


Some of what we discussed in lecture:

→ given some null hypothesis on how data is generated, how surprising is a particular observation of data?



For section today, a mix of things but they're all connected--

→ given data, how do we make a new hypothesis?

→ revisit Bayes Rule

→ estimation of parameters for binomial & normal


→ see how T scores arise from estimation of  $\uparrow$

→ concepts behind p-set

# Parameter Estimation / Inference


Given some data,  $X_1, \dots, X_N$ , we'd like to describe the process that produced the data.

Ex. 1: Weights of Siberian huskies



⇒ what particular  $\mu, \sigma$  describe  $X_1, \dots, X_N$ ?

Ex. 2: # of heads out of  $N$  flips of a coin?



⇒ what particular  $p$  describes # of heads / flips?

Notation:  $\theta$  is a hypothesis on how the data was generated  
↳ particular values of the parameters of the underlying probability distribution.

Ex. for dog weights:  $\theta = \{\mu = 15 \text{ kg}, \sigma = 5 \text{ kg}\}$

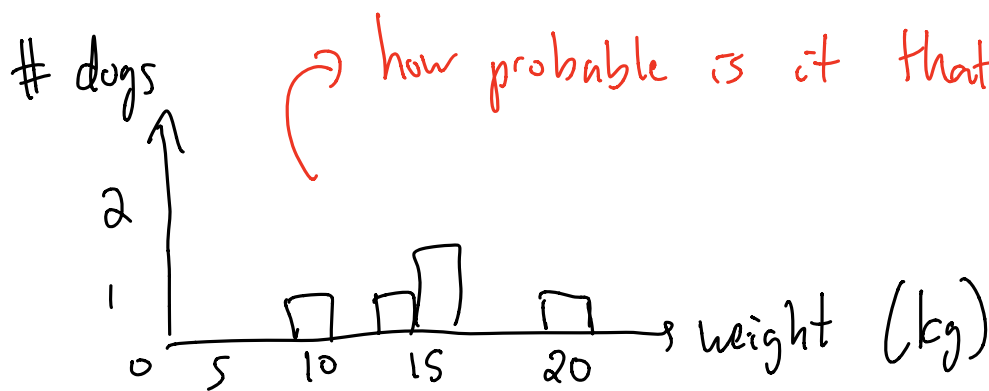
Ex. for an unbiased coin:  $\theta = \{p = 0.5\}$

The big question we'd like to answer:

given data  $D$ , what's the most probable  $\theta$ ?

( $\theta$ : a hypothesis on how data  $D$  was generated)

Ex. I weigh 5 dogs and make a histogram:



↳ how probable is it that  $\mu = 5$ ? meh.  
 $\mu = 15$ ? probable!

Ex. I get 25 heads in 30 flips

↳ "fair" hypothesis: how probable is it that  $p = \mathbb{P}(\text{heads}) = 0.5$ ?

Suppose we had some set of hypotheses,  $\theta_1, \dots, \theta_M$

Given data, how probable is a particular hypothesis  $\theta_k$ ?

likelihood:

how probable is the data  
given our hypothesis is true

prior:

how probable is our  
hypothesis prior to  
seeing data?

$$P(\theta_k | D) = \frac{P(D | \theta_k) P(\theta_k)}{P(D)}$$

posterior:

how probable is our hypothesis  $\theta_k$   
post seeing the data?

marginal:

how probable is the data  
under all possible  
hypotheses?

Marginal recap:


$$P(D) = \sum_{i=1}^M P(D \text{ AND } \theta_i \text{ true}) \leftarrow \text{add over all possible hypotheses}$$

$$= \sum_{i=1}^M P(D | \theta_i) P(\theta_i) \leftarrow P(X \text{ AND } Y) = P(X|Y)P(Y)$$

$$= \int P(D | \theta) P(\theta) d\theta$$

if an  $\infty$  number of  $\theta$ 's

# An example: My Weekly Schedule



$\Theta$ : What day it is. Monday, Thursday, who knows...

Data: how many Zoom meetings I had today.

$\Theta$  generates data, but also data informs  $\Theta$

Ex. Marginal:

$$\begin{aligned} P(3 \text{ meetings}) &= P(3 \text{ meetings AND it's Monday}) \\ &\quad + \dots + \\ &\quad P(3 \text{ meetings AND it's Sunday}) \\ &= \sum_{i=1}^7 P(3 \text{ meetings AND it's } i\text{-th day of week}) \\ &= \sum_{i=1}^7 P(3 \text{ meetings} \mid i\text{-th day}) P(i\text{-th day}) \end{aligned}$$

## Comparing two posteriors / hypotheses

$$P(\text{Monday} \mid 3 \text{ meetings}) = \frac{P(3 \text{ meetings} \mid \text{Mon}) P(\text{Mon})}{P(3 \text{ meetings})}$$

$$P(\text{Thursday} \mid 3 \text{ meetings}) = \frac{P(3 \text{ meetings} \mid \text{Thu}) P(\text{Thu})}{P(3 \text{ meetings})}$$

To compare these two, I can take a ratio.

The denominator cancels out!

Back to our q:

given data  $D$ , what's the most probable  $\theta$ ?

We can scan over a lot of  $\theta$ 's to look

for a particular  $\theta$  w/ the highest  $P(\theta \mid D)$ ,

$$P(\theta \mid D) = \frac{P(D \mid \theta) P(\theta)}{P(D)}$$

~~$P(D)$~~

(can ignore denom when comparing posteriors generally)

If we have some prior beliefs  $P(\theta)$ ,

(it just feels like a Thursday...)

The  $\theta$  that maximizes  $P(D|\theta)P(\theta)$  is the maximum a posteriori (MAP) estimate.

If we further assume uniform priors on  $\theta$ ...

$$P(\theta|D) = \frac{P(D|\theta) \cancel{P(\theta)}}{\cancel{P(D)}}$$


then... the most probable  $\theta$  given the data is the  $\theta$  that maximizes the likelihood

$$P(\theta|D) \propto P(D|\theta)$$


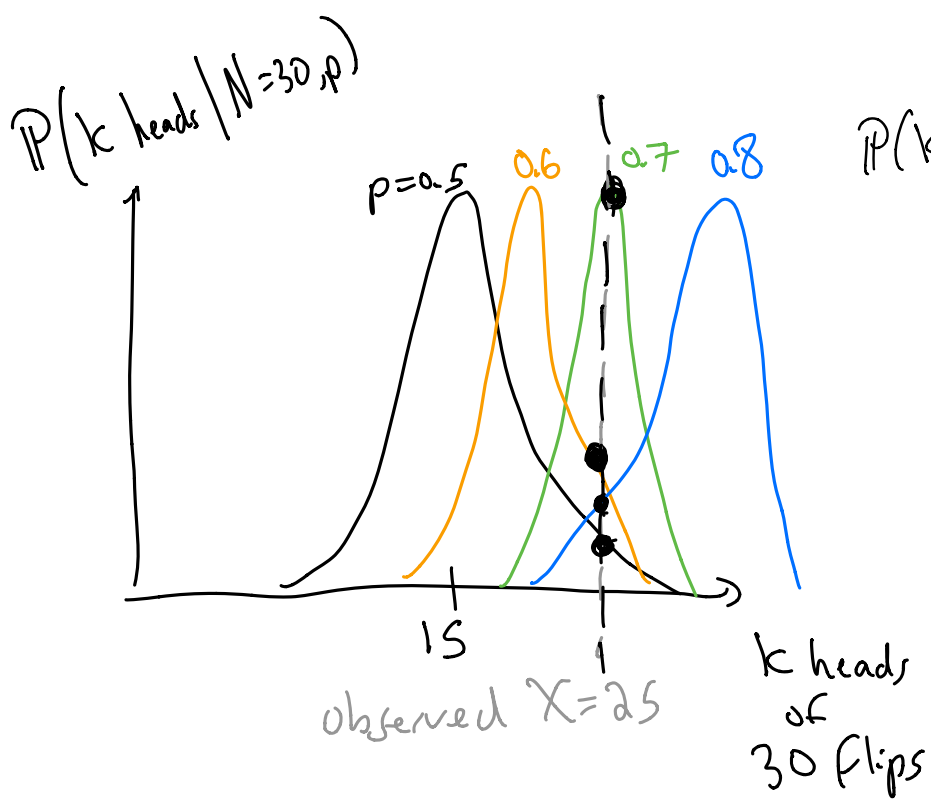
This is Maximum Likelihood Estimation!

# Concept behind likelihood $P(D|\theta)$ w/ binomial

binomial process:  $N$  trials  
 each trial has success prob.  $p$   
 how many successes out of  $N$  trials?



$$P(X = k \text{ successes} \mid N, p) = \binom{N}{k} p^k (1-p)^{N-k}$$



Key point: the likelihood measures overlap of data w/ a data generation/probability process parameterized by  $\theta$ .



## Max likelihood of Binomial

Observed data: 25 heads / 30 flips.

What's the max. likelihood estimate of  $p$ ?

The  $p$  that maximizes the likelihood.

$$P(k | N, p) = L(p) = \binom{N}{k} p^k (1-p)^{N-k}$$

This is a function of  $p$ . We can maximize it

by finding  $p$  s.t.  $\frac{\partial}{\partial p} L(p) = 0$ .

$$\frac{\partial}{\partial p} L(p) = \frac{\partial}{\partial p} \left[ \binom{N}{k} p^k (1-p)^{N-k} \right]$$

gross... chain rule --- take the log!

$$\log L(p) = \log \binom{N}{k} + k \log p + (N-k) \log(1-p)$$

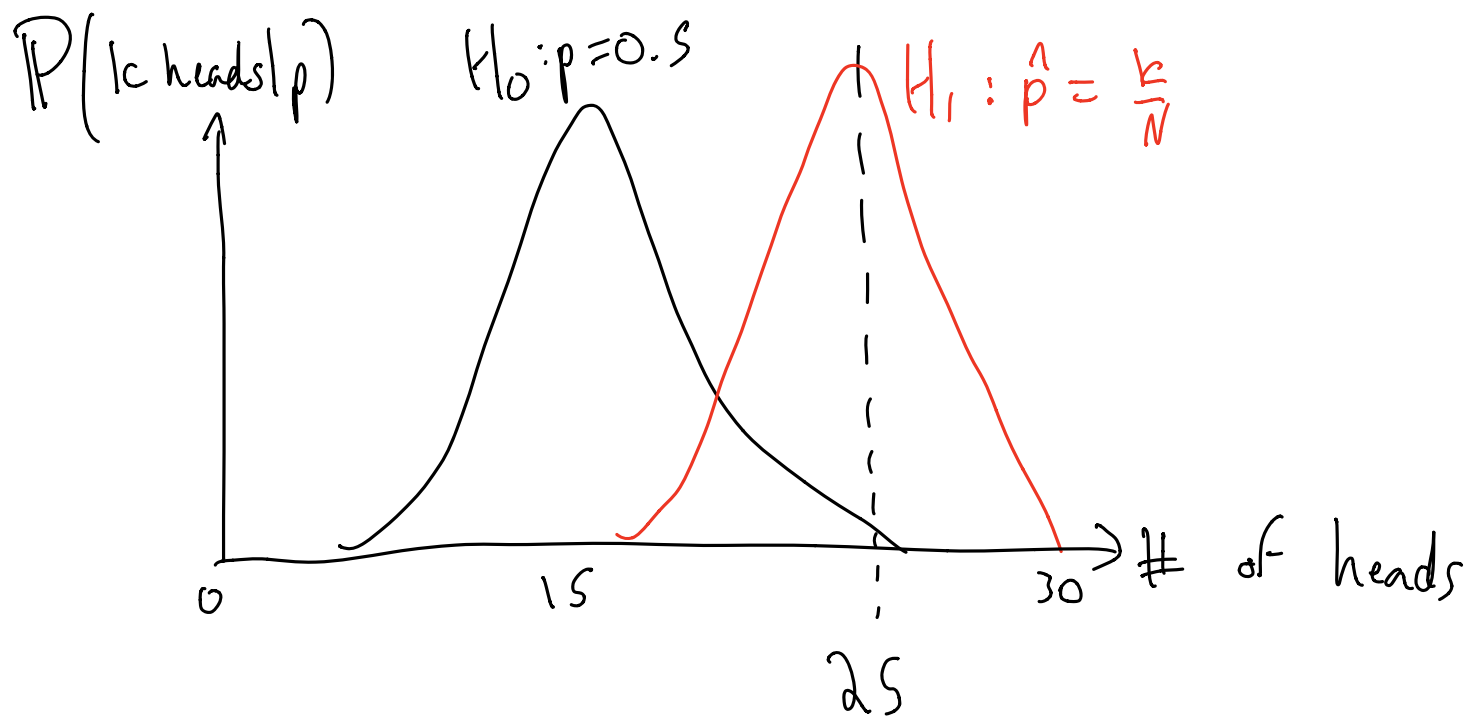
$$\frac{\partial}{\partial p} \log L(p) = \frac{k}{p} - \frac{N-k}{1-p} = 0$$

$$\Rightarrow \frac{k}{p} = \frac{N-k}{1-p}$$

$$\Rightarrow k - kp = Np - kp \Rightarrow$$

$$p = \frac{k}{N}$$

So, given we observe  $k$  heads out of  $N$  flips,  
the max. likelihood estimate of  $p$  is  $\frac{k}{N}$ .



Things to consider:

- is  $H_0$  invalidated?
- does this mean  $H_1$  has to be the correct model?
- what if we had prior beliefs on

$P(\Theta) = P(\text{prob. of heads})?$

(maybe we really trust the coin is fair)

MLE for a normal:  $X_1, \dots, X_n \sim N(\mu, \sigma^2)$

For one data point,

$$\mathbb{P}(X=x \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

For  $n$  data points, write the joint likelihood:

$$\begin{aligned} L(\mu, \sigma^2) &= \mathbb{P}(X_1=x_1, \dots, X_n=x_n \mid \mu, \sigma^2) \\ &= \mathbb{P}(X_1=x_1 \mid \mu, \sigma^2) \cdot \dots \cdot \mathbb{P}(X_n=x_n \mid \mu, \sigma^2) \\ &= \prod_{i=1}^n \mathbb{P}(X_i=x_i \mid \mu, \sigma^2) \\ &= \prod_{i=1}^n (2\pi\sigma^2)^{-\frac{1}{2}} \exp\left(-\frac{(X_i-\mu)^2}{2\sigma^2}\right) \\ &= (2\pi\sigma^2)^{-\frac{n}{2}} \prod_{i=1}^n \exp\left(-\frac{(X_i-\mu)^2}{2\sigma^2}\right) \\ &= (2\pi\sigma^2)^{-\frac{n}{2}} \exp\left(-\sum_{i=1}^n \frac{(X_i-\mu)^2}{2\sigma^2}\right) \end{aligned}$$

Take the log:

$$\log L(\mu, \sigma^2) = -\frac{n}{2} \log(2\pi\sigma^2) - \sum_{i=1}^n \frac{(X_i-\mu)^2}{2\sigma^2}$$

↳ we'll find MLE again by taking derivatives

MLE for  $\mu$ :

$$\frac{\partial}{\partial \mu} \log L(\mu, \sigma^2) = \sum_{i=1}^n \frac{2(X_i - \mu)}{2\sigma^2} = 0$$

$$\Rightarrow \hat{\mu} = \frac{\sum_{i=1}^n X_i}{n}$$

The MLE for  $\mu$  is just the sample mean.

MLE for  $\sigma^2$ :

$$\frac{\partial}{\partial \sigma^2} \log L(\mu, \sigma^2) = -\frac{n}{2\sigma^2} + \sum_{i=1}^n \frac{(X_i - \mu)^2}{2(\sigma^2)^2} = 0$$

Skipping some steps...  $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2$

But wait... if we only observe  $X_1, \dots, X_n$ , we don't know  $\mu$ !

Replacing  $\mu$  w/  $\bar{X}$ :  $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2$

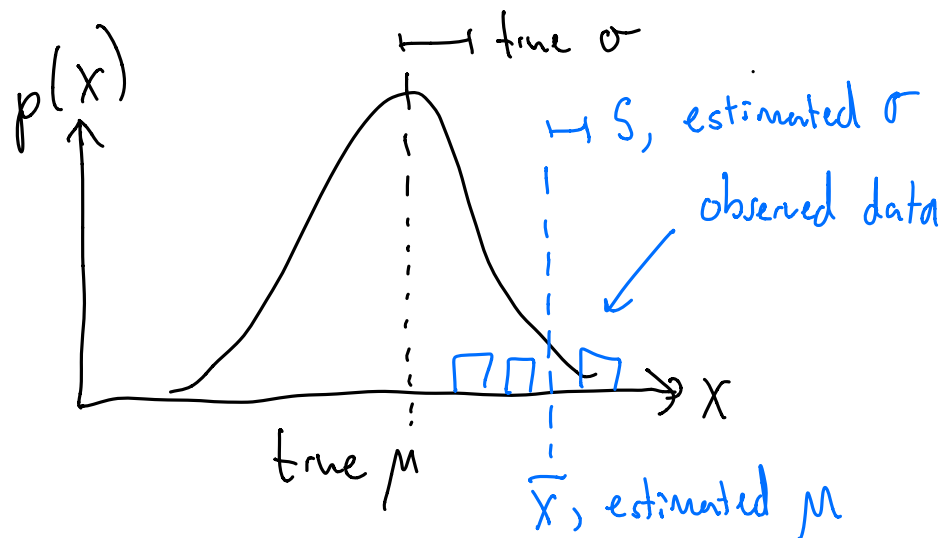
It turns out this is a biased estimator of the population  $\sigma^2$ . (see Bessel's correction)

Unbiased estimate of population variance:

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2$$

# Hypothesis testing on $\bar{X}$

$X_1, \dots, X_n \stackrel{iid}{\sim} N(\mu, \sigma)$ ,  $n$  small



Q: How likely is it that  $\mu$  is the population mean?

Compute distance b/w  $\mu$  and  $\bar{X}$ , scaled by

typical fluctuation in  $\bar{X}$ , which is  $\frac{\sigma}{\sqrt{n}}$ . ('standard error' of the mean)


$\Rightarrow$  Standard score: 
$$\frac{\bar{X} - \mu}{\sigma / \sqrt{n}}$$

If  $\sigma$  is known, this is a Z score, 
$$Z = \frac{\bar{X} - \mu}{\sigma / \sqrt{n}}$$

If  $\sigma$  is unknown, we estimate it with  $S$ ,

this is a T score: 
$$T = \frac{\bar{X} - \mu}{S / \sqrt{n}}$$
, w/  $n-1$  degrees of freedom

Distribution of  
T, Z scores  
(when  $\mu$  really is  
population mean)



⇒ T distribution has fatter tails.

It allows for  $\bar{X}$  to be "far away" from  $\mu$   
because we had to estimate  $\sigma^2$  from the data.

⇒ We could have estimated too small a  $\sigma^2$ , which  
means the T score would be larger than it  
should be → hence more probability of big T scores


⇒ At large  $n$ , we estimate  $\sigma^2$  well,  
T dist converges to Z dist.

So the T dist can arise from estimation of  
an uncertain  $\sigma^2$  given data.

On the p-set we will marginalize over many  
potential  $\sigma^2$ 's.


## Concept for P-Set

There's a true  $N(\mu, \sigma^2)$ ,  
but we only see  
a few data points



⇒ Goal: bet on candidate  $(\mu, \sigma^2)$   
given only a few observed datapoints

How? We'll assume a grid of possible  $(\mu, \sigma)$ :



one of these  $(\mu, \sigma)$  pairs  
generated the data  
we observe,  $X_1, \dots, X_n$

for any candidate  $(\mu, \sigma^2)$   
we can compute  
how probable they are  
given our data...

with a posterior,  $\mathcal{P}(\mu, \sigma^2 | X_1, \dots, X_n)$

$$\begin{aligned}
 & \mathbb{P}(\mu, \sigma^2 \mid X_1, \dots, X_n) \quad \textcircled{1} \\
 &= \frac{\mathbb{P}(X_1, \dots, X_n \mid \mu, \sigma^2) \mathbb{P}(\mu, \sigma^2)}{\mathbb{P}(X_1, \dots, X_n)} \quad \textcircled{2} \\
 & \quad \textcircled{3}
 \end{aligned}$$

① likelihood: (independent observations)

$$\begin{aligned}
 \mathbb{P}(X_1, \dots, X_n \mid \mu, \sigma^2) &= \mathbb{P}(X_1 \mid \mu, \sigma^2) \times \dots \times \mathbb{P}(X_n \mid \mu, \sigma^2) \\
 &= \prod_{i=1}^n \mathbb{P}(X_i \mid \mu, \sigma^2) \\
 &= \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(X_i - \mu)^2}{2\sigma^2}}
 \end{aligned}$$

② prior:  $\mathbb{P}(\mu, \sigma^2) = \mathbb{P}(\mu) \mathbb{P}(\sigma^2)$

③ marginal:  $\mathbb{P}(X_1, \dots, X_n) = \sum_{\mu} \sum_{\sigma} \mathbb{P}(X_1, \dots, X_n \text{ AND } \mu, \sigma)$

$$= \sum_{\mu} \sum_{\sigma} \mathbb{P}(X_1, \dots, X_n \mid \mu, \sigma) \mathbb{P}(\mu, \sigma)$$



## Extra: distribution of a sample mean.

The expectation, or average, of  $\bar{X}$ :

$$\begin{aligned} E(\bar{X}) &= E\left[\frac{1}{N} \sum_{i=1}^N X_i\right] = \frac{1}{N} \sum_{i=1}^N E[X_i] && \text{(mean of sum)} \\ &= \frac{1}{N} (N \mu) && \text{(sum of means)} \\ &= \mu. \end{aligned}$$

So the average  $\bar{X}$  is indeed  $\mu$ , the population average.

Now for the spread in  $\bar{X}$ ? Its variance:

$$\text{Var}(\bar{X}) = \text{Var}\left(\frac{1}{N} \sum_{i=1}^N X_i\right)$$

$$= \frac{1}{N^2} \text{Var}\left(\sum_{i=1}^N X_i\right)$$

$$= \frac{1}{N^2} \sum_{i=1}^N \text{Var}(X_i)$$

$$= \frac{1}{N^2} \sum_{i=1}^N \sigma^2$$

$$= \frac{1}{N^2} N \sigma^2$$

$$= \frac{\sigma^2}{N}$$

$$\rightarrow \text{SD}(\bar{X}) = \sqrt{\frac{\sigma^2}{N}} = \frac{\sigma}{\sqrt{N}}.$$

the more  $N$  we observe, the closer  $\bar{X}$  is to  $\mu$

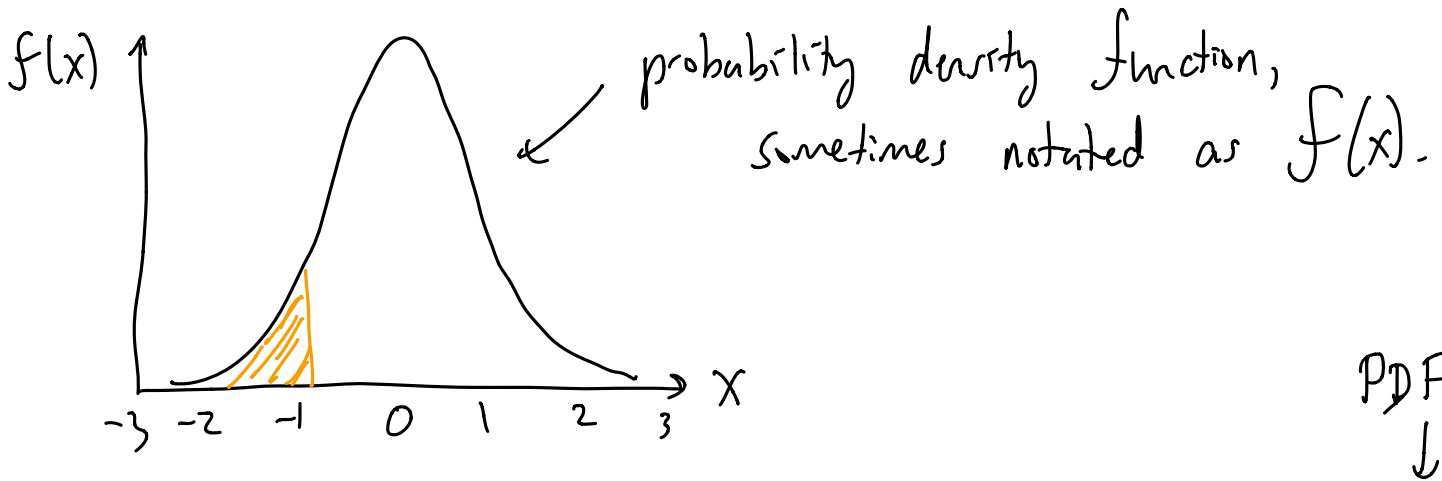
$$\text{Var}(aX) = a^2 \text{Var}(X)$$

independent variables

Each  $X_i \sim N(\mu, \sigma^2)$

# Extra: CDF example

Say we have a normal:




$\mathbb{P}(X \leq -1)$  is the integral from  $-\infty$  to  $-1$  of  $f(x)$

A CDF is an integral from  $-\infty$  to (some place) of  $f(x)$

$$\text{CDF}(x) = \int_{-\infty}^x f(x') dx', \text{ often notated as } F(x).$$

$$\text{So } \mathbb{P}(X \leq -1) = \text{CDF}(-1)$$

The further right we integrate to, we can only add to the integral, so  $\text{CDF}(x)$  cannot decrease w/  $x$ .



Q: draw the CDF of a uniform distribution

(Hint: integrate up to  $x$ )