[SQUEAKING]

[RUSTLING]

[CLICKING]

PETER KEMPTHORNE: All right. So today's topic is stochastic calculus. And this topic is really an engaging one in quantitative finance because it extends ordinary calculus with real variables to calculus that depends upon stochastic processes, and in particular, Brownian motion processes.

And so when we are thinking about how to model the value of a stock or of a derivative, such as a simple option or a more complex option, we can model the dynamics of these assets or derivative contracts as integrals of functions with respect to Brownian motion, where Brownian motion characterizes the underlying dynamics of the asset.

And so it's really neat to see how we can develop stochastic calculus and relate it to what's familiar to us with ordinary real calculus, but see how there are extensions of the formulas that accommodate the randomness in the increments with which we're integrating.

So let's consider just reviewing, first, Brownian motion with drift. We discussed this earlier in our discussion of stochastic processes. And a standard Brownian motion model is one where we have our process X, depending upon the underlying Brownian motion-- standardized Brownian motion B of t, with a drift parameter mu, and a volatility parameter sigma, or variance parameter sigma squared. And the key property of Brownian motion includes independent increments and the variation of increments of the process being proportional to the length of the increment.

So with standard Brownian motion, where we have a drift parameter of 0 and a sigma volatility of 1, the extension of these properties to have essentially mean 0 and standard deviation equal to sigma times t is to have the mean depend upon time, with the drift parameter and the volatility or variance to be proportional to sigma squared times the length of the interval.

Now, in terms of understanding Brownian motion, it's going to be useful for us to focus on the conditional distribution of the value at time t plus delta as a function of the value at time X of t.

And so when we think of a graph of Brownian motion, which is some path, if we fix, at a given value t, the value of X of t-- or maybe I should write this as s-- the value of X at time little t, then if we know the value at time little t, what is the distribution, the conditional distribution into the future?

Well, this increment t plus delta t, delta t in red, we can just expand that out to mu t plus delta t plus sigma B of t plus delta t. And then we can group the known terms on the left-hand side, or the left-hand terms of the second equation, mu t plus sigma B of t. Then we have a term mu delta t, which is a constant, not random. But then we have the increment of the Brownian motion process from time t to time delta t, t plus delta t.

So actually, there's a typo here. That should be B of t plus delta t minus B of t. And so we have that our delta X increment is mu delta t plus sigma delta B of t.

Now, this increment has mean mu delta t, variance sigma squared delta t. And for Brownian motion with drift, this increment is exactly distributed as a normal random variable, with that mean and variance. And what's useful to think about is, what happens as the increment delta t gets infinitesimal?

And so if we look at the infinitesimal increment for delta t, then the expected squared value of the Brownian motion increment is equal to sigma squared delta t plus mu delta t squared. This is actually an exact formula.

But what's important for us is that, as delta t goes to 0, as it becomes infinitesimal, then this squared increment has expectation that's of order delta t. And it actually equals sigma squared delta t plus little o of delta t. So this little o notation means that as delta t goes to 0, this remainder part goes to 0 as well.

And if we consider powers of our increment of delta X, where the power c is greater than 2, then this will also be small o of delta t. So expressions that depend upon second or higher order of delta X will actually be comparable to sigma squared delta t.

So with Brownian motion and drift, we have the notation for a stochastic differential equation, which is simply the increment of X is equal to mu delta t, or dt rather, plus sigma dB of t. And with this notation, the solution turns out to be X of t is equal to the value of X at time 0 plus the integral from 0 to t mu dt plus the integral from 0 to t of sigma dBs.

And this solution seems obvious when we think about how we understand differential equations with real functions, but it does extend to this solution when we have Brownian motion. And so as noted there, this is equal to X0 plus t minus 0 times mu plus sigma times B of t minus B of 0.

So our integration of functions of Brownian motion have this nominal solution, which seems very reasonable. And we'll, in fact, show that this solution is useful and it actually extends to more complex equations.

Now, the more complex differential equations are one where we consider a drift term, mu, that's not a constant. It may vary over time, and it may also vary depending upon the magnitude of the process.

So this kind of an extension is quite general. And we can think about the same extension to the volatility. Rather than being a constant, it might vary with t, and it might vary with Xt.

And with derivative contracts, we have values of contracts which, at any given time, depend on time, maybe time to maturity, and the level of the process. And apart from those two features or properties of the derivative at time little t, we have a valuation model that applies, according to some function. So it will be very useful to have these generalized drift and generalized volatility parameters. Yes.

AUDIENCE: How is the second integral defined? Because, obviously, the first one is like a Riemann integral. What about the second one? Is there an actual formal definition of-- I guess it's an Ito integral.

PETER KEMPTHORNE: Right. And we will get there in a moment. But if we're integrating from-- well, here it's from 0 to t. If we're integrating the process over increments, we're basically going to be summing over the increments, although summing on the infinitesimal timescale over the entire period between 0 and t. And so if we break that sum of infinitesimal increments, we basically get the last value minus the first. And so we'll see that come into play.

So what we will determine with this general setup, though, is that when we have extended drift and volatility, we then will have a solution to this Brownian motion with drift, with generalized drift and volatility, satisfying this equation. So we're going to want to be able to compute these integrals in a consistent way. And the Ito calculus provides us with that. So we have to define Ito integrals.

Well, in defining Ito integrals, we have this notation that we need to be comfortable with. We are going to, first of all, consider a probability process that is defined by this triple. Now, let's see. Have any people here taken advanced probability or measure theory?

Well, we basically have, here, a sample space. And here, we have what's called a sigma field of sets, which are a subset of omega. And then here, we have a probability model, where if we have A, a subset of omega, and it's an element of this sigma field, we have this probability of that set is well defined. So this notation encompasses the generality of probability models that are important to specify and be comfortable with.

Well, with the omega set here, these are a collection of little omegas, which are possible sample paths of this Brownian motion. And we can think of, OK, here's one path, one possible path. Another possible path actually matches the first path up to time t, but then continues on somewhere else.

Another path could be one that is reflected about a given level. We went through this a bit in our discussion of Brownian motion processes. And so we can think of just the collection of all possible paths.

Now, the filtration, Ft-- we think of script Ft, this is going to be the sigma field of events or paths, which are fixed or known at times s less than or equal to t. So with script Ft, the information up to time t-- if this is our path up to time t, the event of following this path up to time t is deterministic if we look back-- deterministic for s less than t-- but then it's random-- random-- for s greater than t.

So we have a value-- or in script Ft, it basically says, OK, we can work with paths that are deterministic before time t but then are random after time t. And so that relates to the concept of being measurable. So we say that events in Ft are measurable with respect to Bs, with the information set up to time t. So if we consider times before t, the path is actually just a constant-- or as a function, a path-valued function that's constant.

Now, if we have this probability measure on omega, the space of all possible paths-- if we consider these a conditional distribution given X at time t, then we're trying to specify the model for all future paths that start at the level X at time t.

And that probability measure over the entire space of paths can also be restricted to the paths up to time little t. So we can basically focus on probability measures that are unconditional but restricted to the information set up to time t.

Now, we'll want to refer to the value of the Brownian motion at time t. And it's important to know that often this notation, Bt, is implicitly the value of a specific path. So to emphasize that, we'll indicate that it's the Brownian motion path indexed by little omega.

And so we then want to consider the Ito integrals with respect to Brownian motion, first where we consider simple functions, real functions, where f of Bs is the value of little f evaluated at the position value of the Brownian motion. And we'll want to consider integrating that function from 0 to T. And so this is our notation.

So with a companion graph, with the same x-axis as our Brownian motion, we could have f being some function which varies. And we want to basically integrate this function f evaluated at the Brownian motion.

So actually, this is actually a wrong graph. I thought we had had the same graph. We basically want for, let's say, x equals Bs of omega, some function. All right.

Well, if we consider this notation to represent our Ito integral, we're going to develop the formulas for these Ito integrals. And importantly, this Ito integral is going to have, basically, a deterministic, simple, real integration term for each path.

But the paths indexed by omega are random. And so this Ito integral is actually going to be a random variable. And as a random variable, if it exists, it'll have an expectation and a variance.

So let's consider a number of special cases that will define Ito integrals. The first case is the simple one, where f of x is equal to 1. So let me just erase this and have the level 1. So we consider f of x equaling 1.

So if we integrate the increments of Brownian motion over from 0 to capital T, let's propose that that's equal to the increment from 0 to-- from 0 to t for each path. So the right-hand side corresponds to the simple integral of this path, Bs of omega, between Bt-- between time 0 and capital T.

And, I guess, it's not in the notes. But if we were to consider the integral of 0 to t dBs of omega, we could think of this as the sum from 0 to n of iT over n. Let's see. Let me do this, if i equals 1 to n, i minus 1 T over n-- the integral from i minus 1 T over n to i times T over n dBs.

So here, I'm thinking of basically dividing this increment from 0 to T into different increments of width capital T over n. And if we approximate-- or each of these integrals here are basically looking at the integral of or the change in Bs from these lower limits to the upper limits. And so the integral of dBs of omega, in a consistent way or well-defined way, is defined by this increment of the whole process.

Now, if that's true, then when we consider the randomness in omega, the expectation of this integral is 0, where B is the standard Brownian motion. And its variance is equal to the length of the interval times the sigma of standard Brownian motion.

Well, if we change our function f to be not the constant 1 over the whole real line but just the indicator of an interval, there's just one between a and b. So we have a and b here. And we have basically 0 up to a. And then it equals 1 up to b, and then it's 0 beyond. Then this Ito integral is essentially looking at the change in Brownian motion from time a to time b.

So no surprise that that's a reasonable definition. Now, from these reasonable definitions, we can consider a second case, where our function f is, again, a function only of time s, but it's given by a step function over time, with steps-- or levels given by the ai's that are constant.

So we go from this single-step function here to considering. basically, different step functions for our function f. Now, these step functions could correspond to, like, the number of shares held of an asset. And so this integral could correspond to the gain or loss if the underlying asset followed Brownian motion.

But here, with this particular step function, we can define the Ito integral to just be the sum over the increments of the step levels times the increment of the process. And so, that's a pretty sensible definition of the integral. And it would apply if we had assets traveling-- values going by a Brownian motion process and, say, shares given by the a constants that were maybe deterministic.

Now, at this point, we're using functions f, which are deterministic functions of time. We can then extend this to random functions. And actually, before I do so, just notice in red, the mean 0 and the variance of this Ito integral of the step function-- the variance is simply the sum of the squares of the levels times the length of the increment over which that level applies.

So now, case 3, things get more interesting. We can allow our levels, ai, to be random variables. And these variables, random variables ai-- so ai equal to ai omega, this is a path. And so we can have the levels vary depending upon what the path value is.

And so we want the ai to depend on-- and, I guess, by depend, I'll say be deterministic-- on script Ft sub i. So if we define our levels to be functions that are measurable in this information set, then that's what we're doing here.

Now earlier, we talked about nonanticipatory variables. And so this is a nonanticipatory variable. It basically says that the level from ti on, for whatever that increment is, is going to be a function of what's happened in the past. So this is like a realizable function f, where when we apply the f function to the Brownian motion increments, we're applying known increments at all times.

And so with this case, we actually want to make an assumption of measurability on that i-th information set. And we want a bounded squared expectation of the level. So if we have those two properties then, we can write out the Ito integral as this formula. And it will have expectation equal to 0 and variance equal to the sum of the expected squared ai's times ti plus 1 minus ti.

So, I guess, rather importantly with these developments so far, we're thinking of the time s ranging from 0 up to time capital T. And we have increments ti, ti plus 1, where this increment is something. For simplicity, we could define this increment to have length, say, t over n, for n increments and have equal increments.

But this formula for the Ito integral is considering integrating the Brownian motion over the interval from 0 to T. So integral from 0 to T, f of s, Bs dBs. This is our integral, our Ito process, I of f. And we can write I of f evaluated at a particular path given by this formula.

All right. Well, an interesting example just from a mathematician's standpoint is, what if we define our levels to be equal to the value of the Brownian motion itself?

So we're integrating B, dB. So we're thinking of this example, so case 3-- special case 3-- integral from 0 to T of B sub s dBs. Well, if we consider discretizing this Ito integral with these random levels ai-- actually, I'm getting ahead of myself. This integral will be a limit of these.

For any fixed set of increments that we apply this random level. given by the value of the Brownian motion on those increments, then for each of these cases with n increments, we get that we have mean 0. And the variance is equal to the sum of the variances of each of the terms. And that turns out to be equal to the sum over the increments of the time Increment t, level ti, times the increment to the next time value.

And this sum of ti times ti plus 1 minus ti is essentially equal or close to the integral from 0 to capital T of sds, which should be 1/2 t squared. Yeah.

AUDIENCE: How do you conclude that the expectation is 0?

PETER KEMPTHORNE: OK. The expectation is 0 because each of these increments here have expectation 0. So we're assuming B of 0 is 0.

AUDIENCE: Yeah. Oh, yeah. But what if ai is not independent? Because when you take the expectations [INAUDIBLE].

PETER KEMPTHORNE: OK. Yep. Yep. Let's see. OK, ai omega is an element of F sub ti. So it's constant on F sub ti. So let's see. So ai of omega and Bti plus 1.

This increment into the future from time ti is independent of the past. And so its expectation, if we take the product of these two and take expectations, we can take the expectation given-- or the expectation over ti times the expectation given-- let's see-- over ti plus 1 given ti.

So we can break down this expectation into its decomposition. And the expectation of this factor, conditionally, will be 0.

AUDIENCE: Yeah, that makes sense.

PETER KEMPTHORNE: So, yep. All right. Well, this case 4 is the case that essentially defines Ito integrals in general, which is, if we consider taking the limit of integrals with respect to step functions with random levels ai, but then allow the increments-- the time steps to get closer and closer to each other, then the limiting value of this Ito integral of the step functions is equal to the Ito integral of f.

So in our special case where f of x equals x, f of the Brownian motion is the Brownian motion itself, we end up getting this formula. And here, we have basically the discretized value of the integral. And then what's of interest is to see how the limit of that equals this formula.

And so to prove this, there's a nice mathematical trick. And so if we look at this term on the right-hand side here, which is an element of the sum, it turns out to be equal to 1/2 the difference in the squared values at the endpoints of the interval minus 1/2 times the difference between ti and ti plus 1 squared.

So this is simple to verify. It's trivial but incredibly useful. And so if we sum over all the time points from 0 to n minus 1, we get the sum of the first terms minus the sum of the second terms.

And so this first term is actually a telescoping sum. When you sum the terms, there's cancelation of the middle value in successive terms. And so we're left at the end with 1/2 the square of the last term here minus the square of the very first term, with i equals 0 here, which is 0.

So this is 1/2 Bt squared minus B0 squared, which is 0. That's the first term. And then for the other, we have minus 1/2 the sum of the squared increments.

And so if we know what the last time interval term is, tn, that's capital T. And then we have the same terms on the right. And what's important to know is that the limit of the sums of these increments converges to T-- capital T.

And this was the property of the total variation of Brownian motion over any increment. If we look at the sum of squared values of increments and take the limit as the number of increments grows large and the partition gets finer and finer, then the limiting value of these sums of increments is equal to capital T. So this is our formula now for the Ito integral of B dBs.

And let's see. So what's really important to highlight is that this is not equivalent to thinking of the Brownian motion process like a real function, where we're integrating BdB or xdx, which is x squared over 2. We actually have this drag on the integral that depends on the variation of the process.

OK. Now, we want to generalize to the general case of Ito integrals. And so what we do is we consider defining the Ito integrals from 0 to T f of omega s dBs. But then we want to basically consider this to be a process Xt of omega. And we want to generalize this to an Xt of omega, where we're integrating just up to time little t.

So what we want to do is, once we have this definition for the Ito integral over a fixed interval from 0 to T, we want to be able to define the Ito integral where we're integrating only up to time little t less than capital T.

And so in order for this kind of function to be well defined, it will turn out that we need some very strong mathematical arguments. And there's a book by Steele, chapter 6, on quantitative finance, where he does go into all the details.

And I have to admit, it's not really a fun read. It's interesting to look and see what's involved so you know how important it is to be a good mathematician, if you want to prove things. But fortunate for us, they're proven well, and we can apply them.

So what do we need? We need that the function little f needs to be measurable, script Fs, the information set up to time s measurable for every time s. And we want the integral of the squared f over time, where that's for each fixed omega, we can integrate the square of f and then take the expectation of that. We want that to be finite.

So this term here, this expectation, is taking the expectation over the probability model given by capital P. And for each omega, we're integrating the squared f over time from 0 to T. And we want that to be finite.

So with this notation, we basically are extending our Ito integral over 0 to capital T, with the Ito integral from 0 to little t. And we can define this by just replacing our f function in the Ito integral I of f by simply the function f times this indicator from 0 to T. So that step basically chops off the integration at time little t and gives us our definition.

Now, let's see. OK. The definition of this extension of our Ito integral over 0 to T to integrating it over 0 to little t's for any little t, less than t, that's relatively straightforward and intuitively obvious. What's not necessarily obvious, though, is that our process Xt is a martingale with respect to this stochastic process model given by the Brownian motions.

So if we look at what it means for a process to be a martingale, if we look at the conditional expectation of the time t plus k value given time t, then it equals the value of the process at time t, little t, plus the integral of the process from time t to t plus k. And we're taking expectations, conditional expectations of that given Xt.

And so the expectation of this future increment from t to t plus k is basically integrating a function times increments of the Brownian motion, where our function values depend on the past, and the increments are independent in the future, with mean 0. So this expectation turns out to be satisfied.

And if we look at the conditional variance of the t plus k value given Xt, then it's equal to the expected square of this f ds. And so when you look at the variance of the discretized sum of increments, you get the discretized sum of variances of those increments, which are f squared. And then the dBs squareds are dt. So this simplifies for the variance to that term.

OK. All of this is fairly-- I know, it's like mathematical results. And it's sort of like, who cares? Well, it turns out that this variance turns out to be a squared norm of our function little f on the space of the probability space of paths cross the time increments.

And so what's-- well, OK. Actually, I'm going to get ahead of myself. But let me get ahead of myself and basically comment that this norm of our function f is going to be a squared norm on the space of Ito-process integrands.

And, let's see. Those of you who have taken, I guess, real analysis or metric spaces, what's an important value of norms? What do we do with norms? Or squared norms?

Well, with norms, we can talk about limits of elements of the space and talk about whether limits exist. And so with this setup, we actually will have the space of Ito processes. And we can calculate the squared norm between different Ito integrals.

And so we can characterize limits in those spaces and be guaranteed that those limits exist. So with this motivation of defining useful norms on different spaces, we have this Ito isometry, which is, if we look at the variance of Ito integrals defined by integrand little f, then those squared norms or the variances, basically, of the Ito processes is equal to the L2 norm of the integrand little f.

And if we write out the definitions of these squared norms, we get these equations. And what's relevant here is being able to express our variance of the Ito integral as the squared norm of the integrand f.

So we have limits and convergences of sequences of Ito integrals. And something that's important mathematically is, if we take the limits of these Ito integrals over increments that are getting-- a grid of increments that's getting finer and finer, do we have a unique limit or not? And so that's what's important. And the squared norms and convergence properties are used to prove that.

So we basically can talk about closed linear subspaces with these norms on the space. So what this does is it allows us to define the distance between different integrands f and the existence of a limiting Ito integral, as we consider increments getting infinitesimal.

All right. Well, we can think of this Ito integral as a Riemann sum, a limiting Riemann sum. And if we have time steps ti, for simplicity, equal to iT over n, we can consider the limiting sum of f of the Brownian motion times the increment of the Brownian motion. And this sum actually converges in probability, as n goes to infinity, to the integral.

So we have a limiting value that exists. Convergence is in probability. And if we consider, I guess, the covariances of the X process given by these Ito integrals from 0 to little t, then these will-- or this Xt will be a mean 0 Gaussian process with independent increments. And the covariance between the process at time s and time t will simply be equal to the integral of the squared f over the times where they both are non-zero.

All right. So we take a partition from 0 to t. We get this result. And we can prove that the limit equals this integral notation. So basically, there's a limit that exists, and this is a notation for that limit. And Steele goes into a lot of the details for that.

Now, what's of interest is that we can now consider some general Ito cases, which will be referred to as Ito's formula, for different kinds of functions f. And so if we consider first the case where f, little f, is a real function mapping reals to reals, then if it has a continuous second derivative, then we can express the function f of Bt in terms of its Taylor approximation, and we get this formula, which is Ito's formula for f of Bt.

And in looking at this expression, what's relevant is that, OK, the first term, f0, is a constant. The second term is a Gaussian process integral. It's an integral of Brownian motion. And then the third term is the integral of the second derivative of f evaluated at the Bs value. And so we have an adaptive drift term.

And we can prove this formula, or outline the proof by saying, OK, let's consider the Taylor approximation of f at x plus delta x and expand it to second-order terms plus a remainder. And then if we take the difference of f at time t with B minus f at time 0 with B, then we're looking at the sum of the increments from time t1 to tn.

And those increments are equal to the sums of the derivative of f times the increment of time plus 1/2 the second derivative of f times the squared increment plus the remainder terms. And if we take the limits of those, we get the integral of f prime dBs plus 1/2 the integral of the second derivative.

So Taylor series gives us this formula. And we can think of developing a notation with Ito calculus, which is that our differential of f at Bs is equal to f prime of Bs dBs plus 1/2 f double prime times Bs times the squared of dB, which is equal to ds, the increment of time.

So this is our Ito's formula shorthand. And this shorthand corresponds in a similar way to derivatives and antiderivatives, where here's the stochastic differential equation, and here's the solution to that stochastic differential equation on the top.

Now, what's really neat and very powerful is that if we want to solve for the integral of a function-- so suppose our objective is to solve for the integral from 0 to T f of Bs dBs. Well, if we have the antiderivative of little f-- so that means that the first derivative of capital F is equal to such that F prime is equal to f-- then, if we choose the antiderivative with f of 0 equal to 0, then we get F of Bt minus F of B0 is equal to this integral of the derivative of f dBs plus 1/2 the integral of the second derivative of capital F, which is simply f prime.

So this formula is what we just derived for Ito's formula. And if we solve for this first term on the right, then that's equal to F of Bt minus 1/2 the integral of the first derivative of f prime of Bs ds.

So what's neat as a technique in stochastic calculus is that if we want to integrate a function little f of Bs dBs, then we can just apply Ito's formula in reverse and characterize it as the antiderivative of little f at Bt minus 1/2 that, defining the antiderivative to have f of 0 equal to 0.

Now, let's apply this just for the case where capital F of x equals x, capital F prime, the derivative is 1, and the second derivative is 0, then the integral of dBs is equal to Bt. And let's see. So basically, the integral of Bs dBs is equal to capital F of Bt, which is Bt, and then minus 1/2 the second derivative-- or the second derivative of f, which is 0. So this gives us this result.

And if we consider capital F equaling 1/2 x squared, then little f is equal to x, and little f prime is equal to 1. And so we get the integral of Bs dBs is f of Bt minus that, which is 1/2 Bt squared minus 1/2 t. So this gives us of a verification of the limiting results we went through before with these Ito integrals.

The next step is to extend Ito's formula to functions of more than one variable. So we basically want to go from f of x, so one variable s. We want to go to two variables functions, f. And we'll write the notation of s comma x, which maps, basically, let's see, the time space cross-- or time range, R plus, cross the real line of x into the reals. And we're going to consider functions little f that are continuous to different orders in terms of the time variable and the space variable.

So we have this notation, C super 1 comma 2, meaning that it's continuous with a first-order derivative, existing and continuous and second-order derivative with respect to x and continuous. And so we have a Taylor series for increments of this function little f, now, of two variables, which has the multidimensional d equals 2 Taylor series formulation.

And it equals the f without the increments on t and x plus delta t times the partial of f with respect to t plus delta x times the partial of f with respect to x. And then we have the second-order partial differential terms. And the only ones that are going to be of interest are the second-order partial of f with respect to x, which will be multiplied by delta x squared. The other second-order terms will actually be of small order delta t.

And with this second-order Taylor approximation, we're able to derive that the function f of t, Bt is equal to the initial value plus an integral over time of a function of the Brownian motion ds plus an integral over increments of the Brownian motion plus 1/2 this integral of the second partial of f with respect to dx squared. And this integral here ends up being multiplied-- or being integrated with respect to the increments of time, because the delta x-squared terms are infinitesimally equivalent to increments ds.

So for this formula, we have here, if we look at the difference in the little f function from time 0 to T, it equals the right-hand side. So this top formula here is simply taking this second-order Taylor series and looking at the increment from 0 to T and what that equals.

And what we get is a sum of two terms, an integral with respect to time of the partial with respect to time plus 1/2 the second derivative of f with respect to x. And with this notation for the value of little f, what's important to notice is that, if the argument in the time integral-- if that's equal to 0, then this will be a martingale because the second term is an integral of terms times increments of the Brownian motion. And the increments of the Brownian motion all have expectation equal to 0.

So we'll have the possibility of this function being a martingale, if that's true. And this integrand equaling 0 is equivalent to this partial differential equation on the little f. So there's this connection between martingales and partial differential equations for f when we're thinking of a process f of t.

So if little f is a martingale, it will turn out that the partial of f with respect to t will equal minus 1/2 d squared f dx squared, which is a PDE condition.

Now the other thing we want actually for this to exist is that the variance of the function, f of T, BT, which is bounded by the squared expectation, is finite. And so this will be finite, the variance, if that second term has finite variance. And so for that to have finite variance, the variance is the integral of the square of that argument df dx, with respect to time.

Now, what's interesting with this Ito formula is we can apply it to a function little f, which is the exponential of sigma x minus sigma squared t over 2. And importantly-- well, let's see. This partial of f with respect to t does equal minus 1/2 sigma squared f. And so this is actually a partial differential equation for little f.

But df dx is equal to sigma f. The second partial of f with respect to x is sigma-squared f. And Ito's formula, if we expand it out, gives us this expression for little f. And this function little f is, in fact, a martingale. If we look at the expected value of f given f of 0, 0, it equals 1.

Now, f of 0, 0 is the exponential of, basically, 0, which is 1. But this function f, little f, if we plug in Bt for x, it's a scaled lognormal random variable. So this little f evaluated at x equal to Bt-- so it's a transformation of the Brownian motion process paths, Bt by this function-- is actually a scaled lognormal random variable.

And so e to the sigma Bt minus 1/2 sigma-squared t is a scaled lognormal. And the expected value of this e to the sigma Bt is actually equal to e to the 1/2 sigma-squared t. We proved that, I believe, in the problem set. And so the expectation of f of t, Bt is going to actually equal 1 for all t. So this little f is, in fact, a martingale.

Now, let's see. With Brownian motion with drift, the ruin problem, we can actually use some Ito-process ideas to deal with this problem, where we assume that Xt is a Brownian motion with drift mu and variance parameter sigma squared. So here's its formula, stochastic differential equation for Xt.

And we're interested in, what's the probability of stopping at A before B, where the stopping time tau is the first time t that hits A or minus B. And this argument is one that's very typical of quantitative finance developments, I guess, by mathematicians. And it's actually a very neat argument, where I don't think I could ever have derived this, but the derivation using martingale theory is really very elegant.

And so let's consider trying to solve for this probability by using martingale theory. So let's consider a martingale. Let's try to define a martingale with value h of Xt defining Mt, where the tau is the same stopping time, and the condition on h is going to be that if we are at the value Xt equals A, then the probability is 1. And if we start at minus B, the probability is 0. So this is fixing the terminal conditions of the function h, when it hits the limits A or minus B.

Well, if we have a martingale that satisfies this function h of Xt, then the expectation of the stopped martingale is simply equal to the probability that X tau equals A. And by the martingale property, the expectation of the stopped time is equal to M0, the initial value of the h function.

Now, the initial value is the value when X is equal to 0. So we're starting at 0 with the X process, and it either moves and hits A or hits minus B.

So in order for this h function to be a martingale, we have to have the PDE condition satisfied on h. So the partial of h with respect to t is equal to minus 1/2 the second partial of h with respect to dx.

And so if this h function, which is the probability of hitting A first, depends just on the value of Xt, then it equals h of mu t plus sigma Bt, substituting in for what Xt is. So we can work with the function little f equal to h of mu t plus sigma x.

And if the partial of f with respect to t, which is, by chain rule, mu h prime of mu t plus sigma x-- if that equals the second partial of f with respect to x, the second partial derivative of h with respect to x is sigma squared h double prime, then this PDE condition says we have this relationship, this differential equation, for h. Although, it's in the first derivative of h. But we can then rewrite this as the second derivative of h evaluated at argument y is minus a constant times the first derivative of h.

And so we can solve that equation by noting that h double prime over h prime is equal to d by dx of log of h prime. So the derivative of the log of h prime is a constant. The solution to that is a linear function. And so we then end up being able to solve for little h of y equaling this formula. So this is the solution to this problem.

All right. Let's see. We're at the end of our lecture today. So next time, we'll finish up these notes talking about standard processes, where our a function for the drift and our b function for the volatilities have to satisfy certain properties.

And if you read ahead in these notes, you'll see that we need the probability of integrating the drift over time has to be finite and absolute integral. And the integral of the squared b's have to also be finite. If those were not finite, then you would, say, have explosive behavior of these integrals, where they won't exist.

So these are reasonable properties to assume for those. And anyway, we'll go through some interesting results here with this particular problem and then turn on to stochastic differential equations next time. All right. Thank you.