Distributions

Hayk Aprikyan, Hayk Tarkhanyan

Distributions

In practice, random variables very often share similar properties: The distributions of their values seem to follow a common pattern, i.e. their PMF/PDFs are similar to each other:

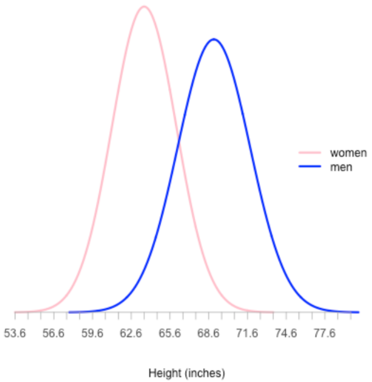

Distributions

In practice, random variables very often share similar properties: The distributions of their values seem to follow a common pattern, i.e. their PMF/PDFs are similar to each other:

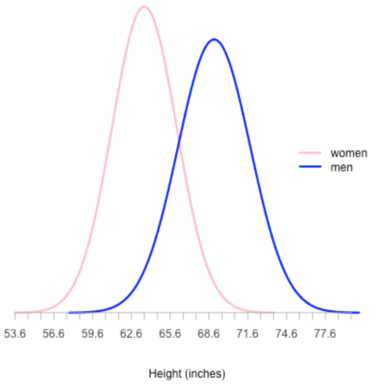

Some of these common patterns are so frequently observed that they have been given specific names.

Bernoulli Distribution

Consider these situations:

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.
- A child is born. The gender at birth is either *male* or *female*.

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.
- A child is born. The gender at birth is either *male* or *female*.

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.
- A child is born. The gender at birth is either *male* or *female*.

In each of these experiments, there are exactly two possible outcomes. We call such experiments *Bernoulli experiments.*

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.
- A child is born. The gender at birth is either *male* or *female*.

In each of these experiments, there are exactly two possible outcomes. We call such experiments *Bernoulli experiments.* Similarly,

Bernoulli Distribution

Consider these situations:

- You take a pass-fail exam. You either *pass* or *fail*.
- You toss a coin. The outcome is either *heads* or *tails*.
- A child is born. The gender at birth is either *male* or *female*.

In each of these experiments, there are exactly two possible outcomes. We call such experiments *Bernoulli experiments.* Similarly,

Definition

If a random variable X has exactly two possible values, say 0 and 1, we say that X follows a *Bernoulli distribution:*

$$\mathbb{P}[X = 1] = p, \qquad \mathbb{P}[X = 0] = 1 - p$$

and write $X \sim \text{Bernoulli}(p)$.

Bernoulli Distribution

$X \sim Bernoulli(p)$

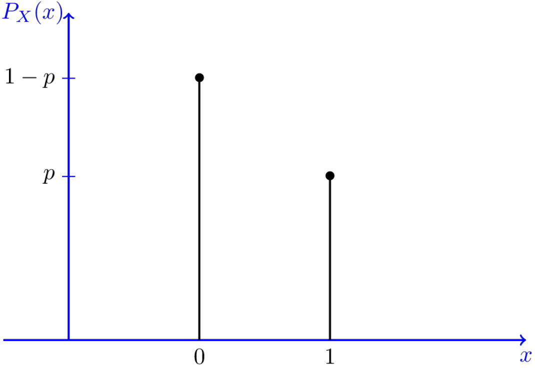

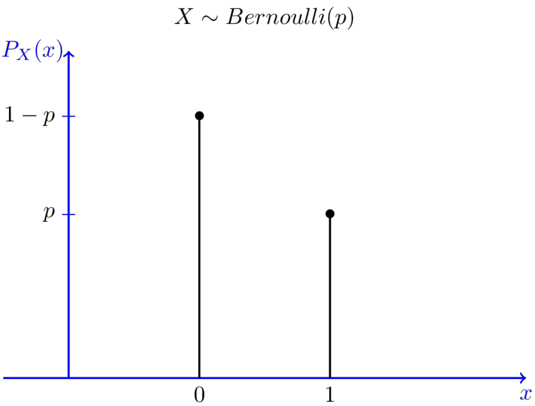

If $X \sim$ Bernoulli(p),

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

Bernoulli Distribution

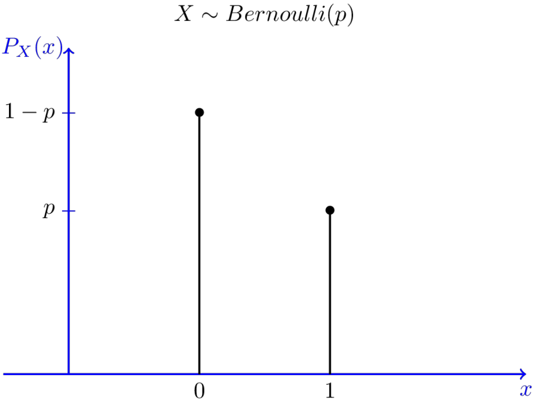

If $X \sim \text{Bernoulli}(p)$,

$$\mathbb{E}[X] = p, \qquad \text{Var}[X] =$$

Bernoulli Distribution

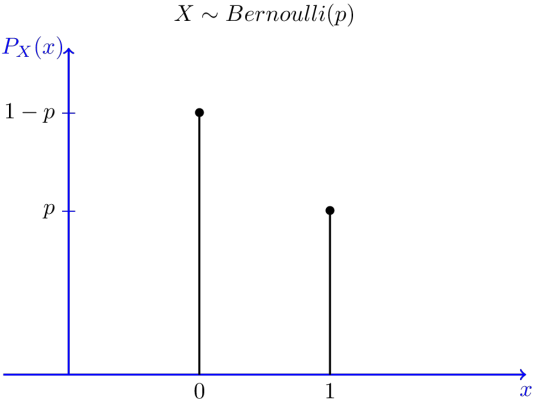

If $X \sim \text{Bernoulli}(p)$,

$$\mathbb{E}[X] = p, \qquad \text{Var}[X] = p(1-p)$$

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car $(X = 1)$ or not $(X = 0)$ with a single ticket, what is the expected value and variance of X?

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car ($X = 1$) or not ($X = 0$) with a single ticket, what is the expected value and variance of X?

What if we buy multiple tickets?

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car ($X = 1$) or not ($X = 0$) with a single ticket, what is the expected value and variance of X?

What if we buy multiple tickets? Let's look at two possible strategies:

- Buy 100 tickets at once, try with each ticket independently.

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car ($X = 1$) or not ($X = 0$) with a single ticket, what is the expected value and variance of X?

What if we buy multiple tickets? Let's look at two possible strategies:

- Buy 100 tickets at once, try with each ticket independently.
- Buy tickets one by one, until you win the car.

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car $(X = 1)$ or not $(X = 0)$ with a single ticket, what is the expected value and variance of X?

What if we buy multiple tickets? Let's look at two possible strategies:

- Buy 100 tickets at once, try with each ticket independently.
- Buy tickets one by one, until you win the car.

Bernoulli Distribution

Example

The probability of winning a car with a single lottery ticket is 0.0001. If X is the random variable telling whether you win the car ($X = 1$) or not ($X = 0$) with a single ticket, what is the expected value and variance of X?

What if we buy multiple tickets? Let's look at two possible strategies:

- Buy 100 tickets at once, try with each ticket independently.
- Buy tickets one by one, until you win the car.

Let's begin with the second strategy.

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car.

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] =$$
$$\mathbb{P}[X = 2] =$$
$$\mathbb{P}[X = 3] =$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] =$$
$$\mathbb{P}[X = 2] =$$
$$\mathbb{P}[X = 3] =$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] =$$
$$\mathbb{P}[X = 3] =$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)$$
$$\mathbb{P}[X = 3] =$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] =$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] = (1 - p)$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] = (1 - p)(1 - p)$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] = (1 - p)(1 - p)p$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] = (1 - p)(1 - p)p$$

so in general, for any natural number k,

$$\mathbb{P}[X = k] = (1 - p)^{k-1}p$$

Geometric Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

How many tickets do we expect to buy until we win the car?

Let X be the **number of tickets** we buy until we win the car. Then

$$\mathbb{P}[X = 1] = p$$
$$\mathbb{P}[X = 2] = (1 - p)p$$
$$\mathbb{P}[X = 3] = (1 - p)(1 - p)p$$

so in general, for any natural number k,

$$\mathbb{P}[X = k] = (1 - p)^{k-1}p$$

In this case, we say that X follows a *geometric distribution*.

Geometric Distribution

Definition

If the PMF of a random variable X has the following form:

$$\mathbb{P}[X = k] = (1-p)^{k-1}p, \qquad k = 1, 2, 3, \ldots$$

we say that X follows a *geometric distribution* with parameter p, and write $X \sim \text{Geo}(p)$.

Geometric Distribution

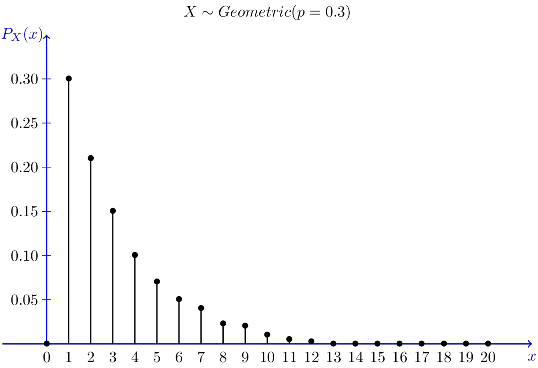

Geometric Distribution

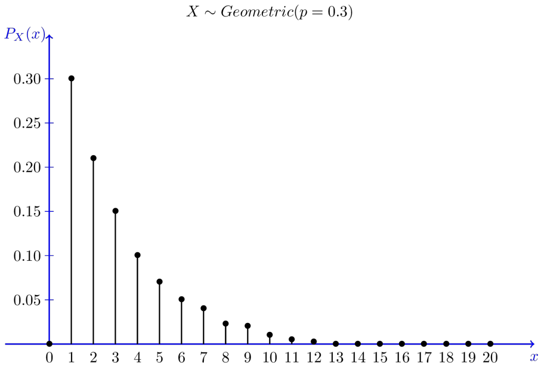

If $X \sim \text{Geo}(p)$,

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

Geometric Distribution

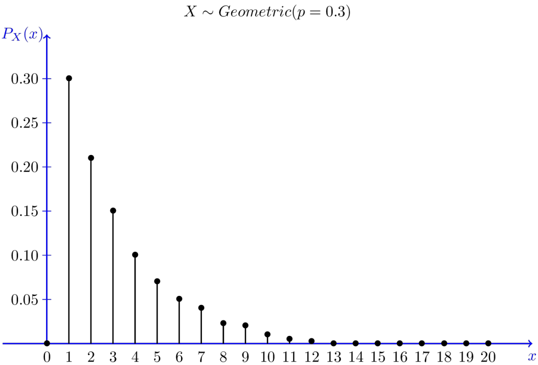

If $X \sim \text{Geo}(p)$,

$$\mathbb{E}[X] = \frac{1}{p}, \qquad \text{Var}[X] = \frac{1-p}{p^2}$$

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Simply because you have lost 400 times in a row, it does not change the probability of winning on the next ticket – it stays the same!

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Simply because you have lost 400 times in a row, it does not change the probability of winning on the next ticket – it stays the same!

Similarly, the probability of losing on the 401st ticket but winning on the 402nd, given that you have lost the first 400 tickets, is the same as losing on the 1st ticket but winning on the 2nd.

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Simply because you have lost 400 times in a row, it does not change the probability of winning on the next ticket – it stays the same!

Similarly, the probability of losing on the 401st ticket but winning on the 402nd, given that you have lost the first 400 tickets, is the same as losing on the 1st ticket but winning on the 2nd.

Moral: Geometric probability does not care about your past losses or wins!

Geometric Distribution

Question

Gevorg and Hayk buy lottery tickets, each with probability of winning p.

1. Is the probability of winning the same for both of them?
2. Is the probability of winning the same for both of them, if we know that before that, Gevorg had bought 400 tickets and never won?

Simply because you have lost 400 times in a row, it does not change the probability of winning on the next ticket – it stays the same!

Similarly, the probability of losing on the 401st ticket but winning on the 402nd, given that you have lost the first 400 tickets, is the same as losing on the 1st ticket but winning on the 2nd.

Moral: Geometric probability does not care about your past losses or wins!

Theorem

Geometric random variables are *memoryless*, i.e. for any m and n,

$$\mathbb{P}[X > m + n \mid X > m] = \mathbb{P}[X > n]$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] =$$

$$\mathbb{P}[X = 1] =$$

$$\mathbb{P}[X = 2] =$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] =$$

$$\mathbb{P}[X = 2] =$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = \quad p$$

$$\mathbb{P}[X = 2] =$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1-p)^n$$

$$\mathbb{P}[X = 1] = p(1-p)^{n-1}$$

$$\mathbb{P}[X = 2] =$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1-p)^n$$

$$\mathbb{P}[X = 1] = np(1-p)^{n-1}$$

$$\mathbb{P}[X = 2] =$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X=0] = (1-p)^n$$

$$\mathbb{P}[X=1] = np(1-p)^{n-1}$$

$$\mathbb{P}[X=2] = \quad p^2$$

$$\mathbb{P}[X=3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = np(1 - p)^{n-1}$$

$$\mathbb{P}[X = 2] = \qquad p^2(1 - p)^{n-2}$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = np(1 - p)^{n-1}$$

$$\mathbb{P}[X = 2] = \binom{n}{2} p^2 (1 - p)^{n-2}$$

$$\mathbb{P}[X = 3] =$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = np(1 - p)^{n-1}$$

$$\mathbb{P}[X = 2] = \binom{n}{2} p^2 (1 - p)^{n-2}$$

$$\mathbb{P}[X = 3] = \binom{n}{3}$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = np(1 - p)^{n-1}$$

$$\mathbb{P}[X = 2] = \binom{n}{2} p^2 (1 - p)^{n-2}$$

$$\mathbb{P}[X = 3] = \binom{n}{3} p^3$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X = 0] = (1 - p)^n$$

$$\mathbb{P}[X = 1] = np(1 - p)^{n-1}$$

$$\mathbb{P}[X = 2] = \binom{n}{2} p^2 (1 - p)^{n-2}$$

$$\mathbb{P}[X = 3] = \binom{n}{3} p^3 (1 - p)^{n-3}$$

Binomial Distribution

Assumptions:

- The probability of winning on each ticket is p.
- Each ticket is independent from the others.

Question

Suppose we buy n tickets. What is the probability of winning exactly k times?

$$\mathbb{P}[X=0] = (1-p)^n$$

$$\mathbb{P}[X=1] = np(1-p)^{n-1}$$

$$\mathbb{P}[X=2] = \binom{n}{2} p^2 (1-p)^{n-2}$$

$$\mathbb{P}[X=3] = \binom{n}{3} p^3 (1-p)^{n-3}$$

$$\mathbb{P}[X=k] = \binom{n}{k} p^k (1-p)^{n-k}$$

Binomial Distribution

Definition

If a random variable X takes values in $\{0, 1, 2, \ldots, n\}$ with probabilities

$$\mathbb{P}[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}$$

Binomial Distribution

Definition

If a random variable X takes values in $\{0, 1, 2, \dots, n\}$ with probabilities

$$\mathbb{P}[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}$$

we say that X follows a *binomial distribution* with parameters n and p, and write $X \sim \mathrm{B}(n, p)$.

Binomial Distribution

Definition

If a random variable X takes values in $\{0, 1, 2, \dots, n\}$ with probabilities

$$\mathbb{P}[X = k] = \binom{n}{k} p^k (1 - p)^{n-k}$$

we say that X follows a *binomial distribution* with parameters n and p, and write $X \sim \mathrm{B}(n, p)$.

Example

- We toss a coin 10 times. If X is the number of heads, what is the distribution of X?

Binomial Distribution

Definition

If a random variable X takes values in $\{0, 1, 2, \ldots, n\}$ with probabilities

$$\mathbb{P}[X = k] = \binom{n}{k} p^k (1-p)^{n-k}$$

we say that X follows a *binomial distribution* with parameters n and p, and write $X \sim \mathrm{B}(n, p)$.

Example

- We toss a coin 10 times. If X is the number of heads, what is the distribution of X?
- The probability of a lightbulb being defective is 0.1. If we randomly buy 50 lightbulbs, what is the probability that exactly 4 are defective?

Binomial Distribution

Definition

If a random variable X takes values in $\{0, 1, 2, \ldots, n\}$ with probabilities

$$\mathbb{P}[X = k] = \binom{n}{k} p^k (1-p)^{n-k}$$

we say that X follows a *binomial distribution* with parameters n and p, and write $X \sim \mathrm{B}(n, p)$.

Example

- We toss a coin 10 times. If X is the number of heads, what is the distribution of X?
- The probability of a lightbulb being defective is 0.1. If we randomly buy 50 lightbulbs, what is the probability that exactly 4 are defective?
- In Yerevan, paying in bus by card fails 30% of the time. If a person takes 5 buses, what is the probability of being able to pay in all buses?

Binomial Distribution

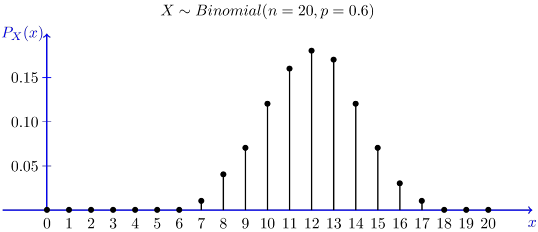

Binomial Distribution

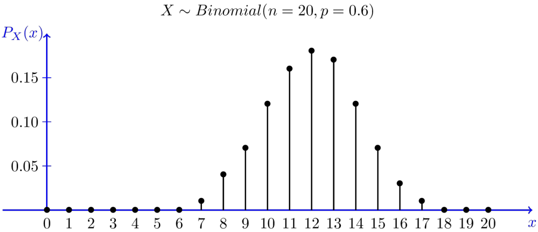

If $X \sim \mathrm{B}(n, p)$,

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

Binomial Distribution

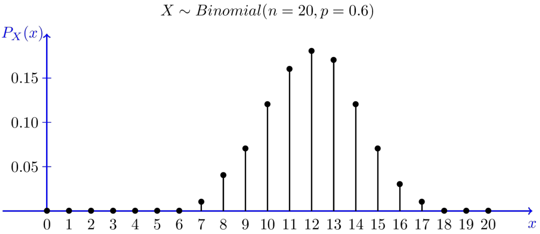

If $X \sim \mathrm{B}(n, p)$,

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

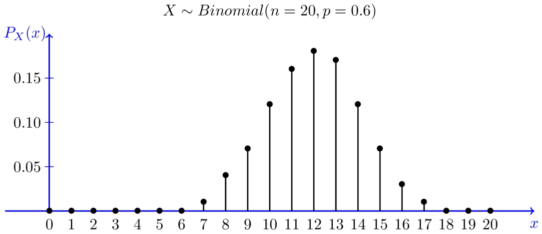

If $X \sim \mathrm{B}(n, p)$,

$$\mathbb{E}[X] = np, \qquad \mathrm{Var}[X] = np(1-p)$$

Poisson Distribution

Let's consider another situation:

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

One way to model this situation is to

- divide the hour into many small intervals, say $n = 3600$,

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

One way to model this situation is to

- divide the hour into many small intervals, say $n = 3600$,

but then

- in each interval, the probability of a customer arriving in each interval would be roughly $p \approx 30/n$.

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

One way to model this situation is to

- divide the hour into many small intervals, say $n = 3600$,

but then

- in each interval, the probability of a customer arriving in each interval would be roughly $p \approx 30/n$.

So the total number of customers can be modeled as a *binomial* random variable $X \sim \mathrm{B}(n, p)$:

$$\mathbb{P}[X = k] = \binom{n}{k} \cdot \left(\frac{30}{n}\right)^k \cdot \left(1 - \frac{30}{n}\right)^{n-k}$$

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

One way to model this situation is to

- divide the hour into many small intervals, say $n = 3600$,

but then

- in each interval, the probability of a customer arriving in each interval would be roughly $p \approx 30/n$.

So the total number of customers can be modeled as a *binomial* random variable $X \sim \mathrm{B}(n, p)$:

$$\mathbb{P}[X = k] = \binom{n}{k} \cdot \left(\frac{30}{n}\right)^k \cdot \left(1 - \frac{30}{n}\right)^{n-k} = \dots$$

Poisson Distribution

Let's consider another situation:

Question

Each day, on average, 30 customers buy coffee from a given Coffee House between 10:00-11:00. Let X be the actual number of customers who buy coffee between 10:00-11:00 on a given day. What is the distribution of X?

One way to model this situation is to

- divide the hour into many small intervals, say $n = 3600$,

but then

- in each interval, the probability of a customer arriving in each interval would be roughly $p \approx 30/n$.

So the total number of customers can be modeled as a *binomial* random variable $X \sim \mathrm{B}(n, p)$:

$$\mathbb{P}[X = k] = \binom{n}{k} \cdot \left(\frac{30}{n}\right)^k \cdot \left(1 - \frac{30}{n}\right)^{n-k} = \ldots \to e^{-30}\frac{30^k}{k!} \qquad \text{as } n \to \infty$$

Poisson Distribution

Definition

If the PMF of a random variable X has the form:

$$\mathbb{P}[X = k] = \frac{\lambda^k}{k!} e^{-\lambda}, \qquad k = 0, 1, 2, \ldots$$

we say that X follows a *Poisson distribution* with parameter λ, and write $X \sim \text{Poisson}(\lambda)$.

Poisson Distribution

Definition

If the PMF of a random variable X has the form:

$$\mathbb{P}[X = k] = \frac{\lambda^k}{k!}e^{-\lambda}, \qquad k = 0, 1, 2, \dots$$

we say that X follows a *Poisson distribution* with parameter λ, and write $X \sim \text{Poisson}(\lambda)$.

The Poisson distribution is one of the most widely used probability distributions – usually used in scenarios where we are counting the *number of occurrences* of certain events in an interval of time or space.

Poisson Distribution

Definition

If the PMF of a random variable X has the form:

$$\mathbb{P}[X = k] = \frac{\lambda^k}{k!}e^{-\lambda}, \qquad k = 0, 1, 2, \ldots$$

we say that X follows a *Poisson distribution* with parameter λ, and write $X \sim \text{Poisson}(\lambda)$.

The Poisson distribution is one of the most widely used probability distributions – usually used in scenarios where we are counting the *number of occurrences* of certain events in an interval of time or space.

Example

- Number of emails received in an hour.

Poisson Distribution

Definition

If the PMF of a random variable X has the form:

$$\mathbb{P}[X = k] = \frac{\lambda^k}{k!} e^{-\lambda}, \qquad k = 0, 1, 2, \ldots$$

we say that X follows a *Poisson distribution* with parameter λ, and write $X \sim \text{Poisson}(\lambda)$.

The Poisson distribution is one of the most widely used probability distributions – usually used in scenarios where we are counting the *number of occurrences* of certain events in an interval of time or space.

Example

- Number of emails received in an hour.
- Number of phone calls received by a call center in a day.

Poisson Distribution

Definition

If the PMF of a random variable X has the form:

$$\mathbb{P}[X = k] = \frac{\lambda^k}{k!}e^{-\lambda}, \qquad k = 0, 1, 2, \ldots$$

we say that X follows a *Poisson distribution* with parameter λ, and write $X \sim \text{Poisson}(\lambda)$.

The Poisson distribution is one of the most widely used probability distributions – usually used in scenarios where we are counting the *number of occurrences* of certain events in an interval of time or space.

Example

- Number of emails received in an hour.
- Number of phone calls received by a call center in a day.
- Number of car accidents on Isakov Ave. in a year.

Poisson Distribution

$X \sim Poisson(\lambda = 5)$

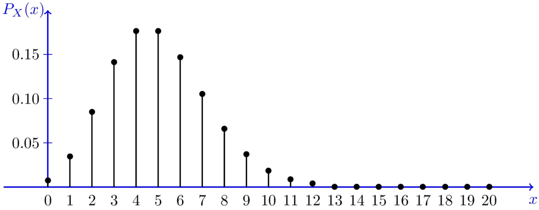

Poisson Distribution

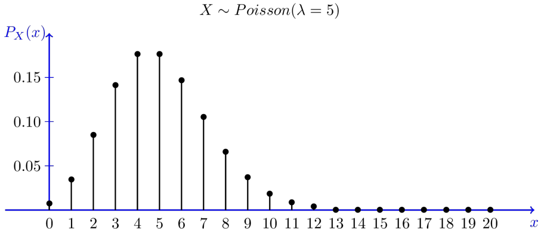

If $X \sim$ Poisson(λ),

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

Poisson Distribution

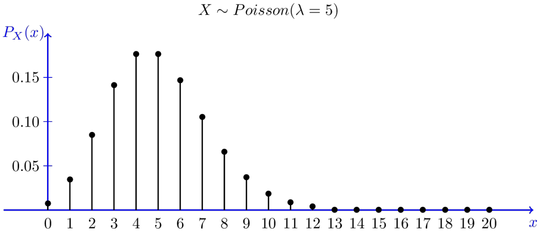

If $X \sim \text{Poisson}(\lambda)$,

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

Poisson Distribution

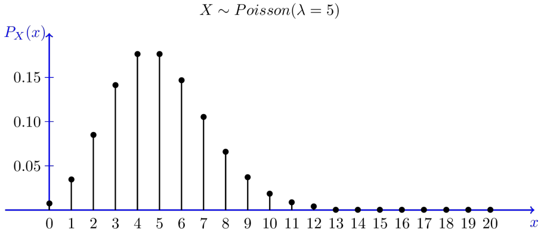

If $X \sim \text{Poisson}(\lambda)$,

$$\mathbb{E}[X] = \lambda, \qquad \text{Var}[X] = \lambda$$

What's the point?

Question

Why is this useful?

What's the point?

Question

Why is this useful?

In practice, statisticians only have very limited data of some N **samples**, they **do not know** the PMFs or PDFs (i.e. distributions) of random variables.

What's the point?

Question

Why is this useful?

In practice, statisticians only have very limited data of some N **samples**, they **do not know** the PMFs or PDFs (i.e. distributions) of random variables.

So, they try to **fit** the data to some known distributions (e.g. Binomial, Poisson, etc.) and use the properties of those (widely known) distributions to make predictions about future samples.

What's the point?

Question

Why is this useful?

In practice, statisticians only have very limited data of some N **samples**, they **do not know** the PMFs or PDFs (i.e. distributions) of random variables.

So, they try to **fit** the data to some known distributions (e.g. Binomial, Poisson, etc.) and use the properties of those (widely known) distributions to make predictions about future samples.

So far, we have considered only discrete RVs. Let's observe some common distributions for continuous RVs.

Uniform Distribution

If we pick a random number X from a given interval, without any number being "more probable" than another, then as we already know,

Uniform Distribution

If we pick a random number X from a given interval, without any number being "more probable" than another, then as we already know,

Definition

If a random variable X takes values from an interval (a, b) with equal probabilities, i.e. if its PDF is constant on (a, b):

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

then we say that X follows a *uniform distribution* on (a, b), and write $X \sim \mathrm{U}(a, b)$.

Uniform Distribution

If we pick a random number X from a given interval, without any number being "more probable" than another, then as we already know,

Definition

If a random variable X takes values from an interval (a, b) with equal probabilities, i.e. if its PDF is constant on (a, b):

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

then we say that X follows a *uniform distribution* on (a, b), and write $X \sim \mathrm{U}(a, b)$.

$$F(x) = \begin{cases} 0 & \text{if } x \leq a \\ \frac{x-a}{b-a} & \text{if } a < x < b \\ 1 & \text{if } x \geq b \end{cases}$$

Uniform Distribution

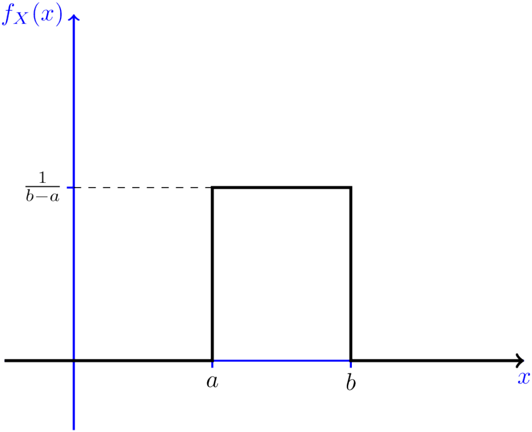

Uniform Distribution

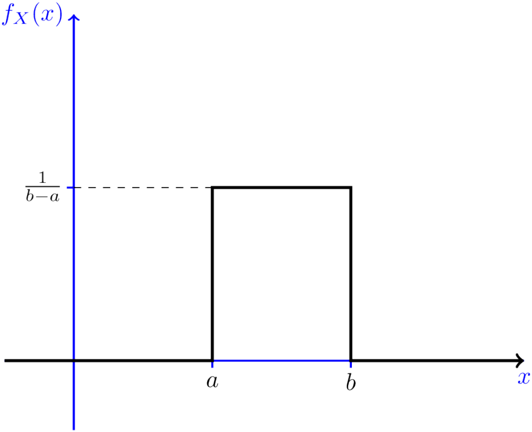

If $X \sim \mathrm{U}(a, b)$,

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

Uniform Distribution

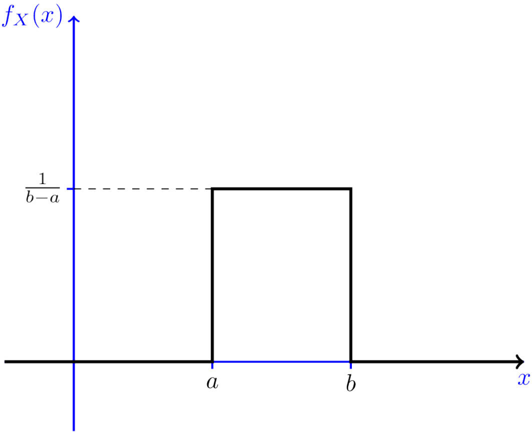

If $X \sim \mathrm{U}(a, b)$,

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

Uniform Distribution

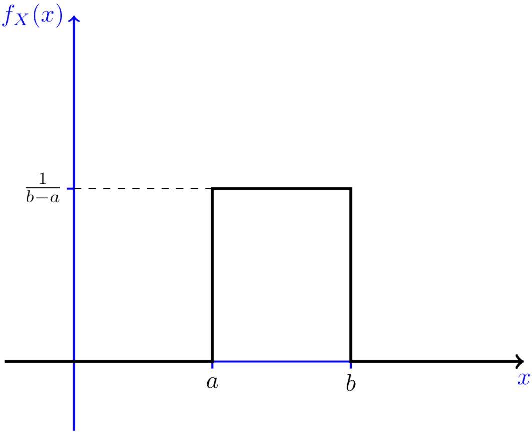

If $X \sim \mathrm{U}(a, b)$,

$$\mathbb{E}[X] = \frac{a+b}{2}, \qquad \mathrm{Var}[X] = \frac{(b-a)^2}{12}$$

Exponential Distribution

Let's look at the continuous analog of the geometric distribution.

If X shows the time until some event occurs, then X often follows an exponential distribution:

Exponential Distribution

Let's look at the continuous analog of the geometric distribution.

If X shows the time until some event occurs, then X often follows an exponential distribution:

Definition

If the PDF of a random variable X has the following form:

$$f(x) = \begin{cases} e^{-\lambda x}\lambda & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases}$$

we say that X follows an *exponential distribution* with parameter $\lambda > 0$, and write $X \sim \text{Exp}(\lambda)$.

Exponential Distribution

Let's look at the continuous analog of the geometric distribution.

If X shows the time until some event occurs, then X often follows an exponential distribution:

Definition

If the PDF of a random variable X has the following form:

$$f(x) = \begin{cases} e^{-\lambda x}\lambda & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases}$$

we say that X follows an *exponential distribution* with parameter $\lambda > 0$, and write $X \sim \text{Exp}(\lambda)$.

Theorem

If $X \sim \text{Exp}(\lambda)$, then X is a **memoryless** random variable:

$$\mathbb{P}[X > x + a \mid X > a] = \mathbb{P}[X > x]$$

Exponential Distribution

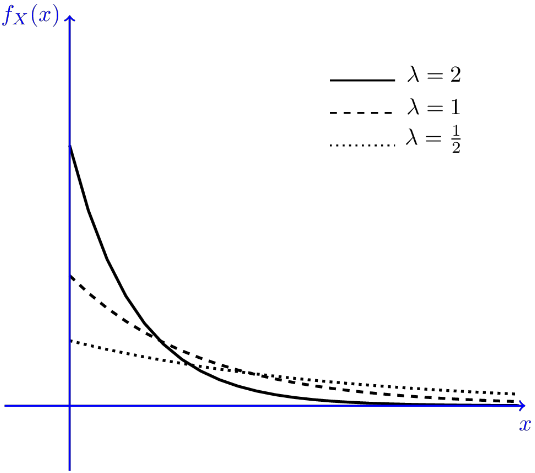

Exponential Distribution

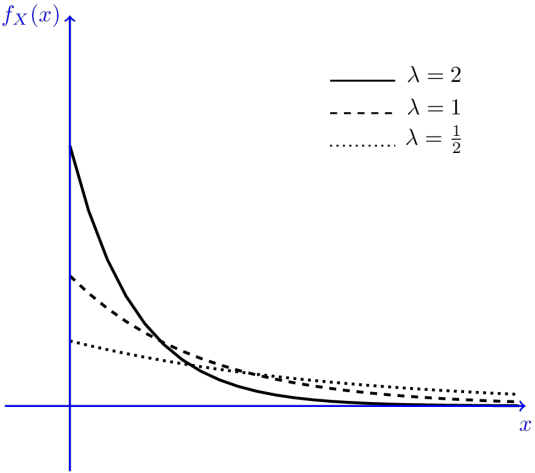

If $X \sim \text{Exp}(\lambda)$,

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

Exponential Distribution

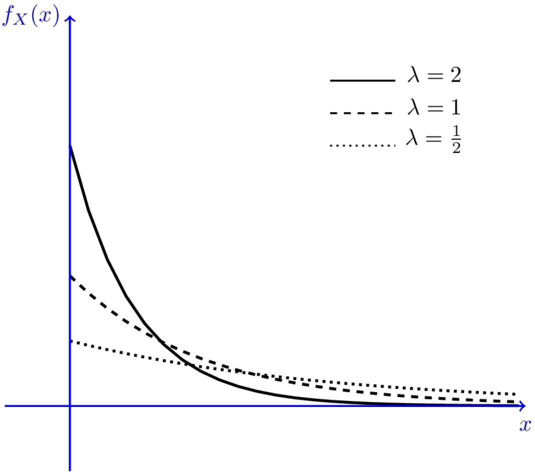

If $X \sim \text{Exp}(\lambda)$,

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

Exponential Distribution

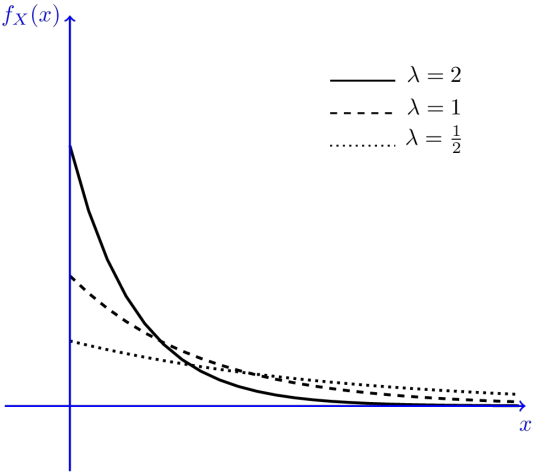

If $X \sim \text{Exp}(\lambda)$,

$$\mathbb{E}[X] = \frac{1}{\lambda}, \qquad \text{Var}[X] = \frac{1}{\lambda^2}$$

Normal Distribution

The *normal* or *Gaussian* distribution is probably the most important probability distribution – with its bell-shaped curve we have seen before.

Many natural phenomena (e.g. heights of people, measurement errors, etc.) follow normal distribution.

Normal Distribution

The *normal* or *Gaussian* distribution is probably the most important probability distribution – with its bell-shaped curve we have seen before.

Many natural phenomena (e.g. heights of people, measurement errors, etc.) follow normal distribution.

Definition

We say that a random variable X follows a *normal distribution* with mean μ and variance σ^2, if its PDF is given by:

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \qquad x \in \mathbb{R} \tag{1}$$

We write $X \sim \mathrm{N}(\mu, \sigma^2)$.

Normal Distribution

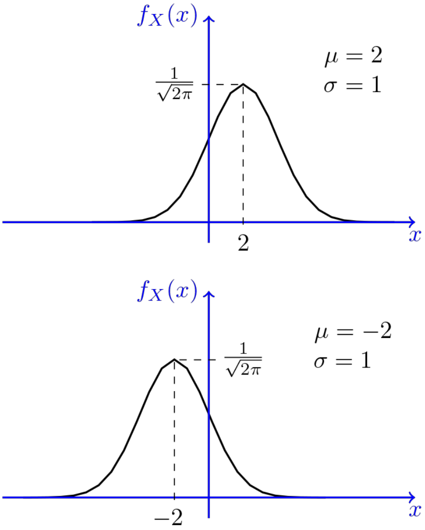

Normal Distribution

If $X \sim \mathsf{N}(\mu, \sigma^2)$,

$$\mathbb{E}[X] = \mu, \qquad \mathrm{Var}[X] = \sigma^2$$

Normal Distribution

If $X \sim \mathsf{N}(\mu, \sigma^2)$,

$$\mathbb{E}[X] = \mu, \qquad \mathsf{Var}[X] = \sigma^2$$

A very important special case is the **standard normal distribution**, where $\mu = 0$ and $\sigma^2 = 1$. We write $Z \sim \mathsf{N}(0, 1)$.

Normal Distribution

If $X \sim \mathrm{N}(\mu, \sigma^2)$,

$$\mathbb{E}[X] = \mu, \qquad \mathrm{Var}[X] = \sigma^2$$

A very important special case is the **standard normal distribution**, where $\mu = 0$ and $\sigma^2 = 1$. We write $Z \sim \mathrm{N}(0, 1)$.

Any normal random variable $X \sim \mathrm{N}(\mu, \sigma^2)$ can be standardized, i.e. converted to a standard normal random variable by doing:

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