Probability Distributions

Loukia Meligkotsidou
Associate Professor of Statistics
Department of Mathematics
National and Kapodistrian University of Athens

Types of Data

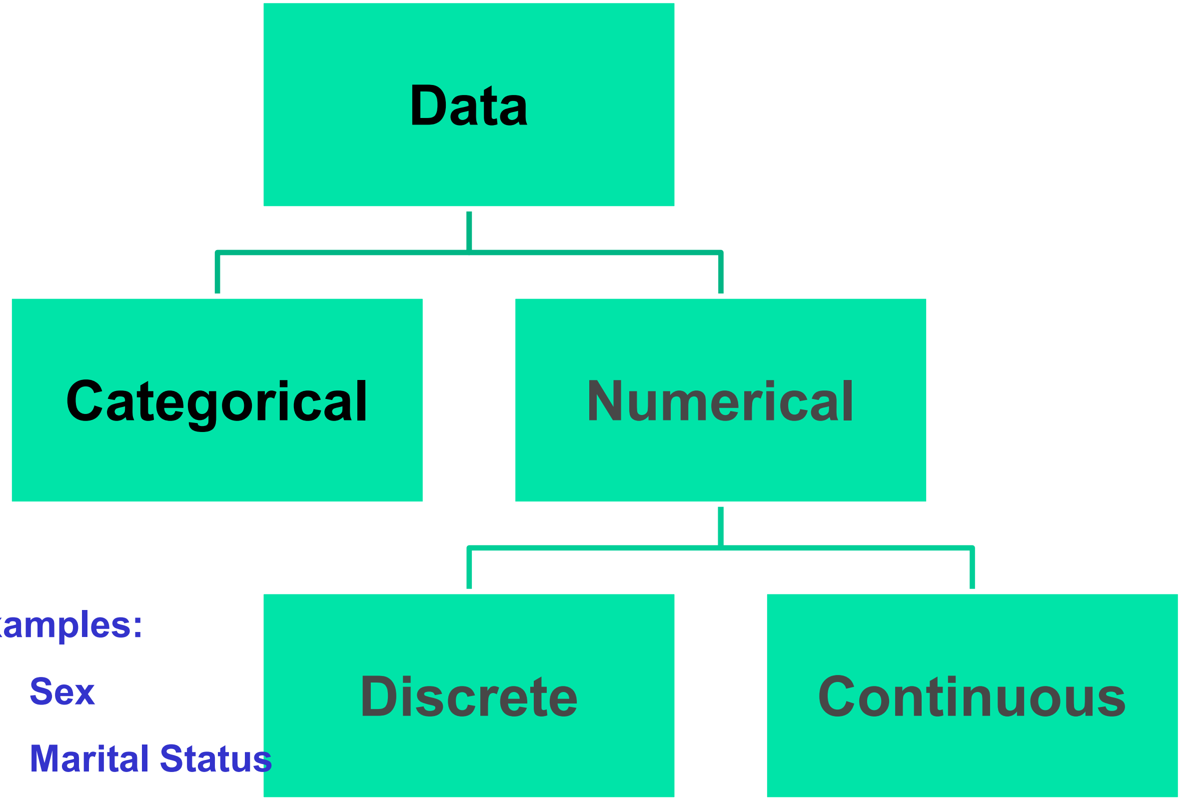

Examples:

- Sex
- Marital Status
- Eye Color

Examples:

- Number of Children
- Number of telephone calls

Examples:

- Weight
- Return of an asset

Random variables and Distributions

- **Random Variable**
  - Represents a possible numerical value from a random experiment

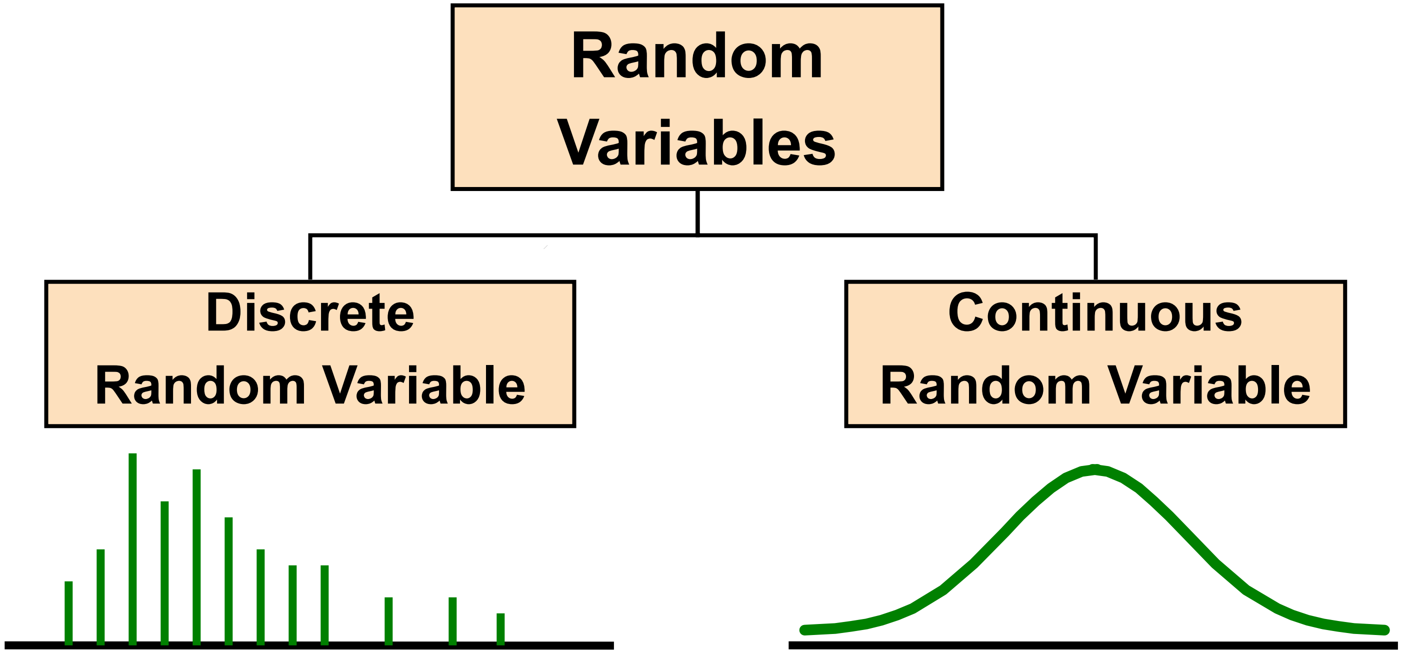

Discrete Random Variables

- Can only take on a countable number of values

Examples:

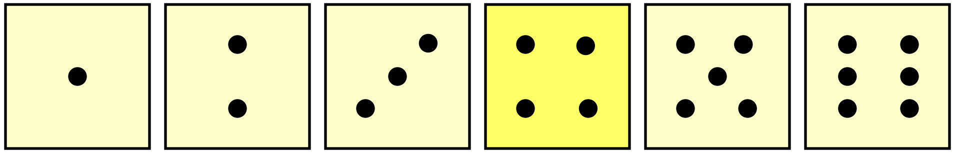

- **Roll a die twice**
  Let X be the number of times 4 comes up
  (then X could be 0, 1, or 2 times)

- **Toss a coin 5 times.**
  Let X be the number of heads
  (then X = 0, 1, 2, 3, 4, or 5)

Discrete Probability Distribution

Experiment: Toss 2 Coins. Let X = # heads.

Show P(x), i.e., P(X = x), for all values of x:

4 possible outcomes

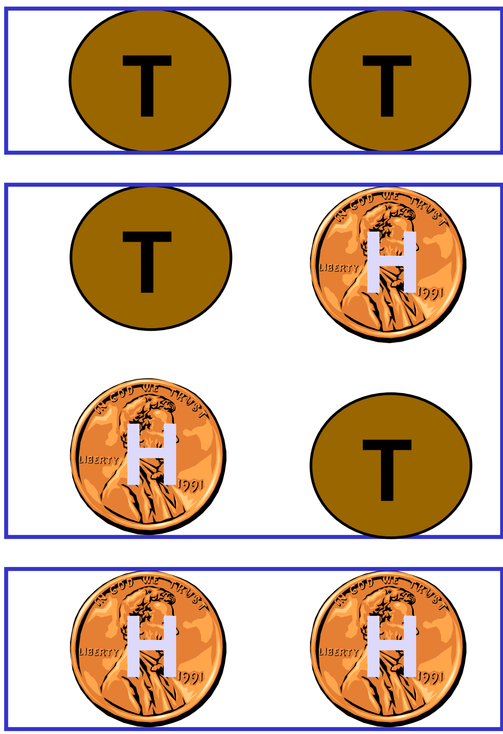

Probability Distribution

x Value	Probability
0	1/4 = .25
1	2/4 = .50
2	1/4 = .25

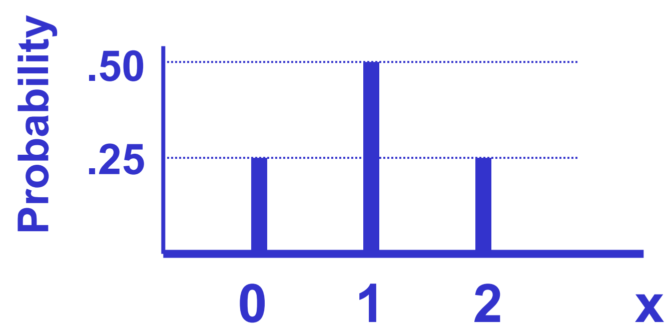

Probability Distribution Required Properties

- $P(x) \geq 0$ for any value of x

- The individual probabilities sum to 1;

$$\sum_{x} P(x) = 1$$

(The notation indicates summation over all possible x values)

Cumulative Probability Function

- The cumulative probability function, denoted $F(x_0)$, shows the probability that X is less than or equal to x_0

$$F(x_0) = P(X \leq x_0)$$

- In other words,

$$F(x_0) = \sum_{x \leq x_0} P(x)$$

- Compute P(x) from F(x), compute F(x) from P(x)

Expected Value

- Expected Value (or mean) of a discrete distribution (Weighted Average)

$$\mu = E(x) = \sum_{x} xP(x)$$

- Example: Toss 2 coins,
  x = # of heads,
  compute expected value of x:

E(x) = (0 x .25) + (1 x .50) + (2 x .25)
= 1.0

x	P(x)
0	.25
1	.50
2	.25

Variance and Standard Deviation

- Variance of a discrete random variable X

$$\sigma^2 = E(X-\mu)^2 = \sum_x (x-\mu)^2 P(x)$$

- Standard Deviation of a discrete random variable X

$$\sigma = \sqrt{\sigma^2} = \sqrt{\sum_x (x-\mu)^2 P(x)}$$

Standard Deviation Example

- Example: Toss 2 coins, X = # heads, compute standard deviation (recall E(x) = 1)

$$\sigma = \sqrt{\sum_x (x-\mu)^2 P(x)}$$

$$\sigma = \sqrt{(0-1)^2(.25)+(1-1)^2(.50)+(2-1)^2(.25)} = \sqrt{.50} = .707$$

Possible number of heads = 0, 1, or 2

Linear Functions of Random Variables

- Let a and b be any constants.

- a) $E(a) = a \qquad \text{and} \qquad Var(a) = 0$

  i.e., if a random variable always takes the value a, it will have mean a and variance 0

- b) $E(bX) = b\mu_X \qquad \text{and} \qquad Var(bX) = b^2\sigma_X^2$

  i.e., the expected value of $b \cdot X$ is $b \cdot E(x)$

Linear Functions of Random Variables

- Let random variable X have mean μ_x and variance $\sigma^2{}_x$
- Let a and b be any constants.
- Let Y = a + bX
- Then the mean and variance of Y are

$$\mu_Y = E(a + bX) = a + b\mu_X$$

$$\sigma^2{}_Y = Var(a + bX) = b^2\sigma^2{}_X$$

- so that the standard deviation of Y is

$$\sigma_Y = |b|\sigma_X$$

Probability Distributions

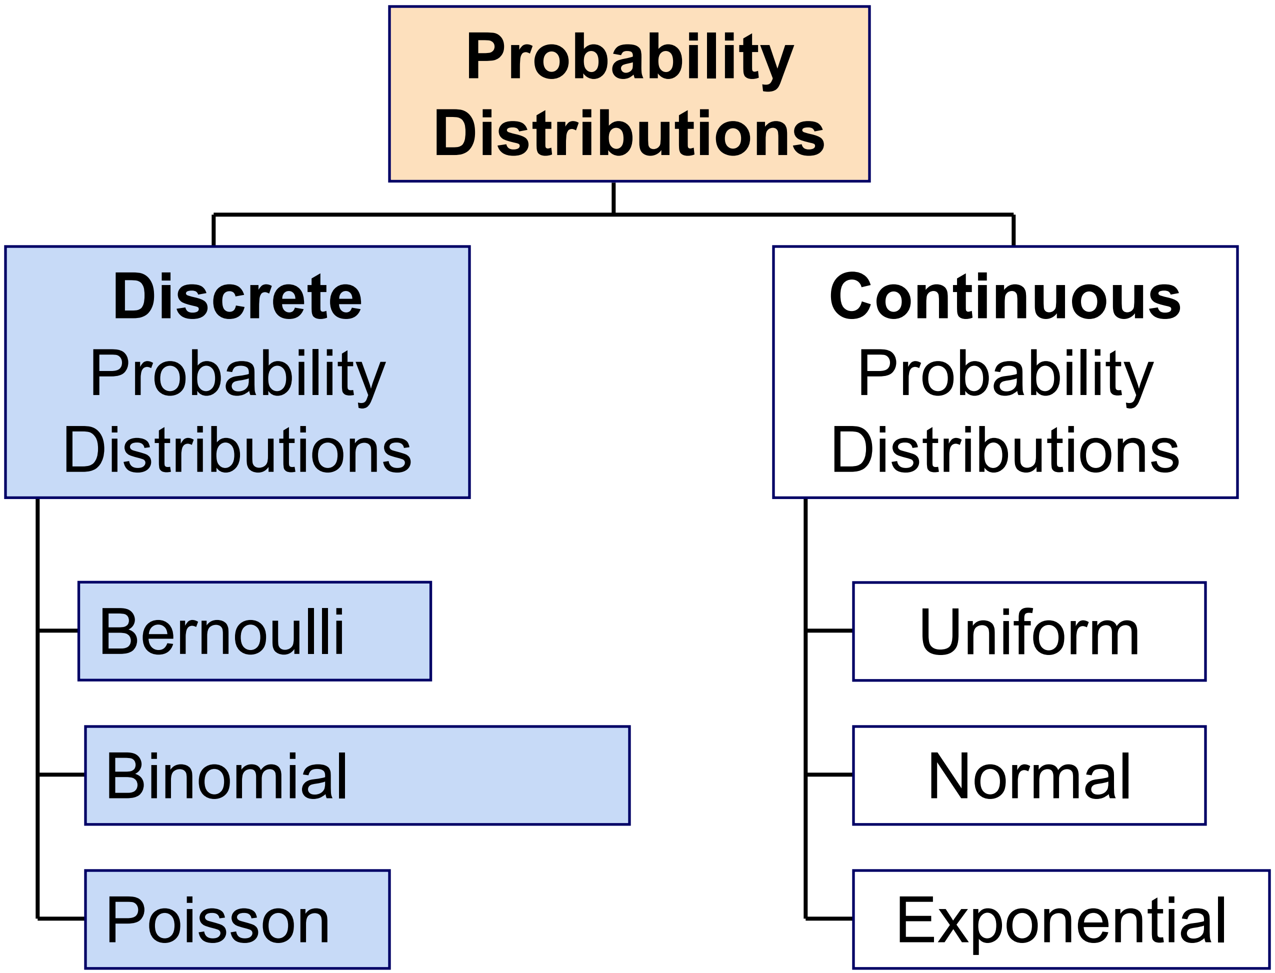

Bernoulli Distribution

- Consider only two outcomes: "success" or "failure"
- Let P denote the probability of success
- Let $1 - P$ be the probability of failure
- Define random variable X:

  x = 1 if success, x = 0 if failure

- Then the Bernoulli probability function is

$$P(0) = (1 - P) \quad \text{and} \quad P(1) = P$$

Bernoulli Distribution Mean and Variance

- The mean is $\mu = P$

$$\mu = E(X) = \sum_{X} xP(x) = (0)(1-P) + (1)P = P$$

- The variance is $\sigma^2 = P(1 - P)$

$$\sigma^2 = E[(X-\mu)^2] = \sum_{X} (x-\mu)^2 P(x)$$

$$= (0-P)^2(1-P) + (1-P)^2 P = P(1-P)$$

Sequences of x Successes in n Trials

- The number of sequences with x successes in n independent trials is:

$$C_x^n = \frac{n!}{x!(n-x)!}$$

Where $n! = n\cdot(n-1)\cdot(n-2)\cdot\ ...\ \cdot 1$ and $0! = 1$

- These sequences are mutually exclusive, since no two can occur at the same time

Binomial Probability Distribution

- A fixed number of observations, n
  - e.g., 15 tosses of a coin
- Two mutually exclusive and collectively exhaustive categories
  - e.g., head or tail in each toss of a coin
  - Generally called "success" and "failure"
  - Probability of success is P , probability of failure is 1 – P
- Constant probability for each observation
  - e.g., Probability of getting a tail is the same each time we toss the coin
- Observations are independent
  - The outcome of one observation does not affect the outcome of the other

Binomial Distribution Formula

$$P(x) = \frac{n!}{x!(n-x)!} P^x (1-P)^{n-x}$$

$P(x)$ = probability of **x** successes in **n** trials,
with probability of success **P** on each trial

x = number of 'successes' in sample, ($x = 0, 1, 2, ..., n$)

n = sample size (number of trials or observations)

P = probability of "success"

Example: Flip a coin four times, let x = # heads:

$n = 4$

$P = 0.5$

$1 - P = (1 - 0.5) = 0.5$

$x = 0, 1, 2, 3, 4$

Example: Calculating a Binomial Probability

What is the probability of one success in five observations if the probability of success is 0.1?

x = 1, n = 5, and P = 0.1

$$P(x=1) = \frac{n!}{x!(n-x)!}P^X(1-P)^{n-X}$$

$$= \frac{5!}{1!(5-1)!}(0.1)^1(1-0.1)^{5-1}$$

$$= (5)(0.1)(0.9)^4$$

$$= .32805$$

Binomial Distribution

- The shape of the binomial distribution depends on the values of P and n

- Here, n = 5 and P = 0.1

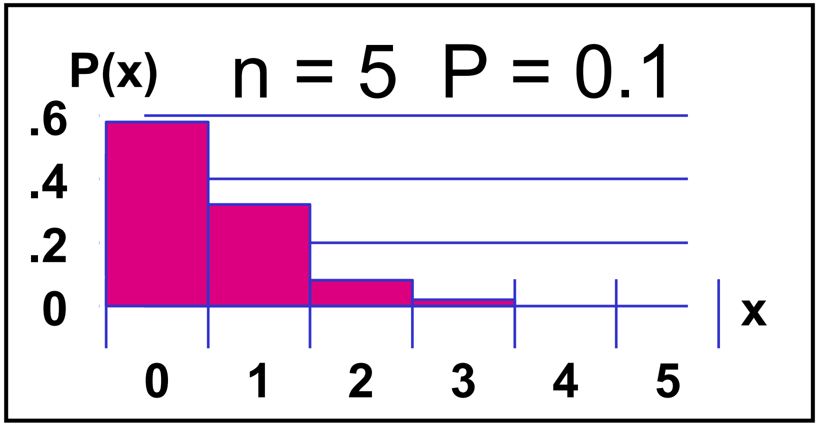

- Here, n = 5 and P = 0.5

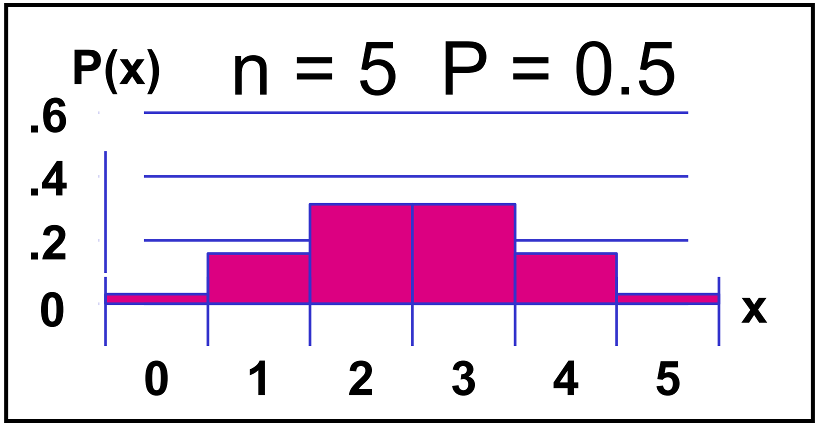

Binomial Distribution Mean and Variance

- Mean

$$\mu = E(x) = nP$$

- Variance and Standard Deviation

$$\sigma^2 = nP(1-P)$$

$$\sigma = \sqrt{nP(1-P)}$$

Where n = sample size

P = probability of success

(1 – P) = probability of failure

The Poisson Distribution

- Apply the Poisson Distribution when:
  - You wish to count the number of times an event occurs in a given continuous interval (of time or space)
  - The probability that an event occurs in one subinterval is very small and is the same for all subintervals
  - The number of events that occur in one subinterval is independent of the number of events that occur in the other subintervals
  - The average number of events per unit is λ

Poisson Distribution Formula

$$P(x) = \frac{e^{-\lambda}\lambda^{x}}{x!}$$

where:

x = number of occurrences per unit

λ = expected number of occurrences per unit

e = base of the natural logarithm system (2.71828...)

Poisson Distribution Characteristics

- Mean

$$\mu = E(x) = \lambda$$

- Variance and Standard Deviation

$$\sigma^2 = E[(X-\mu)^2] = \lambda$$

$$\sigma = \sqrt{\lambda}$$

where λ = expected number of successes per unit

Graph of Poisson Probabilities

Graphically:

λ = .50

X	**λ = 0.50**
0	0.6065
1	0.3033
2	0.0758
3	0.0126
4	0.0016
5	0.0002
6	0.0000
7	0.0000

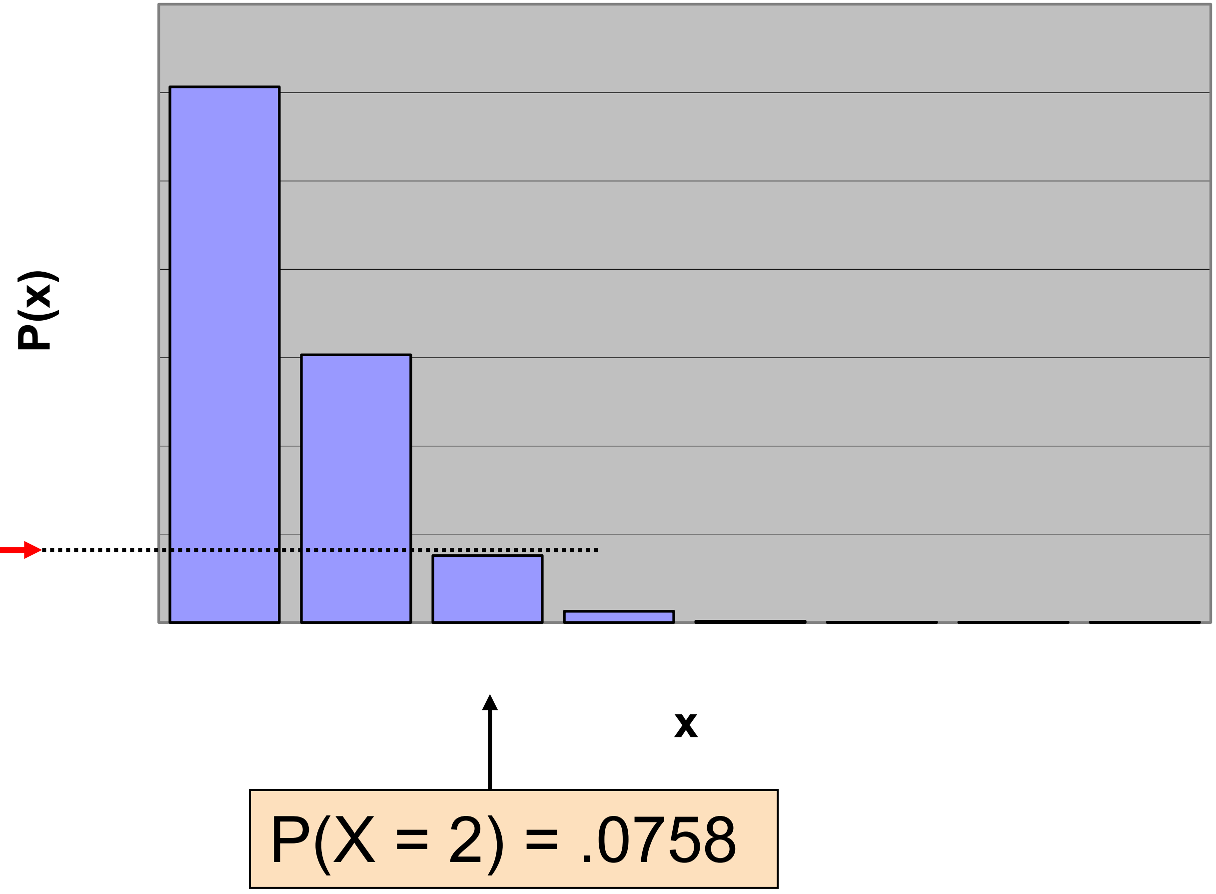

Poisson Distribution Shape

- The shape of the Poisson Distribution depends on the parameter λ :

λ = 0.50

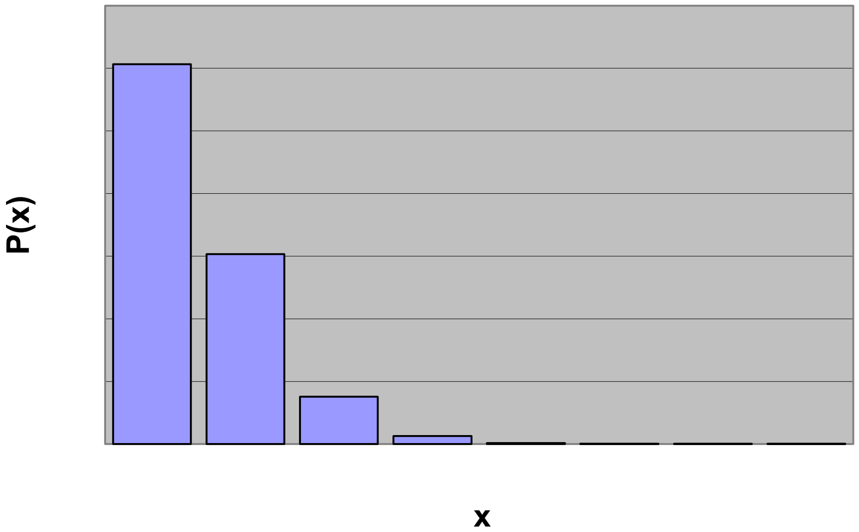

λ = 3.00

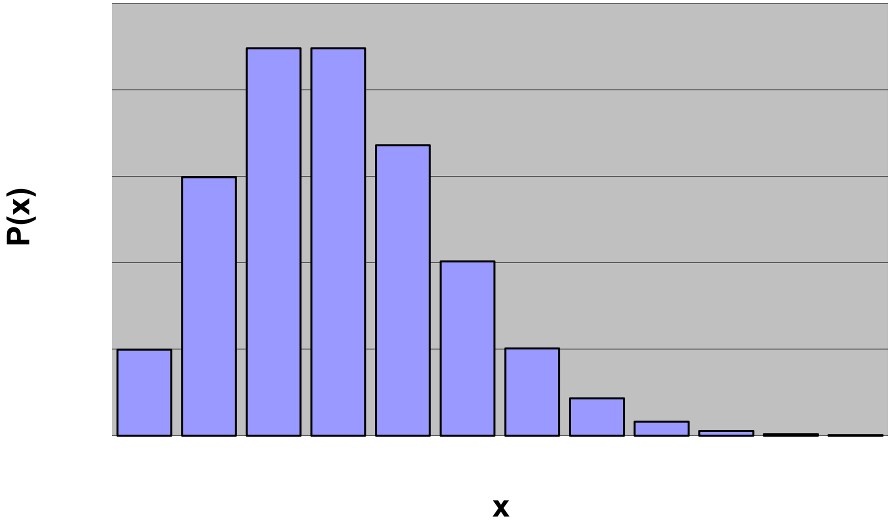

Continuous Probability Distributions

- A continuous random variable is a variable that can assume any value in an interval
  - time required to complete a task
  - height, in inches
  - Return of an asset
- These can potentially take on any value, depending only on the ability to measure accurately.

Cumulative Distribution Function

- The cumulative distribution function, $F(x)$, for a continuous random variable X expresses the probability that X does not exceed the value of x

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

- Let a and b be two possible values of X, with $a < b$. The probability that X lies between a and b is

$$P(a < X < b) = F(b) - F(a)$$

Probability Density Function

The probability density function, f(x), of random variable X has the following properties:

1. f(x) > 0 for all values of x
2. The area under the probability density function f(x) over all values of the random variable X is equal to 1.0, i.e. $\int f(x)dx=1$
3. The probability that X lies between two values is the area under the density function graph between the two values

Probability as an Area

Shaded area under the curve is the probability that X is between a and b

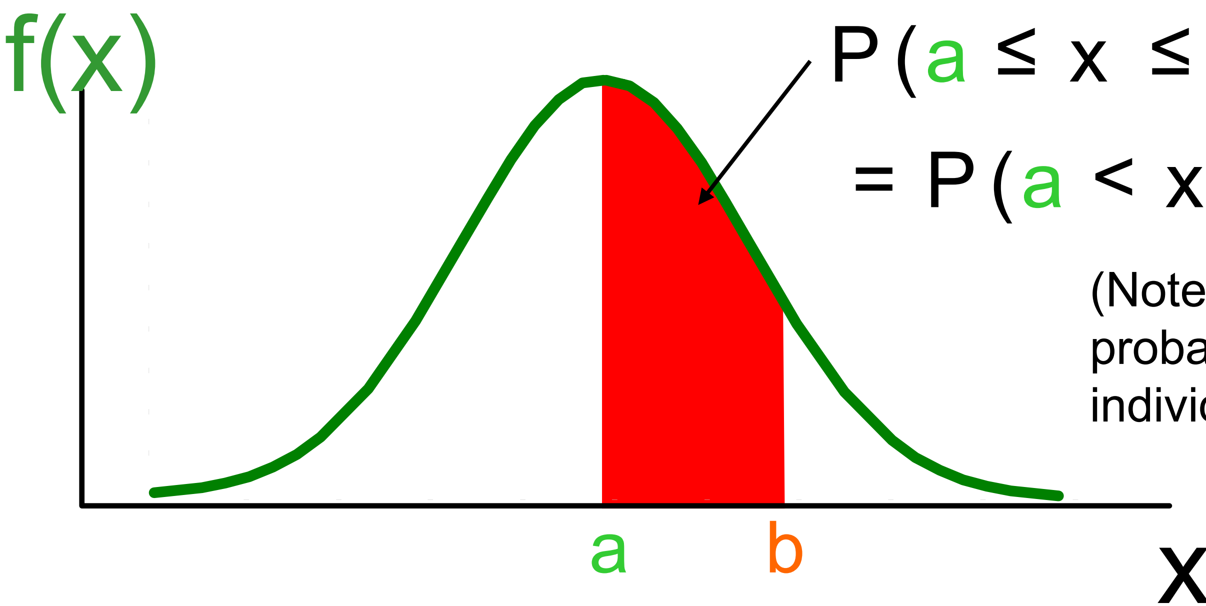

$$P(a \le x \le b) = P(a < x < b)$$

(Note that the probability of any individual value is zero)

Expectations for Continuous Random Variables

- The mean of X, denoted μ_X, is defined as the expected value of X

$$\mu_X = E(X) = \int_x x \cdot f(x)dx$$

- The variance of X, denoted σ_X^2, is defined as the expectation of the squared deviation, $(X - \mu_X)^2$, of a random variable from its mean

$$\sigma_X^2 = E[(X - \mu_X)^2]$$

Linear Functions of Variables

- Let $W = a + bX$, where X has mean μ_X and variance σ_X^2, and a and b are constants
- Then the mean of W is

$$\mu_W = E(a + bX) = a + b\mu_X$$

- the variance is

$$\sigma_W^2 = Var(a + bX) = b^2\sigma_X^2$$

- the standard deviation of W is

$$\sigma_W = |b|\sigma_X$$

Linear Functions of Variables

- An important special case of the previous results is the standardized random variable

$$Z = \frac{X - \mu_X}{\sigma_X}$$

which has a mean 0 and variance 1

The Normal Distribution

- **Bell Shaped**
- **Symmetrical**
- **Mean, Median and Mode are Equal**

Location is determined by the mean, μ

Spread is determined by the standard deviation, σ

The random variable has an infinite theoretical range:
$+\infty$ to $-\infty$

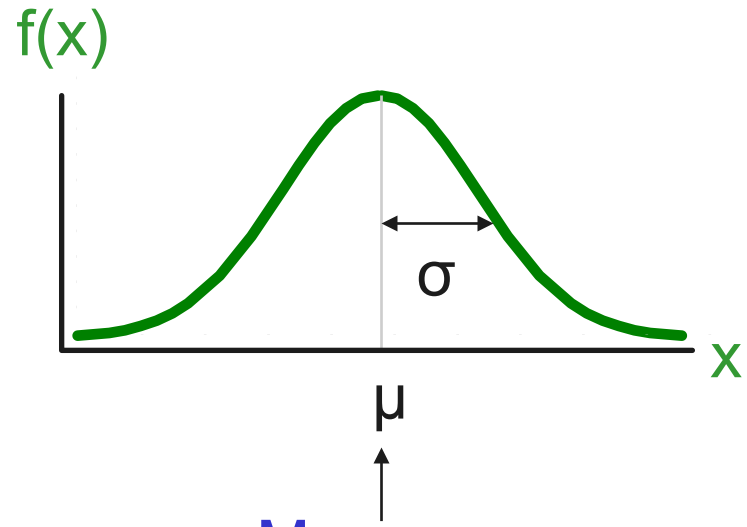

Mean = Median = Mode

The Normal Distribution

- The normal distribution closely approximates the probability distributions of a wide range of random variables
- Distributions of sample means approach a normal distribution given a “large” sample size
- Computations of probabilities are direct and elegant
- The normal probability distribution has led to good business decisions for a number of applications

Many Normal Distributions

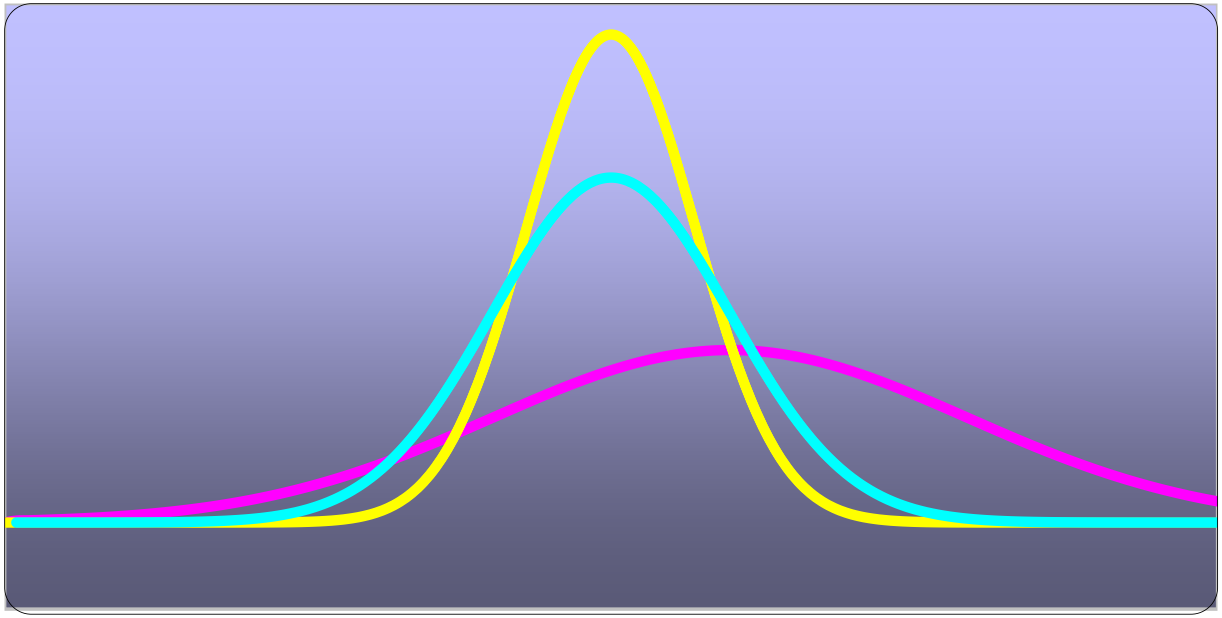

By varying the parameters μ and σ, we obtain different normal distributions

The Normal Distribution Shape

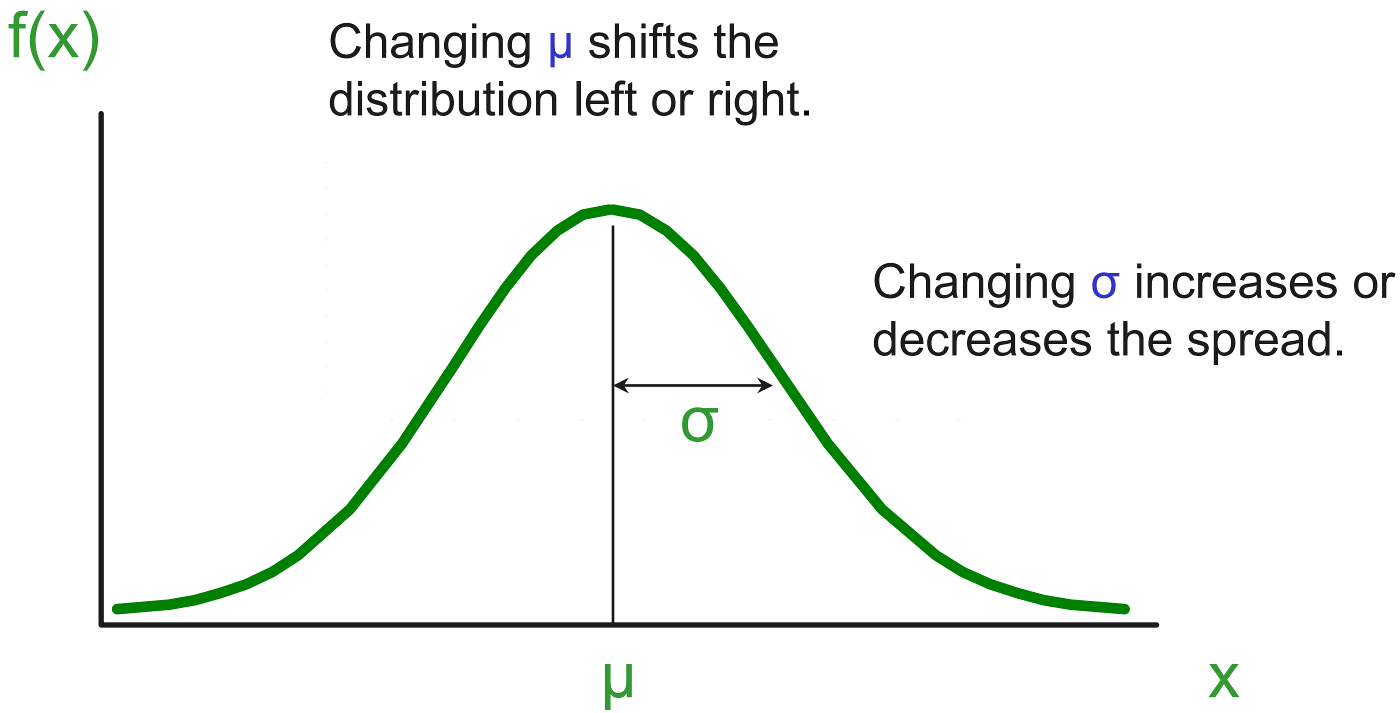

Given the mean μ and variance σ we define the normal distribution using the notation $X \sim N(\mu, \sigma^2)$

The Normal Probability Density Function

- The formula for the normal probability density function is

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma^2}$$

Where e = the mathematical constant approximated by 2.71828
π = the mathematical constant approximated by 3.14159
μ = the population mean
σ = the population standard deviation
x = any value of the continuous variable, $-\infty < x < \infty$

Cumulative Normal Distribution

- For a normal random variable X with mean μ and variance σ^2 , i.e., $X \sim N(\mu, \sigma^2)$, the cumulative distribution function is

$$F(x_0) = P(X \leq x_0)$$

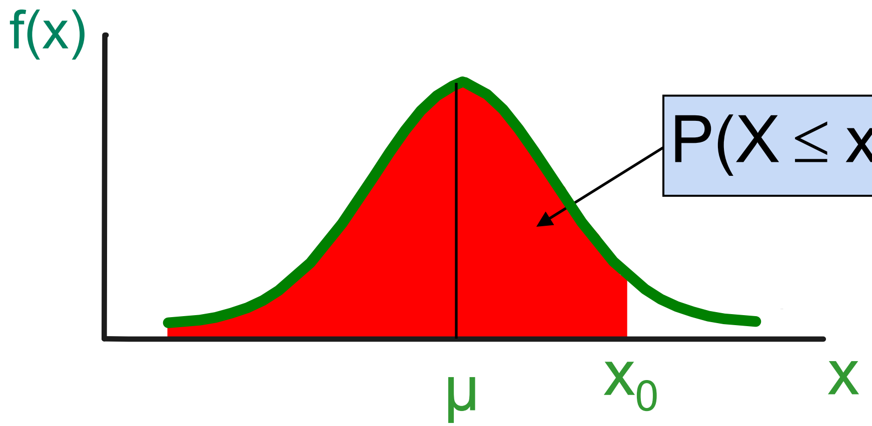

Finding Normal Probabilities

The probability for a range of values is measured by the area under the curve

$$P(a < X < b) = F(b) - F(a)$$

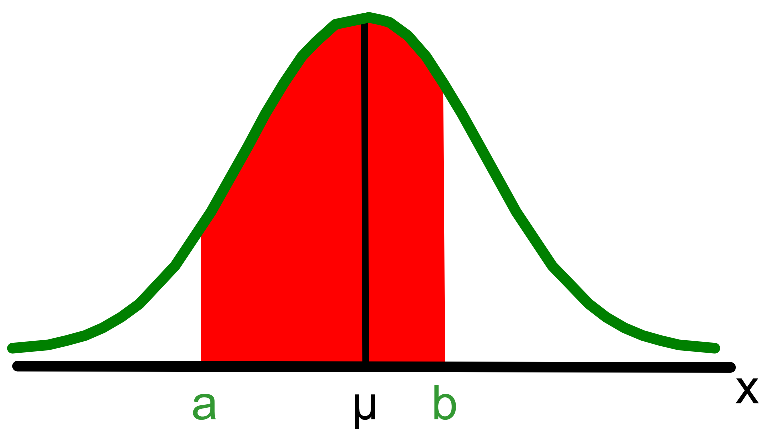

Finding Normal Probabilities

$$F(b) = P(X < b)$$

$$F(a) = P(X < a)$$

$$P(a < X < b) = F(b) - F(a)$$

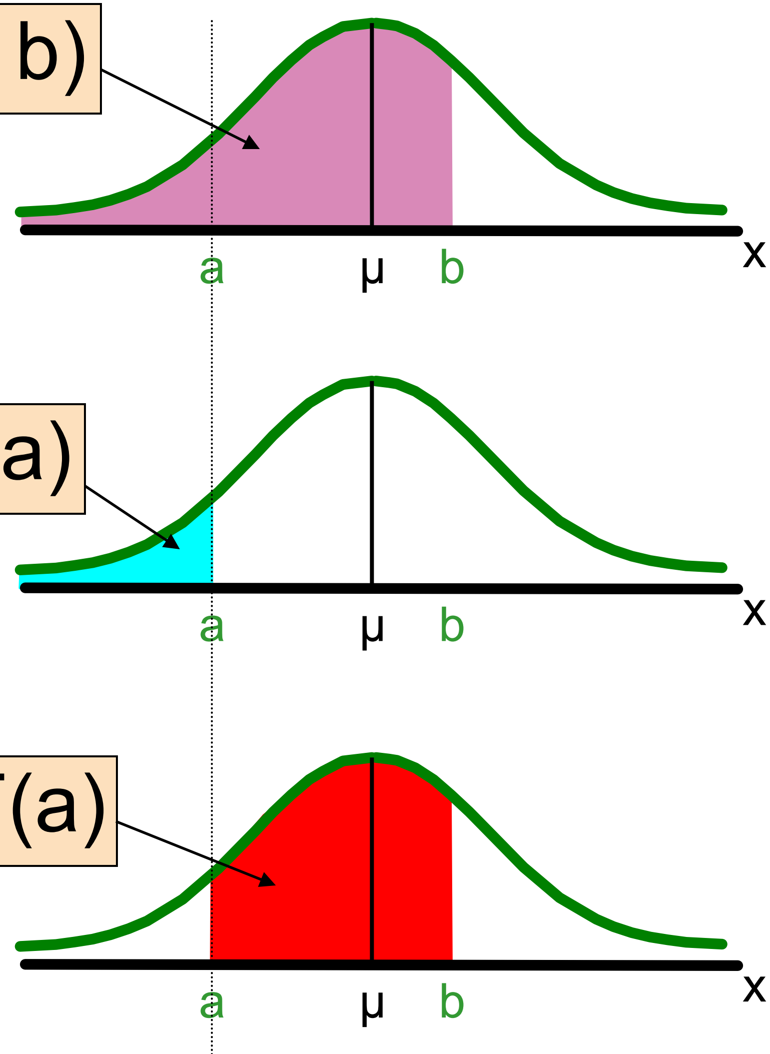

The Standardized Normal

- Any normal distribution (with any mean and variance combination) can be transformed into the standardized normal distribution (Z), with mean 0 and variance 1

$$Z \sim N(0,1)$$

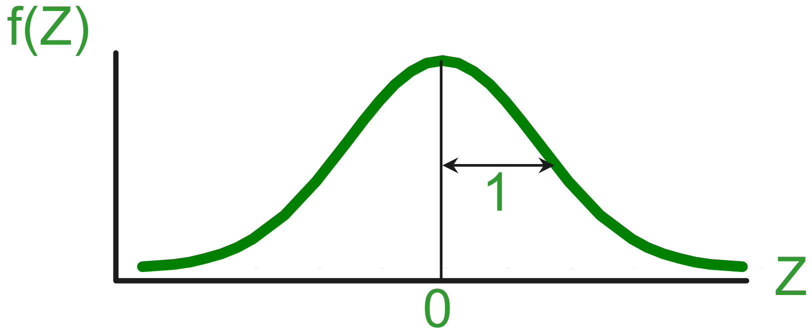

- Need to transform X units into Z units by subtracting the mean of X and dividing by its standard deviation

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

The Standardized Normal Table

- The Standardized Normal table in textbooks shows values of the cumulative normal distribution function

- For a given Z-value a, the table shows F(a) (the area under the curve from negative infinity to a)

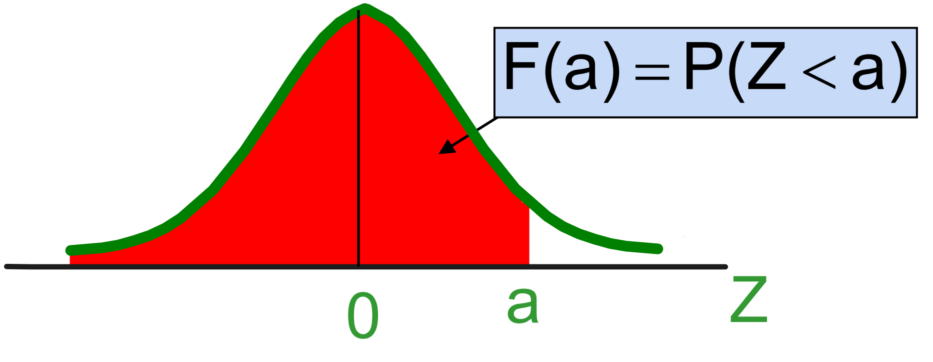

The Standardized Normal Table

- For negative Z-values, use the fact that the distribution is symmetric to find the needed probability:

Example:

P(Z < -2.00) = 1 – 0.9772

= 0.0228