Random Variables

Nate Wells

Math 141, 4/1/22

Outline

In this lecture, we will. . .

Outline

In this lecture, we will...

- Define and explore Random Variables
- Investigate properties of the Normal Distribution

Section 1

Random Variables

Definitions

A **random variable** is a numeric quantity whose value depends on the result of a random process.

Definitions

A **random variable** is a numeric quantity whose value depends on the result of a random process.

- We use capital letters at the end of the alphabet (W, X, Y, Z) to denote random variables.
  - We use lowercase letters (w, x, y, z) to denote the particular values of a random variable

Definitions

A **random variable** is a numeric quantity whose value depends on the result of a random process.

- We use capital letters at the end of the alphabet (W, X, Y, Z) to denote random variables.
  - We use lowercase letters (w, x, y, z) to denote the particular values of a random variable
- We use equations to express events associated to random variables.
  - I.e "$X = 5$" represents the event "The random variable X takes the value 5".

Definitions

A **random variable** is a numeric quantity whose value depends on the result of a random process.

- We use capital letters at the end of the alphabet (W, X, Y, Z) to denote random variables.
  - We use lowercase letters (w, x, y, z) to denote the particular values of a random variable
- We use equations to express events associated to random variables.
  - I.e "$X = 5$" represents the event "The random variable X takes the value 5".
- Events associated to variables have probabilities of occurring.
  - $P(X = 5) = .5$ means X has 50% probability of taking the value 5.

Types of Random Variables

There are two main types of random variables:

1. **Discrete** variables can take only finitely many different values.
2. **Continuous** variables can take values equal to any real number in an interval.

Types of Random Variables

There are two main types of random variables:

1. **Discrete** variables can take only finitely many different values.
2. **Continuous** variables can take values equal to any real number in an interval.

- Examples of discrete variables:
  - The number of credits a randomly chosen Reed student is taking.
  - The number of vegetarians in a random sample of 10 people.
  - The results of a coin flip, where 0 indicates Tails and 1 indicates Heads.

Types of Random Variables

There are two main types of random variables:

1. **Discrete** variables can take only finitely many different values.
2. **Continuous** variables can take values equal to any real number in an interval.

- Examples of discrete variables:
  - The number of credits a randomly chosen Reed student is taking.
  - The number of vegetarians in a random sample of 10 people.
  - The results of a coin flip, where 0 indicates Tails and 1 indicates Heads.
- Examples of continuous variables:
  - The temperature of my office at a particular time of the day.
  - The amount of time it takes a radioactive particle to decay.

Types of Random Variables

There are two main types of random variables:

1. **Discrete** variables can take only finitely many different values.
2. **Continuous** variables can take values equal to any real number in an interval.

- Examples of discrete variables:
  - The number of credits a randomly chosen Reed student is taking.
  - The number of vegetarians in a random sample of 10 people.
  - The results of a coin flip, where 0 indicates Tails and 1 indicates Heads.
- Examples of continuous variables:
  - The temperature of my office at a particular time of the day.
  - The amount of time it takes a radioactive particle to decay.
- Some discrete variables can be well-described by continuous variables:
  - The height of a random person selected from a large population.
  - The proportion of heads in a long sequence of coin flips.

The Distribution of a Random Variable

- Recall that *data* variables have distributions, which tell us. . .
  - the values the variable takes, and the *frequency* the variable takes those values.

The Distribution of a Random Variable

- Recall that *data* variables have distributions, which tell us. . .
  - the values the variable takes, and the *frequency* the variable takes those values.
- But *random* variables also have distributions, which tell us. . .
  - the values the variable can take, and the *probability* the variable takes those values.

The Distribution of a Random Variable

- Recall that *data* variables have distributions, which tell us. . .
  - the values the variable takes, and the *frequency* the variable takes those values.
- But *random* variables also have distributions, which tell us. . .
  - the values the variable can take, and the *probability* the variable takes those values.
- Suppose I play a casino game, where that the amount of money I win (in cents) has the following distribution:

value	1	5	10	25
probability	.3	.4	.2	.1

The Distribution of a Random Variable

- Recall that *data* variables have distributions, which tell us...
  - the values the variable takes, and the *frequency* the variable takes those values.
- But *random* variables also have distributions, which tell us...
  - the values the variable can take, and the *probability* the variable takes those values.
- Suppose I play a casino game, where that the amount of money I win (in cents) has the following distribution:

value	1	5	10	25
probability	.3	.4	.2	.1

- Suppose instead that I have a purse filled with the following 100 coins:

value	1	5	10	25
frequency	30	40	20	10

The Distribution of a Random Variable

- Recall that *data* variables have distributions, which tell us. . .
  - the values the variable takes, and the *frequency* the variable takes those values.
- But *random* variables also have distributions, which tell us. . .
  - the values the variable can take, and the *probability* the variable takes those values.
- Suppose I play a casino game, where that the amount of money I win (in cents) has the following distribution:

value	1	5	10	25
probability	.3	.4	.2	.1

- Suppose instead that I have a purse filled with the following 100 coins:

value	1	5	10	25
frequency	30	40	20	10

- Playing the casino game is very similar to drawing a random coin from the purse.

Visualizing Discrete Distributions

- We often use bar charts to visualize the distribution of discrete random variables.

Visualizing Discrete Distributions

- We often use bar charts to visualize the distribution of discrete random variables.
  - Suppose a fair 6-sided die is rolled 6 times. Let X be the number of 1s rolled. The distribution of X is given by:

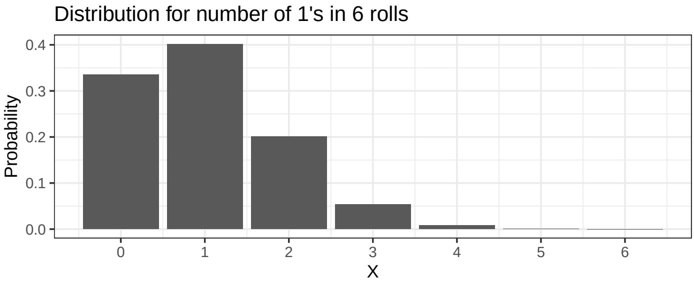

Visualizing Discrete Distributions

- We often use bar charts to visualize the distribution of discrete random variables.
  - Suppose a fair 6-sided die is rolled 6 times. Let X be the number of 1s rolled. The distribution of X is given by:

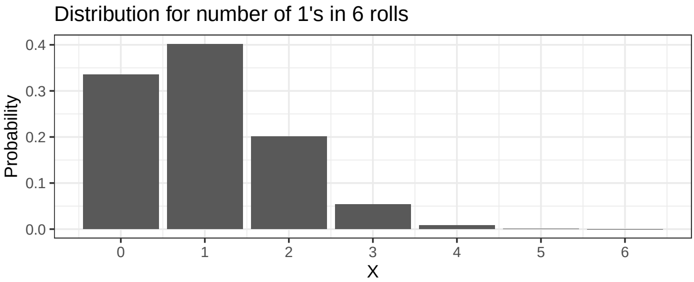

- Heights of bars are probabilities
  - This is analogous to rescaling a histogram to have heights equal to proportions, rather than counts

Expected Value

The **expected value** (or mean) of a discrete random variable X is

$$E[X] = x_1 P(X = x_1) + x_2 P(X = x_2) + \cdots + x_n P(X = x_n) = \sum_{i=1}^{n} x_i P(X = x_i)$$

Expected Value

The **expected value** (or mean) of a discrete random variable X is

$$E[X] = x_1 P(X = x_1) + x_2 P(X = x_2) + \cdots + x_n P(X = x_n) = \sum_{i=1}^{n} x_i P(X = x_i)$$

- The expected value of X is the sum of the value X can take, weighted by the probability it takes those values.

Expected Value

The **expected value** (or mean) of a discrete random variable X is

$$E[X] = x_1 P(X = x_1) + x_2 P(X = x_2) + \cdots + x_n P(X = x_n) = \sum_{i=1}^{n} x_i P(X = x_i)$$

- The expected value of X is the sum of the value X can take, weighted by the probability it takes those values.
- Suppose we have a data set consisting of values $\{1, 1, 2, 2, 2, 2, 3, 4, 5, 5\}$. Let X be a value chosen from this data set randomly. What is the expected value of X?

Expected Value

The **expected value** (or mean) of a discrete random variable X is

$$E[X] = x_1 P(X = x_1) + x_2 P(X = x_2) + \cdots + x_n P(X = x_n) = \sum_{i=1}^{n} x_i P(X = x_i)$$

- The expected value of X is the sum of the value X can take, weighted by the probability it takes those values.
- Suppose we have a data set consisting of values $\{1, 1, 2, 2, 2, 2, 3, 4, 5, 5\}$. Let X be a value chosen from this data set randomly. What is the expected value of X?

$$\begin{aligned} E[X] =& 1P(X = 1) + 2P(X = 2) + 3P(X = 3) + 4P(X = 4) + 5P(X = 5) \\ =& 1\frac{2}{10} + 2\frac{4}{10} + 3\frac{1}{10} + 4\frac{1}{10} + 5\frac{2}{10} = \frac{27}{10} = 2.7 \end{aligned}$$

Expected Value

The **expected value** (or mean) of a discrete random variable X is

$$E[X] = x_1 P(X = x_1) + x_2 P(X = x_2) + \cdots + x_n P(X = x_n) = \sum_{i=1}^{n} x_i P(X = x_i)$$

- The expected value of X is the sum of the value X can take, weighted by the probability it takes those values.
- Suppose we have a data set consisting of values $\{1, 1, 2, 2, 2, 2, 3, 4, 5, 5\}$. Let X be a value chosen from this data set randomly. What is the expected value of X?

$$\begin{aligned} E[X] =& 1P(X = 1) + 2P(X = 2) + 3P(X = 3) + 4P(X = 4) + 5P(X = 5) \\ =& 1\frac{2}{10} + 2\frac{4}{10} + 3\frac{1}{10} + 4\frac{1}{10} + 5\frac{2}{10} = \frac{27}{10} = 2.7 \end{aligned}$$

- But also notice that

$$\begin{aligned} E[X] =& \frac{1}{10}(1 \cdot 2 + 2 \cdot 4 + 3 \cdot 1 + 4 \cdot 1 + 5 \cdot 2) \\ =& \frac{1}{10}(1 + 1 + 2 + 2 + 2 + 2 + 3 + 4 + 5 + 5) \end{aligned}$$

The Law of Large Numbers, again

Previously, we said that by the Law of Large numbers, the proportion of times an outcome occurs in a long sequence of trials is close to the probability for that outcome.

The Law of Large Numbers, again

Previously, we said that by the Law of Large numbers, the proportion of times an outcome occurs in a long sequence of trials is close to the probability for that outcome.

This is a generalization:

Theorem (The Law of Large Numbers)

Let X be a random variable whose value depends on a random experiment. Suppose the experiment is repeated n times and let $\bar{x}_n$ denote the arithmetic mean of the values of X in each trial. As n gets larger, the arithmetic mean $\bar{x}_n$ approaches the expected value $E[X]$ of that variable.

A Roll of the Die

Suppose we roll a fair 6-sided die. What is the expected value of the result?

A Roll of the Die

Suppose we roll a fair 6-sided die. What is the expected value of the result?

- Suppose we roll the same die 1000 times and keep track of the running arithmetic mean of the results. . .

A Roll of the Die

Suppose we roll a fair 6-sided die. What is the expected value of the result?

- Suppose we roll the same die 1000 times and keep track of the running arithmetic mean of the results. . .

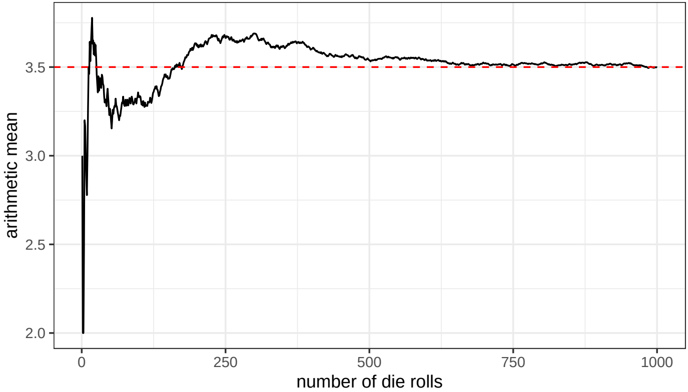

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =&(x_1-\mu)^2P(X=x_1)+(x_2-\mu)^2P(X=x_2)+\cdots+(x_n-\mu)^2P(X=x_n)\\ =&\sum_{i=1}^{n}(x_i-\mu)^2P(X=x_i)\end{aligned}$$

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =& (x_1 - \mu)^2 P(X = x_1) + (x_2 - \mu)^2 P(X = x_2) + \cdots + (x_n - \mu)^2 P(X = x_n) \\ =& \sum_{i=1}^{n} (x_i - \mu)^2 P(X = x_i)\end{aligned}$$

- The variance of X is the sum the squared deviations of X from its mean μ, weighted by the corresponding probabilities.

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =& (x_1-\mu)^2P(X=x_1)+(x_2-\mu)^2P(X=x_2)+\cdots+(x_n-\mu)^2P(X=x_n) \\ =& \sum_{i=1}^{n}(x_i-\mu)^2P(X=x_i)\end{aligned}$$

- The variance of X is the sum the squared deviations of X from its mean μ, weighted by the corresponding probabilities.
  - Variables with low variance tend have values close to the mean, while those with high variance tend to have values farther from the mean.

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =& (x_1-\mu)^2 P(X=x_1) + (x_2-\mu)^2 P(X=x_2) + \cdots + (x_n-\mu)^2 P(X=x_n) \\ =& \sum_{i=1}^{n} (x_i-\mu)^2 P(X=x_i)\end{aligned}$$

- The variance of X is the sum the squared deviations of X from its mean μ, weighted by the corresponding probabilities.
  - Variables with low variance tend have values close to the mean, while those with high variance tend to have values farther from the mean.
- As with data variables, we define the **standard deviation** of a random variable X to be

$$\mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$$

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =& (x_1 - \mu)^2 P(X = x_1) + (x_2 - \mu)^2 P(X = x_2) + \cdots + (x_n - \mu)^2 P(X = x_n) \\ =& \sum_{i=1}^{n} (x_i - \mu)^2 P(X = x_i)\end{aligned}$$

- The variance of X is the sum the squared deviations of X from its mean μ, weighted by the corresponding probabilities.
  - Variables with low variance tend have values close to the mean, while those with high variance tend to have values farther from the mean.
- As with data variables, we define the **standard deviation** of a random variable X to be

$$\mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$$

- We often use σ^2 to denote the variance and σ to denote the standard deviation of a variable.

Variance and Standard Deviation

The **variance** of a discrete random variable X with mean μ is

$$\begin{aligned}\mathrm{Var}(X) =& (x_1-\mu)^2P(X=x_1)+(x_2-\mu)^2P(X=x_2)+\cdots+(x_n-\mu)^2P(X=x_n)\\ =& \sum_{i=1}^{n}(x_i-\mu)^2P(X=x_i)\end{aligned}$$

- The variance of X is the sum the squared deviations of X from its mean μ, weighted by the corresponding probabilities.
  - Variables with low variance tend have values close to the mean, while those with high variance tend to have values farther from the mean.
- As with data variables, we define the **standard deviation** of a random variable X to be

$$\mathrm{SD}(X)=\sqrt{\mathrm{Var}(X)}$$

- We often use σ^2 to denote the variance and σ to denote the standard deviation of a variable.
- Compute the standard deviation for a fair coin flip.

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
    - But for continuous variables, there are too many possible values to provide a meaningful probability for each.

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
    - But for continuous variables, there are too many possible values to provide a meaningful probability for each.
- Instead, we describe the probabilities that continuous variables are in certain ranges of values, specified by a **density curves**

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
    - But for continuous variables, there are too many possible values to provide a meaningful probability for each.
- Instead, we describe the probabilities that continuous variables are in certain ranges of values, specified by a **density curves**
- The **density curve** for a continuous random variable X is the function f so that...

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
  - But for continuous variables, there are too many possible values to provide a meaningful probability for each.
- Instead, we describe the probabilities that continuous variables are in certain ranges of values, specified by a **density curves**
- The **density curve** for a continuous random variable X is the function f so that...
  - The values of the function are always non-negative

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
  - But for continuous variables, there are too many possible values to provide a meaningful probability for each.
- Instead, we describe the probabilities that continuous variables are in certain ranges of values, specified by a **density curves**
- The **density curve** for a continuous random variable X is the function f so that...
  - The values of the function are always non-negative
  - The total area under the function is 1

The Distribution of a Continuous Variable

- Recall: A continuous random variable is one that any value in an interval of real numbers.
- Previously, we described the distribution of *discrete* random variables by listing the probabilities of taking each possible values.
  - But for continuous variables, there are too many possible values to provide a meaningful probability for each.
- Instead, we describe the probabilities that continuous variables are in certain ranges of values, specified by a **density curves**
- The **density curve** for a continuous random variable X is the function f so that...
  - The values of the function are always non-negative
  - The total area under the function is 1
  - The area over any interval is the probability that the variable is in that interval.

Density Curve

- The density curve for the number of seconds T until a radioactive particle decays is:

$$f(t) = e^{-t} \qquad \text{for } t \geq 0$$

Density Curve

- The density curve for the number of seconds T until a radioactive particle decays is:

$$f(t) = e^{-t} \qquad \text{for } t \geq 0$$

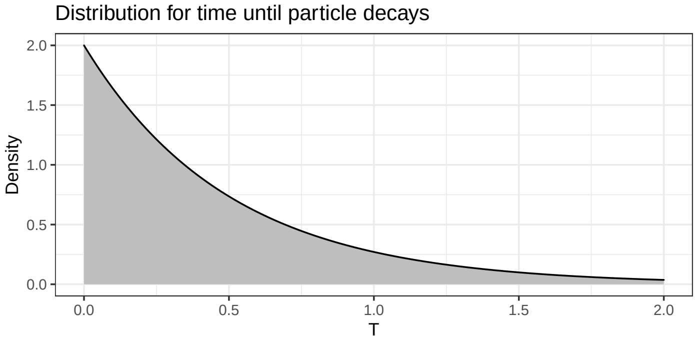

Density Curve

- The density curve for the number of seconds T until a radioactive particle decays is:

$$f(t) = e^{-t} \qquad \text{for } t \geq 0$$

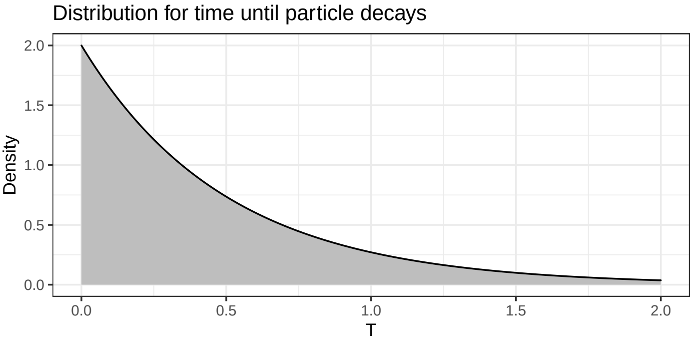

- The probability that it takes between 0.5 and 1.5 seconds to decay is the area under the curve between 0.5 and 1.5. $P(0.5 < T < 1.5) =$

Density Curve

- The density curve for the number of seconds T until a radioactive particle decays is:

$$f(t) = e^{-t} \qquad \text{for } t \geq 0$$

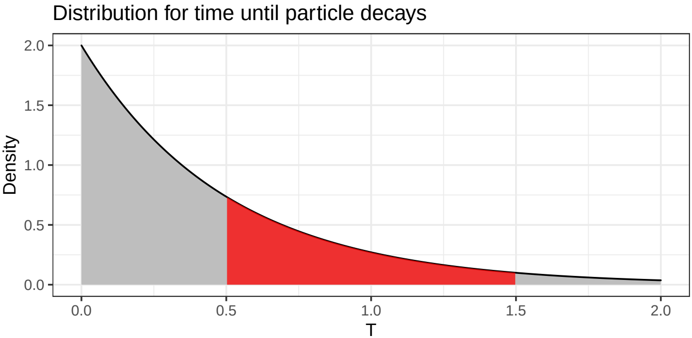

- The probability that it takes between 0.5 and 1.5 seconds to decay is the area under the curve between 0.5 and 1.5. $P(0.5 < T < 1.5) = 0.34$

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.
    - But we cannot use the same definition as before (the sum of values, weighted by the probability of each value)

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.
  - But we cannot use the same definition as before (the sum of values, weighted by the probability of each value)
  - Note that for any real number c, $P(X = c) = 0$. (Why?)

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.
  - But we cannot use the same definition as before (the sum of values, weighted by the probability of each value)
  - Note that for any real number c, $P(X = c) = 0$. (Why?)
- Instead, we use the *integral* from calculus to define the mean and variance:

$$E[X] = \int xf(x)\,dx \qquad \mathrm{Var}(X) = \int (x-\mu)^2 f(x)\,dx \qquad \mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$$

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.
  - But we cannot use the same definition as before (the sum of values, weighted by the probability of each value)
  - Note that for any real number c, $P(X = c) = 0$. (Why?)
- Instead, we use the *integral* from calculus to define the mean and variance:

$$E[X] = \int xf(x)\,dx \qquad \mathrm{Var}(X) = \int (x-\mu)^2 f(x)\,dx \qquad \mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$$

  - These **integrals** are tools to meaningfully average infinitely many values (but we won't compute any integrals in this class)

Mean and Standard Deviation for Continuous Variables

- Just as with discrete random variables, we can define the mean, variance and standard deviations of continuous variables.
  - But we cannot use the same definition as before (the sum of values, weighted by the probability of each value)
  - Note that for any real number c, $P(X = c) = 0$. (Why?)
- Instead, we use the *integral* from calculus to define the mean and variance:

$$E[X] = \int xf(x)\,dx \qquad \mathrm{Var}(X) = \int (x-\mu)^2 f(x)\,dx \qquad \mathrm{SD}(X) = \sqrt{\mathrm{Var}(X)}$$

  - These **integrals** are tools to meaningfully average infinitely many values (but we won't compute any integrals in this class)
- As with discrete variables, the mean of a continuous variables represents its typical value. The standard deviation represents the typical size of deviations from the mean.

Using Densities for Discrete Variables

If a discrete variable takes a large number of values which are close together, we can often approximate it using a continuous variable.

Using Densities for Discrete Variables

If a discrete variable takes a large number of values which are close together, we can often approximate it using a continuous variable.

- Suppose 500 students take a standardized exam, with mean 75 points. The distribution for the score S of a randomly chosen student is:

Using Densities for Discrete Variables

If a discrete variable takes a large number of values which are close together, we can often approximate it using a continuous variable.

- Suppose 500 students take a standardized exam, with mean 75 points. The distribution for the score S of a randomly chosen student is:

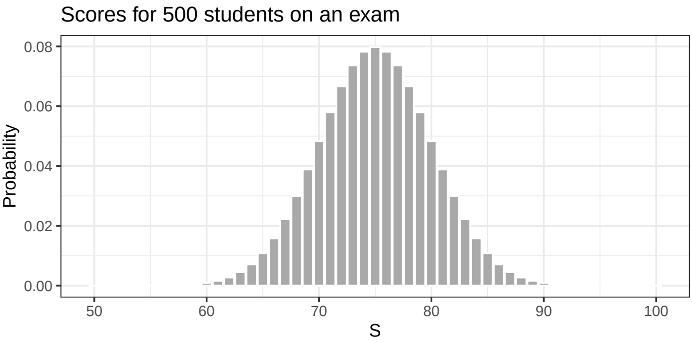

Using Densities for Discrete Variables

If a discrete variable takes a large number of values which are close together, we can often approximate it using a continuous variable.

- Suppose 500 students take a standardized exam, with mean 75 points. The distribution for the score S of a randomly chosen student is:

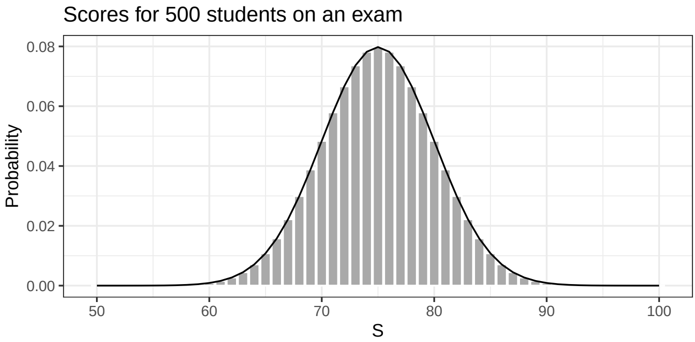

Section 2

The Normal Distribution

The Normal Distribution

- The general Normal density curve with mean μ and standard deviation σ is given by the formula

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma} \qquad \text{Don't memorize this}$$

The Normal Distribution

- The general Normal density curve with mean μ and standard deviation σ is given by the formula

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma} \qquad \text{Don't memorize this}$$

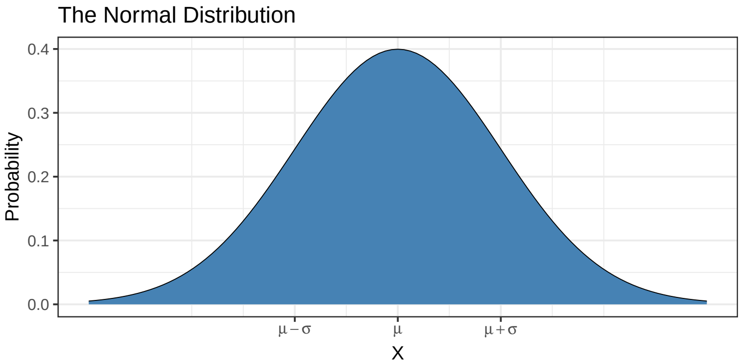

Normal Probabilities

Recall that for a random variable which has a **continuous** distribution, we find probabilities by looking at areas under the density curve.

Normal Probabilities

Recall that for a random variable which has a **continuous** distribution, we find probabilities by looking at areas under the density curve.

Suppose X is Normally distributed with mean 2 and standard deviation 1. What is the probability that X is between 3 and 4?

Normal Probabilities

Recall that for a random variable which has a **continuous** distribution, we find probabilities by looking at areas under the density curve.

Suppose X is Normally distributed with mean 2 and standard deviation 1. What is the probability that X is between 3 and 4?

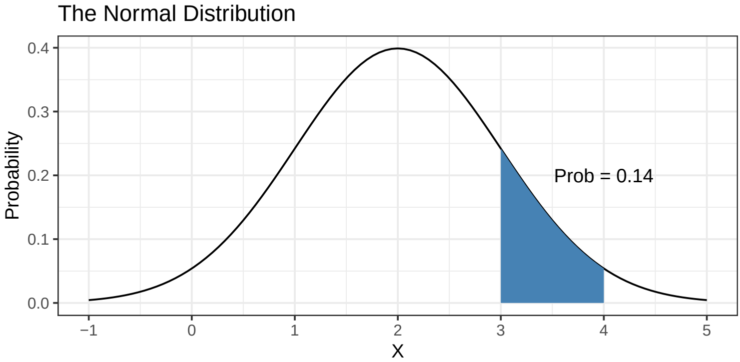

Calculating Normal Probabilities in R

How do we actually find areas under the Normal density curve?

Calculating Normal Probabilities in R

How do we actually find areas under the Normal density curve?

- R has a built-in function for computing cummulative probabilites under Normal densities: `pnorm(q =... , mean =... , sd =... )`

Calculating Normal Probabilities in R

How do we actually find areas under the Normal density curve?

- R has a built-in function for computing cummulative probabilites under Normal densities: `pnorm(q =... , mean =... , sd =... )`
- For example, the following code computes the area left of 1.5 in the Normal distribution with mean 0 and standard deviation 1:

```
pnorm(q =1.5 , mean =0 , sd =1 )
```

```
## [1] 0.9331928
```

Calculating Normal Probabilities in R

How do we actually find areas under the Normal density curve?

- R has a built-in function for computing cummulative probabilites under Normal densities: `pnorm(q =... , mean =... , sd =... )`
- For example, the following code computes the area left of 1.5 in the Normal distribution with mean 0 and standard deviation 1:

```
pnorm(q =1.5 , mean =0 , sd =1 )
```

```
## [1] 0.9331928
```

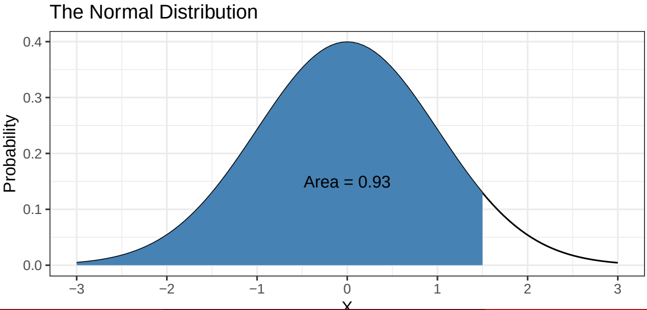

Finding Areas of General Regions

The `pnorm` function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

Finding Areas of General Regions

The `pnorm` function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Finding Areas of General Regions

The pnorm function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Find the area between -.25 and 1.5 under the Normal density with mean 0 and standard deviation 1.

```
pnorm(q =1.5 , mean =0 , sd =1 ) - pnorm(q = -.25 , mean =0 , sd =1 )
```

```
## [1] 0.5318991
```

Finding Areas of General Regions

The pnorm function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Find the area between -.25 and 1.5 under the Normal density with mean 0 and standard deviation 1.

```
pnorm(q =1.5 , mean =0 , sd =1 ) - pnorm(q = -.25 , mean =0 , sd =1 )
```

```
## [1] 0.5318991
```

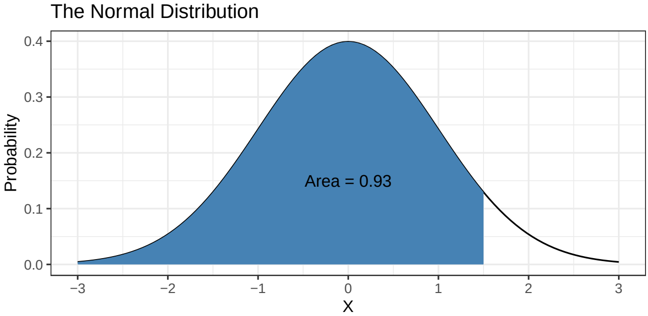

Finding Areas of General Regions

The pnorm function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Find the area between -.25 and 1.5 under the Normal density with mean 0 and standard deviation 1.

```
pnorm(q =1.5 , mean =0 , sd =1 ) - pnorm(q = -.25 , mean =0 , sd =1 )
```

```
## [1] 0.5318991
```

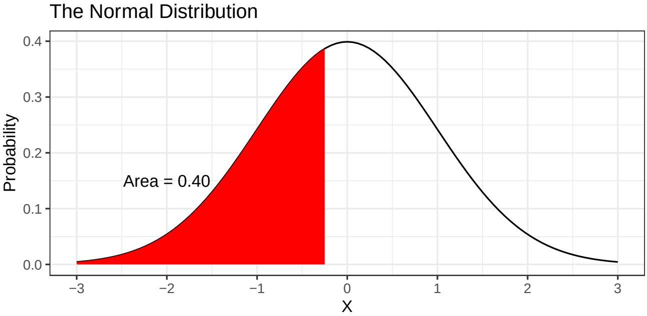

Finding Areas of General Regions under Normal curve

The `pnorm` function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Find the area between -.25 and 1.5 under the Normal density with mean 0 and standard deviation 1.

```
pnorm(q =1.5 , mean =0 , sd =1 ) - pnorm(q = -.25 , mean =0 , sd =1 )
```

```
## [1] 0.5318991
```

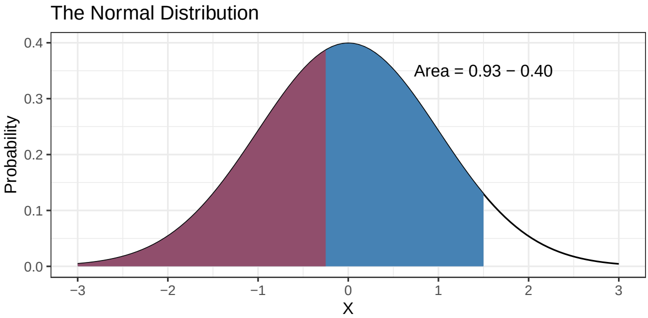

Finding Areas of General Regions

The pnorm function lets us compute cumulative areas (i.e. all area to the left of a given value). But how do we compute the area between two values?

- Answer: By computing two cumulative areas and subtracting the results!

Find the area between -.25 and 1.5 under the Normal density with mean 0 and standard deviation 1.

```
pnorm(q =1.5 , mean =0 , sd =1 ) - pnorm(q = -.25 , mean =0 , sd =1 )
```

```
## [1] 0.5318991
```

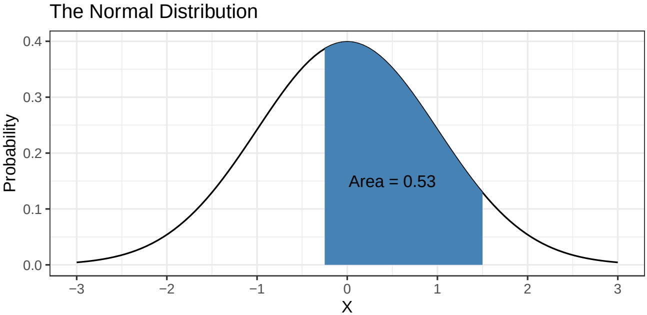

Finding Quantiles

Suppose we instead have the opposite problem: We want to FIND the value of X with a given cumulative area.

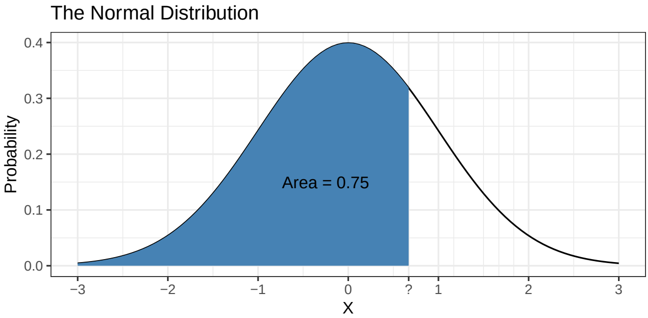

Finding Quantiles

Suppose we instead have the opposite problem: We want to FIND the value of X with a given cumulative area.

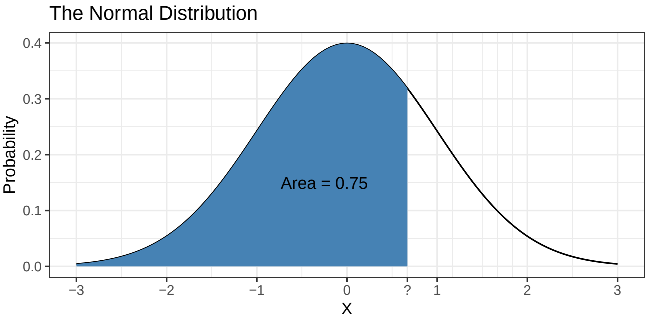

- That is, we want to find the .75 quantile (i.e. the 75th percentile)

Finding Quantiles

Suppose we instead have the opposite problem: We want to FIND the value of X with a given cumulative area.

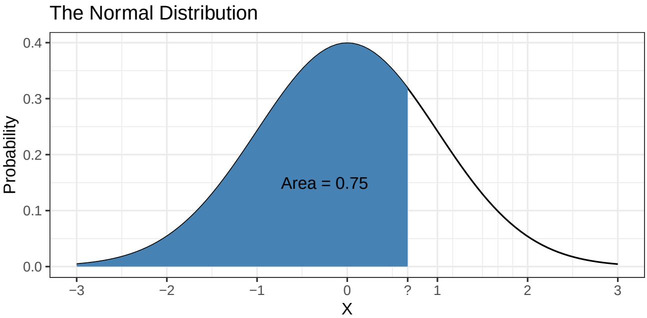

- That is, we want to find the .75 quantile (i.e. the 75th percentile)

R has a built-in function for that too! `qnorm(p =... , mean =... , sd =... )`

Finding Quantiles

Suppose we instead have the opposite problem: We want to FIND the value of X with a given cumulative area.

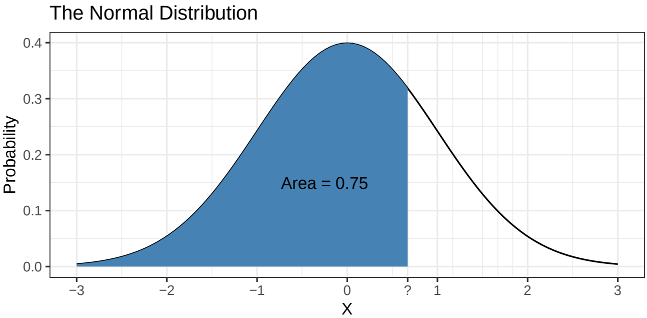

- That is, we want to find the .75 quantile (i.e. the 75th percentile)

R has a built-in function for that too! `qnorm(p =... , mean =... , sd =... )`

```
qnorm(p =.75 , mean =0 , sd =1 )
```

```
## [1] 0.6744898
```

Finding Quantiles

Suppose we instead have the opposite problem: We want to FIND the value of X with a given cumulative area.

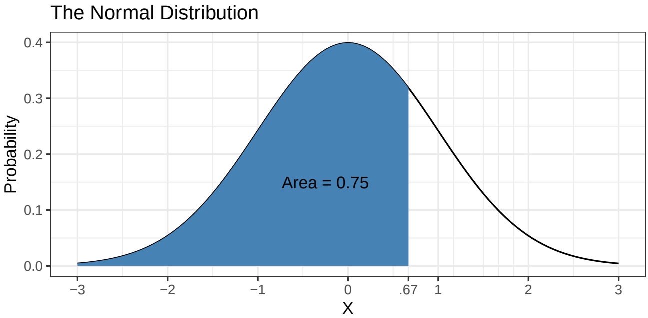

- That is, we want to find the .75 quantile (i.e. the 75th percentile)

R has a built-in function for that too! `qnorm(p =... , mean =... , sd =... )`

```
qnorm(p =.75 , mean =0 , sd =1 )
```

```
## [1] 0.6744898
```

Scale and Translation Invariance

- Consider a Normal variable X with $\mu = 0$ and $\sigma = 1$, and another Normal variable Y with mean $\mu = 2$ and $\sigma = .25$.

Scale and Translation Invariance

- Consider a Normal variable X with $\mu = 0$ and $\sigma = 1$, and another Normal variable Y with mean $\mu = 2$ and $\sigma = .25$.

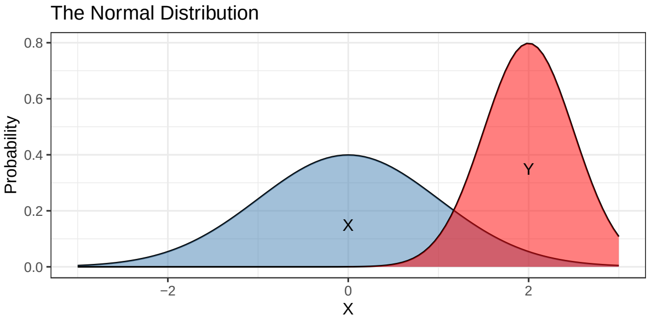

Scale and Translation Invariance

- Consider a Normal variable X with $\mu = 0$ and $\sigma = 1$, and another Normal variable Y with mean $\mu = 2$ and $\sigma = .25$.

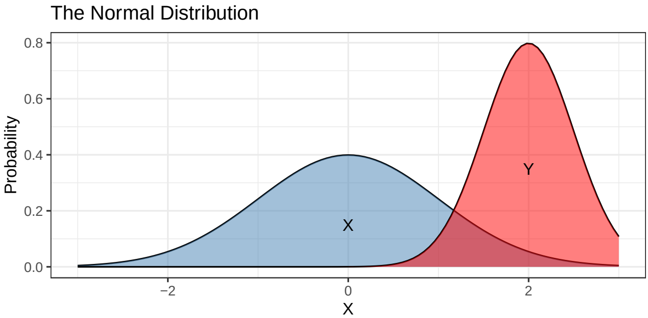

- The distributions for X and Y have different means and different heights and widths. . .
  - But otherwise have identitical shapes!

Scale and Translation Invariance

- Consider a Normal variable X with $\mu = 0$ and $\sigma = 1$, and another Normal variable Y with mean $\mu = 2$ and $\sigma = .25$.

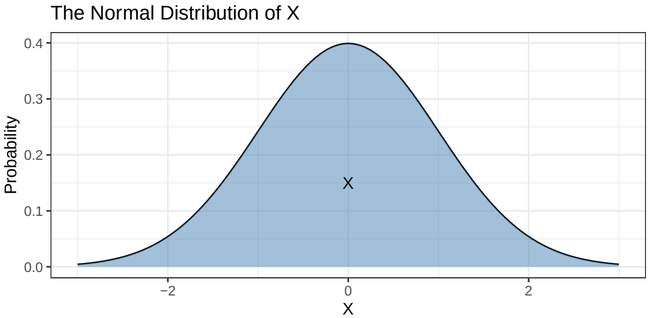

- The distributions for X and Y have different means and different heights and widths. . .
  - But otherwise have identical shapes!

Scale and Translation Invariance

- Consider a Normal variable X with $\mu = 0$ and $\sigma = 1$, and another Normal variable Y with mean $\mu = 2$ and $\sigma = .25$.

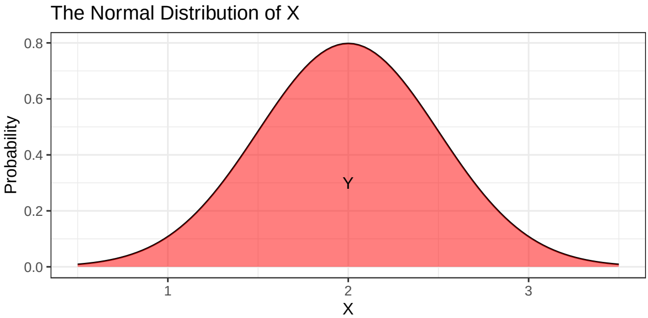

- The distributions for X and Y have different means and different heights and widths. . .
  - But otherwise have identical shapes!

Location-Scale Transformations

The previous example suggest that if we shift and rescale a Normal random variable, we should still get a Normal random variable

Location-Scale Transformations

The previous example suggest that if we shift and rescale a Normal random variable, we should still get a Normal random variable

Theorem

Suppose X is a Normal random variable with mean μ and standard deviation σ. Then $Z = \frac{X-\mu}{\sigma}$ is a Normal random variable with mean 0 and standard deviation 1.

Location-Scale Transformations

The previous example suggest that if we shift and rescale a Normal random variable, we should still get a Normal random variable

Theorem

Suppose X is a Normal random variable with mean μ and standard deviation σ. Then $Z = \frac{X-\mu}{\sigma}$ is a Normal random variable with mean 0 and standard deviation 1.

The Normal variable with mean 0 and standard deviation 1 is given a special name: **the standard Normal**.

Location-Scale Transformations

The previous example suggest that if we shift and rescale a Normal random variable, we should still get a Normal random variable

Theorem

Suppose X is a Normal random variable with mean μ and standard deviation σ. Then $Z = \frac{X-\mu}{\sigma}$ is a Normal random variable with mean 0 and standard deviation 1.

The Normal variable with mean 0 and standard deviation 1 is given a special name: **the standard Normal**.

The process of subtracting off the mean from a random variable and dividing by the standard deviation is called **standardizing.**

Location-Scale Transformations

The previous example suggest that if we shift and rescale a Normal random variable, we should still get a Normal random variable

Theorem

Suppose X is a Normal random variable with mean μ and standard deviation σ. Then $Z = \frac{X-\mu}{\sigma}$ is a Normal random variable with mean 0 and standard deviation 1.

The Normal variable with mean 0 and standard deviation 1 is given a special name: **the standard Normal**.

The process of subtracting off the mean from a random variable and dividing by the standard deviation is called **standardizing.**

It's often useful to standardize a variable so that we only need to consider a single density function (the *standard* Normal density) rather than many (one for each choice of μ and σ)