Basic Stochastic Processes: Financial applications

Lecture 3 (3 December 2020)

Probability theory on uncountable sample spaces

In this lecture we assume that Ω is uncountable (e.g., $\Omega = \mathbb{R}$).

In this case there is no general procedure to construct a probability space, but only an abstract definition.

In particular a probability measure $\mathbb{P}$ on events $A \subseteq \Omega$ is defined only axiomatically by requiring that $0 \leq \mathbb{P}(A) \leq 1$, $\mathbb{P}(\Omega) = 1$ and that, for any sequence of disjoint events $A_1, A_2, \ldots$, it should hold

$$\mathbb{P}(A_1 \cup A_2 \cup \ldots) = \mathbb{P}(A_1) + \mathbb{P}(A_2) + \ldots \, (*)$$

Moreover we do not assume that $\mathbb{P}$ is defined for all events $A \subset \Omega$.

Denote by $\mathcal{F}$ the set of events (i.e., subsets of Ω) which have a well defined probability satisfying the properties above.

The minimial conditions required on $\mathcal{F}$ is that

(i) $\mathcal{F}$ should contain Ω (the "something happens event"),

(ii) the complement of each element A, i.e., $A \in \mathcal{F} \Rightarrow A^c \in \mathcal{F}$ (the "A does not happen event")

(iii) $\mathcal{F}$ must be closed with respect to caountable union (so that the equation (*) above makes sense)

A collections of event satisfying the properties (i), (ii), (iii) is called a σ-**algebra** (or σ-**field**).

Example.

Let $\Omega = \mathbb{R}$. We say that $A \subseteq \mathbb{R}$ is a **Borel** set if it can be written as the union (or intersection) of countably many open (or closed) intervals.

Let $\mathcal{F}$ be the collection of all Borel sets. Let $p : \mathbb{R} \to \mathbb{R}$ be a continuous non-negative function such that

$$\int_{\mathbb{R}} p(\omega) \, d\omega = 1.$$

Then $\mathbb{P} : \mathcal{F} \to [0, 1]$ given by

$$\mathbb{P}(A) = \int_A p(\omega) \, d\omega$$

defines a probability. If $X : \mathbb{R} \to \mathbb{R}$ is a random variables, the expectation of X in this probability measure is given by

$$\mathbb{E}[X] = \int_{\mathbb{R}} X(\omega) p(\omega) \, dx.$$

2

For most applications (and in particular for those in financial mathematics) the knowledge of the full probability space is not necessary.

More precisely, we are only interested in assigning a probability to events of the form $\{X \in I\}$, where X is a random variable on the (abstract) probability space and $I \subset \mathbb{R}$, that is to say, events which can be resolved by one (or more) random variables.

The probability $\mathbb{P}(X \in I)$ can be computed explicitly when X has a density.

Definition 0.3

Let $f_X : \mathbb{R} \to [0, \infty)$ be a continuous function (except possibly on finitely many points). A continuum random variable $X : \Omega \to \mathbb{R}$ is said to have **probability density** f_X if

$$\mathbb{P}(X \in A) = \int_A f_X(x)\, dx,$$

for all Borel sets $A \subseteq \mathbb{R}$.

Note that the density f_X satisfies

$$\int_{\mathbb{R}} f_X(x)\, dx = 1$$

and the **cumulative distribution** $F_X(x) = \mathbb{P}(X \leq x)$ satisfies

$$F_X(x) = \int_{-\infty}^{x} f_X(y)\, dy, \quad \text{for all } x \in \mathbb{R}, \text{ hence } f_X = \frac{dF_X}{dx}.$$

Example.

A random variable $X : \Omega \to \mathbb{R}$ is said to be a **normal** random variable with **mean** $m \in \mathbb{R}$ and **variance** $\sigma^2 > 0$ if it admits the density

$$f_X(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{|x-m|^2}{2\sigma^2}\right).$$

We denote $\mathcal{N}(m, \sigma^2)$ the set of all such random variables.

A variable $X \in \mathcal{N}(0, 1)$ is called a **standard normal** random variable.

The cumulative distribution of standard normal random variables is denoted by $\Phi(x)$ and is called the **standard normal distribution**, i.e.,

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{1}{2}y^2} \, dy.$$

Theorem 0.8

The following holds for all sufficiently regular functions $g : \mathbb{R} \to \mathbb{R}$:

(i) Let $X : \Omega \to \mathbb{R}$ be a random variable with density f_X. Then for all Borel sets $A \subseteq \mathbb{R}$,

$$\mathbb{P}(g(X) \in A) = \int_{x:g(x)\in A} f_X(x) \, dx.$$

(ii) Let $X : \Omega \to \mathbb{R}$ be a random variable with density f_X. Then

$$\mathbb{E}[g(X)] = \int_{\mathbb{R}} g(y) f_X(y) \, dy.$$

Moreover the properties 1,2,3 in Theorem 0.1 still hold on uncountable probability spaces.

By (ii) in Theorem 0.8, the expectation and the variance of a continuum random variable X with density f_X are given by

$$\mathbb{E}[X] = \int_{\mathbb{R}} x f_X(x)\, dx, \quad \mathrm{Var}[X] = \int_{\mathbb{R}} x^2 f_X(x)\, dx - \left(\int_{\mathbb{R}} x f_X(x)\, dx \right)^2.$$

In particular normal variables we obtain

$$X \in \mathcal{N}(m, \sigma^2) \implies \mathbb{E}[X] = m, \quad \mathrm{Var}[X] = \sigma^2.$$

Joint probability density

Definition 0.4

Two continuum random variables $X, Y : \Omega \to \mathbb{R}$ are said to have the **joint probability density** $f_{X,Y} : \mathbb{R}^2 \to [0, \infty)$, if

$$\mathbb{P}(X \in A, Y \in B) = \int_A \int_B f_{X,Y}(x, y)\, dx\, dy,$$

for all Borel sets $A, B \subseteq \mathbb{R}$.

Note that if $f_{X,Y}$ is a joint probability density, then

$$\int_{\mathbb{R}} \int_{\mathbb{R}} f_{X,Y}(x, y)\, dx\, dy = 1.$$

Moreover if we define the **joint cumulative distribution** as $F_{X,Y}(x,y) = \mathbb{P}(X \le x, Y \le y)$ then

$$f_{X,Y}(x,y) = \partial_x \partial_y F_{X,Y}(x,y).$$

When X, Y have the joint density $f_{X,Y}(x,y)$, the random variables X, Y admit the densities

$$f_X(x) = \int_{\mathbb{R}} f_{X,Y}(x,y)\,dy, \quad f_Y(y) = \int_{\mathbb{R}} f_{X,Y}(x,y)\,dx.$$

Example: Jointly normally distributed random variables.

Let $m \in \mathbb{R}^2$ and $C = (C_{ij})_{i,j=1,2}$ be a symmetric, positive definite 2×2 matrix. Two random variables $X_1, X_2 : \Omega \to \mathbb{R}$ are said to be jointly normally distributed with mean m and **covariance matrix** C if they admit the joint density

$$f_{X_1,X_2}(x) = \frac{1}{\sqrt{(2\pi)^2 \det C}} \exp\left(-\frac{1}{2}(x-m)C^{-1}(x-m)\right), \quad \text{for all } x = (x_1, x_2) \in \mathbb{R}^2.$$

The following theorem generalizes Theorem 0.8 in the presence of two variables.

Theorem 0.9

Let $X, Y : \Omega \to \mathbb{R}$ be random variables with joint density f_X and $g : \mathbb{R}^2 \to \mathbb{R}$.

(i) For all Borel sets $A \subseteq \mathbb{R}$ there holds

$$\mathbb{P}(g(X, Y) \in A) = \int_{(x,y):g(x,y)\in A} f_{X,Y}(x, y) \, dx \, dy.$$

(ii) There holds

$$\mathbb{E}[g(X, Y)] = \int_{\mathbb{R}^2} g(x, y) f_{X,Y}(x, y) \, dx \, dy.$$

By (ii) of Theorem 0.9, if X_1, X_2 have the joint density f_{X_1,X_2}, then the covariance of X_1, X_2 can be computed as

$$
\begin{aligned}
\mathrm{Cov}(X_1, X_2) &= \mathbb{E}[X_1 X_2] - \mathbb{E}[X_1]\mathbb{E}[X_2] \\
&= \int_{\mathbb{R}^2} x_1 x_2 f_{X_1,X_2}(x_1, x_2) \, dx_1 \, dx_2 \\
&\quad - \int_{\mathbb{R}^2} x_1 f_{X_1,X_2}(x_1, x_2) \, dx_1 \, dx_2 \int_{\mathbb{R}^2} x_2 f_{X_1,X_2}(x_1, x_2) \, dx_1 \, dx_2.
\end{aligned}
$$

In particular, if X_1, X_2 are jointly normal distributed with mean $m \in \mathbb{R}^2$ and covariance matrix $C = (C_{ij})_{i,j=1,2}$, we find

$$m = (m_1, m_2), \qquad C_{ij} = \mathrm{Cov}(X_i, X_j).$$

The following result on the linear combination of independent normal random variables will play an important role for the project in multi-asset options.

Theorem 0.10

Let $X_1, X_2 \in \mathcal{N}(0,1)$ be independent and $a, b, c, d \in \mathbb{R}$. Then $aX_1 + bX_2 \in \mathcal{N}(0, a^2 + b^2)$. Moreover if

$$Y_1 = aX_1 + bX_2, \quad Y_2 = cX_1 + dX_2,$$

and if the matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is invertible, then Y_1, Y_2 are jointly normally distributed with zero mean and covariant matrix $C = AA^T$.

Stochastic processes. Martingales

Let Ω be an uncountable sample space.

A stochastic process is a one parameter family $\{X(t)\}_{t \geq 0}$ of random variables $X(t) : \Omega \to \mathbb{R}$.

We denote $X(t, \omega) = X(t)(\omega)$.

The parameter t is referred to as the time variable, since this is what it represents in the applications that we have in mind.

For each $\omega \in \Omega$ *fixed*, the function $t \to X(t, \omega)$ is called a path of the stochastic process.

If the paths are all the same for all $\omega \in \Omega$, then we say that $X(t)$ is a deterministic function of time.

Martingales

Martingale stochastic processes play a fundamental role in options pricing theory.

To define martingales on uncountable sample spaces, let $\mathcal{F}_X(t)$ denote the information accumulated by "looking" at the stochastic process up to time t, i.e., the collection of events resolved by $X(s)$ for $0 \leq s \leq t$.

Intuitively, the stochastic process $\{X(t)\}_{t \geq 0}$ is a martingale if, based on the information contained in $\mathcal{F}_X(s)$, our "best estimate" on $X(t)$ for $t > s$ is $X(s)$, i.e., we are not able to estimate whether the process will raise or fall in the interval $[s, t]$ with the information available at time s.

This intuitive definition is encoded in the formula

$$\mathbb{E}[X(t)|\mathcal{F}_X(s)] = X(s), \quad 0 \leq s \leq t,$$

which generalizes the definition of martingales in finite probability theory.

The left hand side is the conditional expectation of $X(t)$ with respect to the information $\mathcal{F}_X(s)$, whose precise definition is not needed here.

It can be shown that martingales have constant expectation:

Brownian motion

Next we recall the definition of the most important of all stochastic processes.

Definition 0.5

A **Brownian motion**, or **Wiener process**, is a stochastic process $\{W(t)\}_{t \geq 0}$ with the following properties:

1. For all[1] $\omega \in \Omega$, the paths are continuous (i.e., $t \to W(t, \omega)$ is a continuous function) and $W(0, \omega) = 0$;

2. For all $0 = t_0 < t_1 < t_2 < \ldots$, the increments

$$W(t_1) = W(t_1) - W(t_0), \ W(t_2) - W(t_1), \ldots,$$

   are independent random variables;

3. The increments are normally distributed, that is to say, for all $0 \leq s < t$,
$$\mathbb{P}(W(t) - W(s) \in A) = \frac{1}{\sqrt{2\pi(t-s)}} \int_A e^{-\frac{y^2}{2(t-s)}} \, dy,$$
   for all Borel sets $A \subseteq \mathbb{R}$.

Important remark:

Since the definition of Brownian motion depends on the probability measure $\mathbb{P}$, then a stochastic process $\{W(t)\}_{t \geq 0}$ which is a Brownian motion in the probability measure $\mathbb{P}$ will in general *not* be a Brownian motion in another probability measure $\widetilde{\mathbb{P}}$. When we want to emphasize that $\{W(t)\}_{t \geq 0}$ is a Brownian motion in the probability measure $\mathbb{P}$, we shall say that $\{W(t)\}_{t \geq 0}$ is a $\mathbb{P}$-Brownian motion.

---

[1]More precisely, for all $\omega \in \Omega$ up to a set of zero probability.

Remark

Letting $s = 0$ in property 3 in the definition we obtain that $W(t) \in \mathcal{N}(0, t)$, for all $t > 0$. In particular, $W(t)$ has zero expectation for all times. It can also be shown that Brownian motions are martingales.

The following result is used a few times in some projects.

Theorem 0.11

Let $g : (0, \infty) \to \mathbb{R}$ be a differentiable function and let

$$X(t) = g(t)W(t) - \int_0^t g'(s)W(s)\, ds.$$

Then

$$X(t) \in \mathcal{N}(0, \Delta(t)), \quad \Delta(t) = \int_0^t g(s)^2\, ds.$$

Remark

By using the formal identity $d(g(t)W(t)) = g'(t)W(t)dt + g(t)dW(t)$, as well as $\int_0^t d(g(s)W(s)) = g(t)W(t)$, we can write the definition of $X(t)$ in the previous theorem as

$$X(t) = \int_0^t g(s)dW(s),$$

which is called **Itô integral** of the deterministic function $g(t)$.

Equivalent probability measures. Girsanov theorem

One further technical complication arising for uncountable sample spaces is the existence of non-trivial events with zero measure, e.g., the event $\{W(t) = 0\}$ that the Brownian motion $W(t)$ takes value zero when $t > 0$.

We shall need to consider the concept of equivalent probability measures:

Definition 0.6 Two probability measure $\mathbb{P}$, $\widetilde{\mathbb{P}}$ on the events $A \in \mathcal{F}$ are said to be equivalent if $\mathbb{P}(A) = 0 \Leftrightarrow \widetilde{\mathbb{P}}(A) = 0$.

Hence equivalent probability measures agree on which events are impossible.

Note that in a finite probability space all probability measures are equivalent, as in the finite case the empty set is the only event with zero probability.

The following important theorem characterizes the relation between equivalent probability measures on uncountable sample spaces and is known as the Radon-Nikodým theorem.

We denote $\mathbb{I}_A$ the **characteristic function** of the set $A \in \mathcal{F}$, i.e., the random variable taking value $\mathbb{I}_A(\omega) = 1$ if $\omega \in A$ and zero otherwise.

Radon-Nikodým theorem

Let $\mathbb{P} : \mathcal{F} \to [0, 1]$ be a probability measure. Then $\widetilde{\mathbb{P}} : \mathcal{F} \to [0, 1]$ is a probability measure equivalent to $\mathbb{P}$ if and only if there exists a random variable $Z : \Omega \to \mathbb{R}$ such that

$$Z > 0 \quad \mathbb{E}[Z] = 1 \quad \text{and} \quad \widetilde{\mathbb{P}}(A) = \mathbb{E}[Z\mathbb{I}_A]$$

Moreover if $\mathbb{P}$ and $\widetilde{\mathbb{P}}$ are equivalent then

$$\widetilde{\mathbb{E}}[X] = \mathbb{E}[ZX]$$

for all random variables $X : \Omega \to \mathbb{R}$.

Example

Assume $\Omega = \mathbb{R}$ and that $\mathbb{P}$ and $\widetilde{\mathbb{P}}$ are defined as in the example on page, namely

$$\mathbb{P}(A) = \int_A p(\omega)\, d\omega, \quad \widetilde{\mathbb{P}}(A) = \int_A \widetilde{p}(\omega)\, d\omega,$$

where A is a Borel set and p, $\widetilde{p}$ are two continuous non-negative functions such that

$$\int_{\mathbb{R}} p(\omega)\, d\omega = \int_{\mathbb{R}} \widetilde{p}(\omega)\, d\omega = 1.$$

Then, according to the Radon-Nikodým Theorem, $\mathbb{P}$ and $\widetilde{\mathbb{P}}$ are equivalent if and only if there exists a function $Z : \mathbb{R} \to \mathbb{R}$ such that $Z > 0$, and

$$\widetilde{\mathbb{P}}(A) = \int_A \widetilde{p}(\omega)\, d\omega = \int_{\mathbb{R}} Z(\omega)\mathbb{I}_A(\omega)p(\omega)\, d\omega = \int_A Z(\omega)p(\omega)\, d\omega.$$

As the equality $\int_A \widetilde{p}(\omega)\, d\omega = \int_A Z(\omega)p(\omega)\, d\omega$ has to be satisfied for all Borel sets $A \subset \mathbb{R}$, then $\widetilde{p}(\omega) = Z(\omega)p(\omega)$ must hold for all $\omega \in \mathbb{R}$.

Theorem 0.13 and Definition

Let $\{W(t)\}_{t \geq 0}$ be a $\mathbb{P}$-Brownian motion. Given $\theta \in \mathbb{R}$ and $T > 0$ define

$$Z_\theta = e^{-\theta W(T) - \frac{1}{2}\theta^2 T}.$$

Then $\mathbb{P}_\theta(A) = \mathbb{E}[Z_\theta \mathbb{I}_A]$, for all Borel sets $A \subseteq \mathbb{R}$, defines a probability measure equivalent to $\mathbb{P}$, which is called **Girsanov's probability** with parameter $\theta \in \mathbb{R}$.

Proof. The proof follows immediately from the Radon-Nikodým Theorem, since the random variable Z_θ satisfies $Z_\theta > 0$ and

$$\mathbb{E}[Z_\theta] = \mathbb{E}[e^{-\theta W(T) - \frac{1}{2}\theta^2 T}] = \int_{\mathbb{R}} e^{-\theta x - \frac{1}{2}\theta^2 T} \frac{e^{-\frac{x^2}{2T}}}{\sqrt{2\pi T}}\, dx = 1,$$

where we used the density of the normal random variable $W(T) \in \mathcal{N}(0, T)$ to compute the expectation of Z_θ in the probability measure $\mathbb{P}$. $\qquad\square$

Note that the Girsanov probability measure $\mathbb{P}_\theta$ depend also on T, but this is not reflected in our notation.

In the following we denote by $\mathbb{E}_\theta[\cdot]$ the expectation computed in the probability measure $\mathbb{P}_\theta$ for $\theta \neq 0$.

When $\theta = 0$ then $\mathbb{P}_\theta = \mathbb{P}$, in which case the expectation is denoted as usual by $\mathbb{E}[\cdot]$.

By Radon-Nikodým Theorem we have $\mathbb{E}_\theta[X] = \mathbb{E}[Z_\theta X]$, for all random variables $X : \Omega \to \mathbb{R}$.

Moreover we now show that $\mathbb{E}_\theta[W(t)] = -\theta t$. In fact by the Radon-Nikodým theorem we have

$$\mathbb{E}_\theta[W(t)] = \mathbb{E}[Z_\theta W(t)] = \mathbb{E}[e^{-\theta W(T) - \frac{1}{2}\theta^2 T} W(t)].$$

Adding and subtracting $W(t)$ in the exponent of the exponential function we have

$$\mathbb{E}_\theta[W(t)] = \mathbb{E}[e^{-\theta(W(T)-W(t)) - \frac{1}{2}\theta^2 T} e^{-\theta W(t)} W(t)] = \mathbb{E}[e^{-\theta(W(T)-W(t)) - \frac{1}{2}\theta^2 T}]\mathbb{E}[e^{-\theta W(t)} W(t)],$$

where in the last step we used that the random variables $X = e^{-\theta(W(T)-W(t)) - \frac{1}{2}\theta^2 T}$ and $Y = e^{-\theta W(t)} W(t)$ are independent (being functions of the independent random variables $W(T) - W(t)$ and $W(t)$).

Using $W(T) - W(t) \in \mathcal{N}(0, T-t)$ and $W(t) \in \mathcal{N}(0, t)$, we can compute the expectations of X and Y as

$$\mathbb{E}[X] = e^{-\frac{1}{2}\theta^2 T} \frac{1}{\sqrt{2\pi(T-t)}} \int_{\mathbb{R}} e^{-\theta x - \frac{x^2}{2(T-t)}}\, dx = e^{-\frac{\theta^2}{2}t},$$

14

$$\mathbb{E}[Y] = \frac{1}{\sqrt{2\pi t}} \int_{\mathbb{R}} e^{-\theta x - \frac{x^2}{2t}} x \, dx = -e^{\frac{\theta^2}{2}t}\theta t.$$

Hence $\mathbb{E}_\theta[W(t)] = \mathbb{E}[X]\mathbb{E}[Y] = -\theta t$, as claimed.

It follows that $\{W(t)\}_{t\geq 0}$ is *not* a $\mathbb{P}_\theta$-Brownian motion, since Brownian motions, by definition, have zero expectation at any time.

Now we can state a fundamental theorem in probability theory with deep applications in financial mathematics, namely Girsanov's theorem. (Actually we consider only a special case of this theorem, which suffices for our purposes)

Girsanov theorem

Let $\{W(t)\}_{t\geq 0}$ be a $\mathbb{P}$-Brownian motion. Given $\theta \in \mathbb{R}$ and $T > 0$, let $\mathbb{P}_\theta$ be the Girsanov probability measure with parameter θ introduced in Theorem **??**. Define the stochastic process $\{W^{(\theta)}(t)\}_{t\geq 0}$ by

$$W^{(\theta)}(t) = W(t) + \theta t.$$

Then $\{W^{(\theta)}(t)\}_{t\geq 0}$ is a $\mathbb{P}_\theta$-Brownian motion.

Note carefully that $\{W^{(\theta)}(t)\}_{t\geq 0}$ is *not* a $\mathbb{P}$-Brownian motion, as it follows by the fact that $\mathbb{E}[W^{(\theta)}(t)] = \theta t$.

In particular, according to the probability measure $\mathbb{P}$, the stochastic process $\{W^{(\theta)}(t)\}_{t\geq 0}$ has a *drift*, i.e., a tendency to move up (if $\theta > 0$) or down (if $\theta < 0$). However in the Girsanov probability this drift is removed, because, as shown before, $\mathbb{E}_\theta[W^{(\theta)}(t)] = \mathbb{E}_\theta[W(t)] + \theta t = 0$.

Multi-dimensional Girsanov theorem

We conclude with a generalization of Girsanov's theorem in the presence of two independent Brownian motions. This generalization is important for the project on multi-asset options

Theorem 0.15 and Definition Let $\{W_1(t)\}_{t\geq 0}$, $\{W_2(t)\}_{t\geq 0}$ be $\mathbb{P}$-independent Brownian motions. Given $\theta = (\theta_1, \theta_2) \in \mathbb{R}^2$ and $T > 0$ define

$$Z_\theta = e^{-\theta_1 W_1(T) - \theta_2 W_2(T) - \frac{1}{2}(\theta_1^2 + \theta_2^2)T}.$$

Then $\mathbb{P}_\theta(A) = \mathbb{E}[Z_\theta \mathbb{I}_A]$ defines a probability measure equivalent to $\mathbb{P}$, which is called Girsanov's probability with parameters $\theta_1, \theta_2 \in \mathbb{R}$.

2-dimensional Girsanov's theorem

Let $\{W_1(t)\}_{t\geq 0}$, $\{W_2(t)\}_{t\geq 0}$ be $\mathbb{P}$-independent Brownian motions. Given $\theta = (\theta_1, \theta_2) \in \mathbb{R}^2$ and $T > 0$, let $\mathbb{P}_\theta$ be the Girsanov probability with parameters θ_1, θ_2 introduced in Theorem **??**. Define the stochastic processes $\{W_1^{(\theta)}(t)\}_{t\geq 0}$, $\{W_2^{(\theta)}(t)\}_{t\geq 0}$ by

$$W_1^{(\theta)}(t) = W_1(t) + \theta_1 t, \quad W_2^{(\theta)}(t) = W_2(t) + \theta_2 t$$

Then $\{W_1^{(\theta)}(t)\}_{t\geq 0}$, $\{W_2^{(\theta)}(t)\}_{t\geq 0}$ are $\mathbb{P}_\theta$-independent Brownian motions.

Black-Scholes options pricing theory

In the binomial model the stock price at time t is a finite random variable $S(t)$. In the Black-Scholes model the stock price is a continuum random variable with image $\text{Im}(S(t)) = (0, \infty)$, namely the **geometric Brownian motion**

$$S(t) = S_0 e^{\alpha t + \sigma W(t)} \quad (GBM)$$

The probability $\mathbb{P}$ with respect to which $\{W(t)\}_{t \geq 0}$ is Brownian motion is the **physical** (or **real-world**) **probability** of the Black-Scholes market.

Moreover α is the **instantaneous mean of log-return**, σ is the **instantaneous volatility** and σ^2 is the **instantaneous variance** of the geometric Brownian motion

The geometric Brownian motion admits the density

$$f_{S(t)}(x) = \frac{H(x)}{\sqrt{2\pi\sigma^2 t}} \frac{1}{x} \exp\left(-\frac{(\log x - \log S(0) - \alpha t)^2}{2\sigma^2 t}\right),$$

where $H(x)$ is the **Heaviside** function. It can be shown that the binomial stock price converges in distribution to the geometric Brownian motion in the time-continuum limit.

The risk-neutral pricing formula in Black-Scholes markets

The purpose of this section is to introduce the definition of Black-Scholes price of European derivatives from a probability theory point of view.

Recall that the probabilistic formulation of the binomial options pricing model is encoded in the risk-neutral pricing formula.

Our goal is to derive a similar risk-neutral pricing formula (at time $t = 0$) for the time-continuum Black-Scholes model.

Motivated by the approach for the binomial model, we first look for a probability measure in which the the discounted stock price in Black-Scholes markets is a martingale (martingale probability measure).

It is natural to seek such martingale probability within the class of Girsanov probabilities $\mathbb{P}_\theta$ equivalent to the physical probability $\mathbb{P}$.

To this purpose we shall need the form of the density function of the geometric Brownian motion in the probability measure $\mathbb{P}_\theta$.

Theorem 0.17

Let $\theta \in \mathbb{R}$, $T > 0$ and $\mathbb{P}_\theta$ be the Girsanov probability measure equivalent to the physical probability $\mathbb{P}$. The geometric Brownian motion has the following density in the probability measure $\mathbb{P}_\theta$:

$$f_{S(t)}^{(\theta)}(x) = \frac{H(x)}{\sqrt{2\pi\sigma^2 t}} \frac{1}{x} \exp\left(-\frac{(\log x - \log S_0 - (\alpha - \theta\sigma)t)^2}{2\sigma^2 t} \right).$$

Proof. Since

$$S(t) = S_0 e^{\alpha t + \sigma W(t)} = S_0 e^{(\alpha - \theta\sigma)t + \sigma W^{(\theta)}(t)}, \quad W^{(\theta)}(t) = W(t) + \theta t$$

and since $\{W^{(\theta)}(t)\}_{t \geq 0}$ is a Brownian motion in the probability measure $\mathbb{P}_\theta$ (see Girsanov's Theorem **??**), then the density $f_{S(t)}^{(\theta)}$ is the same as $f_{S(t)}$ with α replaced by $\alpha - \theta\sigma$. $\square$

Let $\mathbb{E}_\theta[\cdot]$ denote the expectation in the probability $\mathbb{P}_\theta$.

Recall that martingales have constant expectation. Hence in the martingale (or risk-neutral) probability measure the expectation of the discounted value of the stock must be constant, i.e., $\mathbb{E}_\theta[S(t)] = S_0 e^{rt}$. This condition alone suffices to single out a unique possible value of θ.

In fact we now show that the identity $\mathbb{E}_\theta[S(t)] = S_0 e^{rt}$ holds if and only if $\theta = q$, where

$$q = \frac{\alpha - r}{\sigma} + \frac{\sigma}{2}.$$

Proof. Using the density of $S(t)$ in the measure $\mathbb{P}_\theta$ we have

$$\mathbb{E}_\theta[S(t)] = \int_{\mathbb{R}} x f_{S(t)}^{(\theta)}(x)\, dx = \frac{1}{\sqrt{2\pi\sigma^2 t}} \int_0^\infty \exp\left(-\frac{(\log x - \log S_0 - (\alpha - \theta\sigma)t)^2}{2\sigma^2 t}\right) dx.$$

With the change of variable $y = \frac{\log x - \log S_0 - (\alpha - \theta\sigma)t}{\sigma\sqrt{t}}$, $dx = x\sigma\sqrt{t}\, dy$, we obtain

$$\mathbb{E}_\theta[S(t)] = \frac{S_0}{\sqrt{2\pi}} e^{(\alpha - \theta\sigma)t} \int_{\mathbb{R}} e^{-\frac{y^2}{2} + \sigma\sqrt{t}y}\, dy = S_0 e^{(\alpha - \theta\sigma + \frac{\sigma^2}{2})t} \frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} e^{-\frac{(y + \sigma\sqrt{t})^2}{2}}\, dy.$$

As $\frac{1}{\sqrt{2\pi}} \int_{\mathbb{R}} e^{-\frac{x^2}{2}}\, dx = 1$, the result follows. $\square$

Even though the validity of $\mathbb{E}_\theta[S(t)] = S_0 e^{rt}$ is only necessary for the discounted geometric Brownian motion to be a martingale, one can show that the following result holds.

The discounted value of the geometric Brownian motion stock price is a martingale in the probability measure $\mathbb{P}_\theta$ if and only if $\theta = q$

The previous discussion leads us to the following definition.

Definition 0.19

Given $\alpha \in \mathbb{R}$, $\sigma > 0$, $r \in \mathbb{R}$ and $T > 0$, the probability measure

$$\mathbb{P}_q(A) = \mathbb{E}[e^{-\theta W(T) - \frac{1}{2}\theta^2 T}\mathbb{I}_A], \quad q = \frac{\alpha - r}{\sigma} + \frac{\sigma}{2}$$

is called the **martingale probability**, or **risk-neutral probability**, in the interval $[0, T]$ of the Black-Scholes market with parameters α, σ, r.

Remark

In the risk-neutral probability the stock price is given by the geometric Brownian motion

$$S(t) = S(0)e^{(r - \frac{\sigma^2}{2})t + \sigma W^{(q)}(t)},$$

where, by Girsanov's theorem, $W^{(q)}(t) = W(t) + qt$ is a Brownian motion in the risk-neutral probability. This follows by replacing $\alpha = r + q\sigma - \frac{1}{2}\sigma^2$ in (GBM).

Moreover replacing $\theta = q$ in the density $f_{S(t)}^{(\theta)}$ (see Theorem 0.17) we obtain that the geometric Brownian motion has the following density in the risk-neutral probability measure $\mathbb{P}_q$:

$$f_{S(t)}^{(q)}(x) = \frac{H(x)}{\sqrt{2\pi\sigma^2 t}}\frac{1}{x}\exp\left(-\frac{(\log x - \log S_0 - (r - \frac{\sigma^2}{2})t)^2}{2\sigma^2 t}\right).$$

At this point we have all we need to define the Black-Scholes price of European derivatives at time $t = 0$ using the risk-neutral pricing formula.

Definition 0.8

The Black-Scholes price at time $t = 0$ of the European derivative with pay-off Y at maturity T is given by the risk-neutral pricing formula

$$\Pi_Y(0) = e^{-rT}\mathbb{E}_q[Y],$$

i.e., it equals the expected value of the discounted pay-off in the risk-neutral probability measure of the Black-Scholes market.

In the case of standard European derivatives we can use the density of the geometric Brownian motion in the risk-neutral probability measure to write the Black-Scholes price in the following integral form.

Theorem 0.20

For the standard European derivative with pay-off $Y = g(S(T))$ at maturity $T > 0$, the Black-Scholes price at time $t = 0$ can be written as $\Pi_Y(0) = v_0(S_0)$, where S_0 is the price of the underlying stock at time $t = 0$ and $v_0 : (0, \infty) \to \mathbb{R}$ is the **pricing function** of the derivative at time $t = 0$, which is given by

$$v_0(x) = e^{-rT} \int_{\mathbb{R}} g(xe^{(r-\frac{\sigma^2}{2})T+\sigma\sqrt{T}y})e^{-\frac{1}{2}y^2} \frac{dy}{\sqrt{2\pi}}.$$

Proof. Using the density $f_{S(t)}^{(q)}$ for $t = T$ in the risk-neutral pricing formula we obtain

$$\Pi_Y(0) = e^{-rT}\mathbb{E}_q[Y] = e^{-rT}\mathbb{E}_q[g(S(T))] = \int_{\mathbb{R}} g(x)f_{S(T)}^{(q)}(x)\, dx$$

$$= \frac{e^{-rT}}{\sqrt{2\pi\sigma^2 t}} \int_0^\infty \frac{g(x)}{x} \exp\left(-\frac{(\log x - \log S_0 - (r-\frac{\sigma^2}{2})t)^2}{2\sigma^2 t}\right) dx.$$

With the change of variable $y = \frac{\log x - \log S_0 - (\alpha - \theta\sigma)t}{\sigma\sqrt{t}}$ we obtain

$$\Pi_Y(0) = e^{-rT} \int_{\mathbb{R}} g(S_0 e^{(r-\frac{\sigma^2}{2})T + \sigma\sqrt{T}y}) e^{-\frac{1}{2}y^2} \frac{dy}{\sqrt{2\pi}} = v_0(S_0),$$

as claimed. □

Example

For instance, in the case of the European call option with strike K and maturity T, for which the pay-off function is $g(z) = (z - K)_+$, Theorem 0.20 gives

$$\Pi_{\text{call}}(0) = C_0(S_0, K, T), \quad C_0(x, K, T) = x\Phi(d_1) - Ke^{-rT}\Phi(d_2)$$

where Φ is the standard normal distribution and

$$d_2 = \frac{\log \frac{x}{K} + (r - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \quad d_1 = d_2 + \sigma\sqrt{T}.$$

Remark

The risk-neutral pricing formula for $t > 0$ is

$$\Pi_Y(t) = e^{-r(T-t)}\mathbb{E}_q[Y|\mathcal{F}_S(t)],$$

The right hand side is the expectation of the discounted pay-off in the risk-neutral probability measure conditional to the information available at time t, which in a Black-Scholes market is determined by the history of the stock price up to time t.

It can be shown that in the case of the standard European derivative with pay-off $Y = g(S(T))$ at maturity T, the risk-neutral pricing formula entails that the Black-Scholes price at time $t \in [0, T]$ can be written in the integral form

$$\Pi_Y(t) = v(t, S(t)), \quad \text{where} \quad v(t,x) = \frac{e^{-r\tau}}{\sqrt{2\pi}} \int_{\mathbb{R}} g\left(xe^{(r-\frac{\sigma^2}{2})\tau}e^{\sigma\sqrt{\tau}y}\right) e^{-\frac{y^2}{2}} \, dy, \quad \tau = T-t.$$

Hence the pricing function $v(t,x)$ of the derivative at time t is the same as the pricing function v_0 at time $t = 0$ but with maturity T replaced by the time τ left to maturity, which is rather intuitive.

The Monte Carlo method

The Monte Carlo method is, in its simplest form, a numerical method to compute the expectation of a random variable.

Its mathematical validation is based on the **Law of Large Numbers**, which states the following: Suppose $\{X_i\}_{i\geq 1}$ is a sequence of i.i.d. random variables with expectation $\mathbb{E}[X_i] = \mu$. Then the sample average of the first n components of the sequence, i.e.,

$$\overline{X} = \frac{1}{n}(X_1 + X_2 + \cdots + X_n),$$

converges (in probability) to μ as $n \to \infty$.

The law of large numbers can be used to justify the fact that if we are given a large number of independent trials $X_1, \ldots, X_n$ of a random variable X, then

$$\mathbb{E}[X] \approx \frac{1}{n}(X_1 + X_2 + \cdots + X_n).$$

To measure how reliable is the approximation of $\mathbb{E}[X]$ given by the sample average, consider the standard deviation of the trials $X_1, \ldots, X_n$:

$$s_X = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(\overline{X} - X_i)^2}.$$

Viewing $X_1, \ldots, X_n$ as independent copies of X, a simple application of the Central Limit Theorem proves that the random variable

$$\frac{\mu - \overline{X}}{s_X/\sqrt{n}}$$

converges in distribution to a standard normal random variable. We use this result to show that the true value μ of $\mathbb{E}[X]$ has about 95% probability to be in the interval

$$[\overline{X} - 1.96\frac{s}{\sqrt{n}}, \overline{X} + 1.96\frac{s}{\sqrt{n}}].$$

Indeed, for n large,

$$\mathbb{P}\left(-1.96 \leq \frac{\mu - \overline{X}}{s_X/\sqrt{n}} \leq 1.96\right) \approx \int_{-1.96}^{1.96} e^{-x^2/2}\frac{dx}{\sqrt{2\pi}} \approx 0.95.$$

An application to Black-Scholes theory

Using the Monte Carlo method and the risk-neutral pricing formula, we can approximate the Black-Scholes price at time $t = 0$ of the European derivative with pay-off Y and maturity $T > 0$ with the sample average

$$\Pi_Y(0) = e^{-rT}\frac{Y_1 + \ldots Y_n}{n},$$

where $Y_1, \ldots, Y_n$ is a large number of independent trials of the pay-off. Each trial Y_i is determined by a path of the stock price.

Letting $0 = t_0 < t_1 < \cdots < t_N = T$ be a partition of the interval $[0, T]$ with size $t_i - t_{i-1} = h$, we may construct a sample of n paths of the geometric Brownian motion on the given partition with the following simple Matlab function:

```
function Path=StockPath(s,sigma,r,T,N,n)
h=T/N;
W=randn(n,N);
q=ones(n,N);
Path=s*exp((r-sigma^2/2)*h.*cumsum(q')+sigma*sqrt(h)*cumsum(W'));
Path=[s*ones(1,n);Path];
```

Note carefully that the stock price is modeled as a geometric Brownian motion with mean of log return $\alpha = r - \sigma^2/2$, which means that the geometric Brownian motion is risk-neutral. This is of course correct, since the expectation that we want to compute is in the risk-neutral probability measure.

The following Matlab code compute the Black-Scholes price of a call option using the Monte Carlo method. The code also computes the statistical error

$$\text{Err} = 1.96 \frac{s}{\sqrt{n}}$$

of the Monte Carlo price, where s is the standard deviation of the pay-off trials.

```
function [price, conf95]=MonteCarloCall(s,sigma,r,K,T,N,n)
tic
stockPath=StockPath(s,sigma,r,T,N,n);
payOff=max(0,stockPath(N,:)-K);
price=exp(-r*T)*mean(payOff);
conf95=1.96*std(payOff)/sqrt(n);
toc
```

For instance, by running the command

$$[\texttt{price}, \texttt{conf95}] = \texttt{MonteCarloCall}(10, 0.5, 0.01, 10, 1, 100, 100000)$$

we obtain the output

```
price = 1.9976
conf95 = 0.0249
```

The calculation took about half a second.

The exact price for the given call obtained by using the Black-Scholes formula is 2.0144, which lies within the confidence interval $[1.9976 - 0.0249, 1.9976 + 0.0249] = [1.9727, 2.0225]$ of the Monte Carlo price.

Control variate Monte Carlo

The Monte Carlo method just described is also known as **crude** Monte Carlo and can be improved in a number of ways.

In order to reduce the error of the Monte Carlo price, one needs to either

(i) increase the number of trials n or

(ii) reduce the standard derivation s.

As increasing n can be very costly in terms of computational time, the approach (ii) is preferable.

There exist several methods to decrease the standard deviation of a Monte Carlo computation, which are collectively called **variance reduction techniques**. Here we describe the **control variate** method.

Suppose we want to compute $\mathbb{E}[X]$. The idea of the control variate method is to introduce a second random variable Q for which $\mathbb{E}[Q]$ can be computed *exactly* and then write

$$\mathbb{E}[X] = \mathbb{E}[Y] + \mathbb{E}[Q], \quad \text{where } Y = X - Q.$$

Hence the Monte Carlo approximation of $\mathbb{E}[X]$ can now be written as

$$\mathbb{E}[X] \approx \frac{Y_1 + \cdots + Y_n}{n} + \mathbb{E}[Q],$$

where $Y_1, \ldots, Y_n$ are independent trials of the random variable Y.

This approximation improves the crude Monte Carlo estimate (i.e., without control variate) if the sample average estimator of $\mathbb{E}[Y]$ is better than the sample average estimator of $\mathbb{E}[X]$. Because of (**??**), this will be the case if $(s_Y)^2 < (s_X)^2$.

It will now be shown that the latter inequality holds if X, Q *have a positive large correlation.* Letting $X_1, \ldots, X_n$ be independent trials of X and $Q_1, \ldots, Q_n$ be independent trials of Q, we compute

$$(s_Y)^2 = \frac{1}{n-1} \sum_{i=1}^{n} (\overline{Y} - Y_i)^2 = \frac{1}{n-1} \sum_{i=1}^{n} ((\overline{X} - \overline{Q}) - (X_i - Q_i))^2$$
$$= (s_X)^2 + (s_Q)^2 - 2C(X, Q),$$

where $C(X, Q)$ is the sample covariance of the trials $(X_1, \ldots, X_n)$, $(Q_1, \ldots, Q_n)$, namely

$$C(X, Q) = \sum_{i=1}^{n} (\overline{X} - X_i)(\overline{Q} - Q_i).$$

Hence $(s_Y)^2 < (s_X)^2$ holds provided $C(X, Q)$ is sufficiently large and positive (precisely, $C(X, Q) > s_Q/\sqrt{2}$). As $C(X, Q)$ is an unbiased estimator of $\text{Cov}(X, Q)$, then the use of the control variate Q will improve the performance of the crude Monte Carlo method if X, Q have a positive large correlation.

This method is applied in the project on Asian options