Part III

Basic Stochastic Calculus

Michael Boyuan Zhu

UNIVERSITY OF WATERLOO

Fall 2023

Probability Theory

Probability Theory - The Basics

- Financial models in continuous time are modelled using tools from stochastic calculus. This is a branch of mathematics that operates on stochastic processes.
    - Calculus is concerned with analysing functions. Stochastic calculus analyses functions that can be random (stochastic processes).
- It would be too time-consuming to formally develop this theory. Instead, we will rely on some rules to calculate what we want.
- The main topics we want to introduce are (filtered) probability spaces, conditional expectation, and martingales.

Probability Theory - Probability Spaces

- We have seen from the discussion on general market models that random variables are (measurable) functions on a state space Ω. This is in line with the official definition of a probability space from probability theory:

Definition (Probability Space)

A probability space is a triple $(\Omega, \mathcal{F}, \mathbb{P})$, where:

- Ω is a state space (i.e. a set representing future states of the world),
- $\mathcal{F}$ is a σ-algebra on Ω, which represents the amount of information available to us,
- $\mathbb{P}$ is a probability measure.

- We will use the letter X to denote a random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Usually, X will represent the price of a financial instrument.

Probability Theory - Stochastic Processes

Definition (Stochastic Process)

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Then a (continuous) stochastic process is a collection

$$\{X_t : t \in [0, +\infty)\}$$

of random variables on $(\Omega, \mathcal{F}, \mathbb{P})$. We will also use the notation $\{X_t\}_{t\geq 0}$ to denote a stochastic process.

- For a specific state $\omega \in \Omega$, the realization of the stochastic process is

$$\{X_t(\omega)\}_{t\geq 0} \,.$$

- This is a function from $[0, \infty)$ to $\mathbb{R}$, defined by $t \mapsto X_t(\omega)$. This is called the sample path of the process at ω.

Probability Theory - Stochastic Processes

- In other words, each different state of the world $\omega \in \Omega$ produces a different path of the price of the instrument X. A sample path of a stochastic process is given below:

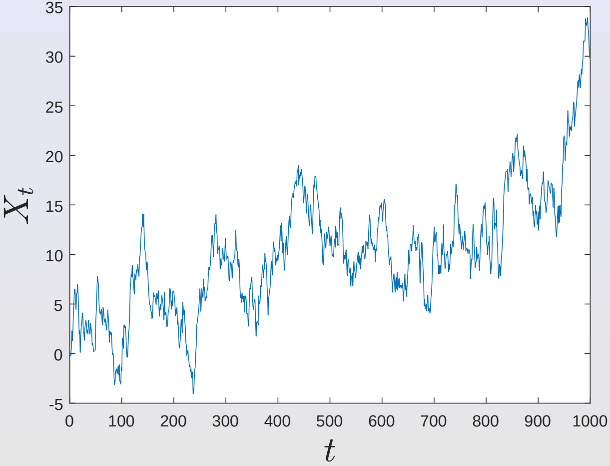

Probability Theory - Stochastic Processes

- Some more sample paths (representing different states of the world ω) are shown below:

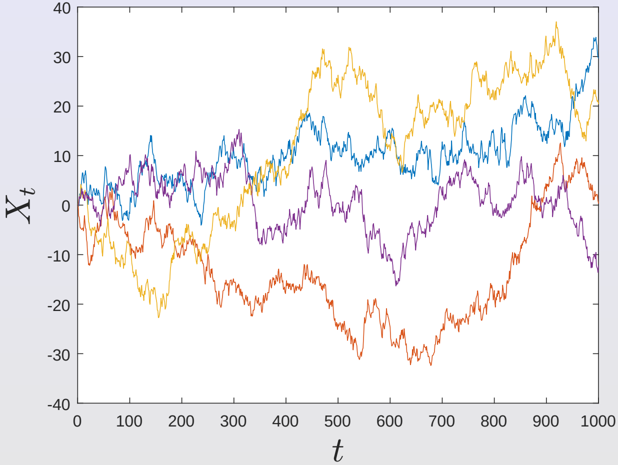

Probability Theory - Stochastic Processes

- We will assume the following on stochastic processes:
- Each sample path is continuous. That is, there are no jumps in sample paths.
- X_t is a continuous random variable for all t.
  - In particular, this means that the state space Ω is a continuum. This is different from the models we saw in Part II, where Ω had finitely many states.

Probability Theory - Information and Filtration

- We have mentioned before that the σ-algebra component of a probability space represents information.
- Suppose $\{X_t\}_{t\geq 0}$ represents the price of some instrument.
- As time passes, we can observe the value of X_t and accumulate more "information". At each time t, define the information set $\mathcal{F}_t$ to represent the information available at time t.

Definition (Filtration)

A filtration is a series of information sets $\{\mathcal{F}_t\}_{t\geq 0}$ such that if $s \leq t$, then

$$\mathcal{F}_s \subseteq \mathcal{F}_t \, .$$

- This condition essentially says that the amount of information available to us becomes larger as time goes on.

Probability Theory - Information and Filtration (extra)

- Some more precise definitions are given in the following for those familiar with measure theory. This material is optional.

Definition (σ-algebra)

A σ-algebra or σ-field $\mathcal{F}$ is a collection of subsets of Ω satisfying:

- $\varnothing \in \mathcal{F}$.
- $\mathcal{F}$ is closed under complements.
- $\mathcal{F}$ is closed under countable unions.

Elements of $\mathcal{F}$ are called events.

- Information sets are σ-algebras. Specifically, $\mathcal{F}_t$ represents all the events that we know have happened or not, at time t.
- A filtration is an increasing collection of σ-algebras.

Probability Theory - Information and Filtration

- We say a random variable X is measurable with respect to $\mathcal{F}$ if information contained in $\mathcal{F}$ can tell us what the realized value of X is.
- We will generally consider filtrations $\{\mathcal{F}_t\}_{t\geq 0}$ that contain the information of the evolution of a price X_t. Then at time t, we would know what the value of X_t is. Therefore, we would expect X_t to be measurable with respect to $\mathcal{F}_t$.

Definition (Adapted Process)

The process $\{X_t\}_{t\geq 0}$ is adapted to the filtration $\{\mathcal{F}_t\}_{t\geq 0}$ if X_t is measurable with respect to $\mathcal{F}_t$.

Probability Theory - Information and Filtration (extra)

- Some more official definitions are as follows:

Definition (Measurability)

A function $X : \Omega \to \mathbb{R}$ is measurable with respect to a σ-algebra $\mathcal{F}$ if for all intervals $I \in \mathbb{R}$, we have

$$X^{-1}(I) \in \mathcal{F},$$

where $X^{-1}(I)$ denotes the pre-image of the set I under X. We say that X is a random variable (on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$) if it is measurable (with respect to $\mathcal{F}$).

- This means that given the information set $\mathcal{F}$, we know whether or not the realized value of X is in the interval I.
- This implies that we would know what the value of X is.

Probability Theory - Information and Filtration (extra)

- We will work with filtered probability spaces in this course.

Definition (Filtered Probability Space)

[Probability Space] A filtered probability space is a quadruple $(\Omega, \{\mathcal{F}_t\}_{t\geq 0}, \mathcal{F}, \mathbb{P})$, where $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space and $\{\mathcal{F}_t\}_{t\geq 0}$ is a filtration. We will also assume that

$$\mathcal{F} = \mathcal{F}_T \, ,$$

for some terminal time T.

- Suppose that X is a random variable on $(\Omega, \mathcal{F}, \mathbb{P})$. That means that it is $\mathcal{F}$-measurable. This suggests that the value of X must be known at time T. Recall that this is true for all contingent claims we saw so far in the course.

Conditional Expectation and Martingales

Conditional Expectation

- Recall that the expectation of a random variable X given another random variable Y is $\mathbb{E}[X|Y]$, which is itself a random variable. We will now introduce the corresponding notion for information sets.

Definition (Conditional Expectation)

Let X be a random variable on $(\Omega, \mathcal{F}, \mathbb{P})$ and let $\mathcal{G} \subseteq \mathcal{F}$ be an information set. Then the conditional expectation of X given the information set $\mathcal{G}$ is denoted by

$$\mathbb{E}[X|\mathcal{G}],$$

which is itself a random variable.

- In essence, $\mathbb{E}[X|\mathcal{G}]$ is our best guess of the value of X given the information that we have ($\mathcal{G}$). Generally, we will consider conditional expectations of the form $\mathbb{E}[X|\mathcal{F}_t]$.

Conditional Expectation (extra)

- The official definition of condition expectation is fairly complicated to prove and interpret. In any case, it is given by the following:

Theorem (Kolmogorov)

Suppose $X \in L^1(\Omega, \mathcal{F}, \mathbb{P})$ *and* $\mathcal{G} \subseteq \mathcal{F}$. *Then there exists a unique,* $\mathcal{G}$*-measurable random variable* $Z \in L^1(\Omega, \mathcal{G}, \mathbb{P})$ *such that for all events* $A \in \mathcal{G}$,

$$\int_A Z \, d\mathbb{P} = \int_A X \, d\mathbb{P} \, .$$

We define $\mathbb{E}[X|\mathcal{G}] = Z$.

Conditional Expectation

The conditional expectation $\mathbb{E}[X|\mathcal{G}]$ exhibits the following:

- $\mathbb{E}[X|\mathcal{G}]$ is $\mathcal{G}$-measurable. If we know the information $\mathcal{G}$, then we should know what the value of our best guess is. Alternatively, we do not need to know anything more than $\mathcal{G}$ in order to give our best guess given $\mathcal{G}$.
- The expectation $\mathbb{E}[X]$ can be interpreted as a conditional expectation with no information given. Formally, $\mathbb{E}[X] = \mathbb{E}[X|\{\varnothing, \Omega\}]$.
- The conditional expectation is linear. That is, $\mathbb{E}[aX + Y|\mathcal{G}] = a\mathbb{E}[X|\mathcal{G}] + \mathbb{E}[Y|\mathcal{G}]$.

Conditional Expectation

The conditional expectation $\mathbb{E}[X|\mathcal{G}]$ exhibits the following:

- The tower rule holds: if $\mathcal{H} \subseteq \mathcal{G}$, then $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{H}] = \mathbb{E}[X|\mathcal{H}]$. This also implies that $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]] = \mathbb{E}[X]$.
- If Y is $\mathcal{G}$-measurable, then $\mathbb{E}[XY|\mathcal{G}] = Y\mathbb{E}[X|\mathcal{G}]$. If we know $\mathcal{G}$, then we know the value of Y, so we can take it out of the expectation as if it were a constant.
- If X is independent of $\mathcal{G}$, then $\mathbb{E}[X|\mathcal{G}] = \mathbb{E}[X]$. This means that $\mathcal{G}$ does not give us any information about X at all.

Martingales

Definition (Martingale)

An adapted process $\{X_t\}_{t\geq 0}$ on a filtered probability space $(\Omega, \{\mathcal{F}_t\}_{t\geq 0}, \mathcal{F}, \mathbb{P})$ is a martingale if

1. $\mathbb{E}[|X_t|] < \infty$ for all t (this is a technical condition),
2. $\mathbb{E}[X_t|\mathcal{F}_s] = X_s$, for all $s, t \geq 0$ such that $s < t$.

- Essentially, a process is a martingale if the best guess of its future value is its current value.
- This also means that the directions of future movements are impossible to forecast. Therefore trajectory must not exhibit any discernible trends or periodicities.

Martingales

Example

We can show that if $\{X_t\}_{t\geq 0}$ is a martingale, then for $s < t$,

$$\mathbb{E}[X_t - X_s|\mathcal{F}_s] = 0\,.$$

This means that the expected change over the interval $[s, t]$, even when given all the information at s, is zero. Hence, a martingale can be understood as the mathematical formalization of a fair game.

$$\begin{aligned}\mathbb{E}[X_t - X_s|\mathcal{F}_s] &= \mathbb{E}[X_t|\mathcal{F}_s] - \mathbb{E}[X_s|\mathcal{F}_s] \\ &= X_s - X_s \\ &= 0\,.\end{aligned}$$

Martingales

- The stochastic process we saw earlier is a martingale. Does it look like there are any discernible trends?

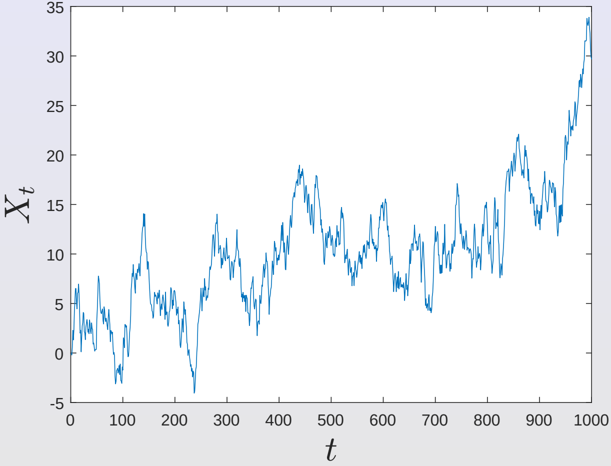

Brownian Motion

Brownian Motion

- The sample paths we saw before are realizations of a very special stochastic process: the one-dimensional Brownian motion.
- Brownian motion is the basis of the most commonly used models of asset prices.
- We will see later that stocks are modeled, roughly speaking, as a function of Brownian motion:

Brownian Motion

Definition (Brownian motion)

A stochastic process $\{W_t\}_{t \geq 0}$ is a (standard, one-dimensional) Brownian motion if it has the following properties.

- $W_0 = 0$.
- Each sample path $t \mapsto W_t(\omega)$ is continuous.
- For each $s < t$, we have $W_t - W_s \sim \mathcal{N}(0, t - s)$.
- It has independent increments. That is, for $0 = t_0 < t_1 < \cdots < t_m$,

$$W_{t_1} - W_{t_0}, W_{t_2} - W_{t_1}, \ldots, W_{t_m} - W_{t_{m-1}}$$

  are independent.

Brownian Motion - A Brief History

1827 BM originates in the work of Scottish botanist **Robert Brown**, who studies the irregular movement of pollen in water.

1900 **Louis Bachelier** introduces it to finance. For this reason, Bachelier is considered by many as the founder of modern mathematical finance.

1905 **Albert Einstein** introduces Brownian motion to physics, which eventually leads to conclusive proof of atomic theory.

1923 **Norbert Wiener** gives the first mathematically rigorous proof of the existence of BM. This is why Brownian motion is sometimes called *Wiener process* and denoted by W.

1934 **Paul Lévy** extends the notion of BM to more general processes.

1940 French army soldier **Vincent Doeblin** develops stochastic calculus: new calculus for stochastic processes.

1942 Independently, **Kiyoshi Itô** develops stochastic calculus.

1955 **Paul Samuelson** introduces the geometric Brownian motion (GBM) as the underlying model for stock prices.

1973 **Fisher Black**, **Myron Scholes** and **Robert Merton** use the GBM for stock prices and come up with the famous option pricing formulas.

Brownian Motion

- Proving that Brownian motion actually exists is a difficult task, and beyond the scope of this course.
- Additionally, we will generally assume that there is a filtration $\{\mathcal{F}_t\}_{t \geq 0}$ for which the Brownian motion $\{W_t\}_{t \geq 0}$ is an adapted process.
- We will also assume that for $s < t$ that the increment $W_t - W_s$ is independent of $\mathcal{F}_s$. In other words, future increments are independent of available information.

Brownian Motion

A Brownian motion $\{W_t\}_{t\geq 0}$ exhibits the following properties:

- $\mathbb{E}[W_t] = 0$ and $\mathrm{Var}(W_t) = t$.
- $\{W_t\}_{t\geq 0}$ is a martingale.
- $\mathrm{Var}(W_t|\mathcal{F}_s) = t - s$ for $s < t$.
- $\mathrm{Cov}(W_t, W_s) = \min\{s, t\}$.
- Sample paths are continuous but nowhere differentiable.

Brownian Motion

Example

When performing calculations involving Brownian motion, it will often be useful to try and express things in terms of increments. For example, consider the following proof that Brownian motion is a martingale:

$$
\begin{aligned}
\mathbb{E}[W_t|\mathcal{F}_s] &= \mathbb{E}[W_t - W_s + W_s|\mathcal{F}_s] \\
&= \mathbb{E}[W_t - W_s|\mathcal{F}_s] + \mathbb{E}[W_s|\mathcal{F}_s] \\
&= \mathbb{E}[W_t - W_s] + W_s \\
&= 0 + W_s \\
&= W_s \, .
\end{aligned}
$$

The Ito Integral

Ito Integral - Motivation

- In mathematical finance, we often find that we need to integrate with respect to random increments.
- In particular, we will look at objects of this form:

$$\int_0^t f(u)\, dX_u \,,$$

where $\{X_t\}_{t\geq 0}$ is a stochastic process. We will often use the notation $f(u) = f_u$ to express functions of time.

- We say that the above expression is the Ito integral of the function f with respect to the process $\{X_t\}_{t\geq 0}$.
  - As a technical detail, we assume that there is a filtration $\{\mathcal{F}_t\}_{t\geq 0}$ such that X_t and $f(t)$ are both adapted processes. This essentially means that we know the values of X_t and $f(t)$ at time t.

Ito Integral - Motivation

- Recall the multiperiod binomial model with time periods $0, h, 2h, \ldots, T$, in which we introduced portfolio strategies $\theta = (\theta_0, \theta_h, \theta_{2h}, \ldots, \theta_{T-h})$. At time t, the portfolio $\theta_t = (\Delta_t, b_t)$ represents the composition of the portfolio for the time period $(t, t+h)$.
- The return of the position in the stock over $(t, t+h)$ is

$$\Delta_t(S_{t+h} - S_t)\,.$$

- Then the total return from the stock over the time period is (with imprecise notation)

$$\sum_t \Delta_t(S_{t+h} - S_t)\,.$$

- This calculation looks like the sum of the areas of a series of rectangles. The height of each rectangle is Δ_t and the width is $S_{t+h} - S_t$.

Ito Integral - Motivation

- In continuous time, we let the length of each period go to zero. Therefore the total return from the stock in a continuous model should look something like:

$$\lim_{h \to 0} \sum_t \Delta_t (S_{t+h} - S_t) \,.$$

- This is a limit of the sum of the areas of rectangles, as the width of each rectangle approaches zero. That is exactly what a Riemann(-Stieltjes) integral is.
- Essentially the only difference here is that S_t is random. The Ito integral can be understood as

$$\int_0^T f(u)\, dX_u = \lim_{h \to 0} \sum_t f(t)(X_{t+h} - X_t) \,.$$

- Note that this is not entirely rigorous, and serves mainly to give some intuition.

Ito Integral - Motivation

- A rigorous and formal construction of the Ito integral is beyond the scope of this course.
- We will assume that the Ito integral is well-defined, and that expressions that look like $\int_0^T f_u \, dX_u$ make sense.
- We will most often look at Ito integrals with respect to Brownian motion:

$$\int_0^T f(u) \, dW_u \, .$$

These Ito integrals are nice because they satisfy a lot of desirable properties.

Ito Integral - Properties

- For each time t and a given stochastic process $\{f_t\}_{t \geq 0}$, define

$$I_t := \int_0^t f_u \, dW_u$$

  to be the Ito integral from time 0 to t, of the process $\{f_t\}_{t \geq 0}$ with respect to a Brownian motion $\{W_t\}_{t \geq 0}$.

- Note that there is randomness in this process. Hence, the Ito integral is a random variable, and $\{I_t\}_{t \geq 0}$ is a stochastic process.

- Recall that we also assume the existence of an underlying filtration $\{\mathcal{F}_t\}_{t \geq 0}$. With respect to this filtration, $\{I_t\}_{t \geq 0}$ is an adapted process.

Ito Integral - Properties

The Ito integral exhibits the following:

- $\{I_t\}_{t\geq 0}$ has continuous sample paths.
- The Ito integral is linear. That is,

$$\int_0^t af_u + g_u \, dW_u = a\int_0^t f_u \, dW_u + \int_0^t g_u \, dW_u \, .$$

- The process $\{I_t\}_{t\geq 0}$ is a martingale. That is, if $s \leq t$,

$$\mathbb{E}[I_t | \mathcal{F}_s] = I_s \, .$$

- Interestingly, the converse statement is also true: every martingale can be written as an Ito integral with respect to Brownian motion. This result is called the Martingale Representation Theorem.

Ito Integral - Properties

Since I_t is a random variable, we can also say a few things about its distribution.

- $\mathbb{E}[I_t] = I_0 = 0$, since the Ito integral is a martingale.
- The variance can be calculated using the following result:

Proposition (Ito Isometry)

$$\mathbb{E}[I_t^2] = \mathbb{E}\left[\left(\int_0^t f_u \, dW_u\right)^2\right] = \mathbb{E}\left[\int_0^t f_u^2 \, du\right].$$

$\Rightarrow$ Hence, $\mathrm{Var}(I_t) = \mathbb{E}[I_t^2] = \mathbb{E}\left[\int_0^t f_u^2 \, du\right]$.

- More generally, if $s \leq t$,

$$\mathbb{E}[I_t | \mathcal{F}_s] = I_s \,,$$

$$\mathrm{Var}(I_t | \mathcal{F}_s) = \mathbb{E}\left[\int_s^t f_u^2 \, du \,\middle|\, \mathcal{F}_s\right].$$

Ito Integral - Properties

- If f_t is deterministic in t, then I_t is normally distributed:

$$I_t \sim \mathcal{N}(0, \mathrm{Var}(I_t))\,.$$

- If f_t is not deterministic in t, then it is not possible to determine the shape of the distribution of I_t in general.

Ito Integral - Properties

Example

We can use the Ito isometry to determine the distribution of W_t, the value of a Brownian motion at time t. Similar to the fundamental theorem of calculus, we can write

$$W_t - W_0 = \int_0^t dW_u \,.$$

Since $W_0 = 0$, by the Ito isometry, we have

$$\mathrm{Var}(W_t) = \mathbb{E}\left[\left(\int_0^t 1\, dW_u\right)^2\right] = \mathbb{E}\left[\int_0^t 1^2\, du\right] = \mathbb{E}[t] = t\,.$$

We can conclude that

$$W_t \sim \mathcal{N}(0, t)\,,$$

which is in line with what we know about the Brownian motion.

Continuous-time Models

Continuous-time Models - Stochastic Differential Equations

- We usually write Ito integrals in differential form, using the following notation:

$$I_t = \int_0^t f_t \, dW_t \iff dI_t = f_t \, dW_t$$

- A differential equation involving Ito integrals, written in differential form, is a stochastic differential equation (SDE).

- You may remember that an ordinary differential equation is an equation where the derivative is taken with respect to one variable. For example, consider the following ODE in differential form:

$$dB_t = rB_t \, dt \iff \frac{dB_t}{dt} = rB_t \, .$$

Continuous-time Models - Differential Form

- You may recall that the solution to this ODE is given by

$$B_t = Ce^{rt} \, ,$$

where C is a constant.

- If we impose the initial condition $B_0 = 1$, then we get $B_t = e^{rt}$. This is the value of 1 invested at a continuous interest rate r.

- These three equations are different notation for the same thing:

$$dB_t = rB_t \, dt$$

$$\frac{dB_t}{dt} = rB_t$$

$$B_T - B_0 = \int_0^T dB_t = \int_0^T rB_t \, dt \, .$$

Continuous-time Models - Differential Form

- Perhaps the nicest notation is the equation in its differential form:

$$dB_t = rB_t \, dt \, .$$

- In this form, we can more easily interpret the meaning of this equation. Suppose that B_t represents the value of an investment in a risk-free bank account.
- Recall that dB_t means, roughly, the change in the value of B_t. Similarly, dt represents the change in time.
- This equation says that the change in value of the investment is the rate r multiplied by the current value of the investment and the change in time.

Continuous-time Models - Ito Processes

- We are finally ready to formulate a model of the market.
- Continuous-time models are formulated in terms of Ito processes:

Definition (Ito Process)

An Ito process $\{X_t\}_{t\geq 0}$ is a stochastic process of the form

$$dX_t = \alpha_t \, dt + \sigma_t \, dW_t \,,$$

such that

- X_0 is a non-random initial condition,
- W_t is a Brownian motion,
- $\{\alpha_t\}_{t\geq 0}$ is an adapted process, called the drift,
- $\{\sigma_t\}_{t\geq 0}$ is also an adapted process, called the diffusion or volatility.

Continuous-time Models - Ito Processes

- A sample path of an Ito process is given below:

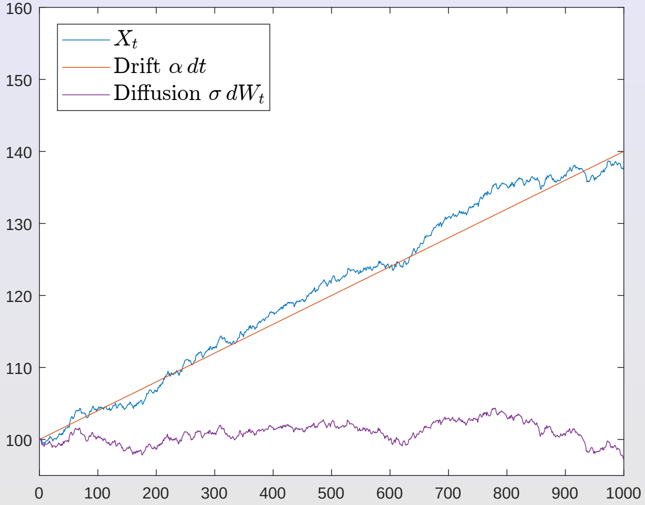

Continuous-time Models - Market Model

- Note that the equation $dX_t = \alpha_t \, dt + \sigma_t \, dW_t$ is just another way of writing

$$X_t = X_0 + \int_0^t \alpha_t \, dt + \int_0^t \sigma_t \, dW_t \, .$$

- The interpretation here is that the change in X at time t is driven by two factors:
  - A factor of α_t multiplied by the change in time,
  - A factor of σ_t multiplied by the change in value of an underlying random process.

Continuous-time Models - Ito Processes

- Consider a market with two assets: a risk-free asset B and a risky asset S. A continuous-time market model is a specification of the behaviour of these assets as Ito processes. These are also called the dynamics of these assets.
- We have seen that a risk-free asset earning continuous interest is specified by

$$dB_t = rB_t\,dt\,.$$

- On the other hand, the risky asset's dynamics are specified more generally, by

$$dS_t = \alpha_t\,dt + \sigma_t\,dW_t\,.$$

Continuous-time Models - Arithmetic Brownian Motion

- One of the simplest Ito processes is the following:

Definition (Arithmetic Brownian Motion)

A stochastic process $\{X_t\}_{t\geq 0}$ is an arithmetic Brownian motion (ABM) if it satisfies the following SDE:

$$dX_t = \alpha\, dt + \sigma\, dW_t\,,$$

where α, σ are constants.

- The example sample path from a few slides ago was an ABM.
- We can show (as an exercise) that under this model, $X_t \sim \mathcal{N}(X_0 + \alpha t, \sigma^2 t)$.

Ito's Lemma and Geometric Brownian Motion

Arithmetic Brownian Motion - Weaknesses

- It turns out that the ABM $dS_t = \alpha\, dt + \sigma\, dW_t$ is not a great model for a stock price S_t.
  1. Since S_t is normally distributed, it could take negative values, which is not realistic.
  2. Additionally, we see that the drift and diffusion do not depend on the level of S_t. This may not be realistic either. We would typically expect that the return would scale with the price of the stock.

- To address the first issue, we could model the instantaneous rate of return as an ABM instead:

$$d\ln(S_t) = \alpha\, dt + \sigma\, dW_t\,.$$

  Since a rate of return can be negative, this is not an issue.

Geometric Brownian Motion

- To address the second issue, we could introduce a factor of S_t in the drift and diffusion terms:

$$dS_t = \alpha S_t \, dt + \sigma S_t \, dW_t \, .$$

- Our goal for the remainder of this section is to show that these two models are actually the same thing.
- This model is the most important model of stock dynamics, and has a special name:

Definition (Geometric Brownian Motion)

A stochastic process $\{X_t\}_{t \geq 0}$ is an geometric Brownian motion (GBM) if it satisfies the following SDE:

$$dX_t = \alpha X_t \, dt + \sigma X_t \, dW_t \, ,$$

where α, σ are constants.

Geometric Brownian Motion - Motivating Example

Example

Recall that the risk-free asset B_t has dynamics given by

$$dB_t = rB_t\,dt\,.$$

Then by applying the chain rule, we have:

$$\begin{aligned}
\frac{d\ln(B_t)}{dt} &= \frac{1}{B_t}\cdot\frac{dB_t}{dt} \\
&= \frac{1}{B_t}\cdot rB_t \\
&= r \\
d\ln(B_t) &= r\,dt\,.
\end{aligned}$$

Geometric Brownian Motion - Motivating Example

Example

Hence, we have shown that the two equations

$$dB_t = rB_t \, dt \, ,$$
$$d \ln(B_t) = r \, dt \, ,$$

are the same model on the risk-free asset.

- However, we used the chain rule to find $d \ln(B_t)$. Can we do the same thing to find $d \ln(S_t)$ if S_t is a GBM?
- The normal chain rule does not work for stochastic processes. We can see that dS_t/dt does not exist, since S_t is not differentiable!

Ito-Doeblin Lemma

- The following result serves as the analogue of the chain rule for stochastic processes.

Theorem (Ito-Doeblin Lemma)

Let X_t be an Ito process with $dX_t = \alpha_t \, dt + \sigma_t \, dW_t$. Suppose $f(t,x)$ is a function such that the partial derivatives $f_t(t,x)$, $f_x(t,x)$, and $f_{xx}(t,x)$ are defined and continuous. Then

$$df(t,X_t) = f_t(t,X_t)\, dt + f_x(t,X_t)\, dX_t + \frac{1}{2} f_{xx}(t,X_t)(dX_t)^2 \, .$$

- In this expression, the first two terms come from the chain rule for total derivatives for regular functions. The third term is new in this case.
- The proof of the Ito-Doeblin Lemma is beyond the scope of this course. A second-order Taylor approximation plays a role, which is where the new term comes from.

Ito-Doeblin Lemma - Cross-Variation Terms

We see in the Ito-Doeblin Lemma that there is a $(dX_t)^2$ term. Officially, this is the cross-variation of X_t with itself, or the quadratic variation of X_t. We will use the following rules to calculate these terms.

1. $dt\, dt = 0$. Very roughly, this is saying that

$$\lim_{h \to 0} \sum_t ((t+h) - t)^2 = 0\,.$$

2. $dt\, dW_t = 0$.

3. $dW_t\, dW_t = dt$.

4. If $W_t^{(1)}$ and $W_t^{(2)}$ are two Brownian motions with correlation ρ, then

$$dW_t^{(1)}\, dW_t^{(2)} = \rho\, dt$$

.

Ito-Doeblin Lemma - Multivariate Case

- The Ito-Doeblin Lemma can be extended to multiple variables. The two-variable case is given below.

Theorem (Multivariate Ito-Doeblin Lemma)

Let X_t and Y_t be Ito processes. Suppose $f(t,x,y)$ is a twice-differentiable function. Then

$$\begin{aligned} df(t, X_t, Y_t) = f_t(t, X_t, Y_t)\, dt + f_x(t, X_t, Y_t)\, dX_t + f_y(t, X_t, Y_t)\, dY_t \\ + \frac{1}{2} f_{xx}(t, X_t, Y_t)(dX_t)^2 + \frac{1}{2} f_{yy}(t, X_t, Y_t)(dY_t)^2 \\ + f_{xy}(t, X_t, Y_t)\, dX_t\, dY_t \,. \end{aligned}$$

Ito-Doeblin Lemma and the GBM

Example (Geometric Brownian Motion)

We can now use the Ito-Doeblin Lemma to determine $d\ln(S_t)$. Suppose that $dS_t = \alpha S_t\,dt + \sigma S_t\,dW_t$. Let $f(t,S) = \ln(S)$.

Then $f_t = 0$, $f_S = \frac{1}{S}$, and $f_{SS} = -\frac{1}{S^2}$. Therefore by the Ito-Doeblin Lemma, we have

$$
\begin{aligned}
d\ln(S_t) &= f_t\,dt + f_S\,dS_t + \frac{1}{2}f_{SS}(dS_t)^2 \\
&= \frac{1}{S_t}\,dS_t - \frac{1}{2S_t^2}(dS_t)^2 \\
&= \frac{1}{S_t}\,(\alpha S_t\,dt + \sigma S_t\,dW_t) - \frac{1}{2S_t^2}(\alpha S_t\,dt + \sigma S_t\,dW_t)^2 \\
&= \alpha\,dt + \sigma\,dW_t - \frac{1}{2S_t^2}\,\sigma^2 S_t^2\,dt \\
&= \left(\alpha - \frac{\sigma^2}{2}\right)\,dt + \sigma\,dW_t\,.
\end{aligned}
$$

Ito-Doeblin Lemma and the GBM

- We have proved the following:

Proposition

The logarithm of a GBM is an ABM. In particular, the following two SDEs have the same solutions:

$$dS_t = \alpha S_t \, dt + \sigma S_t \, dW_t$$

$$d \ln(S_t) = \left(\alpha - \frac{\sigma^2}{2} \right) dt + \sigma \, dW_t$$

GBM - Solution

- In fact, we can get an explicit characterization of the solutions to GBMs.

Proposition (GBM Solution Form)

The unique solution to $dS_t = \alpha S_t\,dt + \sigma S_t\,dW_t$ *is given by*

$$S_t = S_0 e^{\left(\alpha - \frac{\sigma^2}{2}\right)t + \sigma W_t}.$$

Also, for all $s \leq t$,

$$S_t = S_s e^{\left(\alpha - \frac{\sigma^2}{2}\right)(t-s) + \sigma(W_t - W_s)}.$$

GBM - Solution

Proof.

By the result of the previous proposition, we have

$$d\ln(S_t) = \left(\alpha - \frac{\sigma^2}{2}\right) dt + \sigma\, dW_t\,.$$

Integrating this equation from time 0 to t gives

$$\int_0^t d\ln(S_u) = \int_0^t \left(\alpha - \frac{\sigma^2}{2}\right) du + \int_0^t \sigma\, dW_u$$
$$\ln(S_t) - \ln(S_0) = \left(\alpha - \frac{\sigma^2}{2}\right) t + \sigma W_t\,.$$

The result follows by solving for S_t. The second equation follows similarly, by integrating from s to t. □

GBM - Lognormal Distribution

- We can also identify the distribution of the solution to the Geometric Brownian Motion.

Definition (Lognormal Distribution)

We say that a random variable X is lognormally distributed if

$$\ln(X) \sim \mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)\,.$$

We use the notation

$$X \sim log\mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)\,.$$

- Note that the parameters $\tilde{\mu}$ and $\tilde{\sigma}$ are not the mean and variance of X.

GBM - Lognormal Distribution

- If $X \sim log\mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)$, then

$$\mathbb{E}[X] = e^{\tilde{\mu} + \frac{\tilde{\sigma}^2}{2}},$$

$$\mathrm{Var}(X) = (e^{\tilde{\sigma}^2} - 1)e^{2\tilde{\mu} + \tilde{\sigma}^2}.$$

This can be calculated from the moment generating function of the normal distribution.

Proposition

If $\{S_t\}_{t \geq 0}$ is a GBM with $dS_t = \alpha S_t\, dt + \sigma S_t\, dW_t$, then

$$S_t \sim log\mathcal{N}\left(\ln(S_0) + \left(\alpha - \frac{\sigma^2}{2}\right)t, \sigma^2 t\right).$$

GBM - Lognormal Distribution

Proof.

We have seen that

$$\ln(S_t) - \ln(S_0) = \left(\alpha - \frac{\sigma^2}{2}\right) t + \sigma W_t \,.$$

Note that $W_t \sim \mathcal{N}(0, t)$. It follows immediately that

$$\ln(S_t) \sim \mathcal{N}\left(\ln(S_0) + \left(\alpha - \frac{\sigma^2}{2}\right) t, \sigma^2 t\right) \,,$$

which implies the result. □

GBM - Summary

- The model of stock prices as Geometric Brownian Motion seems to be a reasonable model that addresses shortcomings of the ABM.
- The GBM is central to modern asset pricing theory, and is widely used in practice.
- The GBM was the stochastic model used by Black and Scholes to obtain their option pricing formulas. We will derive these formulas in the next part of this course.

GBM - Summary

However, the GBM is not without its drawbacks:

- The GBM does not allow for jumps in the stock prices. Hence, it cannot model the market impact of “black swan events” such as COVID-19. This is addressed by jump-diffusion models.
- The GBM assumes a constant volatility, whereas volatility in practice is stochastic. This is addressed by stochastic volatility models.
- Returns are assumed to be normally distributed, but in reality a distribution with a heavier tail is observed. This is addressed by heavy-tail distributions.

The specifics of these more sophisticated models are beyond the scope of this course.

GBM - Summary

- This is the chart for the price of the S&P 500 over the last 10 years:

GBM - Summary

- This is a sample path of a Brownian motion $dS_t = 0.07S_t\,dt + 0.1S_t\,dW_t$. Does this look like a reasonable model?