\tau = \lim_{T \rightarrow \infty} \tau_T$ is almost surely finite, and $E[\tau] < \infty$, for any starting point $x \in D$.

QUESTION 4: Consider Brownian motion in R^n , starting at a point x with $|x| = b$. Given $r < b$, what is the probability that the path ever enters the ball of radius r centered at 0? Answer: for $n = 1, 2$ this probability is 1. (Interpretation: Brownian motion is "recurrent" in dimensions 1 and 2 – it comes arbitrarily close to any point, infinitely often, regardless of where it starts.) In higher dimensions the situation is different: in dimension $n \geq 3$ the probability of entering the ball of radius r is $(b/r)^{2-n}$. (Interpretation: Brownian motion is "transient" in dimension $n \geq 3$.)

Consider first the case $n \geq 3$. We use the stopping time $\tau_k =$ first exit time from the annulus

$$D_k = \{r < |x| < 2^k r\}.$$

Since D_k is bounded, $E[\tau_k] < \infty$ and we can integrate the stochastic differential equation (6) up to time τ_k . Let's do this with the special choice

$$\phi(y) = |y|^{2-n}.$$

This ϕ solves Laplace's equation $\Delta\phi = 0$ away from its singularity at $y = 0$. (The singularity does not bother us, since we only evaluate ϕ at points $y(s) \in D_k$ and 0 does not belong to D_k .) The analogue of (7) is

$$E \left[|y(\tau_k)|^{2-n} \right] - b^{2-n} = \frac{1}{2} \int_0^{\tau_k} \Delta\phi(y(s)) ds = 0.$$

If p_k is the probability that y leaves the annulus D_k at radius r , and $q_k = 1 - p_k$ is the probability that it leaves the annulus at radius $2^k r$, we have

$$r^{2-n} p_k + (2^k r)^{2-n} q_k = b^{2-n}.$$

As $k \rightarrow \infty$ this gives $p_k \rightarrow (b/r)^{2-n}$, as asserted.

The case $n = 2$ is treated similarly, using

$$\phi(y) = \log y,$$

which solves $\Delta\phi = 0$ in the plane, away from $y = 0$. Arguing as before we get

$$p_k \log r + q_k \log(2^k r) = \log b.$$

As $k \rightarrow \infty$ this gives $q_k \rightarrow 0$. So $p_k \rightarrow 1$, as asserted.

The case $n = 1$ is similar to $n = 2$, using $\phi(y) = |y|$.

Another application: distribution of first arrivals. Consider a scalar diffusion whose drift and volatility are functions of y alone, independent of t :

$$dy = f(y(s))ds + g(y(s))dw.$$

The initial condition is $y(t) = x$. We are interested in the first arrival of $y(s)$ at a given threshold, say $y = 0$. Assume to fix ideas that $x > 0$.

WHAT IS THE DISTRIBUTION OF ARRIVAL TIMES? Let the density of arrival times be $\rho(s)$. Its cumulative distribution function $\int_0^T \rho(s) ds$ is the probability that the first arrival occurs by time T . According to our discussion of the backward Kolmogorov equation, this is $u(x, 0)$ where u solves

$$u_t + f u_x + \frac{1}{2} g^2 u_{xx} = 0 \quad \text{for } x > 0, \quad 0 < t < T \tag{8}$$

with boundary condition

$$u = 1 \quad \text{at } x = 0$$

and final-time condition

$$u = 0 \quad \text{at } t = T.$$

Clearly u depends also on the final time T ; let's make that dependence explicit by writing $u = u(x, t; T)$. Evidently $u(x, 0; T) = \int_0^T \rho(s) ds$, so by differentiation we get

$$\rho(s) = \frac{\partial u}{\partial T}(x, 0; s).$$

For special choices of f and g (e.g. Brownian motion with drift, or lognormal) the PDE (8) can be solved explicitly (we'll discuss how to do this later), yielding an explicit formula for the distribution of first arrivals.

Suppose the mean arrival time is finite.¹ Then we know it should be given by $v(x)$ where $fv_x + \frac{1}{2}g^2v_{xx} = -1$ for $x > 0$ with $v = 0$ at $x = 0$. On the other hand, the mean arrival time is

$$\int_0^\infty s\rho(s) ds = \int_0^\infty s\partial_s u(x, 0; s) ds.$$

Are these apparently different expressions consistent? Yes indeed! To show this, we observe (writing u_s for $\partial_s u(x, 0; s)$, for simplicity) that

$$\int_0^\infty su_s ds = - \int_0^\infty s(1-u)_s ds = \int_0^\infty (1-u) ds$$

by integration by parts, since² $\int_0^\infty \partial_s[s(1-u)] ds = s(1-u)|_0^\infty = 0$. Moreover the function $v(x) = \int_0^\infty (1-u) ds$ clearly satisfies $v = 0$ at $x = 0$, and

$$fv_x + \frac{1}{2}g^2v_{xx} = - \int_0^\infty fu_x + \frac{1}{2}g^2u_{xx} ds.$$

But since f and g are independent of time, $u(x, t; T)$ depends on t and T only through $T - t$, so $\partial u/\partial T = -\partial u/\partial t$. Therefore, using the backward Kolmogorov equation,

$$- \int_0^\infty fu_x + \frac{1}{2}g^2u_{xx} ds = \int_0^\infty u_t ds = - \int_0^\infty \partial_s u(x, 0; s) ds.$$

The last expression is clearly $u(x, 0; 0) - u(x, 0; \infty) = 0 - 1 = -1$. Thus v solves the anticipated PDE.

¹This is *not* true when $f = 0$ and $g = 1$ – see the discussion on page 5, and the problem related to it on HW1. Specifying conditions on f and g that make the mean arrival time finite is a nontrivial task, which we won't attempt to address here.

²We prefer to work with $1 - u$ rather than u , since it *vanishes* at the spatial boundary $x = 0$. Since f and g depend only on position, $u(x, t; T)$ depends only on the “elapsed time” $T - t$. As $T - t \rightarrow \infty$, we expect $w = 1 - u$ to approach a solution of the “stationary” (time-independent) problem $fw_x + \frac{1}{2}g^2w_{xx} = 0$ for $x > 0$, $w = 0$ at $x = 0$. The obvious choice is $w = 0$; with some modest hypotheses on f and g , $w = 0$ is the only solution of this boundary value problem on the half-line $x > 0$ that remains uniformly bounded as $x \rightarrow \infty$. The “integration by parts” calculation that follows uses this fact, and a little more: it uses (more precisely, it assumes) that $s[(1 - u(x, 0, s))] \rightarrow 0$ as $s \rightarrow \infty$. This is closely related to our hypothesis that the mean arrival time is finite. We will simply assume this property holds, without attempting to specify conditions on f and g that guarantee it.

Transition probabilities and the forward Kolmogorov equation. We've shown that when the state evolves according to a stochastic differential equation

$$dy_i = f_i(y, s)ds + \sum_j g_{ij}(y, s)dw_j$$

the expected final position

$$u(x, t) = E_{y(t)=x} [\Phi(y(T))]$$

solves the backward Kolmogorov equation

$$u_t + \sum_i f_i \frac{\partial u}{\partial x_i} + \frac{1}{2} \sum_{i,j,k} g_{ik}g_{jk} \frac{\partial^2 u}{\partial x_i \partial x_j} = 0 \text{ for } t < T, \text{ with } u = \Phi \text{ at } t = T. \quad (9)$$

We can write the backward Kolmogorov equation as

$$u_t + \mathcal{L}u = 0 \quad (10)$$

with

$$\mathcal{L}u = \sum_i f_i \frac{\partial u}{\partial x_i} + \sum_{i,j} a_{ij} \frac{\partial^2 u}{\partial x_i \partial x_j}, \quad (11)$$

where $a_{ij} = \frac{1}{2} \sum_k g_{ik}g_{jk} = \frac{1}{2}(gg^T)_{ij}$.

The solution of the stochastic differential equation is a *Markov process*, so it has a well-defined *transition probability*

$p(z, s; x, t)$ = probability of being at z at time s , given that it started at x at time t .

More precisely: $p(\cdot, s; x, t)$ is the probability density of the state at time s , given that it started at x at time t . Of course p is only defined for $s > t$. To describe a Markov process, p must satisfy the Chapman-Kolmogorov equation

$$p(z, s; x, t) = \int_{R^n} p(z_1, s_1; x, t) p(z, s; z_1, s_1) dz_1$$

for any s_1 satisfying $t < s_1 < s$. Intuitively: the state can get from (x, t) to (z, s) by way of being at various intermediate states z_1 at a chosen intermediate time s_1 . The Chapman-Kolmogorov equation calculates $p(z, s; x, t)$ by adding up (integrating) the probabilities of getting from (x, t) to (z, s) via (z_1, s_1) , for all possible intermediate positions z_1 .

How should we visualize p ? Consider first the case when y is multidimensional Brownian motion. Then $p(\cdot, s; x, t)$ is the density of a Gaussian random variable with mean x and variance $s - t$. The graph of $z \rightarrow p(z, s; x, t)$ always has volume 1 below it (since p is a probability density); as $s \rightarrow \infty$ its maximum value tends to 0 (a Brownian particle diffuses further and further away, on average, as time increases); as $s \rightarrow t$ it becomes infinitely tall and thin (at time $s \approx t$ the Brownian particle is very close to its initial position x). The situation for a general stochastic differential equation is similar: p becomes infinitely tall and

thin, concentrating at $z = x$, as $s \rightarrow t$; and if $gg^T > 0$ then the graph of p keeps spreading as $s \rightarrow \infty$. Of course in the general case p does not describe a Gaussian distribution, and there is no simple formula for the mean or variance – they are simply the mean and variance of $y(s)$.

If the stochastic differential equation does not involve time explicitly, then the transition probability depends only on the “elapsed time”:

if $dy = f(y)dt + g(y)dw$ with f, g depending only on y , then $p(z, s; x, t) = p(z, s - t; x, 0)$.

If the stochastic differential equation does not involve space explicitly, then the transition probability depends only on the “relative position”:

if $dy = f(t)dt + g(t)dw$ with f, g depending only on t , then $p(z, s; x, t) = p(z - x, s; 0, t)$.

The initial position of a Markov process need not be deterministic. Even if it is (e.g. if $y(0) = x$ is fixed), we may wish to consider a later time as the “initial time.” The transition probability determines the evolution of the spatial distribution, no matter what its initial value: if $\rho_0(x)$ is the probability density of the state at time t then

$$\rho(z, s) = \int_{R^n} p(z, s; x, t) \rho_0(x) dx \quad (12)$$

gives the probability density (as a function of z) at any time $s > t$.

The crucial fact about the transition probability is this: it solves the *forward Kolmogorov equation* in s and z :

$$-p_s - \sum_i \frac{\partial}{\partial z_i} (f_i(z, s)p) + \frac{1}{2} \sum_{i,j,k} \frac{\partial^2}{\partial z_i \partial z_j} (g_{ik}(z, s)g_{jk}(z, s)p) = 0 \text{ for } s > t, \quad (13)$$

with initial condition

$$p = \delta_x(z) \text{ at } s = t.$$

We can write the forward Kolmogorov equation as

$$-p_s + \mathcal{L}^*p = 0 \quad (14)$$

with

$$\mathcal{L}^*p = - \sum_i \frac{\partial}{\partial z_i} (f_i p) + \sum_{i,j} \frac{\partial^2}{\partial z_i \partial z_j} (a_{ij} p). \quad (15)$$

Here $a_{ij} = \frac{1}{2}(gg^T)_{ij}$ just as before. The initial condition $p = \delta_x(z)$ encapsulates the fact, already noted, that the graph of $p(\cdot, s; x, t)$ becomes infinitely tall and thin at x as s decreases to t . The technical meaning is that

$$\int_{R^n} p(z, s; x, t) f(z) dz \rightarrow f(x) \text{ as } s \text{ decreases to } t \quad (16)$$

for any continuous f .

Recall that if the initial state distribution is ρ_0 then the evolving distribution is $\rho(z, s) = \int p(z, s; x, t) \rho_0(x) dx$. This function $\rho(z, s)$ automatically solves the forward equation (just bring the derivatives under the integral, and use that p solves it). The initial condition on p is just what we need to have $\rho(z, s) \rightarrow \rho_0(z)$ as $s \rightarrow t$. (Demonstration: multiply (16) by $\rho_0(x)$ and integrate in x to see that

$$\int \rho(z, s) f(z) dz = \int p(z, s; x, t) f(z) \rho_0(x) dz dx \rightarrow \int f(x) \rho_0(x) dx$$

as $s \rightarrow t$. Since this is true for every continuous f , we conclude that $\rho(z, s)$ converges [weakly] to $\rho_0(z)$ as $s \rightarrow t$.)

Please note that the forward Kolmogorov equation describes the probability distribution by solving an initial-value problem, while the backward Kolmogorov equation describes the expected final payoff by solving a final-value problem. Students familiar with pricing options via binomial trees will find this familiar. The stock prices at various nodes of a tree are determined by working forward in time; the option values at various nodes of a tree are determined by working backward in time.

Notice that the forward and backward Kolmogorov equations are, in general, completely different. There is one case, however, when they are closely related: for Brownian motion the forward equation starting at $t = 0$ is

$$p_s - \frac{1}{2} \Delta p = 0 \text{ for } s > 0$$

while the backward equation with final time T is

$$u_t + \frac{1}{2} \Delta u = 0 \text{ for } t < T.$$

In this special case the backward equation is simply the forward equation with time reversed. More careful statement: if $u(x, t)$ solves the backward equation then $\tilde{u}(z, s) = u(z, T - s)$ solves the forward equation, and conversely. This is an *accident*, associated with the self-adjointness of the Laplacian. The situation is different even for Brownian motion with constant drift f : then the forward equation is $p_s + f \cdot \nabla p - \frac{1}{2} \Delta p = 0$, while the backward equation is $u_t + f \cdot \nabla u + \frac{1}{2} \Delta u = 0$, and the two are not equivalent under time-reversal.

Students with a background in physical modeling will be accustomed to equations of the form $v_t = \text{div}(a(x) \nabla v)$. Neither the forward nor the backward Kolmogorov equation has this form. Such equations are natural in physics, but not in problems from control theory and stochastic differential equations.

Application: steady-state distributions. The backward Kolmogorov equation comes up more often than the forward one in finance. But one important application involves the large-time behavior of a diffusion. If $\rho(z, s)$ is the probability density of a diffusion, then evidently $\rho_\infty(z) = \lim_{s \rightarrow \infty} \rho(z, s)$ represents (if it exists) the large-time statistics of the process. For Brownian motion $\rho_\infty = 0$, reflecting the fact that Brownian particles wander a lot. The situation is quite different however for the Ornstein-Uhlenbeck process

$dy = -kyds + \sigma dw$. We expect y to remain near 0 due to the deterministic term $-ky$, which constantly pushes it toward 0. And indeed the steady-state distribution is

$$\rho_\infty(z) = Ce^{-kz^2/\sigma^2}$$

where C is chosen so that ρ_∞ has integral 1. (It's easy to check that this gives a steady-state solution of the forward Kolmogorov equation. In fact this gives the long-time asymptotics of a fairly general initial condition, but this is *not* so obvious.)

This application can be generalized. Consider the stochastic PDE $dy = -V'(y)dt + \sigma dw$. Its deterministic part pushes y toward a local minima of V . If V grows rapidly enough at infinity (so the diffusion is successfully confined, and does not wander off to infinity) then the long-time statistics are described by the steady-state distribution

$$\rho_\infty(z) = Ce^{-2V(z)/\sigma^2}.$$

Testing the plausibility of the forward equation. We will explain presently why the forward equation holds. But first let's get used to it by examining some consequences and checking some special cases. Let $\rho_0(x)$ be the probability density of the state at time 0, and consider

$$\rho(z, s) = \int p(z, s; x, 0)\rho_0(x) dx$$

for $s > 0$. It gives the probability density of the state at time s .

Checking the integral. Since ρ is a probability density we expect that $\int \rho(z, s) dz = 1$ for all s . In fact, from the forward equation

$$\begin{aligned} \frac{d}{ds} \int \rho dz &= \int \rho_s dz \\ &= \int \mathcal{L}^* \rho dz \\ &= 0 \end{aligned}$$

since each term of $\mathcal{L}^* \rho$ is a perfect derivative. (Here and below, we repeatedly integrate by parts, with no "boundary terms" at $\pm\infty$. We are implicitly assuming that ρ and its derivatives decay rapidly as $z \rightarrow \pm\infty$. This is true, provided the initial distribution ρ_0 has this property.)

If the stochastic differential equation has no drift then the expected position is independent of time. In general, $E[y(s)] - E[y(0)] = E \int_0^s f(y(r), r) dr$ since the expected value of the integral dw vanishes. Thus when $f = 0$ the expected position $E[y(s)]$ is constant. Let's prove this again using the forward equation:

$$\begin{aligned} \frac{d}{ds}(\text{expected position}) &= \frac{d}{ds} \int z\rho(z, s) dz \\ &= \int z\rho_s(z, s) dz \\ &= \int z\mathcal{L}^* \rho(z, s) dz \\ &= 0 \quad \text{when } f = 0. \end{aligned}$$

The last step is the result of integration by parts; for example, if y is scalar valued ($dy = g(y, t)dw$) we have

$$\begin{aligned} \int z \mathcal{L}^* \rho dz &= \frac{1}{2} \int z (g^2 \rho)_{zz} dz \\ &= -\frac{1}{2} \int (g^2 \rho)_z dz \\ &= 0. \end{aligned}$$

(As noted above, to justify the integrations by parts one must know that ρ vanishes rapidly enough at spatial infinity.)

The special case $f = \text{constant}$, $g = 0$. If $g = 0$ then we're studying a deterministic motion. If in addition $f = \text{constant}$ then the solution is explicit and very simple: $y(t) = y(0) + ft$. Clearly

$$\text{Prob of being at } z \text{ at time } s = \text{Prob of being at } z - fs \text{ at time } 0,$$

whence

$$\rho(z, s) = \rho_0(z - fs).$$

In particular, $\rho_s + f \cdot \nabla \rho = 0$, which agrees with the forward equation (since f is constant).

Biting the bullet. Enough playing around; let's explain why the forward equation holds. The first main ingredient is the observation that

$$E_{y(t)=x} [\Phi(y(T))] = \int \Phi(z) p(z, T; x, t) dz. \quad (17)$$

We know how to determine the left hand side (by solving the backward equation, with final value Φ at $t = T$). This relation determines the integral of $p(\cdot, T; x, t)$ against any function Φ , for any value of x, t, T . This is a lot of information about p – in fact, it fully determines p . Our task is to make this algorithmic, i.e. to explain how p can actually be computed. (The answer, of course, will be to solve the forward equation in z and s .)

The second main ingredient is the relation between \mathcal{L} and \mathcal{L}^* . Briefly: \mathcal{L}^* is the *adjoint* of \mathcal{L} in the L^2 inner product. Explaining this: recall from linear algebra that if A is a linear operator on an inner-product space, then its adjoint A^* is defined by

$$\langle Ax, y \rangle = \langle x, A^* y \rangle.$$

When working in R^n we can represent A by a matrix, and A^* is represented by the transpose A^T . The situation is similar here, but our inner product space consists of all (square-integrable, scalar-valued) functions on R^n , with inner product

$$\langle v, w \rangle = \int_{R^n} v(x) w(x) dx.$$

We claim that

$$\langle \mathcal{L}v, w \rangle = \langle v, \mathcal{L}^* w \rangle. \quad (18)$$

When y is scalar-valued our claim says that

$$\int_R \left(f v_x + \frac{1}{2} g^2 v_{xx} \right) w dx = \int_R v \left(-(fw)_x + \frac{1}{2} (g^2 w)_{xx} \right) dx.$$

This is a consequence of integration by parts. For example, the first term on the left equals the first term on the right since

$$\int_R [fw] v_x dx = - \int_R [fw]_x v dx.$$

The second term on each side matches similarly, integrating by parts twice. Notice that f and g can depend on time as well as space; it doesn't change the argument. The proof of (18) when y is vector valued is essentially the same as the scalar case.

The third main ingredient is hiding in our derivation of the backward equation. We know from this derivation that

$$E_{y(t)=x} [\phi(y(T), T)] - \phi(x, t) = E_{y(t)=x} \left[\int_t^T (\phi_s + \mathcal{L}\phi)(y(s), s) ds \right] \quad (19)$$

for any function $\phi(y, s)$. Our main use of this relation up to now was to choose ϕ so that the right hand side vanished, i.e. to choose ϕ to solve the backward equation. But we don't have to make such a restrictive choice: relation (19) holds for *any* ϕ .

Let's put these ingredients together. Rewriting (19) using the transition probabilities gives

$$\int_{R^n} \phi(z, T) p(z, T; x, t) dz - \phi(x, t) = \int_t^T \int_{R^n} (\phi_s + \mathcal{L}\phi)(z, s) p(z, s; x, t) dz ds. \quad (20)$$

Using (18) and doing the obvious integration by parts in time, the right hand side becomes

$$\int_t^T \int_{R^n} -\phi p_s + \phi \mathcal{L}^* p dz ds + \int_{R^n} \phi(z, s) p(z, s; x, t) dz \Big|_{s=t}^{s=T}. \quad (21)$$

This is true for *all* ϕ . Since the left hand side of (20) involves only the initial and final times (t and T) we conclude that

$$-p_s + \mathcal{L}^* p = 0.$$

Therefore (20)-(21) reduce to

$$\int_{R^n} \phi(z, t) p(z, t; x, t) dz = \phi(x, t)$$

for all ϕ , which is what we mean by the initial condition " $p = \delta_x$ when $s = t$ ". Done!

The argument is simple; but maybe it's hard to encompass. To recapitulate its essence, let's give a new proof (using the forward equation) of the fact (known via Ito calculus) that

$$u \text{ solves the backward equation} \implies \frac{d}{ds} E[u(y(s), s)] = 0.$$

In fact: if $\rho(z, s)$ is the probability distribution of the state at time s ,

$$\begin{aligned} \frac{d}{ds} E[u(y(s), s)] &= \frac{d}{ds} \int u(z, s) \rho(z, s) dz \\ &= \int u_s \rho + u \rho_s dz \\ &= \int u_s \rho + u \mathcal{L}^* \rho dz \\ &= \int u_s \rho + (\mathcal{L}u) \rho dz \\ &= 0 \end{aligned}$$

using in the last step our hypothesis that u solves the backward equation.

Boundary value problems. The preceding discussion concerned the backward and forward Kolmogorov equations in all space. When working in a bounded domain, the boundary conditions for the forward Kolmogorov equation depend on what the random walk does when it reaches the boundary.

We discuss here just the case of most interest for financial applications: the *absorbing* boundary condition, i.e. a random walk that we track only till it hits the boundary for the first time. (After that time we think of the random walker as disappearing, i.e. being “absorbed” by the boundary.) The corresponding boundary condition for the forward Kolmogorov equation is that the probability density vanish there (since it represents the density of not-yet-absorbed walkers).

Let’s explain briefly why this choice is right. Consider the backward Kolmogorov equation in a bounded domain, with boundary condition $u = 0$:

$$\begin{aligned} u_t + \mathcal{L}u &= 0 \text{ for } x \in D, t < T \\ u(x, T) &= \phi(x) \text{ at } t = T \\ u(x, t) &= 0 \text{ for } x \in \partial D. \end{aligned}$$

We know that

$$u(x, t) = E_{y(t)=x} [\Phi(y(\tau), \tau)]$$

where $\tau = \tau(x)$ is the exit time from D (or T , if the path doesn’t exit by time T) and

$$\Phi = 0 \text{ for } x \in \partial D; \Phi = \phi \text{ at the final time } T.$$

This formula for u can be written as

$$u(x, t) = \int_{R^n} \phi(z) q(z, T; x, t) dz$$

where

$q(z, s; x, t)$ = probability that the diffusion arrives at z at time s , starting from x at time t , without hitting ∂D first.

Our assertion is that $q(z, s; x, t)$ solves the forward Kolmogorov equation for $z \in D$ and $s > t$, with boundary condition $q = 0$ for $z \in \partial D$, and initial condition $q = \delta_x$. The justification is very much like the argument given above for R^n .

One thing changes significantly when we work in a bounded domain: $\int_D q(z, s; x, t) dz < 1$. The reason is that q gives the probability of arriving at z at time s without hitting the boundary first. Thus

$$1 - \int_D q(z, s; x, t) dz = \text{prob of hitting } \partial D \text{ by time } s, \text{ starting from } x \text{ at time } t.$$

Evidently $\int q(z, s; x, t) dz$ is decreasing in time. Let's check this for Brownian motion, for which $q_s - \frac{1}{2}\Delta q = 0$. We have

$$\begin{aligned} \frac{d}{ds} \int_D q(z, s; x, t) dz &= \int_D q_s dz \\ &= \frac{1}{2} \int_D \Delta q dz \\ &= \frac{1}{2} \int_{\partial D} \frac{\partial q}{\partial n} \\ &\leq 0. \end{aligned}$$

The inequality in the last step is elementary: since $q = 0$ at ∂D and $q \geq 0$ in D we have $\partial q / \partial n \leq 0$ at ∂D , where n is the outward unit normal.

Application to the exit time distribution. We used the backward Kolmogorov equation to express the probability that a diffusion reaches a certain threshold before time T (see (8)). The forward Kolmogorov equation gives a very convenient alternative expression for the same quantity. Indeed, if ρ solves the forward Kolmogorov equation in the domain D of interest, with $\rho = 0$ at the boundary and $\rho = \delta_x$ at time 0, then $\int_D \rho(x, T) dx$ gives the probability of surviving till time T . So $1 - \int_D \rho(x, T) dx$ is the probability of hitting the boundary by time T , given that you started at x at time 0. When D is a half-space, this is an alternative expression for $u(x, 0; T)$ defined by (8).