

Lecture 8: From Random Walks to Diffusion

Scribe: Daniel Rudoy*

Division of Engineering and Applied Science, Harvard University

October 3, 2006

In this lecture we provide two examples of problems where reasonable models are given by the deterministic continuum limit of the discrete random walk through which the problem is naturally defined. Once suitably defined, this continuum limit usually takes the form of a partial differential equation (hence deterministic). One interesting use of the continuum limit we will see in upcoming lectures is to understand the behavior of random walks with boundary conditions. To properly incorporate boundary conditions into the calculation of the N -step point mass function of a discrete random walk usually introduces complicated combinatorial calculations. On the other hand, it may be much simpler to solve a differential equation subject to certain boundary conditions in order to answer some questions about the behavior of the random process. However, the applicability of this type of approximation has limits as well. We will explore when these approximations are valid and useful.

1 Example 1: Motion of Bacteria with Flagella

One interesting problem is understanding the motion of *E. Coli* bacteria. The following is a description of how *E. Coli* bacteria move through the use of their flagella that appeared in *Physics Today* [2]: “*Escherichia coli* is a single-celled organism that lives in your gut. It is equipped with a set of rotary motors only 45 nm in diameter. each motor drives a long, thin, helical filament that extends several cell body lengths out into the external medium. The assemblage of motor and filament is called a flagellum. The concerted motion of several flagella enables a cell to swim. A cell can move toward regions that it deems more favorable by measuring changes in the concentrations of certain chemicals in its environment (mostly nutrients), deciding whether life is getting better or worse, and then modulating the direction of rotation of its flagella. Thus, in addition to rotary engines and propellers, *E. colis* standard accessories include particle counters, rate meters, and gear boxes.”

1.1 A Run-and-Tumble Random Walk

When all of the motors spinning the flagella spin in the same direction, they wrap around one another and all the flagella are bundled. This allows the bacterium to travel long distances (compared to its size) in the desired direction and constitutes what is referred to as a “run” mode. On the other hand, when the motors reverse direction, flagella break apart from the bundle, and the bacterium undergoes small random fluctuations (compared to its size) and enters what is referred to as a

*Some additional material by MZB in part 1.

“tumble” mode. Thus the motion of a bacterium can be described by alternating periods of running and tumbling motion for which a reasonable model could be given by a random walk.

In biology, the run/tumble motion is usually approximated more crudely as a simple random walk, where sudden steps of roughly fixed size ($\sigma^2 < \infty$) are made at fixed equal time steps τ with a possible expected drift (or bias) $\vec{\mu}$ in a given direction, as described below. In any case, the details of the walk are neglected by passing to the continuum limit of a diffusion equation, with an additional drift term:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\vec{u}\rho) = D\nabla^2 \rho, \quad (1)$$

where $D = \sigma^2/2d\tau$ and $\vec{u} = \vec{\mu}/\tau$, according to the calculations made in earlier lectures. We have seen that the Green function for this equation approximates represents the “central region” of the PDF of a single bacterium, and thus cannot be trusted for outliers at large deviations from the mean (much farther than $\sigma\sqrt{N}$). The quality of the approximation is also degraded for early times (small N) and large fluctuations (especially diverging higher moments).

In mathematical terms, a more precise model of the run/tumble motion would be a (non-separable) continuous-time random walk of the “leaper” variety. We will encounter such models later in the class. (See 18.366 Lectures 23 and 26 from 2005 or the book of Hughes.) The idea is that the walker moves at constant velocity for a *random* “waiting time” τ between “turning points” where new random direction is chosen. We will see that the crucial aspect to determine the scaling of the walk, and the validity of the diffusion equation approximation, is whether the mean waiting time $\langle \tau \rangle$ is finite. Recent experiments have shown that the run time PDF has a rather fat tail¹, $p(\tau) \propto t^{-3/2}$, which would imply an infinite mean and invalidate the simple analysis leading to the diffusion equation. We will see that fractional derivatives end up replacing normal integer derivatives in the continuum limit in such cases. However, here the PDF tail has a cutoff, beyond which the probability is effectively zero, so that normal diffusion can be recovered at large enough time and length scales. Nevertheless, one should be quite concerned about the validity of continuum modeling based on the diffusion equation, especially for small times when not too many steps have been taken.

1.2 Chemotaxis

The bias in the random walk arises due to the fact that bacteria tend to travel in a specific direction determined, for instance, by a food or chemical gradient. This is the phenomenon of “chemotaxis”. Of course, the bacteria does not actively “steer” itself, but rather it is programmed to lengthen the random time spent in the run mode, if the environment is becoming more favorable. However, it still performs occasionally tumbles, even when it is already climbing a desirable gradient. This allows it to explore its environment more globally, and not get stuck in a locally preferred location.

The chemical attractant itself performs a random walk due to thermal fluctuations (Brownian motion). These are quite small and frequent and well described by the continuum limit at the scale of the bacteria. Therefore, a simple model of the bacterial drift velocity would be to solve a diffusion equation for the concentration C of the attracting chemical, and making the drift proportional to the concentration gradient,

$$\frac{\partial c}{\partial t} = D_c \nabla^2 c, \quad \vec{u} = \chi \nabla c. \quad (2)$$

¹The experimental data can be fit by the Smirnov density, which we will study in later, as the PDF for the first passage time of a normal random walk (or Wiener process). This can be justified by assuming that a large number of molecular “switches” must randomly flip to new position to initiate the tumble mode.

In general, the coefficient $\chi(x, t)$ could depend on position or time, based on how the bacteria react to the attractant in different parts of the system.

More general models of colonies of bacteria are also possible to derive from microscopic random walk type models. For example, in some cases, the bacteria consume and release chemicals which are sensed by other bacteria. This means that an additional reaction-like source term may appear, as in the Keller-Segel Model:

$$\frac{\partial \rho}{\partial t} = D_c \nabla^2 c - \nabla \cdot (\chi \rho \nabla c) \quad (3)$$

$$\frac{\partial c}{\partial t} = D_c \nabla^2 c + k \rho \quad (4)$$

where k is the net production rate of the chemical per bacterium (which can be positive or negative). This is a nonlinear system of partial differential equations, which has interesting behavior, such as spontaneous pattern formation, typical for reaction-diffusion equations.

This example shows the power of continuum modeling to give a reasonable and tractable macroscopic model when the underlying discrete processes become complicated. Here, we actually have *interacting* random walkers, i.e. the bacteria can “talk” to each other via the diffusing chemical signals, which generally leads to *nonlinear diffusion*. Discrete models of interacting stochastic processes commonly arise in statistical physics, and it is generally quite difficult to analyze the models rigorously. However, the continuum limit, even if taken only formally, often suffices for many useful predictions.

2 Example 2: Financial Time Series

In this section, we consider the discrete random walks that arise in the modeling of financial time series along with a deterministic continuum limit of one of the two discrete models. In particular, we are interested in the modeling the time evolution of some financial quantity such as a stock price or interest rate. Much of the discussion is based also on notes from [1].

Bachelier (1900) was probably the first to model the stock market using random walks, as described in his PhD thesis. The random walk takes N steps each of length τ , with $t = N\tau$ equal to the total time required to make N steps. The random walk can be thought of as taking independent displacements over the time interval τ . In various markets that are being considered, τ might have very different lengths, depending on the cost of trading and the logistical delays involved. One way to empirically determine a natural scale for τ is to consider the autocorrelation function of the series such as

$$E(X(t), X(t + t')) = \exp(-t'/\tau) \quad (5)$$

In liquid markets such as stock exchanges, the autocorrelation function decays almost entirely after 20-30 minutes have elapsed, so $\tau = 10$ minutes may be a reasonable time step for a random walk modeling a stock price in such a market.

2.1 Additive Displacements

Random walks that we have considered so far have additive displacements. Moreover, we have seen that CLT type arguments applied to an IID sum for N such displacements characterize the behavior of these walks in the central region, while various types of corrections help to understand what is happening in the tails. For a financial time series, however, this model has a number of limitations:

- Non-negativity. Many of the quantities that we may wish to model are non-negative. For instance prices of goods on a commodities exchange are non-negative. However, it is quite possible that a random walk with additive displacements would take on negative values
- Fluctuations. It is natural that the size of the fluctuations of an asset price are directly related to the magnitude of the price. This is, in fact, confirmed by market observation. However, the size of the additive displacements does not vary with overall magnitude.

Both of the above observations suggest that unless we are looking at a very short interval of time, a multiplicative displacements model may be more appropriate than an additive displacements one.

2.2 Multiplicative Displacements

In previous lectures we have considered processes with additive increments. In this part of the lecture, we consider random walks with multiplicative processes and explore some ways of working with them. A random walk with multiplicative increments is defined by:

$$\frac{X_{N+1}}{X_N} = Y_N \quad (6)$$

One useful instance of this setup is when $Y_N = e^{r_n \tau}$ where r_n is a random variable representing the rate of return on some investment (units of 1/time). The mean and standard deviation of r_n are commonly referred to as the expected rate of return and volatility of X_N . The volatility has units of inverse square root of time. What is important to notice is that while the increments are multiplicative, the logarithm of the random walk is a simple random walk with additive increments. In particular, this yields the following identity:

$$\log\left(\frac{X_N}{X_0}\right) = \sum_{n=1}^N r_n \tau \quad (7)$$

For any $N > 0$, we define a scaled, mean adjusted random variable Z_N by:

$$Z_N = \frac{\log\left(\frac{X_N}{X_0}\right) - N\mu\tau}{\sigma\sqrt{N\tau}} = \frac{\log\left(\frac{X_N}{X_0}\right)}{\sigma\sqrt{N\tau}} \quad (8)$$

where $t = N\tau$ and $\tilde{X}_0 = X_0 e^{\mu N\tau}$. By the Central Limit Theorem we have that Z_N weakly converges to the standard normal distribution with PDF:

$$\phi(Z) = e^{-Z^2/2} / \sqrt{2\pi} \quad (9)$$

in the central region where $Z = O(1)$. This implies that $dZ = dx/x\sigma\sqrt{N\tau}$. This of course means that the limiting distribution of the original process is:

$$P_N(x) = \frac{\exp\left(-\log\left(X_n/\tilde{X}_0\right)/2\sigma^2 N\tau\right)}{x\sqrt{2\pi N\tau\sigma^2}} \quad (10)$$

Setting $t = N\tau$ and rewriting the infinitesimal we have that the limiting PDF of $X(t)$ is given by:

$$\rho(X, t) \sim \frac{\exp\left(-\log\left(X_n/\tilde{X}_0(t)\right)/2\sigma^2 t\right)}{x\sqrt{2\pi\sigma^2 t}} \quad (11)$$

This is the general form of a Log-normal distribution. The moments of $X(t)$ are given by:

$$E(X_N^n) = \tilde{X}_0^n e^{n^2 \sigma^2 t / 2} = X_0^n e^{n \mu t n^2 \sigma^2 t / 2} \quad (12)$$

The first moment here is known as the mean drift. So we have established that a non-linear function of a symmetric random walk, which had no drift, yields a random walk with a noise induced drift. This change of variable from $t = N\tau$ allows us to consider the diffusion equation limit of the previously described random walk. In particular, a standard model of a drifting process as the one we described above is give by the following diffusion equation with a drift term:

$$\frac{\partial \rho}{\partial t} + \bar{\mu} x \frac{\partial \rho}{\partial x} = D x^2 \frac{\partial^2 \rho}{\partial x^2} \quad (13)$$

where $\bar{\mu} = \mu + \sigma^2/2$ and $D = \sigma^2/2\tau$. However, it is important to keep in mind that this approximation is only good in the central region we previously described.

This concludes our second example of passing from a discrete random walk to a continuum approximation. In the next section, we will explore some of more interesting properties of the Log-normal distribution. One striking property is that the Log-normal distribution is not uniquely defined by its moments, that is there are (infinitely many) other distributions with the same moments.

3 Fun with the Lognormal Distribution

In this section, we explore the properties of the Log-normal distribution in a bit more depth. Recall the definition of the scaled lognormal density with subtracted mean:

$$f_0(x) = \frac{1}{\sqrt{2\pi}x} \exp(-(\log(x))^2/2) \quad x \geq 0$$

In particular, we are interested in computing the moments of the Log-normal. We use the fact that if x is a random variable with the standard normal distribution then by the change of variables formula $y = \exp(x)$ has the Log-normal distribution. Thus, the moments can be obtained as follows:

$$\begin{aligned} E[Y^n] &= E[\exp(nx)] = \frac{1}{\sqrt{2\pi}} \int e^{nx} e^{-x^2/2} dx \\ &= \frac{e^{n^2/2}}{\sqrt{2\pi}} \int e^{nx} e^{-(x-n)^2/2} dx = e^{n^2/2} \end{aligned}$$

3.1 Family of distributions with the same moments as the Log-normal

In this section, we explicitly describe a family of distributions such that every member of this family has the same moments as the Log-normal distribution. This counterexample is due to Heyde [4] and the discussion below follows that of Durrett [3].

For any $-1 \leq a \leq 1$, we define the family of distributions:

$$f_a(x) = f_0(x) (1 + a \sin(2\pi \log(x)))$$

To show that f_a has the same moments as f_0 we just need to show that for any $r \in \mathbb{N}$:

$$\int_0^\infty x^r f_0(x) dx - \int_0^\infty x^r f_a(x) dx = \int_0^\infty x^r f_0(x) \sin(2\pi \log(x)) dx = 0$$

To do this integral we change variables by letting $x = \exp(s+r)$. This implies that $s = \log(x) - r$ and $ds = dx/x$. Rewriting the integral yields:

$$\begin{aligned} \int_0^\infty x^r f_0(x) \sin(2\pi \log(x)) dx &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^\infty \exp(rs + r^2) \exp(-(r+s)^2/2) \sin(2\pi s) ds \\ &= \frac{1}{\sqrt{2\pi}} \exp(r^2/2) \int_{-\infty}^{+\infty} \exp(-s^2/2) \sin(2\pi s) ds = 0 \end{aligned}$$

The last equality follows because the integrand is an odd function.

3.2 Discrete distribution with the same moments as the Log-normal

To perhaps make the point even more forcefully, we provide an example of a *discrete* random variable that has the same moments as the Log-normal distribution. Let $a > 0$ and define the random variable Y_a by:

$$P(Y_a = ae^k) = a^{-(k-n)} \exp(-k^2/2)/c_a \quad k \in \mathbb{Z}$$

where c_a is chosen to make the total probability mass sum to 1. To see that this has the same moments as the Log-normal, we multiply the n^{th} moment of Y_a by the reciprocal of the n^{th} moment of the Log-normal:

$$\begin{aligned} \exp(-n^2/2) E(Y_a^n) &= \exp(-n^2/2) \sum_k (ae^k)^n a^{-(k-n)} \exp(-k^2/2)/c_a \\ &= \sum_k a^{-(k-n)} \exp(-(k-n)^2/2)/c_a = 1 \end{aligned}$$

References

- [1] Erik Allen. Applications in finance. *18.366 Course Notes*, March 16, 2005.
- [2] H. C. Berg. Motive behavior of bacteria. *Physics Today*, 53(1):24–29, 2000.
- [3] R. Durrett. *Probability: Theory and Examples*. Duxbury Advanced Series, 2005.
- [4] C. C. Heyde. On a property of the lognormal distribution. *Journal of the Royal Statistical Society B*, 29:392–393, 1963.