Probability II

Meet Patel

Quantitative Finance Club

Quick Review

The set of all possible outcomes of an experiment is called the **sample space**, denoted Ω, and subsets of this set are called **events**, denoted ω. To each $\omega \in \Omega$, we assign a probability $\mathbb{P}(\Omega)$ such that:

Axioms

(a) For any event $\omega \subset \Omega$,

$$0 \leq \mathbb{P}(\omega) \leq 1$$

(b) $\mathbb{P}(\Omega) = 1$

(c) For a collection of mutually exclusive events $\omega_1, \omega_2, \omega_3, \dots,$

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} \omega_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(\omega_i)$$

What is a Random Variable?

Let Ω be the sample space of an experiment. A function $X : \Omega \to \mathbb{R}$ is called a **random variable** of the experiment. If X can take on countably many values, we say that it is **discrete**. If X can take on uncountably many values, we say that it is **continuous**.

"function $X : \Omega \to \mathbb{R}$":
This definition means that random variables are simply functions that take in events as input and output real numbers.

Examples of Random Variables:

(a) A coin toss.
(b) The sum of the outcomes of 5 dice rolls.
(c) The height of a person picked at random from a population.
(d) The price of a given stock one hour into the future.

Discrete Random Variables

If X is a discrete random variable, we can define the **probability mass function (p.m.f.)** of X as:

$$f_X(x) = \mathbb{P}(X = x)$$

We can also define the **cumulative distribution function (c.d.f.)** of X as:

$$F_X(x) = \sum_{a<x} \mathbb{P}(a)$$

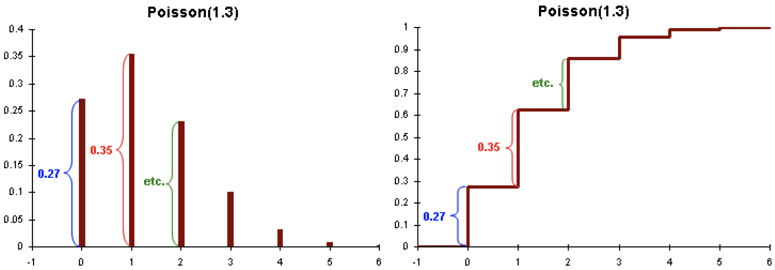

Continuous Random Variables

For continuous random variables, we have to be a little more careful. First, let's define the probability that the random variable will output a value in between two given values, a and b:

$$\mathbb{P}(a \leq X \leq b) = \int_a^b f_X(x)\,\mathrm{d}x$$

The function $f_X(x)$ is called the **probability density function (p.d.f.)**. Now, we can define the **cumulative distribution function (c.d.f.)** of X as: $F_X(x) = \int_{-\infty}^{x} f_X(x)\,\mathrm{d}x$

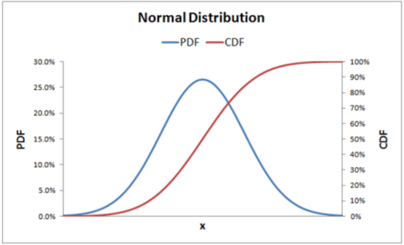

Expected Value

The expected value of a discrete random variable X is given by:

$$\mathbb{E}[X] = \sum x \cdot f_X(x)$$

Similarly, the expected value of a continuous random variable X is given by:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x \cdot f_X(x) \, \mathrm{d}x$$

We can think of the expected value as a probability-weighted average. We can think of this as the "most likely" outcome of the random variable X. Let's do an example.

Expected Value of Uniform Random Variable

The p.d.f. of this random variable is:

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{for } x \in (a, b) \\ 0 & \text{for } x \notin (a, b) \end{cases}$$

So, since this is a continuous random variable, let's plug this into the formula:

$$\begin{aligned} \mathbb{E}[X] &= \int_{-\infty}^{\infty} x \cdot f_X(x)\, \mathrm{d}x = \int_a^b \frac{x}{b-a}\, \mathrm{d}x \\ &= \frac{1}{b-a} \cdot \left(\frac{x^2}{2} \Big|_a^b \right) = \frac{1}{b-a} \left(\frac{b^2 - a^2}{2} \right) \\ &= \frac{1}{b-a} \left(\frac{(b+a)(b-a)}{2} \right) = \frac{b+a}{2} \end{aligned}$$

Variance of a Random Variable

If we are given a random variable X, we now know how to find $\mathbb{E}[X]$, which is essentially an "average value" for the random variable.

Now, we want to answer the question of: "If X were to output another value, how far off would it be from the expected value?"

This is called the **variance** of the random variable, denoted $\mathbb{V}[X]$ or σ_X^2, and the formula for it is:

$$\mathbb{V}[X] = \mathbb{E}\left[(X - \mathbb{E}[X])^2\right]$$

If we are thinking in terms of the graph for a p.d.f., the variance is a measure of how *wide* the distribution is.

The quanitity $\sigma_X = \sqrt{\mathbb{V}[X]}$ is called the **standard deviation** of X.

Normal Random Variables

We say that a random variable X is **normally distributed** with mean μ and variance σ^2, denoted $X \sim N(\mu, \sigma^2)$, if its p.d.f. is:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{ \frac{-1}{2} \left(\frac{x-\mu}{\sigma} \right)^2 \right\}$$

A very useful and common distribution is the *standard* normal distribution,

$N(0, 1)$. This is the normal distribution centered around 0 with variance 1. The p.d.f. of $N(0, 1)$ is:

$$f(x) = \frac{1}{\sqrt{2\pi}} \exp\left\{ \frac{-x^2}{2} \right\}$$

Limit Theorems

Central Limit Theorem: Let $X_1, X_2, \ldots$ be a sequence of independent and identically distributed random variables, each having a mean μ and variance σ^2. Then the distribution of:

$$Z := \frac{(X_1 + X_2 + \cdots + X_n - n\mu)}{\sigma\sqrt{n}}$$

tends to $N(0, 1)$ as $n \to \infty$.

Strong Law of Large Numbers: Let $X_1, X_2, \ldots$ be a sequence of independent and identically distributed random variables, each having a mean $\mu = \mathbb{E}[X_i] < \infty$. Then, as $n \to \infty$,

$$\frac{X_1 + X_2 + \cdots + X_n}{n} \to \mu$$

Recap

In this presentation, we used our knowledge of the fundamental laws of probabilty to discuss *random variables*, the most important of which was the *normal* random variable.

We defined the *expected value* and the *variance* of a random variable.

We concluded by looking at two very important "limit laws": the *Central Limit Theorem* and the *Strong Law of Large Numbers*.