[SQUEAKING]

[RUSTLING]

[CLICKING]

PETER KEMPTHORNE: Today, we're going to talk about stochastic processes, the second set of lectures. And we're going to begin with a discussion of Brownian motion.

Brownian motion is a process that I'm sure most of you have heard of. And it was named after Robert Brown, an English botanist. And this is an interesting case where a scientist observes some dynamics in nature and then is concerned about what mathematical model might describe that behavior.

Turns out that it was, I guess, Albert Einstein who formalized this Brownian motion model, mathematically. And Norbert Wiener, our MIT hero-- those of you familiar with the math department know about the Norbert Wiener room that we have-- and he developed the mathematical foundation for Brownian motion.

Now, let's take a look at a simulation of Brownian motion in two dimensions. So this is the kind of dynamics that Brown was observing. And so in these slides, we're going to start the process at the origin of the space, where the red dot is, and then move around in two dimensions, where each step is a random move in the x direction and a random move in the y direction.

And for this simulation, we use Gaussian distributions. But as I page through these graphs, one can see how a particle moving according to these dynamics will start at the origin and evolve to the final red dot. And if we hold down the graphs and see where it goes, you can see that basically there's purely random dynamics but of a very special kind. And so the special kind is that the movements are jagged and unpredictable in terms of their direction.

Something that's really quite interesting when you just look at this and page through the dynamics here, animating the process, the movements, in fact, have no trend in them to move left or up or down. And so the movements we see basically are movements that are symmetrical in their dynamics. And so that's a very interesting property that we would like to characterize mathematically.

So let's look at how we define Brownian motion in one dimension. So in one dimension, we are simply looking at a graph. Let me find the eraser here. OK.

In one dimension, we have time t. And we're going to think of a process Bt, which will start at 0. And then it's going to evolve over time in such a way that the increments of the process from t to t plus delta t-- so if we have t and t plus delta t here, and we look at B of t here and B of t plus delta t-- this vertical increment is going to be normally distributed for any increment of time delta t.

And the normal distribution for the magnitude of that increment will be a normal random variable with mean 0 and a variance that is proportional to the length of the interval. So as we consider increments, delta t that are larger will have greater variation in the vertical increment that's normally distributed.

Now, in addition to normal increments, we have independent, disjoint increments. So if we have different time points, like t1, t2, and then t3, t4-- if these increments are separate or disjoint from each other, then the change from t1 to t2 and the change from t3 to t4, those will be independent random variables.

Finally, B of t turns out to be a continuous function of time if it satisfies these properties. And so this is a very simple model in terms of a continuous time stochastic process.

Now, the properties of Brownian motion are that it's a Markov process. So if we're looking at the value at B of t plus delta t at this time point, that is equal to the value at any prior time point-- well, at B of t-- plus the increment.

And so given the information set up to time little t, the conditional distribution simply is equal to that starting point, B of t, plus the random increment. So displacements are independent of the past, apart from the last value of the past.

OK. Mean and variance of B of t-- well, because the increments are normally distributed with mean 0, then the expected value of B of t into the future will stay at 0. But the variance of B of t will increase with time, and so the standard deviation of B of t will be proportional-- or will be sigma times the square root of t.

So if we think of drawing curves which correspond to sigma square root of t, this corresponds to 1 standard deviation above the 0 mean for the process.

OK. The covariance function of Brownian motion-- if we want to look at the covariance between two times, s and t-- so suppose here is time s, in here is time t-- then the process paths to time t include the process paths up to time s, but then have this independent increment beyond time s.

And so when we look at the covariance of the process at times s and t, it's equal to the expected value of the product. But we can split B of t into B of s plus the increment, and we simply get sigma squared s for the covariance. So the degree to which the process is correlated with other values of the process, it depends on how much overlap there is in the time intervals corresponding to those two values.

Now, this process that starts at 0, B of 0 or B sub 0 or B of 0 equal to 0, we could start it at any other level and have a process that has similar properties but just has a different mean level from time 0. So we can talk about starting the Brownian motion at any given point.

And so in order to work with these models, we'll find it convenient to introduce the notation for normal random variables. And so if we have Z be a standard normal, that is a normal distribution with mean 0 and standard deviation 1, then the probability density function is given by this formula. I'm sure everyone is comfortable with that. And the cumulative distribution function is simply the integral of the density from minus infinity up to c for phi of c. And this notation of small case phi and uppercase Phi are standard for referring to Gaussian or normal densities in cdf's.

Well, if we have a normal with non-zero mean mu and variance tau squared, or standard deviation tau, then we have the density of this normal Y is a simple function of the density of the standard normal. We simply subtract the mean mu and divide by tau inside phi, and we have a tau denominator factor in front.

And what's particularly convenient at times is that, if we want to calculate the cdf of Y, then this cdf of Y is actually the cdf under the standard normal of c minus mu over tau. So we basically have Y is equal to mu plus sigma or tau z. And so z is equal to Y minus mu over tau. And so we're just substituting in this equation for Y in terms of z.

Now, the probability model of the whole Brownian motion process is given by-- or we can represent it by saying the conditional distribution of the time t value, given the time s value, is normally distributed with mean equal to the time s value, X, and variance sigma squared times the length of the increment.

So its pdf has this form. So the notation here is the density of y at time t, given x, the value at time s. We have this Gaussian density that's centered at x. And the variance of y depends on the increment from s to t multiplied by the sigma squared.

Now, with this probability model, it turns out this density function satisfies the diffusion equation. And the diffusion equation is going to arise again and again in this course, especially when we get into the valuation of derivatives.

And so for now, let's just understand this diffusion equation. It says that the partial derivative of p with respect to time t is equal to a multiple of the second partial derivative of p with respect to y.

So this is easy to verify for s equals 0 and sigma equals 1. And if we look at this set of graphs, here's the density of X, given-- or let's say, here's, essentially, the density of Y represented by X. So the notation symbol here is we're using X now instead of Y. And we're considering the initial value of the Brownian motion being 0.

And if we increase time, then this normal density drops in the middle and it increases far from the middle. And so the change in the density is negative near the middle and positive outside.

And that's because, if we look at the second derivative of this density function, it's positive, then it's negative for the second derivative. The curvature is negative. And then it switches from being a negative curvature to a positive curvature. And those points of inflection with the normal density are actually at plus and minus 1 standard deviation.

So if we increase the time from 1 to 2 to 3, or 2.5 here, you can see how it's basically dropping in the middle of the distribution and it's rising on the other side. So this kind of motion or dynamics satisfies the diffusion equation. And so what's, I think, important is just to have a good sense of what it means when a function such as the density satisfies this diffusion equation.

Now, let's see. This diffusion equation relates to a conservation of probability for every x. So regardless of what the starting value is, we have the probability density integrating to 1 always. And it turns out that this density is a unique solution to this diffusion equation.

So an important issue in the mathematical development framework for models is our solutions to our equations. Do they exist? Are they unique? In fact, for this diffusion equation, we get a unique solution satisfying the properties.

OK. In other disciplines-- well, in physics, there's a heat equation. Let's see. In the dynamics of populations, there's the Fokker-Planck equation and Kolmogorov Forward equation. These arise outside of mathematical Brownian motion as well.

Now, the interesting question perhaps, or at least the interesting question to mathematicians is, how do we really define Brownian motion formally as a stochastic process? And so it turns out that Brownian motion can be defined as the limiting stochastic process corresponding to a random walk.

So let's consider random variables X indexed by integers that are independent and identically distributed. And let's assume that those X's have mean 0 and standard deviation 1. Then, if we look at the random walk with steps X1 to Xn for the n-th position of the random walk, then the limiting distribution of the normalized random walk is actually a standard normal distribution.

So phi of c is our standard normal cumulative distribution function. The standardized random walk will in the limit, as n grows large, have a Gaussian distribution with mean 0 and standard deviation 1.

So what we can do is consider normalizing the random-walk process by having the partial sums of the steps divided by root n. But rather than focusing just on those increments, those integer increments, we consider four times, say, between 0 and 1-- we can consider values for time t between, say, integer k over n and integer k over n plus 1 over n to define a continuous path of Bn of t, given s of n and t.

And so what we then can do is say, let's consider the limit as the number of increments-- number of steps increases arbitrarily. Let's consider the limiting process. And that turns out to be a Brownian motion process.

So with this normalized random-walk process, it turns out that this limiting distribution is-- or this limiting process is invariant to the distribution of the X's. So if we have random walks with independent steps that have mean 0, standard deviation 1, then regardless of what that distribution is, the limiting normalized random walk will be the diffusion process.

So let's take a look at two cases for the steps, Bernoulli X's and Gaussian X's, which of course will be Gaussian. And so here is a cumulation or a random walk with 10 steps. The top panel corresponds to unit plus 1 or minus 1 steps, and the bottom panel corresponds to Gaussian steps.

And if we increase the number of steps and normalize them by dividing by root t, we end up getting these different paths tracing out. And so what's important in this display is that, as we let n get arbitrarily large, the process will be equivalent to Brownian motion.

So we're going to go up to 2,000 steps here. And one can see how this normalized random-walk process evolves. And when we get to the end of the simulation here of 2,000 steps, this is what we have.

And so the key takeaway is that the limiting process will be a Brownian motion, and it doesn't matter what the step distribution is in the limit. And so this implies that this model may be useful in a variety of contexts, where underlying steps of the process increments may have some distribution that's non-Gaussian. But the distribution of increments over long periods of time will then follow the diffusion.

So the key properties of Brownian motion are independent increments. And the variance of those increments are proportional to the time period between the measures.

Now, there's a reflection principle in working with Brownian motion that is very powerful. This is an example where the mathematical theory is a bit tricky. But if you're clever at solving the problem, you may find an elegant, simple solution to the problem.

So let's just understand what it means to be reflecting Brownian motion first. So we consider Brownian motion to be the process B of t. And so this will be one realization of the path perhaps.

And then if we consider the first time the process hits x and define tau equal to the first time B of t is equal to x, then, OK, this is the first time it hits that level x. If we consider reflecting this path about the level at time x going into the future-- so this would be going like this and then up like that--

This is the mathematical formula for doing a reflection of B be about x. And we'll see some examples of this in a minute. But what is important in these reflected paths is that the probability model for the original process and the reflected process, those are identical.

So let's take a look at this reflected Brownian motion simulation. So what I'm doing here is simulating Brownian motion over 10,000 steps and looking at an original process in red, and then reflecting that in green. And because of overlapping colors, I guess, the beginning process at time 0 is the original and the reflected path. So it appears as the reflected.

So once it hits this level x in red, the green reflects about x, the value. And it continues reflecting across the remaining path of the original series.

Now, here's another realization of reflected Brownian motion, has the same property. And here's a third. Basically, with different realizations of Brownian motion, we'll get different paths. And the reflected paths have these same properties.

Now, the interesting question is-- well, it's actually-- let me go back to the notes here. And the interesting question is to ask, what is the maximum of a Brownian motion up to time t? So we can define m of t to be, what is the maximum value that the Brownian motion has realized between time 0 and t?

Well, the probability of the maximum exceeding x, some level, will actually be twice the probability that the ending value of the process exceeds x. So let's take a look again at the reflected paths.

So if we look at the reflected paths about x, then the red and the green paths have the same probabilities of occurring. Let's see. If we look at whether the ending value of the original Brownian motion-- sorry-- or the ending value of a Brownian motion exceeds a given level, then there will be twice that number, twice those realizations for which the maximum exceeded that level.

So for both of these reflected paths, the maximum-- let's see-- the maximum of the path exceeded x. And so looking here, the probability of the maximum exceeding x, well, we can just look at the probability of a path exceeding x at the end point. And the reflected intervals will match those, one for one.

And so we basically have the formula that we have twice the final value of the process exceeding x is equal to the probability that the maximum exceeds x. And we can write out what that probability is in terms of the standard normal cdf.

Now, from this result, we can differentiate this probability with respect to x and get the density of the maximum. And so let's see here.

So here is a graph of the cumulative probability of the maximum exceeding a level x. And this is, basically, we're interested in, what's the probability of exceeding x before time t? And with lower values of x, like 1, the probability is quite high. Let's see. If x is very-- let's see. Or, sorry. Is that x there?

Right. OK. So if the maximum is 3, the probability of exceeding 3 is much lower. The probability of exceeding 2 is in the mid-way there. And what's also very interesting is that we can use these results to calculate the density function for the first hitting time of a given level.

And so importantly, we have the probability of the max up to time t exceeding x is twice the probability that the ending value exceeds x. And we can differentiate this with respect to t. And we get this density function for the density of the stopping time for hitting a given level x.

So let's see here if we have that. So here's the hitting time densities. So with standard Brownian motion, if we're looking at the time to hit the level 1 starting at 0, then there's this nice unimodal distribution for that time, which is relatively short. But if we're looking at hitting times for higher levels, like two or three, then that hitting time distribution spreads out.

So now, let's see. We can talk about some extensions of Brownian motion. One is just reflected Brownian motion around 0.

So if we have the process bouncing off 0 level, we can think of that process being a positive stochastic process. And it's fairly easy to show that the expected value of this reflected Brownian motion is a multiple of root t, square root of 2 over pi, and the variance is 1 minus 2 over pi t.

And we can also consider absorbed Brownian motion. If we consider tau being the first time the level 0-- or level x is hit, then we can define our process being equal to the original Brownian motion B of t, or B at 0 for t greater than tau.

And so with absorbed Brownian motion, these kinds of processes are useful, for example, when we have a Brownian motion-like process for a stock price, which is generally positive, so long as the company is viable. But if the company goes bankrupt, then the price gets absorbed at 0.

Now, we have another extension of Brownian motion, which is called the Brownian Bridge. And so if we consider a Brownian motion process starting at 0, and it goes up to, let's say, time 1, starting at 0, we can define-- we can define the procedure x, which is equal to B of t-- let's see-- B of t minus tB of 1. Let's see. Is that right? Minus t times B of 1.

So if we consider-- let's see-- looking at-- well, let's see. This process, so x of 1 is equal to B of 1 minus 1 times B of 1, which is equal to 0. If we basically draw a line to-- if we consider this line as 1 and this has slope B of 1, then if we subtract tB of 1, we're looking at pinning down the Brownian motion and looking at the distance from a constant slope, tB of 1 and B of t.

And this particular process actually arises in a very interesting context of normalized empirical distribution functions. So what do we want to highlight here?

If we have a sequence of random variables that are uniformly distributed between 0 and 1 and we calculate their empirical distribution function-- so it's the count of the U's less than or equal to t divided by n-- then this empirical distribution function will have expectation equal to t and variance equal to t times 1 minus t over n.

So if we think of these U's-- so we have Ui is distributed uniform from 0 to 1. We have the scale from 0 to 1. And we consider a value Ui as a realization. Then the probability that Ui is less than or equal to t is equal to t.

And the variance of the indicator of Ui being less than or equal to t, this is a Bernoulli outcome, which has variance equal to the probability of occurring times the complement of that. So we get this t times 1 minus t being the variance of-- the indicator of the Ui being less than or equal to t.

And then if we consider the empirical distribution function, which is summing over these indicators, then we get a process which has a limiting value, which is basically the deviation from its mean, scaled up by root n, is normal with mean 0 and variance t times 1 minus t. And this is actually the Brownian Bridge process.

Now, let's see. With these examples where we use uniform random variables and properties of functions of those uniforms, it's important to note that for any continuous distribution, say, Y distributed with some-- this is the cdf of Y. If we calculate F of Y, this happens to be a uniform distribution.

So this is the probability integral transform of any random variable. And so when we have this property for the empirical distribution function of uniforms-- well, of uniforms, if we replace the U's, which are uniform 0, 1 by assumption, with a distribution function F of Y, where we have Y1 to Yn iid, with this cdf F sub Y, then the Ui equal to F of Yi F sub Y of i, these will be iid uniform.

So that's a really nice result. And it leads to issues involved with modeling model selection. If we want to know whether a given an assumed probability model for Y's is true or not, we can do this transformation and test whether the transformed series or transformed sample is consistent with being uniform or not.

All right. Well, with simple Brownian motion, standard Brownian motion, we can consider an extension of that to Brownian motion with drift, mu. So we can consider a drift parameter and a volatility parameter. And such a process will have independent increments, with variances of differences proportional to the time interval between s and t.

And we can work with the Brownian motion with drift using a one-step analysis. And in considering this X process-- so X is equal to t mu plus sigma B of t. So this is X of t.

If we look at the increment ahead delta t from little t, then expanding out the representation of this Brownian motion with drift, we get that it's equal to the Brownian motion at time little t plus mu delta t plus sigma times the increment of the Brownian motion process from t to t plus delta t.

And so increments of X correspond to increments of time plus times the drift rate, mu, plus sigma, the volatility parameter, times the increment of the standard Brownian motion.

So our process X, Brownian with drift, will have non-zero mean increments and increments variances that are proportional to the change in time. And this increment will have an exact normal distribution with these properties.

And what's particularly important is that, as we let delta t, the length of our increments, get infinitesimal, then the expected squared increment of X ends up being equal to-- well, it's exactly equal to the variance of that increment plus the square of the mean. But the square of the mean is a term that is small order of delta t.

So our expected squared increments will be to order delta t equal to sigma squared delta t and then a small order of delta t term after that. Also, if we consider powers of delta X, so our increments to a power, then when our power c is greater than 2-- it's equal to 2 here-- when it's greater than 2, then this will also be a small order delta t.

Now, the reason why we're interested in these properties is because we're going to be working, ultimately, with, say, a value of some derivative that depends upon X of t. So there's basically a derivative value, which depends upon the underlying-- I'll just say the underlying asset price.

And so the dynamics of the derivative prices can be written as a Taylor series. And so we'll want to be able to deal with increments of X and powers of increments of X as our valuation function may be expressed, if it has a Taylor series representation.

All right. So let's consider the problem of Brownian motion underlying a gambler's ruin problem. So let's let X of t be Brownian motion with drift, mu and variance sigma squared. And we'll consider an initial level equal to X that's between two levels, a and b. And we want to consider the stopping time hitting a or b and ask, what's the probability that we hit b first?

So over time t, we have, basically, levels a and b and an initial level X at time 0. And we're interested in hitting a or b so tau is equal to the hit time of a or b. And we're interested in, what's the probability of hitting b first before a?

And this generalizes our problem in the first lecture on stochastic processes, where we had a random walk with unit steps with drift. So there was a bias of maybe a plus 1 versus a minus 1 occurring. And we wanted to solve for, what's the likelihood of hitting a level b before level a?

So with this, we use an infinitesimal one-step analysis. And this solution approach is illustrative of how many problems are solved with stochastic processes.

We can define u of x to be the probability function, which is the probability that the stopped process equals b, given that it starts at X. And so if we consider u of x and write it as the expected value of u of x plus delta X, then if our delta t time increment is very, very short, then this expectation will equal u of x. So the likelihood of hitting b first will be the expectation of u of x, a very short increment into the future.

And so if we take the expectation with respect to the delta X variable, well, delta X is X of t plus delta t minus X of t-- then if we assume that our function u, this probability, is twice differentiable, then the Taylor series about the center X will equal u of x plus delta X times the first derivative plus a half delta X squared times the second derivative, plus small order terms of delta X squared.

And so if we take an expectation of that first Taylor series on both sides, well, the u of x is a constant. And then we have u prime of x, a constant times the expected value of delta X. And then plus a half the expected value of delta x squared times u double prime. And we have mu delta t is the expectation of delta X. Sigma squared delta t plus little o of delta t is the expectation of delta X squared.

So we have this expansion from the Taylor series. And from what we said before, this expected value of u evaluated at x plus delta X should simply equal u of x. And if that's true, then we get this equation for u prime and u double prime.

And if we take that equation and u prime and u double prime and we divide by delta t, we then get 0 is equal to mu u prime of x plus a half sigma squared u double prime of x. And to be really precise, little o factor to make it an equation.

Now, this equation here is a second-order differential equation. And so what are solutions to this second-order differential equation?

This is an interesting challenge. Let me just get some more chalk here. It's maybe been a number of years since people have taken differential equations, right?

But let's consider, let's see, u double prime over u prime is equal to-- let's see. Have I done this right? u double prime divided by u prime is equal to minus 2 mu over sigma squared. And so this is the derivative is equal to the derivative of the log of u prime.

And so if we integrate this, we basically get log of u prime is equal to minus 2 mu over sigma squared x plus a constant. And this is true if mu is not equal to 0. And so we basically get that u prime is equal to e to the minus 2 mu x over sigma squared times a constant.

And what if we integrate both sides of this? We then get u. So what function has a derivative equal to an exponential?

AUDIENCE: The exponential.

PETER KEMPTHORNE: Just the exponential, yeah. So we get u is just going to equal a constant times e to the minus 2 mu over sigma squared x.

And if mu is equal to 0, then we have the log of u prime is equal to 0, or the derivative-- yeah, we basically have the derivative of the log of u prime is equal to 0. So we end up getting-- let's see-- basically a linear function for the solution.

So let me go to the notes here. Well, we basically get u is equal to Ax plus B, if mu is equal to 0, and it's equal to this constant, A times e to the minus 2 mu x over sigma squared. All right. Well, this general solution has to satisfy the initial conditions. u of a is going to equal 0, and u of b is going to equal 1.

So if we start at a, where the-- basically, we stop at a, if we start at a. So u of 0 is going to equal 0. And if we start at b, it's equal to 1. We end up getting this formula for what u is. And so what's interesting to ask is, how does this probability vary with different drifts and volatilities?

OK. So here's the graph of these probabilities of hitting b before a when x-- well, when a is 0. So if we have 0 here and b is equal to 3, then here's the initial value x between 0 and 3. And we have the probabilities of hitting b before a

And if our mu is very high, so the trend is up in the Brownian motion-- so basically, if mu is greater than 0, then we basically have an upward-trending path on average, and we get high probabilities of hitting b first. And as we decrease the drift from 3 down to 0.5, we get this Brown curve.

And the limiting case, where mu is equal to 0, so there's no drift at all, corresponds to the probability given by this line. And so given the starting value x, the closer you are to b, the more likely you will be to hit b first and vice versa with a.

So these probabilities vary with the different model parameters. Here is the corresponding graph when there's a negative drift. So the likelihood of hitting b first is proportional when the drift term is 0-- basically how close you are to b. But as we increase the magnitude of the negative drift, we basically get lower and lower probabilities coming out. So that's that solution.

OK. So let's see. The next comment about Brownian motion is a very important property that the Brownian motion process is not differentiable. So when you look at a Brownian motion path, it looks very jagged. And in fact, that jaggedness is characteristic of there not being a smooth derivative, or it's not being smooth enough to compute a derivative.

And so if we look at the increment of the Brownian motion from t to t plus delta t and have that being normal with mean 0 and variance delta t, the length of that increment-- if we normalize the increment by delta t-- so we're basically trying to look at the slope between t and t plus delta t-- then that has a normal distribution with mean 0, but variance 1 over delta t.

So this normalized increment is going to not be well defined in the limit. It's basically going to be equivalent to-- or the limiting value basically doesn't exist. It basically becomes an explosive distribution. So that's a very important property of Brownian motion.

Something that is a positive property of Brownian motion is the issue of quadratic variation. And so to define quadratic variation, we consider looking at partitions of the time. So we have pi, which is, what, a partition of the interval from 0 to t, say-- or 0 to capital T.

So we have a time scale finishing at T, starting at 0. And we can just draw arbitrary points t, j, j equaling 1 up to n, say, or 1 to n minus 1, that are inside this interval. And they're ordered.

If we consider the partition corresponding to these time points, we can define a magnitude of the partition equal to the maximum of tj minus tj minus 1 over j. And then we can talk about the quadratic variation of the process, which is to look at increments of b over these intervals from between tj minus 1 and tj, square those, and sum the squared of those values.

This is an appropriate term for a quadratic variation. And if we consider the limit as the partition width gets narrower and narrower, then we'll define this limiting value to be the quadratic variation of the Brownian motion process.

So we can define the quadratic variation corresponding to a given partition and then the quadratic variation overall, if we take the limit over partitions that are getting infinitesimally small. And the claim is that this quadratic variation is going to equal capital T, the length of the period, with probability 1.

So what does that mean? That means that for every realization of the process, or almost every realization, this quadratic variation is going to equal a constant.

So here, I'm looking at examples of random walks that are normalized to be between time 0 and 1. So these are the normalized random walks, which are effectively Brownian-- or whose limits are Brownian motion. And in the top panel, I'm drawing the path of the random walk, with Gaussian steps. And in the bottom, I'm drawing the cumulative quadratic variation.

So this is up to 20, 30, 40 steps now. As we increase to 100 steps, this cumulative quadratic variation is getting pretty close to the time t over which we're accumulating the quadratic variation. And as we increase the number of steps from 200 to many more-- 1,000 and then 2,000-- you can see that, basically, the quadratic variation of the process cumulates, basically, to the time period over which that's being cumulated.

So this property is a really nice property. What it highlights is that there are some significant regularities to Brownian motion. And the regularity here is that the variation of the increments normalizes, basically, to the length of the period over which they're accumulated. So the quadratic variation is proportional to the time period.

Now, this idea, or this property of Brownian motion, is one that can be explored with time series that we might think are well-modeled by Brownian motion. So if we were to look at a time series of the log of Pt divided by P0 for Pt equal to 0, 1, up to capital N, basically, for prices of assets and equity markets, this corresponds to the log return from time 0 to t. And that log return from 0 to t might follow a Brownian motion process with drift, positive return. And so we basically can have this process evolving.

If we were to calculate the quadratic variation of this process, then if it followed-- if it followed a Brownian motion process, it would tend to grow according to a straight line, with slope equal to sigma squared.

And in empirical modeling, if we look at this quadratic variation of an asset and see that it's not a straight line-- there may be periods where the slope is low and periods where the slope is higher and it curves around-- a way of trying to model the dynamic volatility could be suggested by studying these plots and using them. So it's a very interesting property.

Now, let's see. Let's see. There's one last note in the slide in the notes. But before doing that, I wanted to highlight a connection of Brownian motion to the Laplace distribution.

And so if we have-- let's see. If we have B of t, which is Brownian motion, and we observe it at-- let's see-- if we look at, basically, x of-- if we observe it at times t equals 0, 1, 2, discrete times, then the delta B of t are distributed as normal with mean-- I guess it's mu t and variance sigma squared times 1, which is delta t. And these are iid.

So we have, basically, time points 0, 1, 2, and so forth. If instead of observing the process at equal increments of time, let's consider the possibility that the time increments are random.

So let's consider B of t for t equal to T1, T2, up to Tn, where the delta T's, the increments of time, are distributed as an exponential random variable, with some mean lambda, say. So we're not observing the process at regular intervals. We're observing it at random intervals.

It turns out, then, that the Brownian motion process has increments that are Laplace distributed. So the delta B of t, then will, in fact, have what's called a Laplace distribution. And so instead of having a bell-shaped curve, which is a Gaussian distribution, it'll be a bilateral, exponential distribution about the mean.

And it turns out that with Standard and Poor's 500 returns, this Laplace model seems to fit the data much better, which leads to the consideration of what are called subordinated stochastic processes, where the time scale is different than our clock time.

All right. Well, let's see. The final slide here for today addresses what are called adapted processes. And so in financial modeling, we can think of there being a stochastic process X of t characterizing relevant dynamics in the system. And we can think of X representing the prices of a single asset or a collection of assets as it varies over time.

Then, as we let time increase, there are events that relate to what has happened so far. And we can think of these script F sub t to be the information on realizations up to time little t. And if we consider script F sub t prime with t prime greater than t, then that corresponds to information about all realizations from t up until t prime. So it's giving us more information about what's happened with the process.

Well, if we want to look at the stochastic modeling of a process Y that depends on trading on assets with dynamics X, then we want our process Y to be adapted to X, which essentially means we can't see the future in trading.

And so our Y of t, which is perhaps the value of trading up until time t, is a function only of the values of the process prior to time t. And so we consider transformations of our original process to the process of, say, trading gains and losses, according to some trading strategy. And these will be well defined if the process is adapted.

So with that, the topics connect with what we had in our previous lecture, which was we can have martingales defined on X, and then we can define martingales relating to Y when we combine adaptive functions to X. And so the martingale theory extends to modeling trading strategies, effectively.

And what's interesting with those techniques is that you end up transforming the process to have a constant mean. Martingales are processes with a constant mean. And so those arguments are applied again and again in solving for their properties. And option pricing can be expressed in this context as well.

All right. We'll finish there for today.