Chapter 8

Continuous Probability Distributions

Probability Density Functions…

Unlike a discrete random variable which we studied in Chapter 7, a ***continuous random variable*** is one that can assume an **uncountable** number of values.

→ We cannot list the possible values because there is an infinite number of them.

→ Because there is an infinite number of values, the probability of each individual value is virtually 0.

Point Probabilities are Zero

➔Because there is an infinite number of values, the probability of each individual value is virtually 0.

Thus, we can determine the probability of a ***range of values*** only.

E.g. with a **discrete** random variable like tossing a die, it is meaningful to talk about P(X=5), say.

In a **continuous** setting (e.g. with time as a random variable), the probability the random variable of interest, say task length, takes **exactly** 5 minutes is infinitesimally small, hence $P(X=5) = 0$.

It is meaningful to talk about $P(X \le 5)$.

Probability Density Function…

A function f(x) is called a ***probability density function*** (over the range $\mathbf{a} \le \mathbf{x} \le \mathbf{b}$ if it meets the following requirements:

1) $f(x) \ge 0$ for all **x** between **a** and **b**, and

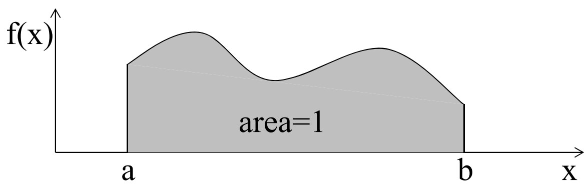

2) The total area under the curve between **a** and **b** is 1.0

Uniform Distribution...

Consider the ***uniform probability distribution*** (sometimes called the *rectangular probability distribution*).

It is described by the function:

$$\boxed{f(x) = \frac{1}{b-a}, \quad where \quad a \le x \le b}$$

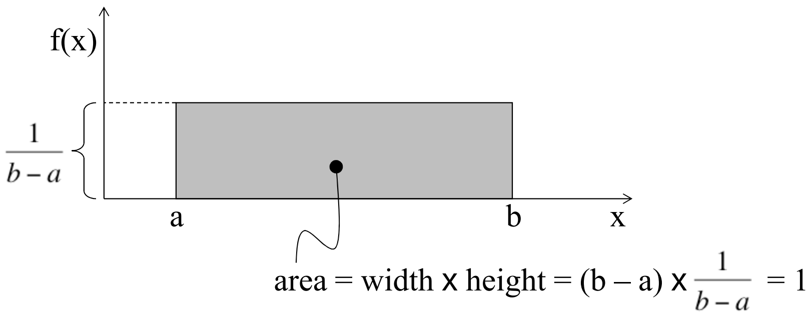

Example 8.1(a)...

The amount of gasoline sold daily at a service station is uniformly distributed with a minimum of 2,000 gallons and a maximum of 5,000 gallons.

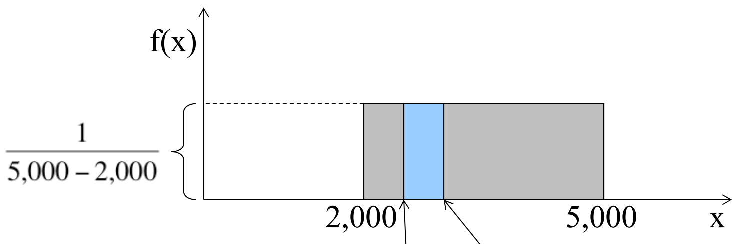

Find the probability that daily sales will fall between 2,500 and 3,000 gallons.

Algebraically: what is **P(2,500 ≤ X ≤ 3,000)** ?

Example 8.1(a)...

$\mathbf{P(2{,}500 \leq X \leq 3{,}000)} = (3{,}000 - 2{,}500) \times \frac{1}{3{,}000} = .1667$

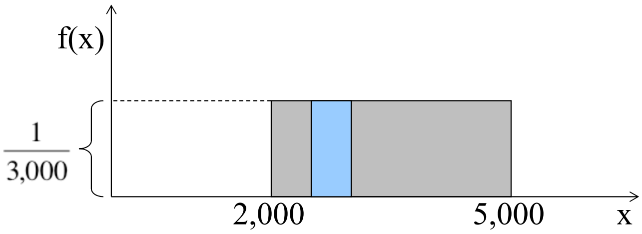

"there is about a 17% chance that between 2,500 and 3,000 gallons of gas will be sold on a given day"

Example 8.1(b)...

The amount of gasoline sold daily at a service station is uniformly distributed with a minimum of 2,000 gallons and a maximum of 5,000 gallons.

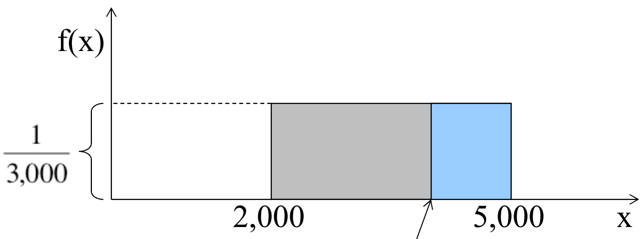

What is the probability that the service station will sell <u>at least</u> 4,000 gallons?

Algebraically: what is $\mathbf{P(X \geq 4{,}000)}$?

Example 8.1(b)...

P(X ≥ 4,000) = (5,000 – 4,000) x $\frac{1}{3,000}$ = .3333

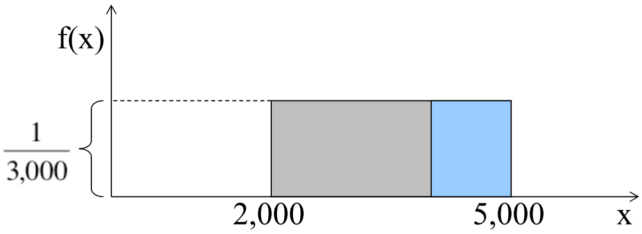

"There is a one-in-three chance the gas station will sell more than 4,000 gallons on any given day"

Example 8.1(c)...

The amount of gasoline sold daily at a service station is uniformly distributed with a minimum of 2,000 gallons and a maximum of 5,000 gallons.

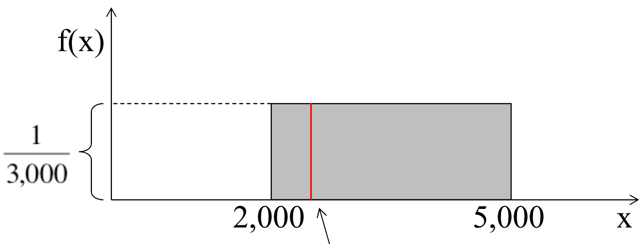

What is the probability that the station will sell exactly 2,500 gallons?

Algebraically: what is **P(X = 2,500)** ?

Example 8.1(c)...

P(X = 2,500) = (2,500 – 2,500) x $\frac{1}{3{,}000}$ = 0

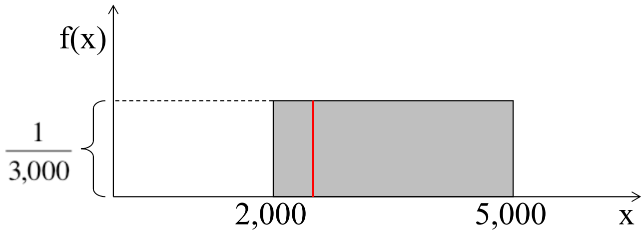

"The probability that the gas station will sell <u>exactly</u> 2,500 gallons is zero"

The Normal Distribution...

The ***normal distribution*** is the most important of all probability distributions. The probability density function of a ***normal random variable*** is given by:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \qquad -\infty < x < \infty$$

It looks like this:

Bell shaped,

Symmetrical around the mean μ …

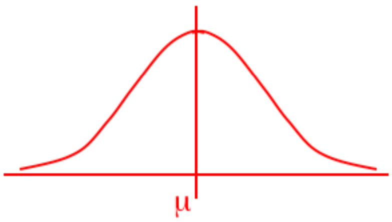

The Normal Distribution...

Important things to note:

The normal distribution is fully defined by two parameters: its standard deviation and mean

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \qquad -\infty < x < \infty$$

The normal distribution is bell shaped and symmetrical about the mean

Unlike the range of the uniform distribution ($a \le x \le b$)
Normal distributions ***range from minus infinity to plus infinity***

Standard Normal Distribution...

A normal distribution whose mean is zero and standard deviation is one is called the ***standard normal distribution***.

$$f(x) = \frac{1}{1\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-0}{1}\right)^2} \qquad -\infty < x < \infty$$

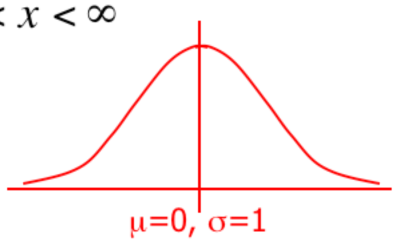

As we shall see shortly, any normal distribution can be ***converted*** to a standard normal distribution with simple algebra. This makes calculations much easier.

Normal Distribution...

The normal distribution is described by two parameters: its mean μ and its standard deviation σ. Increasing the mean *shifts* the curve to the right…

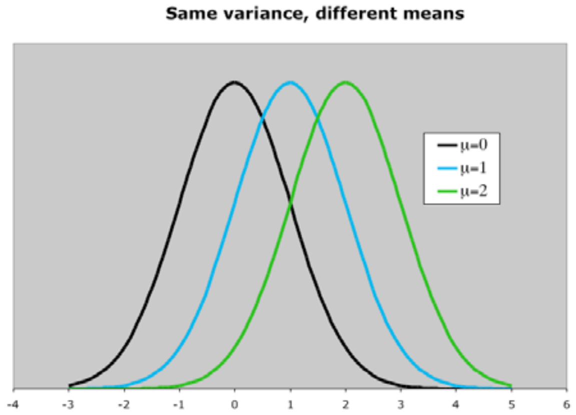

Normal Distribution...

The normal distribution is described by two parameters: its mean μ and its standard deviation σ. Increasing the standard deviation "*flattens*" the curve…

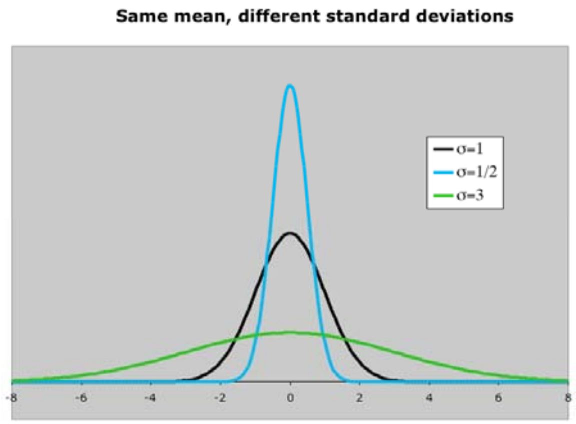

Calculating Normal Probabilities…

We can use the following function to convert any normal random variable to a **standard** normal random variable…

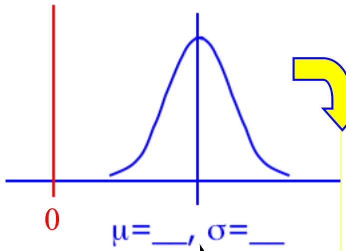

$$Z = \frac{X - \mu}{\sigma}$$

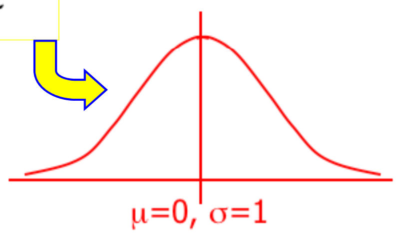

Some advice: always draw a picture!

Calculating Normal Probabilities…

We can use the following function to convert any normal random variable to a **standard** normal random variable…

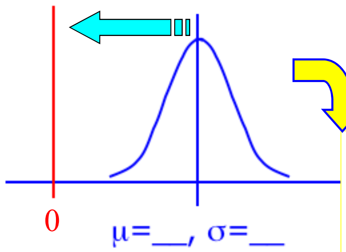

This shifts the mean of X to zero…

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

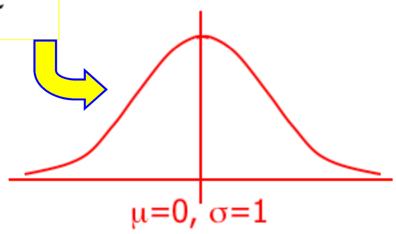

Calculating Normal Probabilities...

We can use the following function to convert any normal random variable to a **standard** normal random variable…

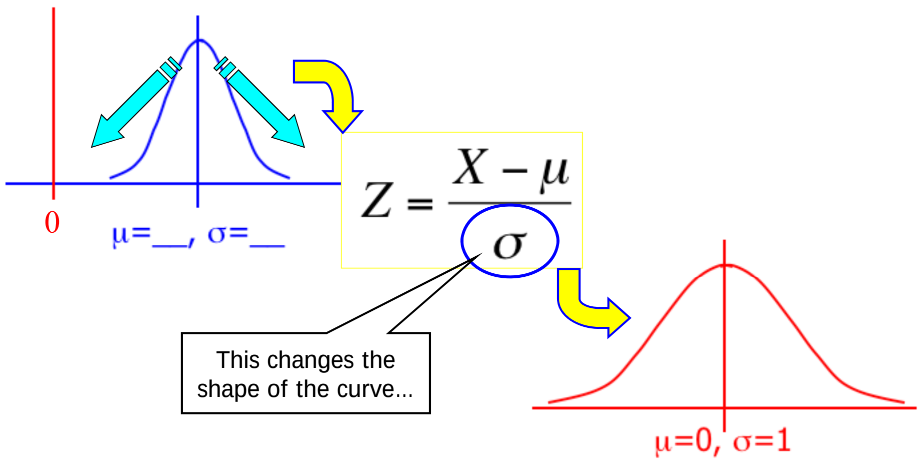

Example 8.2...

Suppose that at another gas station the daily demand for regular gasoline is normally distributed with a mean of 1,000 gallons and a standard deviation of 100 gallons.

The station manager has just opened the station for business and notes that there is exactly 1,100 gallons of regular gasoline in storage.

The next delivery is scheduled later today at the close of business. The manager would like to know the probability that he will have enough regular gasoline to satisfy today's demands.

Example 8.2...

The demand is normally distributed with mean $\mu = 1{,}000$ and standard deviation $\sigma = 100$. We want to find the probability

$$P(X < 1{,}100)$$

Graphically we want to calculate:

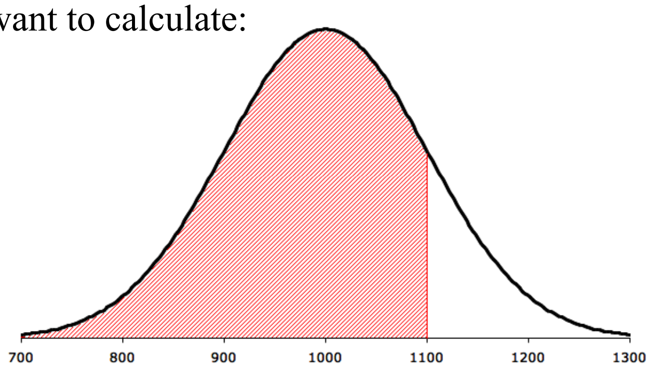

Example 8.2...

The first step is to standardize X. However, if we perform any operations on X we must perform the same operations on 1,100. Thus,

$$P(X < 1{,}100) = P\left(\frac{X-\mu}{\sigma} < \frac{1{,}100 - 1{,}000}{100}\right) = P(Z < 1.00)$$

Example 8.2...

The figure below graphically depicts the probability we seek.

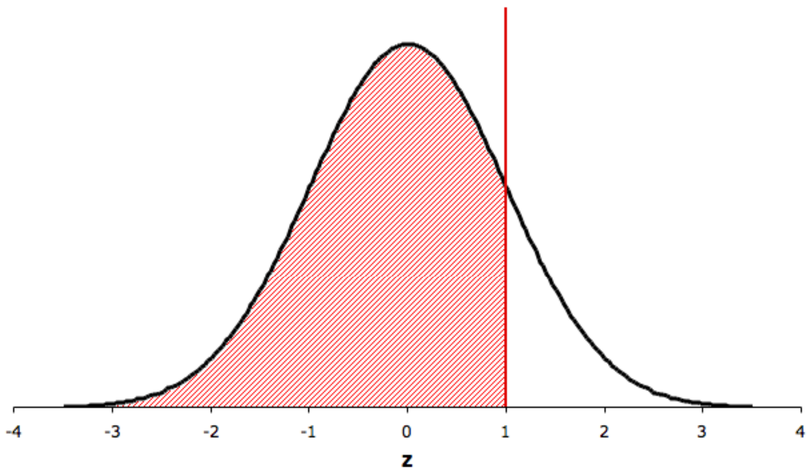

Example 8.2...

The values of Z specify the location of the corresponding value of X.

A value of Z = 1 corresponds to a value of X that is 1 standard deviation above the mean.

Notice as well that the mean of Z, which is 0 corresponds to the mean of X.

Example 8.2...

If we know the mean and standard deviation of a normally distributed random variable, we can always transform the probability statement about X into a probability statement about Z.

Consequently, we need only one table, Table 3 in Appendix B, the standard normal probability table.

Table 3...

This table is similar to the ones we used for the binomial and Poisson distributions.

That is, this table lists cumulative probabilities $P(Z < z)$ for values of z ranging from -3.09 to $+3.09$

Table 3...

Suppose we want to determine the following probability.

$P(Z < -1.52)$

We first find -1.5 in the left margin. We then move along this row until we find the probability under the .02 heading. Thus,

$P(Z < -1.52) = .0643$

Table 3...

$P(Z < -1.52) = .0643$

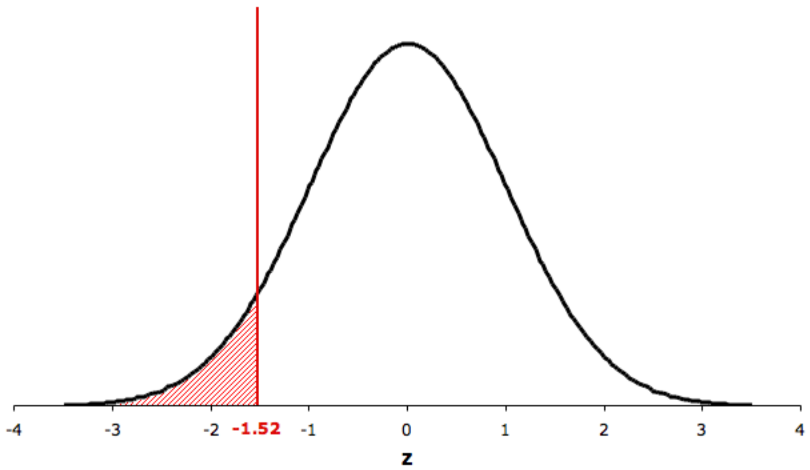

Table 3...

As was the case with Tables 1 and 2 we can also determine the probability that the standard normal random variable is greater than some value of z.

For example, we find the probability that Z is greater than 1.80 by determining the probability that Z is less than 1.80 and subtracting that value from 1.

Applying the complement rule we get

$$P(Z > 1.80) = 1 - P(Z < 1.80) = 1 - .9641 = .0359$$

Table 3...

$P(Z > 1.80) = 1 - P(Z < 1.80) = 1 - .9641 = .0359$

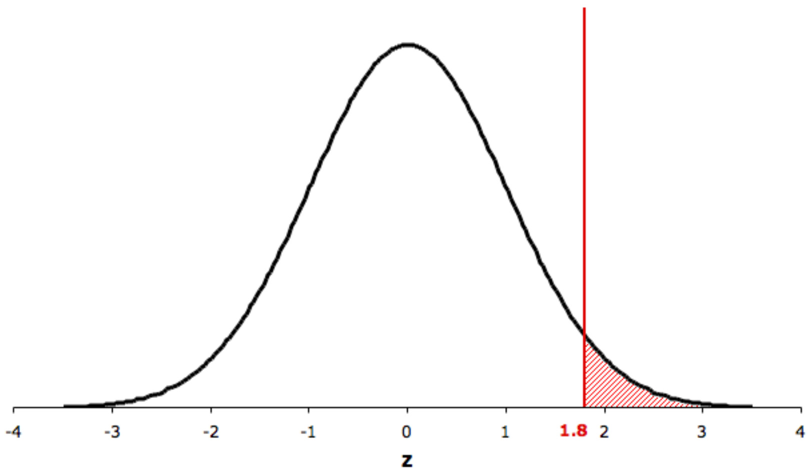

Table 3...

We can also easily determine the probability that a standard normal random variable lies between 2 values of z. For example, we find the probability

$$P(-1.30 < Z < 2.10)$$

By finding the 2 cumulative probabilities and calculating their difference. That is

$$P(Z < -1.30) = .0968$$

and

$$P(Z < 2.10) = .9821$$

Hence,

$$P(-1.30 < Z < 2.10) = P(Z < 2.10) - P(Z < -1.30)$$
$$= .9821 - .0968 = .8853$$

Table 3...

$P(-1.30 < Z < 2.10) = .8853$

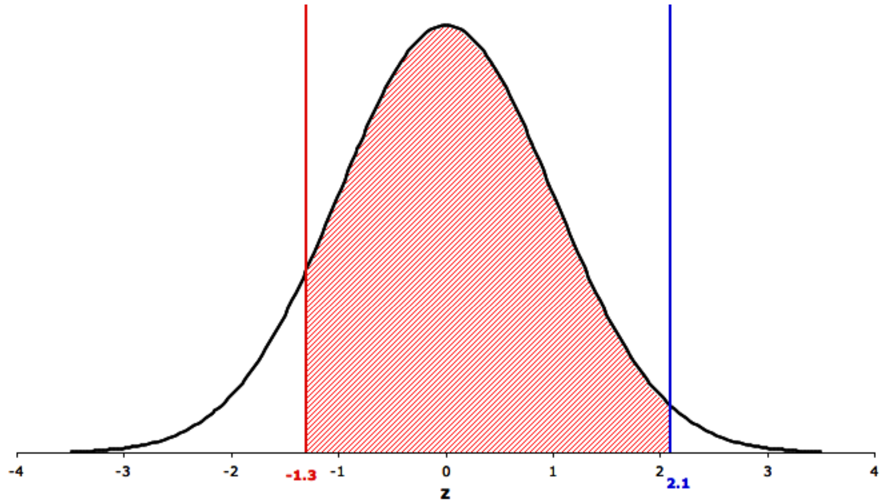

Table 3...

Notice that the largest value of z in the table is 3.09, and that $P(Z < 3.09) = .9990$. This means that

$$P(Z > 3.09) = 1 - .9990 = .0010$$

However, because the table lists no values beyond 3.09, we approximate any area beyond 3.10 as 0. That is,

$$P(Z > 3.10) = P(Z < -3.10) \approx 0$$

Table 3...

Recall that in Tables 1 and 2 we were able to use the table to find the probability that X is *equal* to some value of x, but that we won't do the same with the normal table.

Remember that the normal random variable is continuous and the probability that a continuous random variable it is equal to any single value is 0.

Example 8.2...

Finally returning to Example 8.2, the probability that we seek is

$$P(X < 1{,}100) = P(Z < 1.00) = .8413$$

Example 8.2...

$P(X < 1{,}100) = P(Z < 1.00) = .8413$

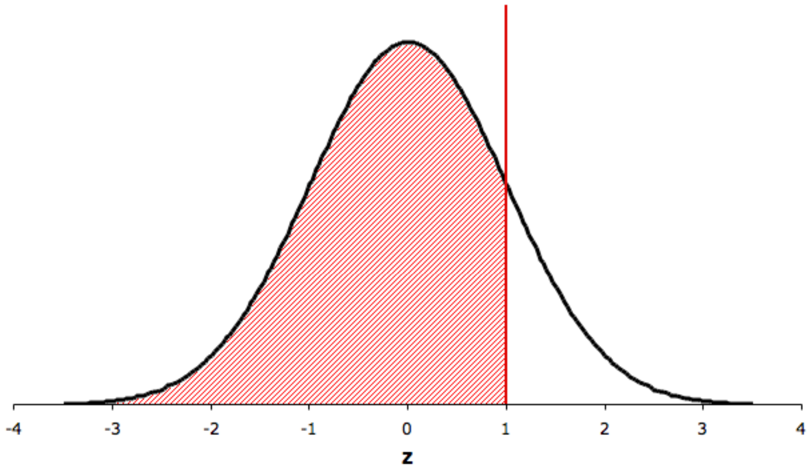

APPLICATIONS IN FINANCE: Measuring Risk

In Section 7.4 we developed an important application in finance where the emphasis was placed on reducing the variance of the returns on a portfolio. However, we have not demonstrated why risk is measured by the variance and standard deviation. The following example corrects this deficiency.

Example 8.3

Consider an investment whose return is normally distributed with a mean of 10% and a standard deviation of 5%.

a. Determine the probability of losing money.

b. Find the probability of losing money when the standard deviation is equal to 10%.

Example 8.2

a The investment loses money when the return is negative. Thus we wish to determine

$$P(X < 0)$$

The first step is to standardize both X and 0 in the probability statement.

$$P(X < 0) = P\left(\frac{X-\mu}{\sigma} < \frac{0-10}{5}\right) = P(Z < -2.00)$$

Example 8.2

From Table 3 we find

$$P(Z < -2.00) = .0228$$

Therefore the probability of losing money is .0228

Example 8.2

b. If we increase the standard deviation to 10% the probability of suffering a loss becomes

$$P(X < 0) = P\left(\frac{X-\mu}{\sigma} < \frac{0-10}{10}\right)$$

$$= P(Z < -1.00)$$

$$= .1587$$

Finding Values of Z...

Often we're asked to find some value of Z for a given probability, i.e. given an area (A) under the curve, what is the corresponding value of z (z_A) on the horizontal axis that gives us this area? That is:

$$\mathbf{P(Z > z_A) = A}$$

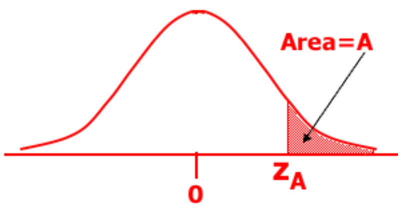

Finding Values of Z...

What value of z corresponds to an area under the curve of 2.5%? That is, what is $z_{.025}$?

(1 – A) = (1–.025) = .9750

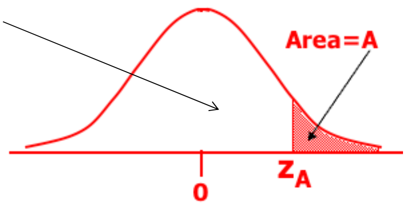

If you do a "reverse look-up" on Table 3 for .9750, you will get the corresponding $z_A = 1.96$

Since $P(z > 1.96) = .025$, we say: $\mathbf{z_{.025} = 1.96}$

Exponential Distribution...

Another important continuous distribution is the ***exponential distribution*** which has this probability density function:

$$f(x) = \lambda e^{-\lambda x}, \quad x \geq 0$$

Note that $x \geq 0$. Time (for example) is a non-negative quantity; the exponential distribution is often used for time related phenomena such as the length of time between phone calls or between parts arriving at an assembly station.

For the exponential random variable

$$\mu = \sigma = \frac{1}{\lambda}$$

Exponential Distribution...

The exponential distribution depends upon the value of λ

Smaller values of λ "flatten" the curve:

(E.g. exponential distributions for

$\lambda = .5, 1, 2$)

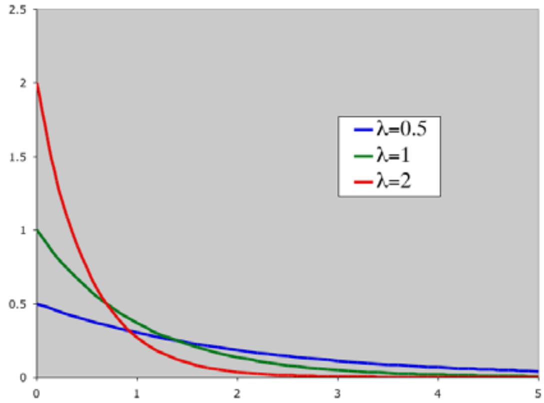

Exponential Distribution...

If X is an exponential random variable, then we can calculate probabilities by:

$$P(X > x) = e^{-\lambda x}$$

$$P(X < x) = 1 - e^{-\lambda x}$$

$$P(x_1 < X < x_2) = P(X < x_2) - P(X < x_1) = e^{-\lambda x_1} - e^{-\lambda x_2}$$

Example 8.6...

The lifetime of an alkaline battery (measured in hours) is exponentially distributed with $\lambda = .05$

Find the probability a battery will last between 10 & 15 hours…

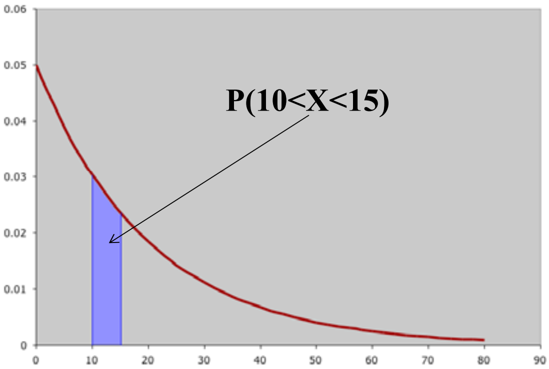

P(10<X<15)

$$P(10 \le X \le 15) = e^{-.05(10)} - e^{-.05(15)}$$
$$= e^{-.5} - e^{-.75} = .6065 - .4724$$
$$= .1341$$

"There is about a 13% chance a battery will only last 10 to 15 hours"

Other Continuous Distributions...

Three other important continuous distributions which will be used extensively in later sections are introduced here:

Student ***t*** Distribution,

Chi-Squared Distribution, and

F Distribution.

Student *t* Distribution...

Here the letter ***t*** is used to represent the random variable, hence the name. The density function for the Student ***t*** distribution is as follows…

$$f(t) = \frac{\Gamma[(\nu+1)/2]}{\sqrt{\nu\pi}\,\Gamma(\nu/2)}\left[1+\frac{t^2}{\nu}\right]^{-(\nu+1)/2}$$

ν (nu) is called ***the degrees of freedom***, and

Γ (Gamma function) is Γ(k)=(k-1)(k-2)…(2)(1)

Student *t* Distribution...

Much like the standard normal distribution, the Student ***t*** distribution is "mound" shaped and symmetrical about its mean of zero:

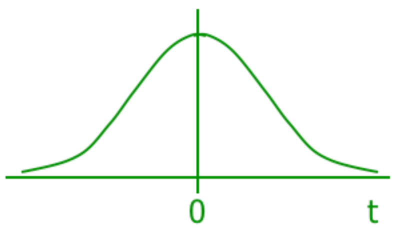

The mean and variance of a Student t random variable are

$E(t) = 0$

and

$V(t) = \dfrac{\nu}{\nu - 2}$ for ν > 2.

Student *t* Distribution...

In much the same way that μ and σ define the normal distribution, ν, the degrees of freedom, defines the Student *t* Distribution:

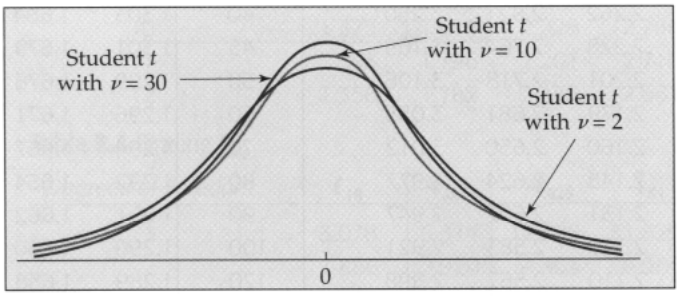

Figure 8.24

As the number of degrees of freedom increases, the *t* distribution approaches the standard normal distribution.

Determining Student *t* Values...

The student *t* distribution is used extensively in statistical inference. Table 4 in Appendix B lists values of $t_{A,\nu}$

That is, values of a Student *t* random variable with ν degrees of freedom such that:

$$P(t > t_{A,\nu}) = A$$

The values for A are pre-determined "critical" values, typically in the 10%, 5%, 2.5%, 1% and 1/2% range.

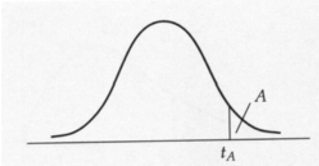

Using the t table (Table 4) for values…

For example, if we want the value of t with 10 degrees of freedom such that the area under the Student t curve is .05:

Area under the curve value (t_A) : COLUMN

$t_{.05,10}$

$t_{.05,10}$=1.812

Degrees of Freedom : ROW

DEGREES OF FREEDOM	$t_{.100}$	$t_{.050}$	$t_{.025}$	$t_{.010}$	$t_{.005}$
1	3.078	6.314	12.706	31.821	63.657
2	1.886	2.920	4.303	6.965	9.925
3	1.638	2.353	3.182	4.541	5.841
4	1.533	2.132	2.776	3.747	4.604
5	1.476	2.015	2.571	3.365	4.032
6	1.440	1.943	2.447	3.143	3.707
7	1.415	1.895	2.365	2.998	3.499
8	1.397	1.860	2.306	2.896	3.355
9	1.383	1.833	2.262	2.821	3.250
10	1.372	1.812	2.228	2.764	3.169
11	1.363	1.796	2.201	2.718	3.106

χ^2 Chi-Squared Distribution…

The chi-squared density function is given by:

$$f(\chi^2) = \frac{1}{\Gamma(\nu/2)}\frac{1}{2^{\nu/2}}(\chi^2)^{(\nu/2)-1}e^{-\chi^2/2} \qquad ,\chi^2 > 0$$

As before, the parameter ν is the number of degrees of freedom.

$$E(\chi^2) = \nu$$

$$V(\chi^2) = 2\nu$$

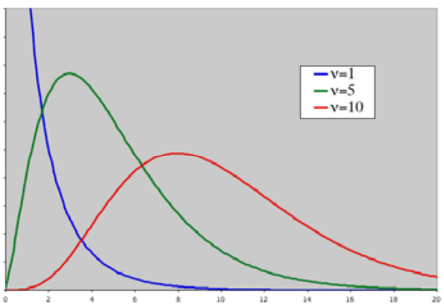

Figure 8.27

χ^2 Chi-Squared Distribution...

Notes:

The chi-squared distribution is not symmetrical

The square, i.e. χ^2 , forces non-negative values (e.g. finding P(χ^2< 0) is illogical).

Table 5 in Appendix B makes it easy to look-up probabilities of this sort, i.e.

P($\chi^2 > \chi^2_A$) = A:

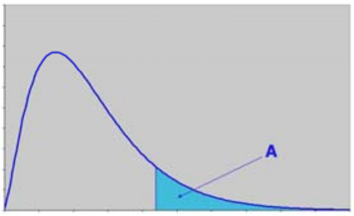

χ^2_A

χ^2 Chi-Squared Distribution...

For probabilities of this sort:

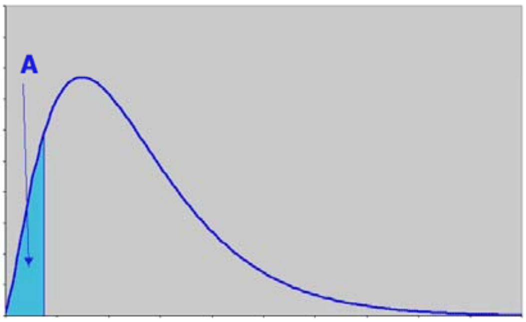

χ^2_{1-A}

We use 1–A, i.e. we are determining

P($\chi^2 < \chi^2_{1-A}$) = A

For Example...

To find the point in a chi-squared distribution with 8 degrees of freedom, such that the area to the right is .05,

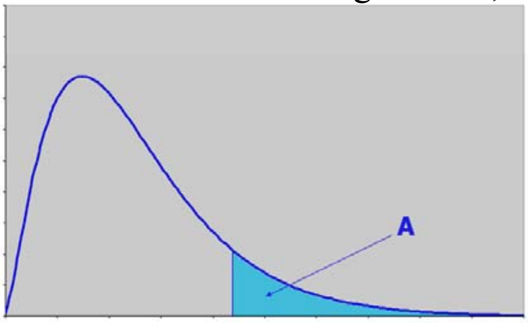

Look up the intersection of the 8 d.f. row with the $\chi^2_{.050}$ column, yielding a value of 15.5

For Example...

To find the point in a chi-squared distribution with 8 degrees of freedom, such that the area to the left is .05,

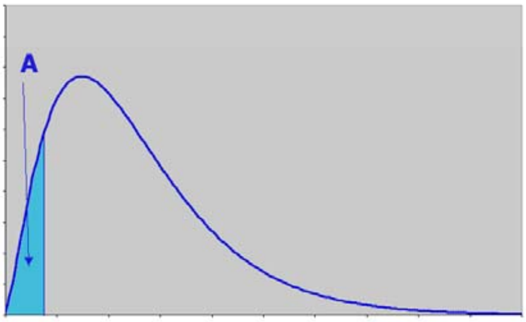

Look up the intersection of the 8 d.f. row with the $\chi^2_{.950}$ column, yielding a value of 2.73

For Example...

This makes sense:

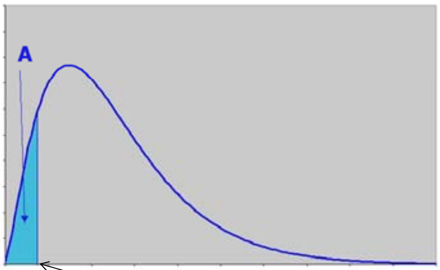

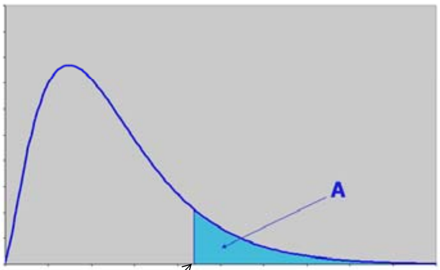

$\chi^2_{.950}$ =2.73 $\chi^2_{.050}$ =15.5

Remember the χ^2 axis starts and zero and increases!

F Distribution...

The ***F*** density function is given by:

$$f(F) = \frac{\Gamma\left(\frac{\nu_1+\nu_2}{2}\right)}{\Gamma\left(\frac{\nu_1}{2}\right)\Gamma\left(\frac{\nu_2}{2}\right)}\left(\frac{\nu_1}{\nu_2}\right)^{\frac{\nu_1}{2}} \frac{F^{\frac{\nu_1-2}{2}}}{\left(1+\frac{\nu_1 F}{\nu_2}\right)^{\frac{\nu_1+\nu_2}{2}}}$$

F > 0. Two parameters define this distribution, and like we've already seen these are again ***degrees of freedom***.

ν_1 is the "numerator" degrees of freedom and

ν_2 is the "denominator" degrees of freedom.

F Distribution…

The mean and variance of an ***F*** random variable are given by:

$$E(F) = \frac{\nu_2}{\nu_2 - 2} \qquad \nu_2 > 2$$

and

$$V(F) = \frac{2\nu_2^2(\nu_1 + \nu_2 - 2)}{\nu_1(\nu_2 - 2)^2(\nu_2 - 4)} \qquad \nu_2 > 4$$

The F distribution is similar to the χ^2 distribution in that its starts at zero (is non-negative) and not symmetrical.

Determining Values of *F*…

For example, what is the value of F for 5% of the area under the right hand “tail” of the curve, with a numerator degree of freedom of 3 and a denominator degree of freedom of 7?

Solution: use the F look-up (Table 6)

There are different tables for different values of A. Make sure you start with the correct table!!

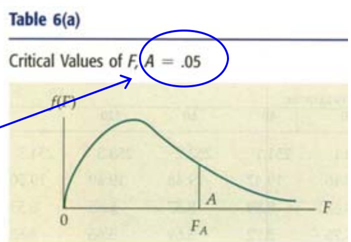

$F_{.05,3,7}$

$F_{.05,3,7}$=4.35

Denominator Degrees of Freedom : ROW

Numerator Degrees of Freedom : COLUMN

Determining Values of *F*...

For areas under the curve on the left hand side of the curve, we can leverage the following relationship:

$$F_{1-A,\nu_1,\nu_2} = \frac{1}{F_{A,\nu_2,\nu_1}}$$

Pay close attention to the order of the terms!