Bernoulli Random Variables

Recall that a *Bernoulli* random variable is a random variable whose only possible values are 0 and 1.

Bernoulli Random Variables

Recall that a *Bernoulli* random variable is a random variable whose only possible values are 0 and 1.

In general we can consider a Bernoulli random variable to be the result of an experiment with two outcomes, which for convenience we will label "success" and "failure"

Bernoulli Random Variables

Recall that a *Bernoulli* random variable is a random variable whose only possible values are 0 and 1.

In general we can consider a Bernoulli random variable to be the result of an experiment with two outcomes, which for convenience we will label "success" and "failure"

As before we define the Bernoulli random variable X by agreeing to assign the value of 1 to X if the result of the experiment is "success", and zero if the result is "failure":

$$X = \begin{cases} 1 & \text{if the outcome of the experiment is "success"} \\ 0 & \text{if the outcome of the experiment is "failure"} \end{cases}$$

Bernoulli Random Variables

To be consistent with the Kolmogorov probability axioms the probability of "success" must be a number p between zero and one (inclusive), and the probability of "failure", which is the compliment of "success", must be $1 - p$.

Bernoulli Random Variables

To be consistent with the Kolmogorov probability axioms the probability of "success" must be a number p between zero and one (inclusive), and the probability of "failure", which is the compliment of "success", must be $1 - p$.

This results in the following probability mass function $f(x)$ which we will refer to as the *Bernoulli distribution*:

$$f(x) = P(X = x) = \begin{cases} p & \text{if} \quad x = 1 \\ 1 - p & \text{if} \quad x = 0 \end{cases}$$

Bernoulli Random Variables

To be consistent with the Kolmogorov probability axioms the probability of "success" must be a number p between zero and one (inclusive), and the probability of "failure", which is the compliment of "success", must be $1-p$.

This results in the following probability mass function $f(x)$ which we will refer to as the *Bernoulli distribution*:

$$f(x) = P(X = x) = \begin{cases} p & \text{if} \quad x = 1 \\ 1-p & \text{if} \quad x = 0 \end{cases}$$

Most of the discrete probability distributions we will now consider are related to the Bernoulli distribution.

Discrete Distributions

Now consider a series of *independent* experiments, each of which produces a Bernoulli random variable with probability of success p (p is the same for all of the trials)

Discrete Distributions

Now consider a series of *independent* experiments, each of which produces a Bernoulli random variable with probability of success p (p is the same for all of the trials)

The following discrete probability distributions arise from this model:

If the number of trials n is fixed in advance, the number of successes X has a **binomial** distribution

Discrete Distributions

Now consider a series of *independent* experiments, each of which produces a Bernoulli random variable with probability of success p (p is the same for all of the trials)

The following discrete probability distributions arise from this model:

If the number of trials n is fixed in advance, the number of successes X has a **binomial** distribution

If trials continue indefinitely until the first success is obtained, the number of failures obtained X has a **geometric** distribution.

Discrete Distributions

Now consider a series of *independent* experiments, each of which produces a Bernoulli random variable with probability of success p (p is the same for all of the trials)

The following discrete probability distributions arise from this model:

If the number of trials n is fixed in advance, the number of successes X has a **binomial** distribution

If trials continue indefinitely until the first success is obtained, the number of failures obtained X has a **geometric** distribution.

If trials continue indefinitely until the r^{th} success is obtained, the number of failures obtained X has a **negative binomial** distribution.

Discrete Distributions

Note that the geometric distribution is a special case of the negative binomial distribution, with $r = 1$.

Discrete Distributions

Note that the geometric distribution is a special case of the negative binomial distribution, with $r = 1$.

Unfortunately, different authors define the random variable X in the negative binomial (and geometric) distribution in different ways.

Discrete Distributions

Note that the geometric distribution is a special case of the negative binomial distribution, with $r = 1$.

Unfortunately, different authors define the random variable X in the negative binomial (and geometric) distribution in different ways.

However, the characterization as a sequence of Bernoulli trials that ends at the r^{th} success is common to all definitions.

Discrete Distributions

Note that the geometric distribution is a special case of the negative binomial distribution, with $r = 1$.

Unfortunately, different authors define the random variable X in the negative binomial (and geometric) distribution in different ways.

However, the characterization as a sequence of Bernoulli trials that ends at the r^{th} success is common to all definitions.

That said, you should be prepared to encounter a different definition of X (and a different, but equivalent pmf)if you look at a different text.

Discrete Distributions

The other related distribution we will consider is the **Poisson** distribution.

Discrete Distributions

The other related distribution we will consider is the **Poisson** distribution.

The Poisson is a limiting form of the binomial distribution that you get if you let n become very large and the probability of success p very small, but always keep $np = \lambda$ the same.

Discrete Distributions

The other related distribution we will consider is the **Poisson** distribution.

The Poisson is a limiting form of the binomial distribution that you get if you let n become very large and the probability of success p very small, but always keep $np = \lambda$ the same.

Another way to say this is that we take binomial random variables with larger and larger n, but we keep the *expected number of successes* $np = \lambda$ the same for all of them.

Discrete Distributions

The other related distribution we will consider is the **Poisson** distribution.

The Poisson is a limiting form of the binomial distribution that you get if you let n become very large and the probability of success p very small, but always keep $np = \lambda$ the same.

Another way to say this is that we take binomial random variables with larger and larger n, but we keep the *expected number of successes* $np = \lambda$ the same for all of them.

The limit of the distribution of such a sequence of random variables as $n \to \infty$ is a Poisson.

The Binomial Distribution

The binomial experiment consists of:

- n independent Bernoulli trials are performed
- The random variable X is the sum of the results (i.e., the number of successes)
- The probability of success p is the same for all trials

The Binomial Distribution

The binomial experiment consists of:

- n independent Bernoulli trials are performed
- The random variable X is the sum of the results (i.e., the number of successes)
- The probability of success p is the same for all trials

The probability mass function (pmf) $f(x)$ is:

$$f(x) = P(X = x) = b(x; n, p) = \binom{n}{x} p^x (1-p)^{n-x}, \quad x = 0, 1, 2, \ldots,$$

The Binomial Distribution

It is not obvious, but if you sum the values of $f(x)$ over all values from zero to n, the sum is one.

$$\sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \; = \; 1$$

The Binomial Distribution

It is not obvious, but if you sum the values of $f(x)$ over all values from zero to n, the sum is one.

$$\sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \;=\; 1$$

One way to make this clear is to consider the algebraic identity

$$(x+y)^n = \sum_{i=0}^{n} \binom{n}{i} x^i y^{n-i}$$

The Binomial Distribution

If we let x be the probability of success p and y the probability of failure $1-p$, on substitution we get

$$[p+(1-p)]^n = 1^n = 1 = \sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \quad 0 \le p \le 1$$

The Binomial Distribution

If we let x be the probability of success p and y the probability of failure $1-p$, on substitution we get

$$[p+(1-p)]^n = 1^n = 1 = \sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \quad 0 \leq p \leq 1$$

For any distribution, the cumulative distribution function (cdf) $F(x)$, is always defined by

$$F(x) = P(X \leq x)$$

The Binomial Distribution

If we let x be the probability of success p and y the probability of failure $1-p$, on substitution we get

$$[p+(1-p)]^n = 1^n = 1 = \sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \quad 0 \leq p \leq 1$$

For any distribution, the cumulative distribution function (cdf) $F(x)$, is always defined by

$$F(x) = P(X \leq x)$$

For the binomial distribution, this gives:

$$F(x) = \sum_{i=0}^{x} \binom{n}{i} p^i (1-p)^{n-i}$$

The Binomial Distribution

For the binomial distribution, there is no simple expression for $F(x)$

The Binomial Distribution

For the binomial distribution, there is no simple expression for $F(x)$

Values of $F(x)$ for the binomial can be obtained from:

- Tables (See table A.1 in the appendix)
- Spreadsheets: $= BINOMDIST(x, n, p, TRUE)$
- R: $pbinom(x, n, p)$

The Binomial Distribution

For the binomial distribution, there is no simple expression for $F(x)$

Values of $F(x)$ for the binomial can be obtained from:

- Tables (See table A.1 in the appendix)
- Spreadsheets: $= BINOMDIST(x, n, p, TRUE)$
- R: $pbinom(x, n, p)$

Example: A fair coin is tossed 10 times. What is the probability that 7 or fewer heads turn up?

The Binomial Distribution

For the binomial distribution, there is no simple expression for $F(x)$

Values of $F(x)$ for the binomial can be obtained from:

- Tables (See table A.1 in the appendix)
- Spreadsheets: $= BINOMDIST(x, n, p, TRUE)$
- R: $pbinom(x, n, p)$

Example: A fair coin is tossed 10 times. What is the probability that 7 or fewer heads turn up?

We want $P(X \leq 7)$, the probability that a binomial experiment with 10 trials and probability of success 0.5 produces 7 or fewer "successes".

The Binomial Distribution

If you are using a spreadsheet, enter:

$$= BINOMDIST(7, 10, 0.5, TRUE)$$

The Binomial Distribution

If you are using a spreadsheet, enter:

$= BINOMDIST(7, 10, 0.5, TRUE)$

If you are using R, enter:

$pdist(7, 10, 0.5)$

The Binomial Distribution

If you are using a spreadsheet, enter:

$= BINOMDIST(7, 10, 0.5, TRUE)$

If you are using R, enter:

$pdist(7, 10, 0.5)$

If you are using Table A.1, look under $n = 10$ on page 664, in the row with $x = 7$ and column with $p = 0.50$

The Binomial Distribution

If you are using a spreadsheet, enter:

$= BINOMDIST(7, 10, 0.5, TRUE)$

If you are using R, enter:

$pdist(7, 10, 0.5)$

If you are using Table A.1, look under $n = 10$ on page 664, in the row with $x = 7$ and column with $p = 0.50$

All of these should give the value $F(7) = .945$

The Binomial Distribution

If you are using a spreadsheet, enter:

$= BINOMDIST(7, 10, 0.5, TRUE)$

If you are using R, enter:

$pdist(7, 10, 0.5)$

If you are using Table A.1, look under $n = 10$ on page 664, in the row with $x = 7$ and column with $p = 0.50$

All of these should give the value $F(7) = .945$

This means that if we toss a fair coin 10 times, the probability of 7 or fewer heads is .945

The Binomial Distribution

If you are using a spreadsheet, enter:

$= BINOMDIST(7, 10, 0.5, TRUE)$

If you are using R, enter:

$pdist(7, 10, 0.5)$

If you are using Table A.1, look under $n = 10$ on page 664, in the row with $x = 7$ and column with $p = 0.50$

All of these should give the value $F(7) = .945$

This means that if we toss a fair coin 10 times, the probability of 7 or fewer heads is .945

If we repeat the experiment, tossing the coin 10 times, over and over, the *proportion* of all of the replications of the experiment that have 7 or fewer heads will approach .945.

The Binomial Distribution

Example: Suppose every time the Red Sox play the Yankees, the probability that the Red Sox win is 0.6.

If they play 7 games, what is the probability that the Red Sox win 5 or fewer?

The Binomial Distribution

Example: Suppose every time the Red Sox play the Yankees, the probability that the Red Sox win is 0.6.

If they play 7 games, what is the probability that the Red Sox win 5 or fewer?

If we assume that each game is an independent Bernoulli trial with probability of "success" equal to 0.6, then the number of games the Red Sox win will have a binomial distribution with $n = 7$ and $p = 0.6$.

The Binomial Distribution

We want to find the probability that the Red Sox win 5 or fewer,

$$P(X \leq 5) \; = \; F(5)$$

The Binomial Distribution

We want to find the probability that the Red Sox win 5 or fewer,

$$P(X \leq 5) \;=\; F(5)$$

Since there is no simple formula for F for a binomial distribution, we have to use one of the methods listed earlier

The Binomial Distribution

We want to find the probability that the Red Sox win 5 or fewer,

$$P(X \leq 5) \;=\; F(5)$$

Since there is no simple formula for F for a binomial distribution, we have to use one of the methods listed earlier

In R, enter *pdist(5,7,0.6)*

The Binomial Distribution

We want to find the probability that the Red Sox win 5 or fewer,

$$P(X \leq 5) \; = \; F(5)$$

Since there is no simple formula for F for a binomial distribution, we have to use one of the methods listed earlier

In R, enter *pdist(5,7,0.6)*

The result should be 0.841

The Binomial Distribution

Example: A baseball player has a $.300$ batting average.

If the player gets to bat five times in a game, what is the probability that he gets one hit or less:

The Binomial Distribution

Example: A baseball player has a $.300$ batting average.

If the player gets to bat five times in a game, what is the probability that he gets one hit or less:

We'll assume a binomial distribution with $n = 5$ and $p = 0.300$, then we want $F(1) = P(X \leq 1)$:

In R enter: *pdist(1,5,0.300)*

The Binomial Distribution

Example: A baseball player has a $.300$ batting average.

If the player gets to bat five times in a game, what is the probability that he gets one hit or less:

We'll assume a binomial distribution with $n = 5$ and $p = 0.300$, then we want $F(1) = P(X \leq 1)$:

In R enter: *pdist(1,5,0.300)*

The result is 0.528, so in games where a $.300$ hitter bats five times, more than 50 percent of the time they get one hit or less.

The Binomial Distribution

Example: The probability that it rains on a given weekend is 0.20.

In a month with four weekends, what is the probability that two or fewer are rainy?

The Binomial Distribution

Example: The probability that it rains on a given weekend is 0.20.

In a month with four weekends, what is the probability that two or fewer are rainy?

Assume a binomial distribution with $n = 4$ and $p = 0.2$.

In R enter: *pdist(2,4,0.20)*

The Binomial Distribution

Example: The probability that it rains on a given weekend is 0.20.

In a month with four weekends, what is the probability that two or fewer are rainy?

Assume a binomial distribution with $n = 4$ and $p = 0.2$.

In R enter: *pdist(2,4,0.20)*

The result is 0.9728,

The Binomial Distribution

Example: If

$$F(x) = P(X \leq x)$$

is the probability of the event A="x or fewer successes", the **compliment** of this event A' is "more than x successes"

Recall that the probability of the compliment A' is always $1 - P(A)$.

If the chance of rain on a weekend is 0.2 and there are four weekends in a month, what is the probability that it rains on more than 2 weekends?

The Binomial Distribution

Example: If

$$F(x) = P(X \leq x)$$

is the probability of the event A="x or fewer successes", the **compliment** of this event A' is "more than x successes"

Recall that the probability of the compliment A' is always $1 - P(A)$.

If the chance of rain on a weekend is 0.2 and there are four weekends in a month, what is the probability that it rains on more than 2 weekends?

As before, assume a binomial distribution with $n = 4$ and $p = 0.2$.

In R enter: *1-pdist(2,4,0.20)*

The Binomial Distribution

Example: If

$$F(x) = P(X \leq x)$$

is the probability of the event A="x or fewer successes", the **compliment** of this event A' is "more than x successes"

Recall that the probability of the compliment A' is always $1 - P(A)$.

If the chance of rain on a weekend is 0.2 and there are four weekends in a month, what is the probability that it rains on more than 2 weekends?

As before, assume a binomial distribution with $n = 4$ and $p = 0.2$.

In R enter: *1-pdist(2,4,0.20)*

The result is 0.0272,

The Binomial Distribution

The expected value of a binomial random variable $E(X)$ is:

$$E(X) = \sum_{x=0}^{n} x \cdot f(x)$$

The Binomial Distribution

The expected value of a binomial random variable $E(X)$ is:

$$E(X) = \sum_{x=0}^{n} x \cdot f(x)$$

$$E(X) = \sum x \cdot \binom{n}{x} p^x (1-p)^{n-x} \ = \ np$$

The Binomial Distribution

To find the variance $V(X)$ of a binomial random variable, first we find $E(X^2)$:

$$E(X^2) = \sum_{x=0}^{n} x^2 \cdot f(x)$$

The Binomial Distribution

To find the variance $V(X)$ of a binomial random variable, first we find $E(X^2)$:

$$E(X^2) = \sum_{x=0}^{n} x^2 \cdot f(x)$$

$$E(X^2) = \sum x^2 \cdot \binom{n}{x} p^x (1-p)^{n-x} \; = \; n^2p^2 - np^2 + np$$

The Binomial Distribution

To find the variance $V(X)$ of a binomial random variable, first we find $E(X^2)$:

$$E(X^2) = \sum_{x=0}^{n} x^2 \cdot f(x)$$

$$E(X^2) = \sum x^2 \cdot \binom{n}{x} p^x (1-p)^{n-x} \;=\; n^2p^2 - np^2 + np$$

Then

$$V(X) = E(X^2) - [E(X)]^2 = n^2p^2 - np^2 + np - n^2p^2$$

and

$$V(X) \;=\; np(1-p)$$

The Binomial Distribution

Now we will perform some numerical experiments.

First generate a sample of $1,000,000$ observations for a binomial experiment with $n = 6$ trials and probability of success $p = 0.4$:

x<-rbinom(1000000,6,0.4)

The Binomial Distribution

Now we will perform some numerical experiments.

First generate a sample of $1,000,000$ observations for a binomial experiment with $n = 6$ trials and probability of success $p = 0.4$:

x<-rbinom(1000000,6,0.4)

Now plot a histogram of the results:

hist(x)

The Binomial Distribution

Now we will perform some numerical experiments.

First generate a sample of $1,000,000$ observations for a binomial experiment with $n = 6$ trials and probability of success $p = 0.4$:

x<-rbinom(1000000,6,0.4)

Now plot a histogram of the results:

hist(x)

To get a table of the results enter

table(x)

The Binomial Distribution

Now we will perform some numerical experiments.

First generate a sample of $1,000,000$ observations for a binomial experiment with $n = 6$ trials and probability of success $p = 0.4$:

x<-rbinom(1000000,6,0.4)

Now plot a histogram of the results:

hist(x)

To get a table of the results enter

table(x)

The results should look something like:

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Now compare the frequencies to the probabilities.

First compute the probability that $X = 0$:

dbinom(0,5,0.4)

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Now compare the frequencies to the probabilities.

First compute the probability that $X = 0$:

dbinom(0,5,0.4)

The result should be something like

[1] 0.07776

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Now compare the frequencies to the probabilities.

First compute the probability that $X = 0$:

dbinom(0,5,0.4)

The result should be something like

[1] 0.07776

To get the probability that $X = 1$ enter

dbinom(1,5,0.4)

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Now compare the frequencies to the probabilities.

First compute the probability that $X = 0$:

dbinom(0,5,0.4)

The result should be something like

[1] 0.07776

To get the probability that $X = 1$ enter

dbinom(1,5,0.4)

This time the results should look something like:

[1] 0.2592

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Next compute the probability that $X = 2$:

dbinom(2,5,0.4)

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Next compute the probability that $X = 2$:

dbinom(2,5,0.4)

The result should be something like

[1] 0.3456

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Next compute the probability that $X = 2$:

dbinom(2,5,0.4)

The result should be something like

[1]   0.3456

To get the probability that $X = 5$ enter

dbinom(1,5,0.4)

The Binomial Distribution

0	1	2	3	4	5
77647	258841	346623	230275	76253	10361

Next compute the probability that $X = 2$:

dbinom(2,5,0.4)

The result should be something like

[1]   0.3456

To get the probability that $X = 5$ enter

dbinom(1,5,0.4)

This time the results should look something like:

[1]   0.01024

The Binomial Distribution

The expected value $E(X)$ in this case is:

$$E(X) = np = 5 \cdot 0.4 = 2$$

The Binomial Distribution

The expected value $E(X)$ in this case is:

$$E(X) \; = \; np \; = \; 5 \cdot 0.4 \; = \; 2$$

To compute the sample mean $\overline{x}$, enter

mean(x)

The Binomial Distribution

The expected value $E(X)$ in this case is:

$$E(X) = np = 5 \cdot 0.4 = 2$$

To compute the sample mean $\overline{x}$, enter

mean(x) The result should be something like

[1] 1.999759

The Binomial Distribution

The variance $V(X)$ in this case is:

$$V(X) \;=\; np(1-p) \;=\; 5 \cdot 0.4 \cdot 0.6 \;=\; 1.2$$

The Binomial Distribution

The variance $V(X)$ in this case is:

$$V(X) \;=\; np(1-p) \;=\; 5 \cdot 0.4 \cdot 0.6 \;=\; 1.2$$

To compute the sample variance s^2, enter

var(x)

The Binomial Distribution

The variance $V(X)$ in this case is:

$$V(X) = np(1-p) = 5 \cdot 0.4 \cdot 0.6 = 1.2$$

To compute the sample variance s^2, enter

var(x) The result should be something like

[1] 1.197966

The Binomial Distribution

At a certain intersection, the probability that a car goes stright through is 0.8.

If we observe 15 cars, what is the probability that 10 or fewer go straight through?

The Binomial Distribution

At a certain intersection, the probability that a car goes stright through is 0.8.

If we observe 15 cars, what is the probability that 10 or fewer go straight through?

Enter *pbinom(10,15,0.8)*

The Binomial Distribution

At a certain intersection, the probability that a car goes stright through is 0.8.

If we observe 15 cars, what is the probability that 10 or fewer go straight through?

Enter *pbinom(10,15,0.8)* The result should be .164

The Binomial Distribution

92% of a certain airline's flights arrive on time.

On a day when the airline operates 30 flights, what is the probablility that more than 27 arrive on time?

The Binomial Distribution

92% of a certain airline's flights arrive on time.

On a day when the airline operates 30 flights, what is the probablility that more than 27 arrive on time?

Enter *1-pbinom(27,30,0.92)*

The Binomial Distribution

92% of a certain airline's flights arrive on time.

On a day when the airline operates 30 flights, what is the probablility that more than 27 arrive on time?

Enter *1-pbinom(27,30,0.92)* The result should be .565