

Variance + Discrete Zoo I

CSE 312 26Su

Lecture 10

Fun fact!

Over a large number of trials, the average result will get closer to the expected value. This principle is why casinos make consistent profits. Even though individual results can vary widely, the “law of large numbers” ensures that the casino’s profit prevails over time.

Announcements

HW2 grades are out (regrade request closes Saturday night)

HW3 due tonight

No HW released tonight! (HW4 released next Wednesday)

Announcements

midterm ref. sheet

Quiz 3 Random Variables on Friday

Lectures 7-9 content

LOE

Midterm next week on **Wednesday 7/22**

See information on website

~ 5 questions

T/F/short answer, counting, probability, conditional probability, random variables

Where Are We?


Last time:

- Expectation
- Linearity of Expectation

Today:

- Recap LOE
- Independence of RVs
- Expectation of a function
- Variance
- Zoo of Discrete Random Variables

RV



Recap LOE

Rotating the table

 n people are sitting around a circular table. There is a name tag in each place. Nobody is sitting in front of their own name tag.

Rotate the table by a random number k of positions between 1 and $n-1$ (equally likely)

X is the number of people that end up in front of their own name tag. **Find** $\mathbb{E}[X]$.

Rotating the table

n people are sitting around a circular table. There is a name tag in each place. Nobody is sitting in front of their own name tag.

Rotate the table by a random number k of positions between 1 and $n-1$ (equally likely)

X is the number of people that end up in front of their own name tag. Find $\mathbb{E}[X]$.

→ Decompose: Define X_i as follows: indicator rv

$$X_i = \begin{cases} 1 & \text{if person } i \text{ sits in front of their own name tag} \\ 0 & \text{otherwise} \end{cases}$$

$\mathbb{P}(\text{event}) = \mathbb{E}[X_i]$

$$X = \sum_{i=1}^n X_i$$

Rotating the table

n people are sitting around a circular table. There is a name tag in each place. Nobody is sitting in front of their own name tag.

Rotate the table by a random number k of positions between 1 and $n-1$ (equally likely)

X is the number of people that end up in front of their own name tag. Find $\mathbb{E}[X]$.

Decompose: Define X_i as follows:

$$X_i = \begin{cases} 1 & \text{if person } i \text{ sits in front of their own name tag} \\ 0 & \text{otherwise} \end{cases} \quad X = \sum_{i=1}^n X_i$$

LOE:

it's OR if X_i is depend

$$\mathbb{E}[X] = \mathbb{E}[\sum_{i=1}^n X_i] = \underbrace{\sum_{i=1}^n \mathbb{E}[X_i]}_{\text{sum}}$$

Conquer:

$$\mathbb{E}[X_i] = \underbrace{P(X_i = 1)} = \frac{1}{n-1}$$

$$\mathbb{E}[X] = \underbrace{n} \cdot \mathbb{E}[X_i] = \frac{n}{n-1}$$

Rotating the table

n people are sitting around a circular table. There is a name tag in each place. Nobody is sitting in front of their own name tag.

Rotate the table by a random number k of positions between 1 and $n-1$ (equally likely)

X is the number of people that end up in front of their own name tag. **Find** $\mathbb{E}[X]$.

Decompose: Define X_i as follows:

$$X_i = \begin{cases} 1 & \text{if person } i \text{ sits in front of their own name tag} \\ 0 & \text{otherwise} \end{cases} \quad X = \sum_{i=1}^n X_i$$

LOE:

$$\mathbb{E}[X] = \mathbb{E}[\sum_{i=1}^n X_i] = \sum_{i=1}^n \mathbb{E}[X_i]$$

Techniques For Finding Expectation

1. Definition of Expectation

When the support is small enough and we can find the PMF of X


$$\mathbb{E}[X] = \sum_{k \in \Omega_X} k \cdot \mathbb{P}(X = k)$$



2. Linearity of Expectation


When we can break X into a sum
Break into indicator RVs if it's some kind of count

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$

3. LOTUS (Law of the unconscious statistician)

When it's a function of X (e.g, X^2)

We'll see this in a couple slides!



Independence of Random Variables

Independence of events

Recall the definition of independence of **events**:

Independence

Two events A , B are independent if
 $\mathbb{P}(A \cap B)$ = $\mathbb{P}(A)$ · $\mathbb{P}(B)$



$$\mathbb{P}(A|B) = \mathbb{P}(A)$$

$$\mathbb{P}(B|A) = \mathbb{P}(B)$$

“knowing whether one event occurred doesn’t tell us anything about whether the other event occurred”

Independence of Random Variables

That's for events...what about random variables?

Independence (of random variables) ★

X and Y are independent if for all $k \in \Omega_X, \ell \in \Omega_Y$
 $\mathbb{P}(X = k, Y = \ell)$ = $\mathbb{P}(X = k)$ $\mathbb{P}(Y = \ell)$

just an event

We'll often use commas instead of $\cap$ symbol to save space.

$$P(A \cap B) = P(A, B)$$

"knowing the value of one random variable doesn't tell us anything about what the value of the other might be"

Independence of Random Variables

The "for all values" is important.

$$P(S = 7 \cap R = 5) = P(S = 7)P(R = 5)$$

↓
We say that the event "the sum is 7" is independent of "the red die is 5"

What about S = "the sum of two dice" and R = "the value of the red die"

Independence of Random Variables

The “for all values” is important.

We say that the event “the sum is 7” is independent of “the red die is 5”

What about S = “the sum of two dice” and R = “the value of the red die”

NOT independent.

$$\mathbb{P}(S = 2, R = 5) \neq \mathbb{P}(S = 2)\mathbb{P}(R = 5) \text{ (for example)}$$

0

Independence of Random Variables

↓ ↓
Flip a coin independently $2n$ times.

Let X be "the number of heads in the first n flips."

Let Y be "the number of heads in the last n flips."

X and Y are independent.

Mutual Independence for RVs

A little simpler to write down than for events

Mutual Independence (of random variables)

$X_1, X_2, \dots, X_n$ are mutually independent if for all $x_1, x_2, \dots, x_n$
 $\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{P}(X_1 = x_1) \mathbb{P}(X_2 = x_2) \cdots \mathbb{P}(X_n = x_n)$

DON'T need to check all subsets for random variables...

But you do need to check all values (all possible x_i) still.

Techniques For Finding Expectation

1. Definition of Expectation

When the support is small enough and we can find the PMF of X

$$\mathbb{E}[X] = \sum_{k \in \Omega_X} k \cdot \mathbb{P}(X = k)$$

2. Linearity of Expectation

*all X, Y
regardless of
indp.*


When we can break X into a sum
Break into indicator RVs if it's some kind of count

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$

3. LOTUS (Law of the unconscious statistician)

When it's a function of X (e.g, X^2)

We'll see this now!



Expectation of a *Function* of a Random Variable

Expectation

Expectation

The “expectation” (or “expected value”) of a random variable X is:

$$\mathbb{E}[X] = \sum_{k \in \Omega_X} \underline{k \cdot \mathbb{P}(X = k)}$$
$$\mathbb{E}[X] = \sum_{\omega \in \Omega} X(\omega) \cdot \mathbb{P}(\omega)$$

Intuition: The weighted average of values X could take on.
Weighted by the probability you actually see them.

What about $\mathbb{E}[g(X)]$? (e.g., $\mathbb{E}[X^2]$, $\mathbb{E}[2^X]$)

Applying **functions** on a random variable(s).

$$\underline{g(X) = 2X + 3}$$

$$g(X) = X^2$$

$$g(X) = 2^X$$

$$g(X, Y) = X + Y$$

Still gives us a random variable!

Given an outcome, these functions give you a number.

They're functions from Ω $\rightarrow \mathbb{R}$. That's the definition of a random variable!

What about $\mathbb{E}[g(X)]$? (e.g., $\mathbb{E}[X^2]$, $\mathbb{E}[2^X]$)

What if we want to find the expected value of some function of X ?

Let's say we want to find $\mathbb{E}[X^2]$. Is $\mathbb{E}[X^2] = (\mathbb{E}[X])^2$?

What about $\mathbb{E}[g(X)]$? (e.g., $\mathbb{E}[X^2]$, $\mathbb{E}[2^X]$)

What if we want to find the expected value of some function of X ?

Let's say we want to find $\mathbb{E}[X^2]$. Is $\mathbb{E}[X^2] = (\mathbb{E}[X])^2$?

Not necessarily! For example,

If we have a random variable X that following the PMF:

$$p_X(k) = \begin{cases} 0.5 & \underline{k = 1} \\ 0.5 & k = -1 \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{E}[X] = 0.5 \cdot 1 + 0.5 \cdot -1 = \underline{0}$$

$$\underline{\mathbb{E}[X]^2 = 0}$$

What about $\mathbb{E}[g(X)]$? (e.g., $\mathbb{E}[X^2]$)

What if we want to find the expected value of some function of X ?

Let's say we want to find $\mathbb{E}[X^2]$. Is $\mathbb{E}[X^2] = (\mathbb{E}[X])^2$?

Not necessarily! For example,

If we have a random variable X that follows the PMF:

$$p_X(k) = \begin{cases} 0.5 & k = 1 \\ 0.5 & k = -1 \\ 0 & \text{otherwise} \end{cases} \quad p_{X^2}(k) = \mathbb{P}(X^2 = k) = \begin{cases} 0.5 & k = 1^2 \\ 0.5 & k = (-1)^2 \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{E}[X] = 0.5 \cdot 1 + 0.5 \cdot (-1) = 0$$

$$\mathbb{E}[X^2] = 0.5 \cdot 1^2 + 0.5 \cdot (-1)^2 = 1$$

$$\mathbb{E}[X]^2 = 0$$

What about $\mathbb{E}[g(X)]$? (e.g., $\mathbb{E}[X^2]$)

What if we want to find the expected value of some function of X ?

Let's say we want to find $\mathbb{E}[X^2]$. Is $\mathbb{E}[X^2] = (\mathbb{E}[X])^2$?

Not necessarily! For example,

If we have a random variable X that follows the PMF:

$$p_X(k) = \begin{cases} 0.5 & k = 1 \\ 0.5 & k = -1 \\ 0 & \text{otherwise} \end{cases} \quad p_{X^2}(k) = \mathbb{P}(X^2 = k) = \begin{cases} 1 & k = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{E}[X] = 0.5 \cdot 1 + 0.5 \cdot (-1) = 0$$

$$\mathbb{E}[X^2] = 0.5 \cdot 1^2 + 0.5 \cdot (-1)^2 = 1$$

$$\mathbb{E}[X]^2 = 0$$

Expectation of $g(X)$

“Law of the unconscious statistician” (LOTUS)

The “expectation” (or “expected value”) of $g(X)$ is:

$$\underline{\mathbb{E}[g(X)]} = \sum_{k \in \Omega_X} \underline{g(k)} \cdot \underline{\mathbb{P}(X = k)}$$

Exact same as formula for $E[X]$, but we **apply the function on each of the values in the support of X** (the corresponding probabilities are the same)

$$E[X^2] = \sum_{k \in \Omega_X} k^2 \cdot P(X = k)$$

Expectation of $g(X)$

$$\mathbb{E}[(X+Y)^2] \quad \underline{Z=2XY}$$
$$\mathbb{E}[X^2 + \underbrace{2XY}_{\text{circled}} + Y^2]$$

“Law of the unconscious statistician” (LOTUS)

The “expectation” (or “expected value”) of $g(X)$ is:

$$\mathbb{E}[g(X)] = \sum_{k \in \Omega_X} g(k) \cdot \mathbb{P}(X = k)$$

if X, Y are idpt
 $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$

Exact same as formula for $E[X]$, but we **apply the function on each of the values in the support of X** (the corresponding probabilities are the same)

$$\mathbb{E}[X+Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$

What if $g(X)$ is a linear function? E.g., $g(X, Y) = X + Y$

Techniques For Finding Expectation

1. Definition of Expectation

When the support is small enough and we can find the PMF of X

$$\mathbb{E}[X] = \sum_{k \in \Omega_X} k \cdot \mathbb{P}(X = k)$$

2. Linearity of Expectation

When we can break X into a sum
Break into indicator RVs if it's some kind of count

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$$


3. LOTUS (Law of the unconscious statistician)

When it's a function of X (e.g, X^2)

$$\mathbb{E}[g(X)] = \sum_{k \in \Omega_X} g(k) \cdot \mathbb{P}(X = k)$$

Where are we?

A *random variable* is a way to summarize what outcome you saw.




→ **Support** (set of values RV can take),

→ **Probability Mass Function & Cumulative Distribution Function** describe probabilities

→ The **expectation** of a RV is its average value. (summarizes a RV)

> Expectation is **linear**: $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$.

- $X + Y$ is a random variable – it's a function that outputs a number given an outcome (or, here, a combination of outcomes).



Variance



Variance


Another one number summary of a random variable.

But wait, we already have expectation, what's this for?

Consider these two games

Would you be willing to play these games?


Game 1: I will flip a fair coin; if it's heads, I pay you \$1. If it's tails, you pay me \$1. Let X_1 be your profit if you play game 1




Consider these two games

Would you be willing to play these games?

Game 1: I will flip a fair coin; if it's heads, I pay you \$1. If it's tails, you pay me \$1. Let X_1 be your profit if you play game 1



Game 2: I will flip a fair coin; if it's heads, I pay you \$10,000. If it's tails, you pay me \$10,000. Let X_2 be your profit if you play game 2.



Both games are "fair" ($\mathbb{E}[X_1] = \mathbb{E}[X_2] = 0$)

What's the difference

Expectation tells you what the average will be...

But it doesn't tell you how "extreme" your results could be.

Nor how likely those extreme results are.

Game 2 has many (well, only) very extreme results.

In expectation they "cancel out" but if you can only play once...

...it would be nice to measure that.

Designing a Measure – Try 1

Let's measure how far all the outcomes are away from the center, weighted on how likely they are

$$\sum_{\omega} (\mathbb{P}(\omega) \cdot (X(\omega) - \mathbb{E}[X]))$$

What happens with Game 1?

$$\frac{1}{2} \cdot (1 - 0) + \frac{1}{2} \cdot (-1 - 0)$$
$$\frac{1}{2} + (-\frac{1}{2}) = 0$$

What happens with Game 2?

$$\frac{1}{2} \cdot (100000 - 0) + \frac{1}{2} \cdot (-100000 - 0)$$
$$50000 + (-50000) = 0$$

Designing a Measure – Try 2

How do we prevent cancelling? Squaring the distances makes everything positive.

$$\sum_{\omega} (\mathbb{P}(\omega) \cdot \underbrace{(X(\omega) - \mathbb{E}[X])^2}_{\text{This is the variance!}})$$

What happens with Game 1?

$$\frac{1}{2} \cdot (1 - 0)^2 + \frac{1}{2} \cdot (-1 - 0)^2$$
$$\frac{1}{2} + \frac{1}{2} = 1$$

What happens with Game 2?

$$\frac{1}{2} \cdot (100000 - 0)^2 + \frac{1}{2} \cdot (-100000 - 0)^2$$
$$\underline{5,000,000,000 + 5,000,000,000} = 10^{10}$$

Why Squaring

Why not absolute value? Or fourth power?

- > Squaring is nicer algebraically. 
- > Our goal with variance was to talk about the spread of results.
Squaring makes extreme results even more extreme. 
- > Fourth power over-emphasizes the extreme results (for our purposes).

Variance

$$E[X] = c$$

$$E[(X - E[X])^2] = E[X^2 - 2E[X]X + E[X]^2]$$

Use LOTUS to compute $E[X^2]$ ($g(x) = x^2$)

Variance

The variance of a random variable X is

$$\begin{aligned}\text{Var}(X) &= \sum_{\omega} \mathbb{P}(\omega) \cdot (X(\omega) - E[X])^2 = E[(X - E[X])^2] \\ &= E[X^2] - E[X]^2\end{aligned}$$

$$= E[X^2] - E[X]^2$$

$$E[2E[X] - X]$$

$\downarrow$
 $\downarrow$
 $\downarrow$
 $E[X]$

The first form forms are the definition. The last one is an algebra trick.

Intuition: Variance is a quantity that measures on average *how "far" the random variable is from its expectation*

- > higher values means values are very spread out
- > smaller values means values are closer and tend to be closer to expectation

Proof of Calculation Trick

$$\mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2 - 2X\mathbb{E}[X] + (\mathbb{E}[X])^2] \text{ expanding the square}$$
$$=$$

Proof of Calculation Trick

$$\begin{aligned}\mathbb{E}[(X - \mathbb{E}[X])^2] &= \mathbb{E}[X^2 - 2X\mathbb{E}[X] + (\mathbb{E}[X])^2] \text{ expanding the square} \\ &= \mathbb{E}[X^2] - \mathbb{E}[2X\mathbb{E}[X]] + \mathbb{E}[(\mathbb{E}[X])^2] \text{ linearity of expectation.} \\ &= \mathbb{E}[X^2] - 2\mathbb{E}[X]\mathbb{E}[X] + \mathbb{E}[(\mathbb{E}[X])^2] \text{ linearity of expectation.} \\ &= \mathbb{E}[X^2] - 2\mathbb{E}[X]\mathbb{E}[X] + (\mathbb{E}[X])^2 \text{ expectation of a constant is the constant} \\ &= \mathbb{E}[X^2] - 2(\mathbb{E}[X])^2 + (\mathbb{E}[X])^2 \\ &= \mathbb{E}[X^2] - (\mathbb{E}[X])^2\end{aligned}$$

$$\text{So } \text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2.$$

Variance of a die

$$\checkmark E[X] = 3.5$$

Let X be the result of rolling a fair die.

$$\begin{aligned}\text{Var}(X) &= \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[(X - 3.5)^2] \\ &= \frac{1}{6}(1 - 3.5)^2 + \frac{1}{6}(2 - 3.5)^2 + \frac{1}{6}(3 - 3.5)^2 + \frac{1}{6}(4 - 3.5)^2 + \frac{1}{6}(5 - 3.5)^2 + \frac{1}{6}(6 - 3.5)^2 \\ &= \frac{35}{12} \approx 2.92.\end{aligned}$$

$$\text{Or } \mathbb{E}[X^2] - (E[X])^2 = \sum_{k=1}^6 \frac{1}{6} \cdot k^2 - 3.5^2 = \frac{91}{6} - 3.5^2 \approx 2.92$$

Variance of n Coin Flips

Flip a coin n times, where it comes up heads with probability p each time (independently). Let X be the total number of heads.

We saw last time $\mathbb{E}[X] = np$.

$$X_i = \begin{cases} 1 & \text{if flip } i \text{ is heads} \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{E}[X] = \mathbb{E}[\sum_{i=1}^n X_i] = \sum_{i=1}^n \mathbb{E}[X_i] = \sum_{i=1}^n p = np.$$

Variance of n Coin Flips

Flip a coin n times, where it comes up heads with probability p each time (independently). Let X be the total number of heads.

What about $\text{Var}(X)$

$$\begin{aligned}\mathbb{E}[(X - \mathbb{E}[X])^2] &= \sum_{k \in \Omega_X} \mathbb{P}(X = k) (k - np)^2 \\ &= \sum_{k=0}^n \binom{n}{k} \cdot p^k (1-p)^{n-k} \cdot (k - np)^2\end{aligned}$$




LOTS

Algebra time?

Variance Adds If Independent

If X and Y are independent then
 $\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$

Are the X_i independent? Yes!

In this problem X_i is independent of X_j for $i \neq j$ where

$$X_i = \begin{cases} 1 & \text{if flip } i \text{ was heads} \\ 0 & \text{otherwise} \end{cases}$$

Variance Adds If Independent

$$X_i = \begin{cases} 1 & \text{event} \\ 0 & \text{ow.} \end{cases}$$

$$X_i^2 = \begin{cases} 1 \\ 0 \end{cases}$$

$$\text{Var}(X) = \text{Var}(\sum_{i=1}^n X_i) = \sum_{i=1}^n \text{Var}(X_i)$$

X_i are independent

What's the $\text{Var}(X_i)$?

$$E[X_i^2] = E[X_i] \text{ indicator RV}$$

$$\text{Var}(X_i) = E[X_i^2] - E[X_i]^2$$

$$= E[X_i] - p^2$$

$$= p - p^2$$

$$= p(1 - p)$$

Variance Adds If Independent

$$\text{Var}(X) = \text{Var}(\sum_{i=1}^n X_i) = \sum_{i=1}^n \text{Var}(X_i)$$

What's the $\text{Var}(X_i)$?

$p(1 - p)$.


$$\text{Var}(X) = \sum_{i=1}^n p(1 - p) = np(1 - p).$$


Expectation and Variance aren't everything

Alright, so expectation and variance is everything right?

No!

Flip a fair coin 3 times indep. Count heads.


PMF 1 with $E=3/2$, $Var=3/4$



→ support for X

Flip a biased coin (prob heads=2/3) until heads. Count flips.

PMF 2 with $E=3/2$, $Var=3/4$



A PMF or CDF *does* fully describe a random variable.



Useful Facts

-

Make a prediction

$$\text{Var}(X+2)$$

How should $\text{Var}(X + c)$ relate to $\text{Var}(X)$ if c is a constant?

How should $\text{Var}(aX)$ relate to $\text{Var}(X)$ if a is a constant?

$$\text{Var}(5X)$$

Facts About Variance

$$\text{Var}(X + c) = \text{Var}(X)$$

Proof:

$$\begin{aligned}\text{Var}(X + c) &= \mathbb{E}[(X + c)^2] - \mathbb{E}[X + c]^2 \\&= \mathbb{E}[X^2] + \mathbb{E}[2Xc] + \mathbb{E}[c^2] - (\mathbb{E}[X] + c)^2 \\&= \mathbb{E}[X^2] + 2c\mathbb{E}[X] + c^2 - \mathbb{E}[X]^2 - 2c\mathbb{E}[X] - c^2 \\&= \mathbb{E}[X^2] - \mathbb{E}[X]^2 \\&= \text{Var}(X)\end{aligned}$$

Intuition: Adding a constant just shifts the distribution – the spread says the same

Facts about Variance

$$\text{Var}(aX) =$$

Facts about Variance

$$\begin{aligned}\text{Var}(aX) &= a^2 \text{Var}(X) \\ &= \mathbb{E}[(aX)^2] - (\mathbb{E}[aX])^2 \\ &= a^2 \mathbb{E}[X^2] - (a\mathbb{E}[X])^2 \\ &= \underline{a^2} \mathbb{E}[X^2] - \underline{a^2} \mathbb{E}[X]^2 \\ &= a^2 (\mathbb{E}[X^2] - \mathbb{E}[X]^2)\end{aligned}$$

$$\mathbb{E}[X^2] - \mathbb{E}[X]^2$$

$$\boxed{a^2 (\text{Var}(X))}$$

Intuition: Multiplying by a positive constant makes it more spread out

Summary of *Expectation* and *Variance*

Expectation ("on average")

$$\mathbb{E}[X] = \sum_{k \in \Omega_X} k \cdot \mathbb{P}(X = k)$$

$$> \mathbb{E}[g(X)] = \sum_{k \in \Omega_X} g(k) \cdot \mathbb{P}(X = k)$$

Linearity of Expectation (LoE)

$$\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y] \text{ (always true)}$$

Shifting and Scaling:

$$\mathbb{E}[aX + b] = a\mathbb{E}[X] + b$$

If X and Y are independent:

$$\mathbb{E}[X \cdot Y] = \mathbb{E}[X]\mathbb{E}[Y]$$

Variance ("how spread out")

$$\text{Var}(X) = \mathbb{E}[X^2] - \mathbb{E}[X]^2$$

Linearity of Variance


only if X and Y are independent

Shifting and Scaling:

$$\text{Var}(aX + b) = a^2 \text{Var}(X)$$

If X and Y are independent:

$$\text{Var}(X + Y) = \text{Var}(X) + \text{Var}(Y)$$



Discrete Random Variable Zoo



Discrete Random Variable Zoo

There are common **patterns** of experiments:

- > Flip a [fair/unfair] coin [blah] times and count the number of heads.
- > Flip a [fair/unfair] coin until the first time that you see a heads
- > Draw a uniformly random element from [set]
- > Define an indicator random variable for [event]

...

Instead of calculating the support, PMF, CDF, expectation, variance every time, why not calculate it once and look it up every time?

for example...

E.g.

- the number of coin tosses till we get the first heads
- the number of games we need to play till we first win
- the number of dice rolls till the first 6

all follow the same format as,

- the number of trials till we get the first success

The discrete zoo defines some common patterns and gives us the PMF, expectation, and variance, so we don't have to compute it every time!

What's our goal?

Your goal is NOT to memorize these facts (it'll be convenient to memorize some of them, but don't waste time making flash cards). Everything is on Wikipedia anyway. Everyone checks Wikipedia when they forget these.



Clément Canonne

@ccanonne_



Are you trying to shame me, @Google?

[https://en.wikipedia.org > wiki > Binomial_distribution](https://en.wikipedia.org/wiki/Binomial_distribution) ▼

Binomial distribution - Wikipedia

In probability theory and statistics, the **binomial distribution** with parameters n and p is the discrete probability distribution of the number of successes ...

[Negative binomial distribution](#) · [Poisson binomial](#) · [Binomial test](#) · [Beta-binomial](#)

[You've visited this page many times](#). Last visit: 9/07/21

9:08 PM · Aug 11, 2021


What's our goal?

Your goal is NOT to memorize these facts (it'll be convenient to memorize some of them, but don't waste time making flash cards). Everything is on Wikipedia anyway. Everyone checks Wikipedia when they forget these.

While learning about this zoo, we will also get the chance too:

0. Introduce one new distribution we haven't seen at all (next time).
1. Practice expectation, variance, etc. for ones we have gotten hints of.
2. Review the first half of the course with some probability calculations.

Zoo!



$X \sim \text{Unif}(a, b)$

$$p_X(k) = \frac{1}{b - a + 1}$$

$$\mathbb{E}[X] = \frac{a + b}{2}$$

$$\text{Var}(X) = \frac{(b - a)(b - a + 2)}{12}$$

$X \sim \text{Ber}(p)$

$$p_X(0) = 1 - p;$$

$$p_X(1) = p$$

$$\mathbb{E}[X] = p$$

$$\text{Var}(X) = p(1 - p)$$

$X \sim \text{Bin}(n, p)$

$$p_X(k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

$$\mathbb{E}[X] = np$$

$$\text{Var}(X) = np(1 - p)$$

$X \sim \text{Geo}(p)$

$$p_X(k) = (1 - p)^{k-1} p$$

$$\mathbb{E}[X] = \frac{1}{p}$$

$$\text{Var}(X) = \frac{1 - p}{p^2}$$

$X \sim \text{NegBin}(r, p)$

$$p_X(k) = \binom{k-1}{r-1} p^r (1 - p)^{k-r}$$

$$\mathbb{E}[X] = \frac{r}{p}$$

$$\text{Var}(X) = \frac{r(1 - p)}{p^2}$$

$X \sim \text{HypGeo}(N, K, n)$

$$p_X(k) = \frac{\binom{K}{k} \binom{N-K}{n-k}}{\binom{N}{n}}$$

$$\mathbb{E}[X] = n \frac{K}{N}$$

$$\text{Var}(X) = \frac{K(N - K)(N - n)}{N^2(N - 1)}$$

$X \sim \text{Poi}(\lambda)$

$$p_X(k) = \frac{\lambda^k e^{-\lambda}}{k!}$$

$$\mathbb{E}[X] = \lambda$$

$$\text{Var}(X) = \lambda$$

Scenario: Uniform

You want an integer in some range, with each integer equally likely.

$$k \in [a, b], k \in \mathbb{Z}$$

Familiar example:

Outcome of rolling a fair die (or draw a random integer from $1, \dots, n$)
equally likely to be an integer between 1 and 6

Discrete Uniform Distribution

$$\underline{X \sim \text{Unif}(a, b)}$$

← parameters

X is a uniformly random integer between a and b (inclusive)

Parameter a is the minimum value in the support, b is the maximum value in the support.

$$\text{PMF: } p_X(k) = \frac{1}{b-a+1} \text{ for } k \in \mathbb{Z}, \underline{a \leq k \leq b}$$

$$\text{CDF: } \underline{F_X(k)} = \frac{k-a+1}{b-a+1} \text{ for } k \in \mathbb{Z}, a \leq k \leq b. \quad \mathcal{P}(X \leq k) = F_X(k)$$

$$\text{Expectation: } \mathbb{E}[X] = \frac{a+b}{2}$$

$$\text{Variance: } \text{Var}(X) = \frac{(b-a)(b-a+1)}{12}$$

Situation: Bernoulli

Do I get a success in one trial with probability p of "success"?

Familiar examples:

You flip a biased coin (once) and want to record whether its heads.

An indicator random variable that is either 1 or 0

$$X_i = \begin{cases} 1 & \text{successful event} \\ 0 & \text{otherwise} \end{cases}$$

success = heads

Bernoulli Distribution

$$\underline{X} \sim \underline{\text{Ber}(p)}$$

X is the indicator random variable that the trial was a success.

Parameter p is probability of success on the trial.

$$\text{PMF: } \underline{p_X(0)} = 1 - p, \underline{p_X(1)} = p$$

$$\text{CDF: } \underline{F_X(k)} = \begin{cases} 0 & \text{if } k < 0 \\ 1 - p & \text{if } 0 \leq k < 1 \\ 1 & \text{if } k \geq 1 \end{cases}$$

all $\mathbb{R}$? $k \in \mathbb{N}$ $\Omega_X = \{0, 1\}$

$\leftarrow P(X \leq 1) = P(X \leq \underline{1.2})$

$$\text{Expectation: } \mathbb{E}[X] = p = P(X=1)$$

$$\text{Variance: } \text{Var}(X) = p(1 - p)$$

Bernoulli Distribution Examples

one trial

Did a particular bit get written correctly on the device?

Did you guess right on a multiple-choice test?

Did a server in a cluster fail?

Did a disk drive crash?

Did buying a particular share of a stock pay off?

Did a user like or dislike a YouTube video? (Back in ye olden days, at least)

Does a user click an ad?

In each of these, we find the probability of success on a trial, define a RV as an indicator for whether the event occurred, say the RV follows a bernoulli distribution with relevant and we have the PMF, $E[X]$ and $\text{Var}(X)$!

Situation: Binomial

How many success did you see in n independent trials, where each trial has probability p of success? ★

Familiar example:

You flip a coin n times independently, each with a probability p of coming up heads. How many heads are there?

X - # of successes

Binomial Distribution

$$X \sim \text{Bin}(n, p)$$

X is the number of successes across n independent trials.

n is the number of independent trials.

p is the probability of success for one trial.

$$\text{PMF: } p_X(k) = \binom{n}{k} p^k (1-p)^{n-k} \text{ for } k \in \{0, 1, \dots, n\}$$

CDF: F_X is ugly.

$$\text{Expectation: } \mathbb{E}[X] = np$$

$$\text{Variance: } \text{Var}(X) = \underline{np(1-p)}$$

all n successful

↓

$k=3$

↓

$$P(X \leq k)$$

↑
0 successes

$$\begin{aligned} &= P(X=0) + \\ &\quad P(X=1) + \\ &\quad P(X=2) + \\ &\quad P(X=3) \end{aligned}$$

Binomial Distribution *Examples*

How many bits were written correctly on the device?

How many questions did you guess right on a multiple-choice test?

How many servers in a cluster failed? P

How many keys went to one bucket in a hash table?

Example: Unpopular Donuts

$$X \sim \text{Bin}(50, 0.2)$$

A donut shop serves 50 people a day and serves a *mango chili lime* donut. The probability that a customer chooses this donut is 0.2. All customers' choices are independent of each other.

$$X \sim \text{Bin}(n=50, p=0.2)$$

What is the probability that exactly 10 people choose this flavor? ✓

$$\underline{\underline{P(X=10)}}$$

What is the probability that at least 3 people choose this flavor?

Example: Unpopular Donuts

A donut shop serves 50 people a day and serves a *mango chili lime* donut. The probability that a customer chooses this donut is 0.2. All customers' choices are independent of each other.

What is the probability that exactly 10 people choose this flavor?


$$\text{PMF: } p_X(k) = \binom{n}{k} p^k (1 - p)^{n-k} \text{ for } k \in \{0, 1, \dots, n\}$$

What is the probability that at least 3 people choose this flavor?

Example: Unpopular Donuts

A donut shop serves 50 people a day and serves a *mango chili lime* donut. The probability that a customer chooses this donut is 0.2. All customers' choices are independent of each other.

What is the probability that exactly 10 people choose this flavor?

$X \sim$ number of people who choose this flavor. $X \sim \text{Bin}(50, 0.2)$

$$\mathbb{P}(X = 10) = \binom{50}{10} 0.2^{10} (1 - 0.2)^{50-10}$$

What is the probability that at least 3 people choose this flavor?

$$\begin{aligned} \mathbb{P}(X \geq 3) &= 1 - \mathbb{P}(X < 3) = 1 - (\mathbb{P}(X = 0) + \mathbb{P}(X = 1) + \mathbb{P}(X = 2)) \\ &= 1 - ((\binom{50}{0} 0.2^0 (0.8)^{50-0} + \binom{50}{1} 0.2^1 (0.8)^{50-1} + \binom{50}{2} 0.2^2 (0.8)^{50-2}) \end{aligned}$$

$\mathbb{P}(X=0)$ $\mathbb{P}(X=1)$ $\mathbb{P}(X=2)$

Situation: Geometric

How many *independent* trials are needed until the first success?

Familiar Example:

You flip a coin (which comes up heads with probability p) until you get a heads. How many flips did you need?

X = # flips to & first success
including

T T T T H
↑

$x = 5$

Geometric Distribution

$$X \sim \text{Geo}(p)$$

X is the number of trials needed to see the first success.

p is the probability of success for one trial.

$$\text{PMF: } p_X(k) = (1-p)^{k-1}p \text{ for } k \in \{1, 2, 3, \dots\}$$

$$F_X(k) = 1 - (1-p)^k \text{ for } k \in \mathbb{N}$$

$$P(X \leq k) = 1 - P(X > k)$$

$$\mathbb{E}[X] = \frac{1}{p}$$

$$\text{Var}(X) = \frac{1-p}{p^2}$$

Geometric Distribution *Examples*

How many bits can we write *before one is incorrect*?

How many questions do you have to answer *until you get one right*?

How many times can you run an experiment *until it fails for the first time*?

Geometric: Analysis

$$E[X] = \frac{1}{p}$$
$$\text{Var}(X) = \frac{1-p}{p^2}$$

Both the expectation and variance are new to us.

The derivations of both are uninformative

Every derivation I've ever seen has wild algebra tricks.

Geometric: Expectation

$$\begin{aligned}\mathbb{E}[X] &= \sum_{k=1}^{\infty} \underbrace{k(1-p)^{k-1}p}_{\substack{\text{red arrow from } \mathbb{P}(X=k) \\ \text{red arrow from } \mathbb{E}[X]}} \\ &= \underbrace{p}_{\text{red underline}} \sum_{k=1}^{\infty} \underbrace{k(1-p)^{k-1}}_{\text{red underline}} = \underbrace{p}_{\text{red underline}} \cdot \underbrace{\frac{1}{p^2}}_{\text{red underline}} = \underbrace{\frac{1}{p}}_{\text{red underline}}.\end{aligned}$$

Intuition: Smaller p means longer wait

$$\begin{aligned}\text{Var}(X) &= \mathbb{E}[X^2] - (\mathbb{E}[X])^2 \\ &= \frac{2-p}{p^2} - \frac{1}{p^2} = \frac{1-p}{p^2}\end{aligned}$$

(Red arrows point from the p^2 terms in the final expression to the p terms in the denominator of the first expression.)

Intuition: for small p lots of variance
(might have to wait a long time, and it's
variable)

For large p very little variance (for $p = 1$
there's no variation at all!)

Geometric Property

Geometric random variables are called "memoryless"

Suppose you're flipping coins (independently) until you see a heads.

The first three came up tails.

How many flips are left until you see the first heads?

It's another independent copy of the original!

The coin "forgot" it already came up tails 3 times.

Formally...

Let X be the number of flips needed, Y be the flips after the third.

$$\mathbb{P}(Y = \underline{k} | X \geq 3) = ? \quad \rightarrow \quad \frac{\mathbb{P}(Y = k \cap X \geq 3)}{\mathbb{P}(X \geq 3)} = \frac{(1-p)^{\cancel{k-3}-1} p}{(1-p)^3}$$

...

$E[Y \mid \underline{3 \text{ times}}]$
is something else

Which is $p_X(k)$.

$$= (1-p)^{k-1} p$$

Formally...

Let X be the number of flips needed, Y be the flips after the third.

$$\mathbb{P}(Y = k | X \geq 3) = \mathbb{P}(Y = k \cap X \geq 3) / \mathbb{P}(X \geq 3)$$

$$\frac{(1-p)^{k+3-1}p}{(1-p)^3}$$

$$= (1-p)^{k-1}p$$

Which is $p_X(k)$.

Summary

Today we covered the following random variables from the zoo:

- • **Bernoulli:** Whether there is success in one trial
 $\text{Ber}(p)$ is 1 with probability p and 0 otherwise
- **Binomial:** Number of successes in n independent trials
 $\text{Bin}(n, p)$ - n independent trials, probability p of success on each trial
- **Geometric:** Number of trials until the first success
 $\text{Geo}(p)$ - probability p of success on a single trial
- **Uniform:** Every integer between a and b are equally likely
 $\text{Unif}(a, b)$

More on Friday ☺

