BNAD 276
Lecture 6
Continuous Probability Distributions

Phuong Ho

May 14, 2017

Outline

Outline

- While a discrete random variable only takes clearly separated values, a continuous random variable can take any value in an interval.
- The followings are the example of continuous random variables.

Experiment	RV	Possible Values for the RV
Operate a bank	Time between customer arrivals	$x \geq 0$
Fill a soft drink can (Max=12 ounces)	Number of ounces	$0 \leq x \leq 12$
Construct a New library	Percentage of completion after 6 months	$0 \leq x \leq 100$

Continuous Random Variables

- One difference between a discrete RV and a continuous RV is the values that each RV can take.
- However, **the more important difference** is **the way we calculate probabilities**.
- Recall that for a discrete RV, X, we have the probability function, or **probability mass function**, $f(x)$.
- If we want to calculate $Prob(X = x)$, we can directly use $f(x)$.

e.g If we want to get $Prob(X = 2)$, we can get this by calculating $f(2)$.

Probability Density Function

- However, a continuous random variable does not have such a probability function.
- Instead, it has the **Probability Density Function**, $f(x)$.
- Though the **Probability Density Function** has the same notation ($f(x)$) as the **Probability Function** for a discrete random variable, these two are TOTALLY DIFFERENT!!
- A **probability density function**, $f(x) \neq P(X = x)$
- Instead, the area under the graph of $f(x)$ corresponding to a given interval does provide the probability that the continuous random variable X takes a value **in that interval**.

Example. Probability density function

Suppose we have a probability density function, $f(x)$, that has the following graph.

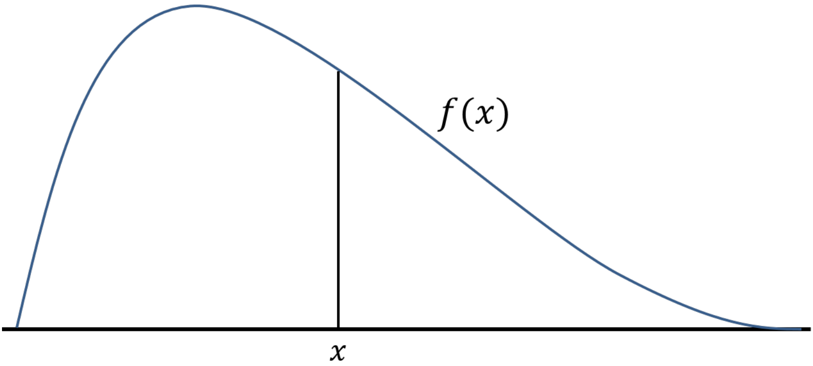

Example cont'd

If we calculate the shaded area below,

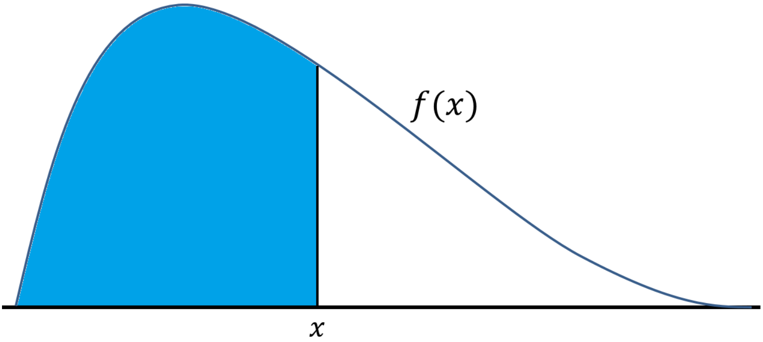

the calculated area will give us $Prob(X \leq x)$.

Continuous Probability Distributions

- There are two things we need to keep in mind when we work with a continuous random variable:
  1. $f(x)$ is NOT probability function but a probability density function. Thus, $f(1)$ or $f(100)$ do NOT have any meaning.
  2. Any probability that X takes a specific value is 0.
     i.e. $Prob(X = x) = 0.$
     because the area under the graph of $f(x)$ at any specific point/value is 0.

- Again, with continuous random variable, we're using a Probability Density Function, $f(x)$.
- Whenever we talk about the probability of a continuous random variable, it is an **AREA** under the graph of $f(x)$ (!!NOT the value of $f(x)$!!).

Requirements for the Probability Density Function

- As we saw before, an area under $f(X)$ will be a probability.
- Thus, it has 2 requirements to satisfy.

Requirements for the Probability Density Function

1. The whole area under $f(x)$ should be 1, i.e. $\int f(x)dx = 1$
2. $0 \leq f(x) \leq$ 1 for all x

- The range of x at which $f(x) > 0$ is called the **support** of the the random variable X, (or of pdf $f(x)$).
- Comment: the first requirement is analogous to the requirement that the sum of all probability should be equal to 1.

Outline

(Continuous) Uniform Distribution

- The uniform distribution is the simplest pdf of a continuous random variable.
- The name comes from the shape of the probability density function.

Uniform Probability Density Function

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

where a and b define the support of the random variable X.

- The support of the uniform variable X in this definition is the interval $[a, b]$.

Uniform Distribution

If we draw the probability density function for a uniform distribution, we will have the following graph.

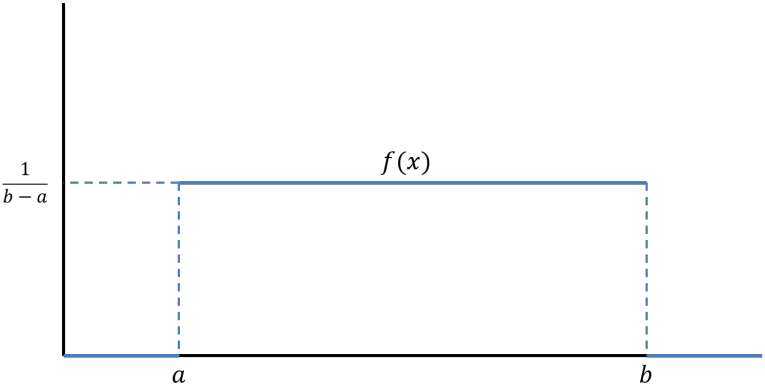

Uniform Distribution, cont'd

- For the relevant values of X (i.e. $a \leq X \leq b$), it has a uniformly horizontal line.
  - This is why this distribution is referred to a uniform distribution.
- The height of the horizonal line above the relevant values of X is $\frac{1}{b-a}$.
  - It comes from the requirement that the whole area below $f(x)$ should be equal to 1.
- For any other values outside of the range from a to b, $f(x)$ is equal to 0. In mathematical language, we write:

$$f(x) = 0 \; \forall x < a \text{ and } x > b$$

Recall. Area as a Measure of Probability

Area as a Measure of Probability

- As we saw before, we need to calculate an **AREA** under $f(x)$ to get a probability.
- When we work with a continuous random variable, an area under $f(x)$ is interpreted as a measure of probability.

Eg. Consider the random variable X representing the flight time of an airplane from Tucson to Los Angeles. It usually takes time from 120 min. to 140 min. Let's assume that X follows a uniform distribution.

Example cont'd

- First, we need to figure out the probability density function of X.
- We assumed that X follows a uniform distribution and we also know the support of X is $[120, 140]$ (from 120 min. to 140 min.)
- Given this information, we can infer the $f(x)$ for X as follows.

$$f(x) = \begin{cases} \frac{1}{140-120} = \frac{1}{20} & \text{for } 120 \leq x \leq 140 \\ 0 & \text{otherwise,} \end{cases}$$

- Given this, what is the probability that $X = 121$?
- Given this, what is the probability that $X = 140$?

Example cont'd

- As said before, the **probability** that X takes a **specific value** (even when the value is in the relevant range) is always **zero.**
- Thus, a meaningful question should be "what is the probability that values of X is in a range?"
- Say, what is the probability that $120 \leq X \leq 130$?
- To answer to this question, we need to calculate the area below $f(x)$.

Example cont'd

- The following is the graph of $f(x)$ in this example.

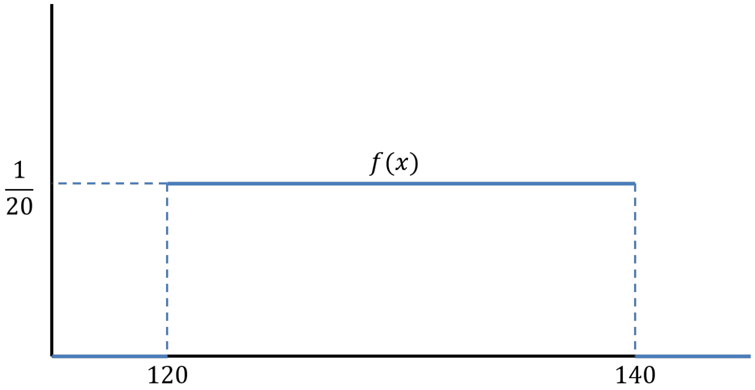

Example cont'd

To calculate the probability that $120 \leq X \leq 130$, we need to calculate the following shaded area.

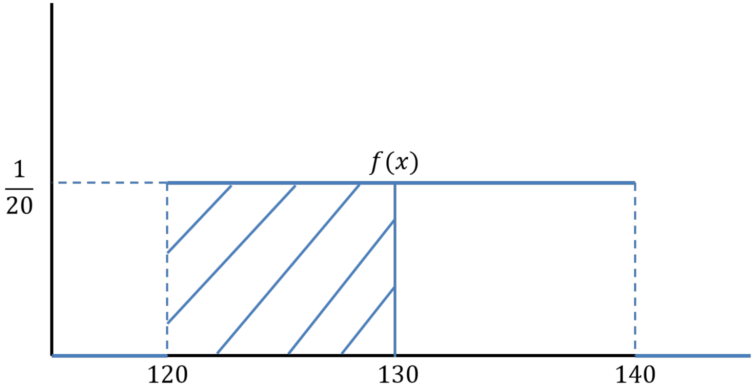

Example cont'd

- The area we want to calculate is a rectangle and we know the height and width of the rectangle.
- The height directly comes from the probability density function and the width is specified by the probability that we are interested in (i.e $P(120 \leq X \leq 130)$).
- Thus, $P(120 \leq X \leq 130) = (130 - 120) \times \frac{1}{20}$.
- What is $P(130 \leq X \leq 150)$?
- What is $P(125 \leq X \leq 1135)$?

Expected Value (Mean) and Variance of a Uniform Distribution

- Since a discrete random variable has its mean and variance, a continuous random variable also has its mean and variance.
- The mean and variance of a continuous random variable involve its probability density function.
- The way to calculate the mean and variance of a continuous random variable is analogous to the ones of a discrete random variable. However, continuous ones involve higher math skills (calculating integration).
- We will only look at the formula for the mean and variance for a random variable that follows a uniform distribution.

Expected Value (Mean) and Variance of a Uniform Distribution

Suppose that X follows a uniform distribution with the following pdf:

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

Then,

Expected Value (Mean) and Variance of a Uniform Distribution

$$E(X) = \mu = \frac{a+b}{2}$$

$$Var(X) = \sigma^2 = \frac{(b-a)^2}{12}$$

Expected Value (Mean) and the Variance of the Uniform Distribution

Note that when X follows a uniform distribution and its range of relevant values is from a to b,

- $E(X)$ is the midpoint between a and b.
- $Var(X)$ is the squared length of the support, $(b-a)^2$, over 12.
  - The number 12 comes from when we calculate $Var(X)$ by using the integration.
- Standard deviation:

$$\sigma = \frac{b-a}{2\sqrt{3}}$$

- The standard deviation σ_X is proportional to the length of the support.

Example. Uniform Random Variable

Recall the previous example of flight time. We have X representing the flight time between Tucson and Los Angeles.

- Since the relevant range of X is from 120 to 140, the mean of X is the midpoint between 120 and 140:

$$E(X) = \frac{120 + 140}{2} = 130,$$

- The variance of X is:

$$Var(X) = \frac{(140 - 120)^2}{12} = 1.667$$

- The standard deviation of X is

$$SD(X) = \sigma_X = \sqrt{Var(X)} = \sqrt{\sigma^2} = \sqrt{1.667} = 1.29$$

Exercise 1

A random variable X is known to be uniformly distributed between 1.0 and 1.5.

a. Show the graph of the probability density function.

b. Compute $Prob(X = 1.25)$

c. Compute $Prob(1.0 \leq X \leq 1.25)$

d. Compute $Prob(1.20 \leq X \leq 150)$

e. Compute $E(X)$ and $Var(X)$.

Outline

Normal Distribution

- The most important continuous probability distribution is the normal distribution.
- It is widely assumed in many contexts.

  e.g. Studies about heights or weights of people
- It is also important in which it is intensively used in the statistical inference.

Normal Distribution

- Suppose that our random variable X follows a normal distribution.
- When X follows a normal distribution, it can take any value between $-\infty$ and $+\infty$. That is, the support of X is $[-\infty, \infty]$
- X is a continuous random variable.
- Thus, it should have a probability density function for X.
- We have the following probability density function for a normal distribution.

Normal Density Function. Mean and Variance

Normal Density Function

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

where μ = mean, σ = s.d., $\pi = 3.14$, and $e = 2.72$.

- The first thing we need to notice is that once $f(x)$ is given, we can immediately know what is the mean and variance of X.

E.g. Suppose our X follows a normal distribution with the following $f(x)$.

$$f(x) = \frac{1}{2\sqrt{2\pi}} e^{-\left(\frac{1}{2}\right)\frac{(x-3)^2}{2^2}},$$

- What is μ and σ of X?

Shape of the Normal Density Function

Normal Density Function

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

where μ = mean, σ = s.d., $\pi = 3.14$, and $e = 2.72$.

- The second thing we need to know is how $f(x)$ looks like.
- Let's draw $f(x)$ step by step.
- First recall that $k^{-2} = \frac{1}{k^2}$ when k is a number.
- Now look at the part, $e^{-\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}$.
- We can rewrite this part as $\frac{1}{e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}}$.

Shape of the Normal Density Function. Mathematical Reasoning

Thus, we have

$$
\begin{aligned}
f(x) &= \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}} \\
&= \left(\frac{1}{\sigma\sqrt{2\pi}}\right)\left(\frac{1}{e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}}\right)
\end{aligned}
$$

- What is $f(x)$ if $x = \infty$?
- What is $f(x)$ if $x = -\infty$?
- What is the maximum value of $f(x)$?

Shape of the Normal Density Function. Mathematical Reasoning

- If X takes a huge value close to ∞, the denominator in
$$\left(\frac{1}{e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}}\right)$$
will be huge or really close to ∞.
- Then, $f(x)$ will be really close to zero.
- If X takes a huge negative value close to $-\infty$, the denominator in
$$\left(\frac{1}{e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}}\right)$$
will be huge or really close to ∞.
- Thus, $f(x)$ will be really close to zero too.

Shape of the Normal Density Function. Mathematical Reasoning

- The first finding is that as X takes large positive or negative value, $f(x)$ will be close to zero.
- Note that the largest value of $f(x)$ will be obtained when $e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}$ is smallest.
- The smallest value of $e^{\left(\frac{1}{2}\right)\frac{(x-\mu)^2}{\sigma^2}}$ will be 1 when $x = \mu$.
- Thus, $f(x)$ will have the largest value when $x = \mu$.
- The second finding is that $f(x)$ will have the largest value when $x = \mu$.

Shape of the Normal Density Function. Mathematical Reasoning

- The last thing we need to notice is that $f(x)$ is symmetric around $x = \mu$.

e.g. Thus, $f(\mu+2) = f(\mu-2)$

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

- Thus, the graph of $f(x)$ will be symmetric.
- Now, summarize all information we have got.

Important Characteristics of the Shape of a Normal Density Function

- $f(x)$ has the largest value at $x = \mu$.
- $f(x)$ is decreasing as x is getting far from μ.
- $f(x)$ is be close to zero when x takes huge positive or negative values.
- $f(x)$ is symmetric around $x = \mu$.

Shape of a Normal Density Function

Normal Density Function

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

where μ = mean, σ = s.d., $\pi = 3.14$, and $e = 2.72$.

The whole area below $f(x)$ will be 1.

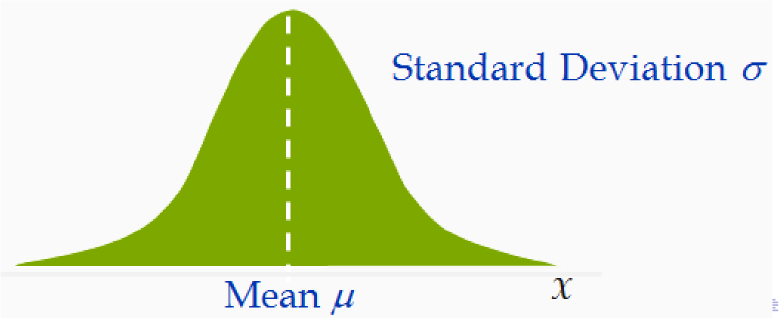

Change in the mean

1. It is centered and symmetric around its mean μ.
   Mean $=$ Median $=$ Mode $= \mu$
   μ can be any value. For example, it can be -10, 0, or 15.

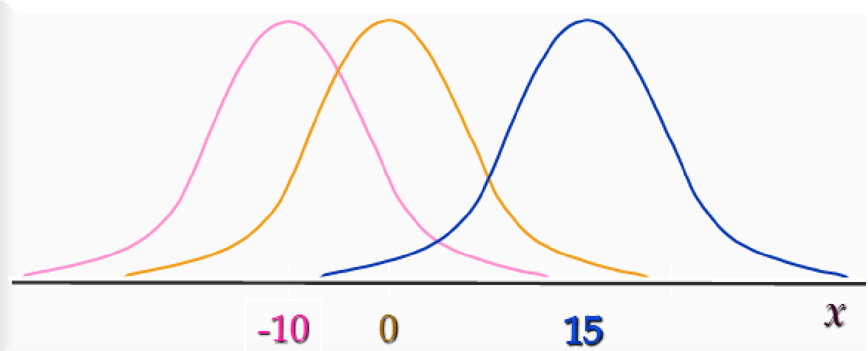

Change in the standard deviation

- The standard deviation determines how flat and wide the normal density function is.

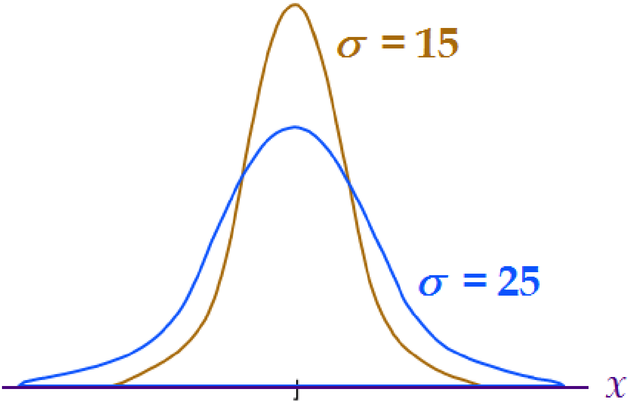

Calculating Probabilities of an Area

- So fare we didn't discuss any thing about how to calculate probabilities that X is in an interval when X follows a normal distribution.
- To calculate probabilities that X is in an interval, we first should know a special normal distribution which is called **the Standard Normal Distribution**

Standard Normal Distribution

Standard Normal Distribution

A random variable X that has a normal density function with $\mu = 0$ and $\sigma = 1$ is said to follow the standard normal distribution.

Standard Normal Density Function

$$f(z) = \frac{1}{\sqrt{2\pi}} e^{-\left(\frac{1}{2}\right)z^2},$$

- The standard normal distribution is nothing but a normal distribution with $\mu = 0$ and $\sigma = 1$.
- It is popular to use Z for the standard normal random variable and z for the possible values that Z can take.

Standard Normal Density Distribution

- The graph of $f(z)$ is as follows. It is centered around $\mu = 0$.

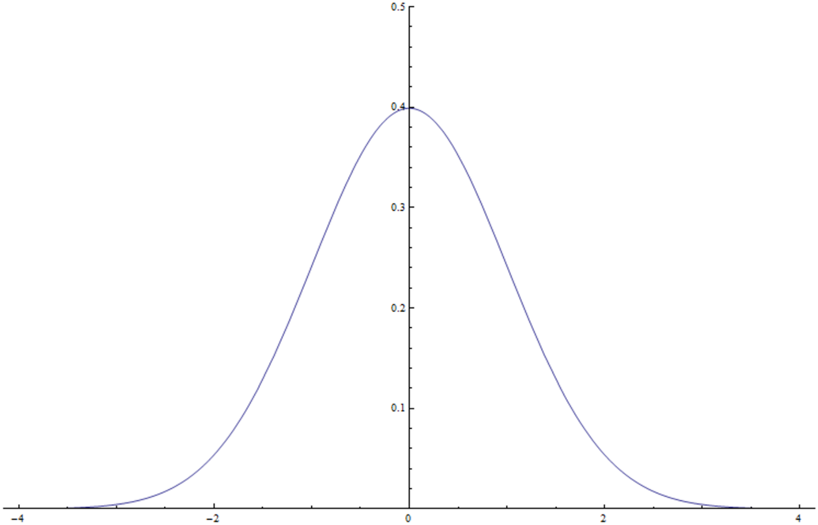

Standard Normal Distribution. Probability of an Area

- Let's calculate the probability that Z takes values less than 1.
- Graphically, the probability we are interested in is the shaded area in the following graph.

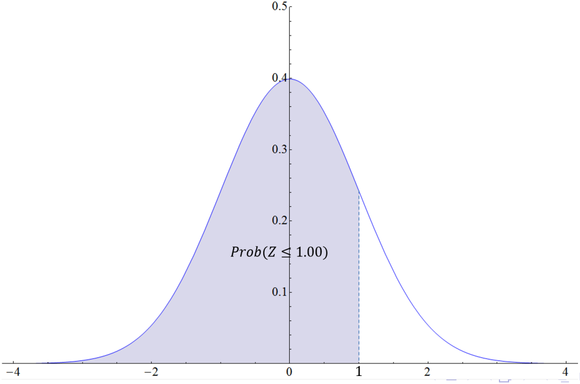

Standard Normal Distribution. Probability of an Area

- However, we have an area which is not square or triangle.
- Thus, it is impossible to calculate this area without higher math skills.
- Usually, we are given a table that summarizes numerous probabilities of different areas.
- The table is called Z-table, and it contains numerous probabilities that Z is less than different values z.
- The Z-table contains cumulative probabilities such as $P(Z \leq z_0)$ for some z_0.

Standard Normal Distribution. Z-table

- The following table is the part of the whole table that we need to use to find $P(Z \leq 1.00)$.

z	.00	0.01	.02	$\cdots$	0.05
-.5	.3085	.3959	.3915	$\cdots$	.2912
.	.	.	.	$\cdots$	.
1.0	.8413	.8438	.8461	$\cdots$	.8531
.	.	.	.	$\cdots$	.
1.5	.9332	..9343	,9357	$\cdots$	.9394

- All the probabilities in the above is a probability that Z is less than some number between $-\infty$ and ∞, (i.e. $P(Z \leq$ a number$)$).

z	.00	0.01	.02	$\cdots$	0.05
-.5	.3085	.3959	.3915	$\cdots$	.2912
.	.	.	.	$\cdots$	.
1.0	.8413	.8438	.8461	$\cdots$	.8531
.	.	.	.	$\cdots$	.
1.5	.9332	..9343	,9357	$\cdots$	.9394

- Let's find $P(Z \leq 1.00)$.
- At the first column, there are numbers up to the first decimal place.
- The first row in the table indicates the second decimal place of a number that in the first column.
- We are interested in 1.00. Start from 1.0 at the first column and stop at the second column.
- The number we found is .8413. And, it is the probability we want to calculate, $P(Z \leq 1.00)$.

z	.00	0.01	.02	$\cdots$	0.05
-.5	.3085	.3959	.3915	$\cdots$	.2912
.	.	.	.	$\cdots$	.
1.0	.8413	.8438	.8461	$\cdots$	.8531
.	.	.	.	$\cdots$	.
1.5	.9332	..9343	,9357	$\cdots$	.9394

- Let's find $P(Z \leq -0.55)$.
- The part up to the first decimal place in -0.55 is -0.5. Find -.5 in the first column.
- Proceed to the right until we hit the column of 0.05, and stop there. The number in that cell is $P(Z \leq -0.55)$.
- Thus, $P(Z \leq -0.55) = 0.2912$.

- Sometimes, we want to get a probability that Z is between two numbers such as $P(-0.50 \leq Z \leq 1.5)$.
- Then, how to calculate this kind of probability?
- Graphically, what we want to find is blue shaded area in the following graph.

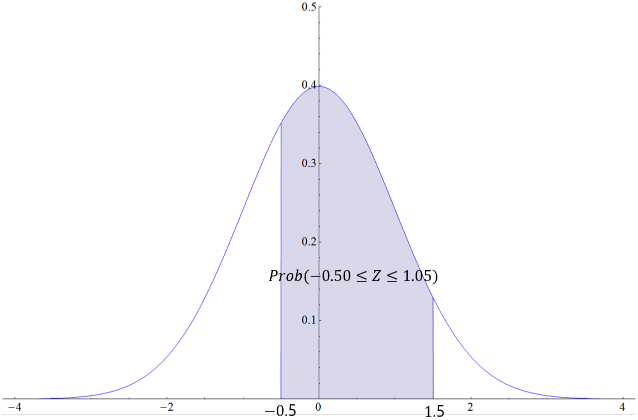

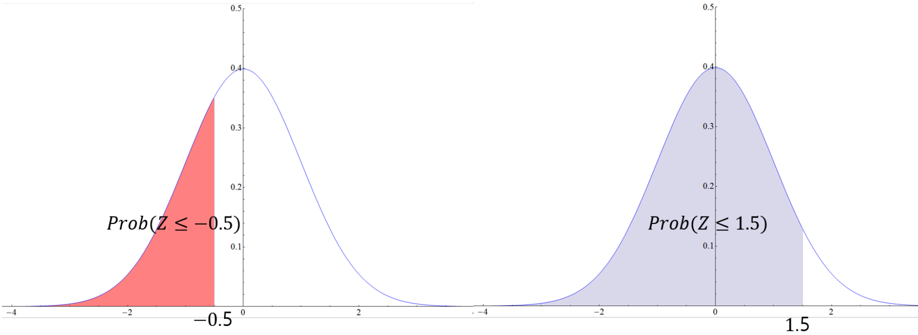

- The blue shaded area is nothing but $P(Z \leq 1.5)$
- The red shaded area is nothing but $P(Z \leq -0.50)$
- If we subtract the red shaded area from the blue shaded area, we get the area we wanted to have.
- Thus, $P(-0.50 \leq Z \leq 1.5) = P(Z \leq 1.5) - P(Z \leq -0.50)$.

- Finally, we could want to find a probability that Z is greater than a number such as $P(Z > 1.05)$.
- Graphically, $P(Z > 1.05)$ is the blue shaded area in the following graph.

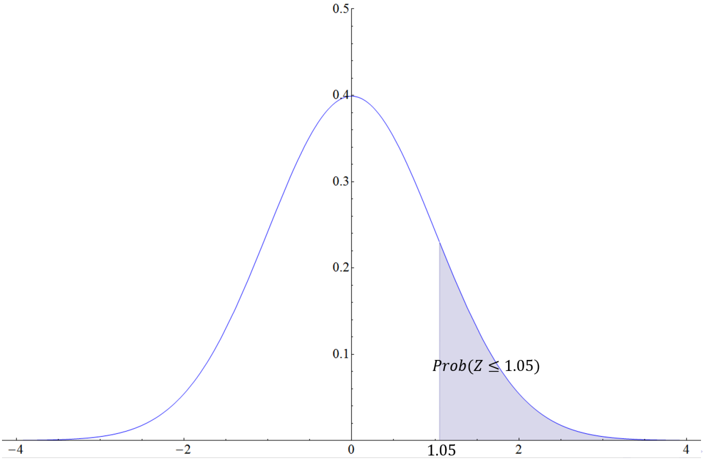

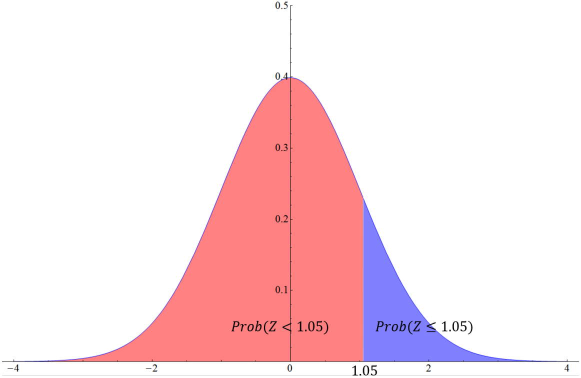

- We know that the whole area below $f(x)$ should be equal to 1.
- We also know the red shaded area is $P(Z \leq 1.05)$.
- Thus,

$$\begin{aligned} P(Z \geq 1.05) &= \text{The blue shaded area} \\ &= \text{1-the red shaded area} \\ &= 1 - P(Z \leq 1.05) \end{aligned}$$

Calculate Probabilities for Any Normal Distribution

- As explained before, we have infinitely many normal distribution as μ and σ can be any value.
- We can't have infinitely many tables. In fact, the table for the standard normal distribution (Z-table) is the only table we have.
- Thus, we need to use Z-table to calculate probabilities for any normal distributions, which means we need to know how to convert a normal random variable, X, to the standard normal random variable Z.

Convert to the Standard Normal Random Variable

Convert to the Standard Normal Random Variable

Suppose we have a random variable X that follows a normal distribution with mean μ and standard deviation σ. Then,

$$Z = \frac{X - \mu}{\sigma}$$

will follow the standard normal distribution.

- By subtracting the μ from X, we will have mean zero.
- By dividing $(X - \mu)$ by σ, the standard deviation of Z will be 1.
- Thus, Z will follow the standard normal distribution.

Calculate Probabilities for Any Normal Distribution

- Suppose X follows a normal distribution with μ and σ.
- Then, what is $P(-10 \leq X \leq 20)$?
- At first, we need to convert X to Z so that we can use Z-table.

$$\begin{aligned} P(-10 \leq X \leq 20) &= P(-10 - \mu \leq X - \mu \leq 20 - \mu) \\ &= P\left(\frac{-10 - \mu}{\sigma} \leq \frac{X - \mu}{\sigma} \leq \frac{20 - \mu}{\sigma}\right) \\ &= P\left(\frac{-10 - \mu}{\sigma} \leq Z \leq \frac{20 - \mu}{\sigma}\right) \end{aligned}$$

- Now we can calculate $P\left(\frac{-10-\mu}{\sigma} \leq Z \leq \frac{20-\mu}{\sigma}\right)$ and it is equal to $P(-10 \leq X \leq 20)$ by the construction in the above.

Example

Suppose that X follows the normal distribution with $\mu = 36500$ and $\sigma = 5000$. What is $P(30100 \leq X \leq 40000)$?

$$\begin{aligned} P(30100 \leq & X & \leq 40000) \\ & = & P(30100 - 36500 \leq X - 36500 \leq 40000 - 36500) \\ & = & P\left(\frac{30100 - 36500}{5000} \leq \frac{X - 36500}{5000} \leq \frac{40000 - 36500}{5000}\right) \\ & = & P(-1.28 \leq Z \leq 0.70) \end{aligned}$$

Outline

Exercise 2

The driving distance for the top golfers on the PGA tour is between 284 and 310 yard. Assume that the driving distance for these golfers is uniformly distributed over this interval.

a. Give a mathematical expression for the probability density function of the driving distance.

b. What is the probability that the driving distance for one of these golfers is less than 290 yard?

c. What is the probability that the driving distance for one of theses golfers is between 290 and 305 yard?

d. What is the probability that the driving distance for one of these golfers is at least 300 yard?

e. What is the expected driving distance of one of these golfers?

Exercise 3

Present graph of the standard normal density function with shaded area that represents each probability below. Then compute each probability.

a. $P(Z \leq 1.5)$

b. $P(Z \leq 1)$

c. $P(1 \leq Z \leq 1.5)$

d. $P(0 \leq Z \leq 2.5)$

Exercise 4

The time needed to complete a final exam in a particular college course is normally distributed with a mean of 80 min. and a standard deviation of 10 min. Answer the following questions.

a. What is the probability of completing the exam in one hour or less?

b. What is the probability that a student will complete the exam in more than 60 min but less than 75 min?

Exercise 5

The average stock price for companies making up the S&P 500 is $ 30, and the standard deviation is $ 8. Assume the stock prices are normally distributed.

a. What is the probability a company will have a stock price of at least $ 40?

b. What is the probability a company will have a stock price no higher than $20?

c. How high does a stock price have to be to put a company in the top 10%?