Welcome!

I wrote these notes based on the feedback from my students….almost 100% of them said they just don't have the time to read the entire curriculum, so they wished they had a set of notes that they could rely on to give them a solid knowledge base, show them how the concepts might be applied on the exam and guidance on what topics to touch lightly and what topics to study hard….

So that is exactly what I did, and you are now looking at an excerpt from the study notes that will allow you to:

1) Learn the complex financial concepts required to answer CFA exam questions.
2) Expand your knowledge base, layering in extra marks to get you over the passing score.
3) Retain your knowledge right up to the exam day.

Best of luck with your studies!

Prof. Brian Gordon, CFA, CFP, CIM, MBA, FCSI

P.S.   To get the greatest benefit out of using the notes, I suggest that you print the CFA Curriculum Learning Outcome Statements (LOS) and use them as a guide as you go through the notes….you will see my bolded titles follow along with the LOS telling you exactly what you need to know.

Level 1

Study Session 3

Quantitative Methods (QM)

READINGS

1. Common Probability Distributions

2. Sampling and Estimation

3. Hypothesis Testing

4. Technical Analysis

- Expect to see 28 questions that relate to QM. SS 2 and SS 3

- 14 questions from SS2 and 14 questions from SS 3

If you can handle the following items from SS 3, you should be in good shape for the exam.

1. Define standard normal distribution

2. Calculate probabilities using standard normal distribution

3. Calculate and interpret Roy's safety first ratio

4. Steps in a hypothesis test

5. 1 tailed vs. 2 tailed tests

6. Type I vs. Type II error

7. Identify the appropriate test statistic to use

8. Calculate t stat and degrees of freedom

9. Calculate and interpret a confidence interval

10. Be able to describe biases in estimation

11. Binomial distribution calculation

12. Be able to calculate and describe tracking error

13. Be able to calculate and interpret an F test

14. Know the assumptions of technical analysis

15. Know the basic characteristics of the different charts

16. DO NOT memorize the list of technical trading rules!

Common Probability Distributions

Professor's Comment:
Try not to get lost in the technical nature of the material, the symbols and formulas can be very intimidating. I doubt you will be presented with any formulas on the exam, so all you need is to have a high level understanding of the concepts that follow. I suggest that you focus on recognizing the information given in questions as your signal of what to do….are you ever going to be using this stuff after the Level 1 exam???

Basic Concepts to Know:
A ***random variable*** is just something that has possible values that are not known with exact certainty. Our quest is to use statistical techniques to narrow down a range (ie. create a distribution of possible outcomes) to give us a better idea of what the value could be.

Formally, a ***probability distribution*** is defined as a list of the probabilities of the possible outcomes of a random variable.

A ***discrete random variable*** has exact possible values that we can count, such as 1, 2, 3, or 4.

A ***continuous random variable*** includes all values that fit in between the discrete values in some interval, for example, 1, 2, 3, 4 would also include 1.1, 1.2, 1.3, etc. but these values are not given in our list, so we cannot count all the outcomes.

Bottom Line:
A continuous random variable differs from a discrete random variable in that it takes on an *uncountable* infinite number of possible outcomes.

Quick Example:
1.   The number of stocks in a portfolio or index would *most likely* be an example of:

      A.   a discrete random variable
      B.   a continuous random variable
      C.   a probability distribution

2.   Rate of return on a portfolio or index would *most likely* be an example of:

      A.   a discrete random variable
      B.   a continuous random variable
      C.   a probability distribution

Solutions:
1. A   This is an example of a specific countable number of items.
2. B   This is an example of an uncountable infinite number of possibilities that can take on any value within the range.

Possible Outcomes of a Specified Discrete Random Variable

All we are talking about here is "…what is the probability of a variable being equal to a specific number…" The probability is calculated as:

$$p(x) = P(X = x)$$

Where: $p(x)$ is the probability that $X = x$
the probability must be between $0 \leq p(x) \leq 1$
the sum of all probabilities, $p(x) = 1$

Quick Example:

Table 1

Outcome	1	2	3	4
Probability	10%	30%	40%	?

1.   What is the probability that $X = 4$?

2.   Based on the data in Table 1, **calculate** $P(X = 2)$:

3.   Based on the data in Table 1, **calculate** $P(X = 2 \text{ or } X = 3)$

Solutions:

1.   The sum of the probabilities must equal 100% (or 1 if using decimals), so the probability that $X = 4$ is 100% - 10% - 30% - 40% = 20%.

2.   Just read the probability from the table, so the probability that the variable is equal to 2 is 30%.

3.   Reading the probabilities from the table, we get 30% + 40% = 70%.

Cumulative Distribution Function

All random variables (discrete and continuous) have a cumulative distribution function. It is a function giving the *probability* that the random variable X is *less than or equal* to x, for every value x.

For a discrete random variable, the cumulative distribution function is found by summing up the probabilities.

Quick Example:

Table 1

Outcome	1	2	3	4
Probability	10%	30%	40%	20%

1.   Based on the data in Table 1, calculate $P(X \leq 1)$:

2.   Based on the data in Table 1, calculate $P(X \leq 2)$:

3.   Based on the data in Table 1, calculate $P(X \leq 3)$:

Solutions:

1.   This asks us for all the values equal to 1 and less than 1. In this case there is only 1 possible outcome that satisfies this description and the probability is 10%.

2. $P(X \leq 2) = 10\% + 30\% = 40\%$.

3. $P(X \leq 3) = 10\% + 30\% + 40\% = 80\%$

Other Definitions to Know:
When the outcomes for discrete random variables all have the same chance of occurring, we call them *discrete uniform random variables*.

If a random variable has just two possible outcomes, success or failure, then it is referred to as a *Bernoulli random variable*:

Probability of success => $P(X = 1) = p$
Probability of failure => $P(X = 0) = 1 - p$

Where p is greater than 0 and less than 1.

We use this type of distribution to value options (**LINK:** SS 17 Binomial Option Pricing Model)

A *binomial random variable* is a random variable that counts the number of successes in a sequence of independent Bernoulli trials with fixed probability of success.

Binomial Distribution and Probabilities
You will need to recognize the "binomial signals" in the question to know what to do. Common signals include phrases like, "probability of success" or "chance of an up movement".

More formally:
The *binomial distribution* describes the behavior of a variable if the following conditions apply:

1. The number of observations n is fixed.
2. Each observation is independent.
3. Each observation represents one of two outcomes ("success" or "failure").
4. The probability of "success" p is the same for each outcome.

You need to know the following:
The discrete probability distribution of the number of successes in a sequence of n independent yes/no experiments, each of which yields success with probability p:

Mean $= \mu = np$

Variance $= \sigma^2 = np(1-p)$

To calculate the probability for a binomial distribution, use the following formula:

$P(r) = [\, n \, C \, r \,] \, (P)^r \, (1 - P)^{n-r}$

Binomial Learning Examples:

1.   Over the last 10 years, a company's earnings increased from the previous year seven times and decreased from the previous year three times. You decide to model the number of earnings increases over the next decade as a binomial random variable. Based on this information the expected number of yearly earnings increases and variance of the number of yearly earnings increases over the next 10 years are respectively:

      A.   7.0 and 1.44
      B.   7.0 and 2.10
      C.   6.3 and 2.10

2.   Suppose a day trader makes 12 trades in day, 3 are profitable and 9 are not. **Calculate** the probability of 3 or fewer profitable trades per day assuming the probability of a profitable trade is 50%.

Solutions:

1.B $E(x) = np = 10 \times 0.7 = 7$

      $\sigma^2 = np(1 - p) = 10 \times 0.7 \times (1 - 0.7) = 2.1$

2.   Picking up on the language "3 or fewer" is the key to answering this question correctly:

      $F(3) = p(3) + p(2) + p(1) + p(0)$
         $= 0.0537 + 0.0161 + 0.000293 + 0.000244 = 7.03\%$

      $P(0) = (12 \text{ C } 0)(0.5)^0 (1 - 0.5)^{(12-0)} = 0.000244$
      $P(1) = (12 \text{ C } 1)(0.5)^1 (1 - 0.5)^{(12-1)} = 0.000293$
      $P(2) = (12 \text{ C } 2)(0.5)^2 (1 - 0.5)^{(12-2)} = \quad 0.0161$
      $P(3) = (12 \text{ C } 3)(0.5)^3 (1 - 0.5)^{(12-3)} = \quad \underline{0.0537}$
      $P(3 \text{ or less}) = \qquad\qquad\qquad\qquad 0.0703 \quad$ or 7.03%

Continuous Uniform Distribution

The *continuous uniform distribution* is the probability distribution of random number selection from the continuous interval between **a** and **b**.

A *uniform distribution* is a distribution that has constant probability:

$$p(x) = \frac{1}{b - a}$$

Where b is the upper limit and a is the lower limit of the interval

The cumulative distribution is calculated as:

$$P(x) = \frac{x - a}{b - a}$$

Where x is the limit you are interested in finding
x must be between a and b

The mean (or expected value) $= \mu = \dfrac{b + a}{2}$

The variance $= \sigma^2 = \dfrac{(b - a)^2}{12}$

Quick Example:

A random variable follows a continuous uniform distribution over the interval 30 to 90.

1.   **Calculate** the expected value for the interval:

2.   **Calculate** the variance for the interval:

3.   **Calculate** the probability of an outcome falling between 40 and 55:

Solutions:

1. The expected value is simply the average of the interval end points: $(90+30)/2 = 60$
2. Variance $(90 - 30)^2 / 12 = 300$
3. $P(x) = (55 - 40) / (90 - 30) = 0.25$

The **probability distribution** of a discrete random variable is a list of probabilities associated with each of its possible values. It is also sometimes called the ***probability function***.

Uniform Probability Distribution Example:

Table 1

X = x	Probability Function $P(x) = P(X = x)$	Cumulative Distribution Function $F(x) = P(X \leq x)$
1	0.125	0.125
2	0.125	0.250
3	0.125	0.375
4	0.125	0.500
5	0.125	0.625
6	0.125	0.750
7	0.125	0.875
8	0.125	1.000

Based on the information given in Table 1, calculate the following:

a) $P(2 \leq X \leq 4)$

b) $P(X \leq 4)$

Solutions:

a) $P(2 \leq X \leq 4) = 0.125 + 0.125 + 0.125 = 0.375$

b) $P(X \leq 4) = 0.5$

Tracking error is the total return on a portfolio (gross of fees) less the total return on the benchmark. It is used to evaluate the performance of Passive Managers (Indexers!)….whose goal is to simply match the index return!

For a Passive Manager, the GOAL is to minimize the tracking error, NOT to generate a positive ALPHA by outperforming the index.
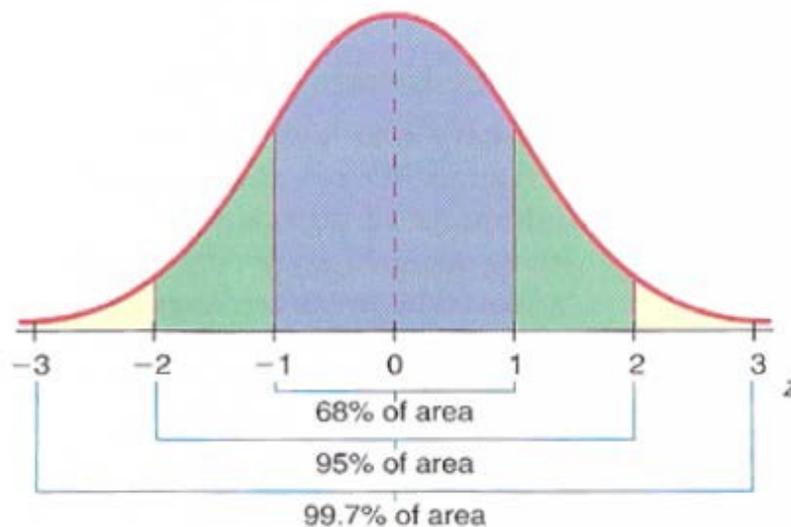
Normal Distribution

A distribution is said to be *"normal"* if it exhibits a symmetrical pattern, has one mode, extends to infinity on either side and includes 100% of all possible outcomes (area under the curve equals 100%)

If the mean and standard deviation are known, then you can draw conclusions about any data point!

Rules of Thumb (Empirical Rules)

Approximately 68% of the data lie within +/- 1 Standard Deviation from the Mean
Approximately 95% of the data lie within +/- 2 Standard Deviations from the Mean
Approximately 99% of the data lie within +/- 3 Standard Deviations from the Mean

Normal Distribution Examples:

1. If common stocks have an expected return of 13% with a standard deviation of 20% and expected returns are normally distributed, the probability of experiencing a return greater than 33% is *closest to*:

   A.   8%
   B.   16%
   C.   34%

2. If an investment strategy has an expected annual return of 15 percent and an annual standard deviation of 15 percent and investment returns are normally distributed, in any one year there is:

   A.   about a 68% chance that the actual return will fall within the range of -15% to 15%.
   B.   about a 32% chance that the actual return will exceed 30%.
   C.   about a 68% chance that the actual return will fall within the range of 0% to 30%.

Solutions:

1. Calculate z score: (33% - 13%) / 20% = 1

   Look up the z-table to get the probability of 16%. Or use your rules of thumb to answer this question!

2.C   Normally Distributed: ± 1 σ, 68% observations, symmetrical about the mean.

Standard Normal Distribution

Standard normal distribution is defined as a normal distribution where the mean is 0 and the standard deviation is 1 "unit".

Probability numbers derived from the standard normal distribution are used to help standardize a random variable, that is, express that number in terms of how many standard deviations it is away from its mean.

Standardizing a random variable is done by subtracting the expected value from the observed value, and then dividing the result by the standard deviation (σ). The result is a standard normal random variable which is denoted by the letter Z:

$$Z = \frac{\text{Observed} - \text{Expected}}{\text{Standard Deviation}}$$

The Z stat tells us the cumulative area under the Normal Curve.

For the CFA exam:

$Z = 1.65$, denotes 95% cumulative area under the curve, it is also used as the "reliability factor" for the 90% confidence interval

$Z = 1.96$, denotes 97.5% cumulative area under the curve, it is also used as the "reliability factor" for the 95% confidence interval

Z Table Learning Example:

It is unlikely that you will need to use a Z-Table on the exam. Most likely you will just use the rules of thumb that you have memorized!

However, just to make sure, this quick example will show you how to use it:

Z-Table

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06
-1.9	0.0287	0.0281	0.0274	0.0268	0.0262	0.0256	0.0250
-1.6	0.0548	0.0537	0.0526	0.0516	0.0505	0.0495	0.0485
-1.0	0.1587	0.1562	0.1539	0.1515	0.1492	0.1469	0.1446
-0.5	0.3085	0.3050	0.3015	0.2981	0.2946	0.2912	0.2877
-0.4	0.3446	0.3409	0.3372	0.3336	0.3300	0.3264	0.3228
0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4801	0.4761
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750

A Portfolio Manager is reviewing a client account that has an expected return of 18% and a standard deviation of returns of 5%. Assuming that the portfolio returns are normally distributed, **calculate** the following:

a)   the probability of earning a return less than 15.35%

b)   the probability of earning a return less than 23%

c)   the probability of earning a return greater than 18%

Solutions:

a)   First calculate the Z-stat:

Z =   Actual – Expected = 15.35% - 18% = -0.53
      Standard deviation   5%

Second, look up the area under the curve using the Z-stat of -0.53:

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06
-1.9	0.0287	0.0281	0.0274	0.0268	0.0262	0.0256	0.0250
-1.6	0.0548	0.0537	0.0526	0.0516	0.0505	0.0495	0.0485
-1.0	0.1587	0.1562	0.1539	0.1515	0.1492	0.1469	0.1446
-0.5	0.3085	0.3050	0.3015	0.2981	0.2946	0.2912	0.2877
-0.4	0.3446	0.3409	0.3372	0.3336	0.3300	0.3264	0.3228
0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4801	0.4761

Third, interpret what the Z Table tells us:

There is approximately 29.81% chance of earning less than 15.35%

b)   First calculate the Z-stat:

Z =   Actual – Expected = 23% - 18% = 1
      Standard deviation   5%

Second, look up the area under the curve using the Z-stat of 1:

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06
0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4801	0.4761
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515

Third, interpret what the Z Table tells us:

There is approximately 84.13% chance of earning less than 23%

c)   OK…you've got it now…repeat the pattern:

1) Calculate the Z-stat: (18 – 18) / 5 = 0
2) Look up area under the curve: 0.5
3) Interpret the results: there is a 50% probability

A **Univariate distribution** describes a single random variable (such as one specific stock), while a **Multivariate distribution** specifies probabilities related to a group of random variables (such as a portfolio of stocks).

Shortfall risk is the risk that your portfolio value will fall below some minimum acceptable level over a specific period of time.

Roy's Safety First Ratio (SFR) is used as a measure of shortfall risk. The higher the SFR, the lower the shortfall risk, conversely, the lower the SFR, the higher the shortfall risk.

The SFR is calculated as:

$$SFR = \frac{\text{Expected Portfolio Return} - \text{Minimum Required Return}}{\text{Portfolio Standard Deviation}}$$

Roy's Safety First Ratio Example:

1.   A Portfolio Manager collects the following information regarding two Defined Benefit Pension Funds:

	Miller Corp. Pension Fund	CAAT Pension Fund
Expected Return	10%	20%
Standard Deviation of Return	10%	20%
Additional Data: Risk free rate of return is 5% Pension discount rate (minimum required return) is 6%		

On the basis of Roy's Safety First measure (SFR) which of the following statements is *most likely* correct?

A.   The risk of shortfall is the same for both Pensions since they both have the same minimum required return.

B.   The risk of shortfall is lower the CAAT Pension because its Roy's Safety First Measure of 0.70 is better than the SFR of 0.40 for the Miller Corp. Pension.

C.   The risk of shortfall is lower for the Miller Corp. Pension because its Roy's Safety First Measure of 0.50 is better than the SFR of 0.75 for the CAAT Pension.

2. You are looking at a client's portfolio that has an expected return(Rp) of 10% and a total risk level (SD) of 10%. If the client must earn at least 5% to achieve his investment goals, **calculate** the probability that the client's portfolio will fall below his minimum required return. (use the Z table below)

Z-Table

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06
-1.9	0.0287	0.0281	0.0274	0.0268	0.0262	0.0256	0.0250
-1.6	0.0548	0.0537	0.0526	0.0516	0.0505	0.0495	0.0485
-1.0	0.1587	0.1562	0.1539	0.1515	0.1492	0.1469	0.1446
-0.5	0.3085	0.3050	0.3015	0.2981	0.2946	0.2912	0.2877
-0.4	0.3446	0.3409	0.3372	0.3336	0.3300	0.3264	0.3228
0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4801	0.4761
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750

Solutions:

1.B Roy's Safety First Ratio for a portfolio is $SFR_p = [E(R_p) – R_{min}]/\sigma_p$. Note that larger values are better here because they indicate that the minimum return is a larger number of standard deviations away from the expected return, hence returns below the required minimum are less likely to occur.

Miller Corp. Pension SFR = [10 – 6]/10 = 0.40

CAAT Pension SFR = [20 – 6]/20 = 0.70

Hence the CAAT Pension has a lower risk of shortfall on the basis of the SFR criterion.

2. $SFR_A = (10-5)/10 = 0.5$
    $P(R_A < 5\%) = N(-0.5) = 0.3085$

We would interpret this result as there being a 30.85% chance of not meeting the 5% minimum required return

Normal vs. Lognormal Distributions
Normal distribution for stock returns is just an approximation!

If you know the Mean and Standard Deviation of the distribution, the Normal distribution can be defined.

The Log-normal *Density* distribution is skewed to the right with a fatter tail, whereas the Normal distribution is symmetrically distributed about the mean.

A lognormal distribution has two distinct properties:
1) it is always positive (bounded on the left by zero)
2) it is skewed to the right

Prices for financial assets are often found to be "lognormally" distributed.

The lognormal and normal distributions are related:
If a random variable is "lognormally" distributed, then its natural log, (ln) is normally distributed.

Continuously Compounded Rate of Return

$$r^c = \ln(1 + HPR) \qquad \text{or} \qquad r^c = \ln(S_1/S_0)$$

$$HPR = e^{(rc)(T)} - 1$$

Price Relative

Quick Example:
1. Hemma Litchin paid $38.50 for one share of GIL Metals, A NYSE listed security. Exactly one year later she sold her share for a profit of $4.43. Based on this information, Hemma would be *least likely* to conclude:

    A.   her holding period return was 11.5%
    B.   the continuously compounded rate of return is 10.89%
    C.   the price relative is 8.69

Solution:
1.C   HPR = (End – Beg) / Beg = ($38.50 + $4.43 - $38.50) / $38.50 = 0.115 or 11.5%
      $r^c = \ln(1 + HPR) = \ln(1 + 0.115) = 10.89\%$
      Price Relative = S_1/S_0 = ($38.50 + $4.43) / $38.50 = 1.11

Monte Carlo simulation is a complex, forward looking, mathematical technique used to calculate the likelihood of an event occurring given hundreds or thousands of possible outcomes.

Historical simulation uses past data to predict the likelihood of a future outcome occurring.

Sampling and Estimation

Sampling Error is the difference between the sample mean and the population mean

Random sampling – a process whereby data points are selected without any pre-ordered method, which permits the analyst to draw out *unbiased conclusions* from the sample to represent the entire population.

Stratified sampling – a process whereby the population set is divided into "layers" or "strata" defined by specific characteristics and then the analyst selects a few random data points from each layer to draw out conclusions about the entire population.

Bottom Line:
Compared to stratified sampling, random sampling is a less precise method, which results in a larger sampling error.

Time-series data – a process whereby the analyst uses historical data to draw out conclusions about possible future performance.

Cross-sectional data – refers to comparing data across different groups at the same point in time to draw out conclusions about relative value.

The **central limit theorem** tells us that the sample mean for large sample sizes will be distributed normally regardless of the distribution of the underlying population. This allows us to construct confidence intervals and run hypothesis tests.

Standard Error of the Sample Mean (Se)
The standard error is often used to provide an indication of the size of uncertainty in an estimation and is used to calculate **confidence intervals:**

$$Se = \sigma / \sqrt{n}$$

Confidence Intervals
In general, you calculate a confidence interval as follows:

Expected Value +/- ( ) ( )

"Reliability Factor" or "Critical Value"   Standard Deviation or Standard Error (for the exam it could be: t, z, 1.96 or 1.65)
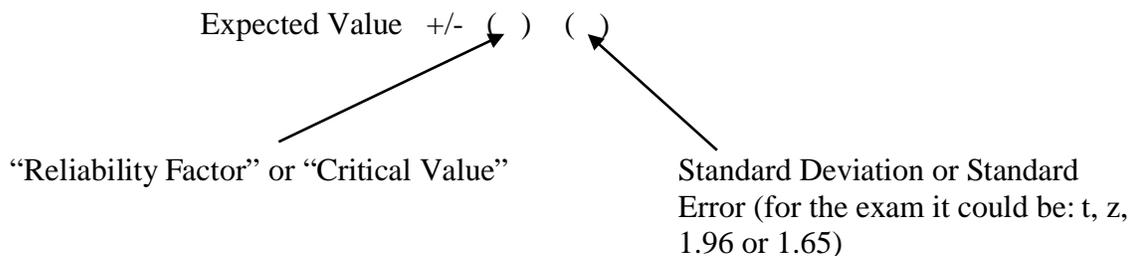
Note: a larger sample size decreases the width of a confidence interval.

Confidence Interval Examples:

1. Given a data series that is normally distributed with a mean of 100 and a standard deviation of 10, *approximately* 95 percent of the numbers in the series will fall within:

   A.   70 to 130
   B.   90 to 110
   C.   80 to 120

2. Given a data series that is normally distributed with a mean of 100 and a standard deviation of 10, *approximately* 90 percent of the numbers in the series will fall within:

   A.   90 to 110
   B.   83.5 to 116.4
   C.   80.4 to 119.6

3. Three securities had the following mean monthly returns, standard deviations of return over the period January 2013 to December 2015.

Security	Mean Return	Standard Deviation
ABX	2.5%	3.1%
BCE	1.1%	2.5%
CCI	0.8%	2.1%

   Assuming that security CCI's returns are normally distributed, the upper and lower limits of the 90 percent confidence interval are *closest to*:

   A.   4.916% and -3.316%
   B.   4.265% and -2.665%
   C.   1.377% and 0.222%

4. Given a sample of 16 data points that is normally distributed with a mean of 18% and a standard deviation of 5%, about 95 percent of the observations will lie between:

   A.   8% and 28%
   B.   15.5% and 20.4%
   C.   8.2% and 27.8%

Solutions:

1.C   About 95% of the numbers in the series will fall within, 100 +/- (2)(10), which is 80 to 120.

2.B   In this question you must use the exact reliability factor 1.645 because the answers are so close:

Upper Limit = 100 + (1.645 x 10) = 116.4
Lower Limit = 100 - (1.645 x 10) = 83.5

3.C   The tricky part about this question is to be able to identify if we are given a sample or population. According to CFA Institute, this is an example where we are given the population, so the upper and lower limits are:

Upper Limit = 0.8 + (1.645 x 2.1) = 4.265
Lower Limit = 0.8 - (1.645 x 2.1) = -2.665

If you chose A, you used 1.96, which is the reliability factor for the 95% confidence interval and if you chose C, you used standard error instead of standard deviation.

4.B   Notice in this question we are given the sample size and that is the reason we are using standard error and not standard deviation to find the upper and lower limits of the confidence interval.

95% CI = Mean $\pm$ (t)(SD/$\sqrt{n}$) = 18 $\pm$ (1.96)(5/$\sqrt{16}$) , thus upper and lower bounds are 15.5% and 20.4%

Student's t-distribution

Small samples of data do not usually follow the normal distribution pattern exactly, so the t-distribution is used to account for the difference. As the sample size is increased, the t-distribution more closely approximates the standard normal distribution.

Bottom Line:

With a small sample size and an unknown population variance, we use the t-distribution to provide the critical values or *"reliability factors"* to construct a confidence interval or to run a hypothesis test.

$$t = \frac{\text{Observed value} - \text{Expected value}}{\sigma / \sqrt{n}}$$

Degrees of freedom = n – 1, where n = sample size

t-Test Learning Example:

Critical Values for Student's t-Distribution		
Degrees of Freedom	Area in Upper Tail	
	5%	2.5%
28	1.701	2.048
29	1.699	2.045
30	1.697	2.042

Assume that monthly returns are normally distributed with a mean of 1% and a sample standard deviation of 4%. The population standard deviation is unknown and the sample size is 29. Given the information in the table above, the lower and upper bounds of the 95% confidence interval are, respectively:

A.   -0.26% and 2.26%
B.   -0.45% and 2.45%
C.   -0.52% and 2.52%

Solution:

C   Since the population standard deviation is unknown we must use the t-distribution. A 95% confidence interval means there is 2.5% in the upper tail and 2.5% in the lower tail, so we are looking at the 2.5% column in the table.

The degrees of freedom (df) = n – 1 = 29 – 1 = 28, so the critical t-value is 2.048.

95% CI = Mean ± (t)(SD/$\sqrt{n}$)

95% CI = 1 ± (2.048)(4/$\sqrt{29}$), thus upper and lower bounds are 2.52% and -0.52%

Data-mining involves searching through data over and over again until you find a pattern that supports your thoughts.

Sample selection bias refers to the practice of choosing specific data that supports your initial thoughts and excluding data that does not.

In the context of the CFA exam, **survivorship bias** refers to the situation where poor performing investment firms, hedge funds or companies go out of business leaving only the good performing ones. This leads to over-estimation of returns and under-estimation of risk.

Look-ahead bias refers to the situation where an analyst uses data known today, in the current period, that would not have been known in the past (ie. during the period of study) to draw conclusions about that past period.

Time period bias refers to the situation where an analyst selects a specific time interval (usually a short time period) that supports his/her initial thoughts. If a longer period of time were to be used, then the conclusions drawn would most likely be different than those of the analyst.

Hypothesis Testing

A **hypothesis** is a statement about one or more populations of data.

Null Hypothesis: It is a hypothesis which states that there is no difference between the procedures and is denoted by H_0. It is always the null hypothesis that is tested, i.e., we want to either accept or reject the null hypothesis because we have information only for the null hypothesis.

Alternative Hypothesis: It is a hypothesis which states that there is a difference between the procedures and is denoted by H_A.

Steps:
State the hypothesis
Select the test statistic
Significance level
Decision rule   nice to know, but not needed for the exam!
Collect data
Make the decision
Act on your decision

One Tail vs. Two Tail Tests (Link to significance level below)

This is a 2 tail test:

H_0: x = 0 vs. H_A: x ≠ 0   The statement must say "equal to, not equal to"

This is a 1 tail test:

H_0: x ≤ 0 vs. H_A: x > 0   To recognize look for statements such as "greater than" or "at least" or "less than"

Decision Rules

t-value:
If the calculated t-value, (Actual – Expected)/Se, is greater than the reliability factor (from the table or rule of thumb), you can reject the null hypothesis (and in the CFA world accept the alternative…yes I just said that!)

If the calculated t-value is less than the reliability factor, you cannot reject the null hypothesis.

p-value:
The p-value represents the lowest level of significance at which you can reject the null hypothesis, so:

- If the p-value is less than the significance level (alpha), you can reject the null hypothesis.
- If the p-value is greater than the significance level, you cannot reject the null hypothesis

Significance Level (Alpha α)
This is your risk level, it is usually 10%, 5% or 1%...this will be given on the exam. You should be able to make the link between significance level and reliability factors:

Two Tail Tests:

Significance Level or Alpha	Reliability Factor or Critical Value
5%	1.96
10%	1.65
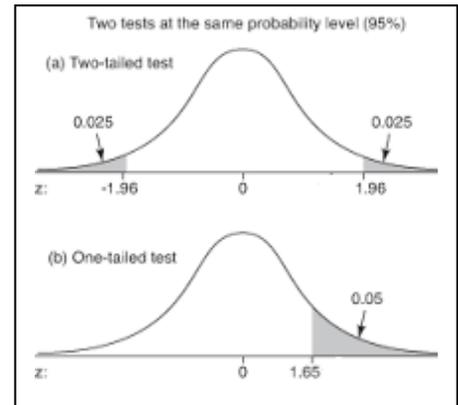
One Tail Tests:

Significance Level or Alpha	Reliability Factor or Critical Value
5%	1.65

Two tests at the same probability level (95%)
(a) Two-tailed test
0.025   0.025
z: -1.96   0   1.96

(b) One-tailed test
0.05
z: 0   1.65

Type I vs. Type II Errors:

The probability of rejecting the null hypothesis when it is in fact true is referred to as a *Type I error*. The probability of a Type I error is equal to the significance level.

A *Type II error* occurs when you do not reject the null hypothesis when it is in fact false.

The *power of a test* is the probability of correctly rejecting the null hypothesis and it is equal to one minus the probability of a Type II error.

Bottom Line:
It is unlikely that you will have to perform a complete hypothesis test on the exam, you will most likely need to interpret the results given all the information.

You should be able to link the significance of the one tail vs. two tail test to the concept of selecting the correct reliability factor!

Statistical Result vs. Economically Meaningful Result
Statistical significance refers to the interpretation of our hypothesis test – it's just the math – do we accept or reject. Economic significance refers to the relative size (or value) of our conclusion…does it actually make a difference in the real world or is it so small that it is inconsequential?

Parametric vs. Nonparametric Tests
Parametric tests are based on many restrictive assumptions, while nonparametric relax these assumptions or simply do not make any assumptions at all.

Nonparametric tests are used when data do not have assumptions about distribution, when data are given in ranks, or when the hypothesis we are addressing does not concern a parameter.

Hypothesis (t-test) Learning Example:

Assume normal distribution of returns and a calculated t-value of 1.75 apply in all cases.

1. Given the following hypothesis, at a significance level (alpha) of 5%, what would an analyst be *most likely* to conclude?

$$H_0: x = 0 \text{ vs. } H_A: x \neq 0$$

    A.   Reject the null
    B.   Fail to reject the null
    C.   Reject the alternative

2. Given the following hypothesis, at a significance level (alpha) of 5%, what would an analyst be *most likely* to conclude?

$$H_0: x \leq 0 \text{ vs. } H_A: x > 0$$

    A.   Reject the null
    B.   Fail to reject the null
    C.   Reject the alternative

3. Given the following hypothesis, at a significance level (alpha) of 10% what would an analyst be *most likely* to conclude?

$$Ho: x = 0 \qquad Ha: x \neq 0$$

    A.   Reject the null
    B.   Fail to reject the null
    C.   Reject the alternative

Solutions:

1.B   This is a two tail so the critical value (reliability factor) at a 5% significance level is 1.96. Since the calculated t-value is less than the critical value, we cannot reject the null and conclude x is equal to zero.

2.A   This is a one tail so the critical value (reliability factor) at a 5% significance level is 1.645. Since the calculated t-value is greater than the critical value, we reject the null and conclude x is greater than zero.

3.A   This is a two tail so the critical value (reliability factor) at a 10% significance level is 1.645. Since the calculated t-value is greater than the critical value, we reject the null and conclude x is not equal to zero.

Which Test Statistic Should You Use?

The LOS asks us to know which test statistic to use for various test, so here are your simple rules:

- **_t-stat_** is used to test the mean and it is used when the population variance is unknown.
- **_z-stat_** is used to test the mean and it is used when the population variance is known.
- **_F-stat_** is used to test the variance of two populations (or samples)
- **_Chi Squared_** is used to test the variance of one population (or sample) vs. a hypothesized value.

We already know how to use a t-stat and z-stat to test the mean, now we need to focus on the F-stat and Chi Squared:

The F-stat is used to test the equivalence of the variance of two populations or samples:

$$H_0: S_1^2 = S_2^2 \qquad H_A: S_1^2 \neq S_2^2$$

$$\text{Calculated F-stat} = \frac{S^2}{S^2} \quad \Leftarrow \text{put the larger variance on top!}$$

Decision Rule:

- If the calculated F-stat is greater than the critical value (which will be given), reject H_0 and conclude the variances are not the same.
- If the calculated F-stat is less than the critical value, do not reject H_0 and conclude the variances are the same.

The Chi Squared is used to test the variance of one population or sample vs. a hypothesized value:

$$H_0: S^2 = S_0^2 \qquad H_A: S^2 \neq S_0^2 \quad \Leftarrow \text{be aware this could be} <, >, \leq \text{or} \geq$$

$$\text{Calculated } X^2 \text{ stat} = \frac{(n-1)S^2}{S_0^2}$$

n = number of observations
S^2 = variance of the population (or sample)
S_0^2 = hypothesized variance

Decision Rule:

- If the calculated X^2 stat is greater than the critical value (which will be given), reject H_0.
- If the calculated X^2 stat is less than the critical value, do not reject H_0.

F Test Learning Example:

Rita Heatherton, is analyzing the variance of returns on the TSX index. She gathers data for a period prior to October 2014 and data for a period after October 2014. The information is summarized below:

Time Period	n	Average Daily Return	Variance of Returns
Before October 2014	120	1.55	27.9
After October 2014	120	1.43	13.7

Determine if the variance of returns has changed after the market activity in October 2014 assuming a 5% significance level and a critical value of 1.63.

Solution:

State our hypothesis: $H_0: S^2_{before} = S^2_{after}$ $\qquad$ $H_A: S^2_{before} \neq S^2_{after}$

Calculate the F-stat: $F = \dfrac{27.9}{13.7} = 2.03$

Make your decision: Since the calculated F is greater than the critical value, we reject H_0 and conclude that the variance has changed.

Chi Squared Learning Example:

Over a 30 year period, the standard deviation of annual returns on a portfolio you are analyzing was 18% a year. You want to see whether this record is sufficient evidence to support the conclusion that the portfolio's underlying variance of returns was less than the variance of the benchmark, which was 625. Assuming a critical value of 10.117 applies, you decide to test this.

Solution:

Sine we are testing the variance vs. the hypothesized value of 625, we must use the Chi Square test.

State your hypothesis: $\qquad$ $H_0: \sigma^2 \geq 625$ $\qquad$ $H_A: \sigma^2 < 625$

Calculate the Chi Squared: $X^2 = \dfrac{(n-1)S^2}{\sigma^2} = \dfrac{(30-1)18^2}{625} = 15.03$

Make your decision: $\qquad$ Since the calculated X^2 is greater than the critical value, we reject the null and conclude the portfolio variance is less than 625.

Basic Hypothesis Testing Review Questions:

1. If the *p*-value for a hypothesis test is 0.07 and the chosen level of significance, alpha = 0.05, then an Analyst conducting a hypothesis test would be *most likely* to conclude:

    A. accept the alternative hypothesis.
    B. there is not enough information given to make a decision.
    C. do not reject the null hypothesis.

2. For the following hypothesis:

    $$H_0: a = b; \qquad H_A: a \neq b;$$

    The t-statistic is 2.1. If the reliability factor is 1.66, then an Analyst would be *most likely* to conclude:

    A. reject null hypothesis because a = b
    B. reject alternative hypothesis because a = b
    C. accept the alternative hypothesis because a is significantly different from b

3. You conduct a two-tailed t-test to determine whether a sample mean involving 100 observations differs from a theoretical mean of zero. The computed t-statistic is 2.90. Using a 5 percent significance level, which one of the following conclusions is the most appropriate to reach?

    A. Reject the null hypothesis and accept the alternative hypothesis that the sample mean is significantly different from zero.
    B. Reject the alternative hypothesis that the sample mean is significantly different from zero.
    C. Refrain from drawing a conclusion because the number of observations is insufficient.

Solutions:

1.C The p-value represents the lowest level of significance (a.k.a alpha) at which you can reject the null hypothesis. Rule: if p-value is < significance level, then reject null hypothesis.

2.C Since the calculated t-stat is greater than the critical t-stat, we can reject the null hypothesis. (accepting the alternative is very poor language and technically incorrect, but that is the CFA for you!)

3.A With a 5% significance level, the critical value for a two tailed test from the t-table is approximately 1.98. The calculated t stat is greater than the critical value, therefore we can reject the null hypothesis. (accepting the alternative is very poor language and technically incorrect, but that is the CFA for you!)

Warning!
If you are maxed out on formulas, I suggest that you skip this next section and go on to the Technical Analysis reading. The logic is you have a good chance of forgetting what you are about to read and you have a high probability of making a mistake on the exam anyway. The other, and more significant reason is, you might confuse the foundation hypothesis testing concepts you just learned with these new ones and end up getting everything wrong!

Hypothesis Test Concerning the Equality of Means
Here we want to see if the mean of two samples is different (or alternatively the same, higher or lower).

In this context, typical null and alternative hypothesis are:

1) $H_0: \mu_1 - \mu_2 = 0$ <= says no difference between μ_1 and μ_2
     $H_a: \mu_1 - \mu_2 \neq 0$ <= says μ_1 is different than μ_2

2) $H_0: \mu_1 - \mu_2 \leq 0$
     $H_a: \mu_1 - \mu_2 > 0$ <= says μ_1 is greater than μ_2

3) $H_0: \mu_1 - \mu_2 \geq 0$
     $H_a: \mu_1 - \mu_2 < 0$ <= says μ_1 is less than μ_2

If the samples are deemed to be independent use the following two tests:

1) If the sample variances are assumed to be equal:
If the question tells you that the population variances are assumed to be equal, then you need to calculate the t-stat to make your decision about the null hypothesis. Look for this language or variances that are very close to each other. You will need to calculate the following three items:

1) Pooled Estimator of Common Variance

$$S_p{}^2 = \frac{(n_1 - 1)S_1{}^2 + (n_2 - 1)S_2{}^2}{n_1 + n_2 - 2}$$

Where: n_1 and n_2 are the sample sizes
$S_1{}^2$ and $S_2{}^2$ are the sample variances

2) t-statistic

$$t = \frac{(x_1 - x_2) - (\mu_1 - \mu_2)}{\left(\dfrac{S_p{}^2}{n_1} + \dfrac{S_p{}^2}{n_2} \right)^{1/2}}$$

Where: x_1 and x_2 are the sample means
n_1 and n_2 are the sample sizes
$S_p{}^2$ pooled estimator of common variance

3) Degrees of Freedom
You will need the df to look up the "rejection points" on the t-table.

$$df = n_1 + n_2 - 2$$

2) If the sample variances are assumed not to be equal:

If the question tells you that the population variances are not equal, then you will still need to calculate the t-stat to make your decision about the null hypothesis. It may be referred to as the ***"approximate t-test"***. Look for this language or <u>variances that are very far apart from each other</u>. You will need to calculate the following two items:

1) t-statistic

$$t = \frac{(x_1 - x_2) - (\mu_1 - \mu_2)}{\left(\dfrac{S_1^2}{n_1} + \dfrac{S_2^2}{n_2} \right)^{1/2}}$$

Where: x_1 and x_2 are the sample means
n_1 and n_2 are the sample sizes
S_1^2 and S_2^2 are the sample variances

2) "modified" estimate of the degrees of freedom

$$df = \frac{\left(\dfrac{S_1^2}{n_1} + \dfrac{S_2^2}{n_2} \right)^2}{\dfrac{\left(\dfrac{S_1^2}{n_1} \right)^2}{n_1} + \dfrac{\left(\dfrac{S_2^2}{n_2} \right)^2}{n_2}}$$

Exam tip:

You may not need to do the complex calculation for the degrees of freedom or even find the rejection points on a t-table…use some common sense…if the t-stat is a large number, say 3 or higher…you can safely assume that the null hypothesis can be rejected.

Learning Example:

An analyst gathers the following data from two independent samples drawn from normally distributed populations:

	Size	Sample mean	Sample standard deviation
Sample 1	40	350	33
Sample 2	22	210	35

Assuming that the variances of the underlying populations are equal, answer the following questions:

a)   Write the null and alternative hypothesis testing no difference between μ_1 and μ_2:

b)   **Calculate** the pooled estimate of the sample variance:

c)   **Calculate** the degrees of freedom:

d)   **Calculate** the t-statistic:

e)   Assuming rejection values of –2 and +2, should the null hypothesis be rejected or not rejected?

f)   What adjustments to the process would the analyst have to make if the sample variance are not assumed to be equal?

Solutions:

a) $H_0: \mu_1 - \mu_2 = 0$
$Ha: \mu_1 - \mu_2 \neq 0$

In this case the analyst is testing whether or not the sample means are different and if the difference is statistically significant.

b) $S_p^2 = \dfrac{(n_1 - 1)S_1^2 + (n_2 - 1)S_2^2}{n_1 + n_2 - 2} = \dfrac{(40 - 1)33^2 + (22 - 1)35^2}{40 + 22 - 2} = 1{,}136.6$

c) $df = n_1 + n_2 - 2 = 40 + 22 - 2 = 60$

d) $t = \dfrac{(x_1 - x_2) - (\mu_1 - \mu_2)}{\left(\dfrac{S_p^2}{n_1} + \dfrac{S_p^2}{n_2}\right)^{1/2}} = \dfrac{(350 - 210) - 0}{\left(\dfrac{1136.6}{40} + \dfrac{1136.6}{22}\right)^{1/2}} = 15.64$

e) The t-value of 15.64 is significant at the 5% level, so we can reject the null hypothesis and conclude that the sample means are statistically different.

f) There would not be a pooled estimator of common variance, instead we would use the variances for each sample in the t-statistic calculation, as follows:

$t = \dfrac{(x_1 - x_2) - (\mu_1 - \mu_2)}{\left(\dfrac{S_1^2}{n_1} + \dfrac{S_2^2}{n_2}\right)^{1/2}} = \dfrac{(350 - 210) - 0}{\left(\dfrac{33^2}{40} + \dfrac{35^2}{22}\right)^{1/2}} = 15.38$

And the analyst would need to calculate a "modified" estimate of the degrees of freedom, as follows:

$df = \dfrac{\left(\dfrac{S_1^2}{n_1} + \dfrac{S_2^2}{n_2}\right)^2}{\dfrac{\left(\dfrac{S_1^2}{n_1}\right)^2}{n_1} + \dfrac{\left(\dfrac{S_2^2}{n_2}\right)^2}{n_2}} = \dfrac{\left(\dfrac{33^2}{40} + \dfrac{35^2}{22}\right)^2}{\dfrac{\left(\dfrac{33^2}{40}\right)^2}{40} + \dfrac{\left(\dfrac{35^2}{22}\right)^2}{22}} = 43.1$

Exam tip:
If the calculations are blowing you away…focus on understanding how to answer questions like part e). I think may be given all the numbers and just have to put them together to draw the correct conclusion.

Hypothesis Test Concerning the Mean Difference

You will use the test described here if the question asks for a *"paired comparisons test"* or says you are dealing with *"paired observations"*.

Paired observations are observations that are dependent because they have something in common and a paired comparisons test is a statistical test for differences in dependent items.

By calculating a standard error based on differences, the *t*-test presented below takes account of correlation between the observations.

Notation to Know:

x_a and x_b = two random variables

$d_i = x_{ai} - x_{bi}$ = the difference between two paired observations

μ_d = population mean difference

μ_{d0} = hypothesized value for the population mean difference (most often equal to zero)

If the samples are deemed to be dependent (not independent) use the following test:

You will do a t-test, remember the basics, for a t-test you need the observed value, the expected value and the standard error. In this case, you will need to calculate the following five items:

1) Sample Mean Difference

$$\bar{d} = \frac{\sum d_i}{n}$$

Where: n is the number of pairs of observations

d_i is the difference between tow paired observations

2) Sample Standard Deviation (Variance)

$$S_d = \sqrt{S_d^2} = \sqrt{\frac{\sum (d_i - \bar{d})^2}{n-1}}$$

3) Standard Error of the Mean Difference

$$S_{\bar{d}} = \frac{S_d}{\sqrt{n}}$$

4) t-statistic

$$t = \frac{\bar{d} - \mu_{d0}}{S_{\bar{d}}}$$

5) Degrees of Freedom

$$df = n - 1$$

Learning Example:
An analyst summarizes the return data for a Korean mutual fund and Canadian mutual fund, based on the last 24 months. The results are presented in Table 1.

Table 1

	Mean Return	Standard Deviation
Korean Fund	15%	25%
Canadian fund	13%	20%

Excerpt from t-table

Degrees of Freedom	p = 0.05	p = 0.025
22	1.717	2.074
23	1.714	2.069
24	1.711	2.064

Based on the data above answer the following questions:

a) In your own words briefly describe what the hypothesis H_0: $\mu_d = 0$ H_a: $\mu_d \neq 0$ means:

b) **Calculate** the sample mean difference:

c) **Calculate** the standard error of the mean difference:

d) Regarding the hypothesis in part a, what would the analyst be *most likely* to conclude assuming,

   i) a 5% significance level

   ii) a 10% significance level

ExamSuccess.ca 2016
Exam Preparation & Investment Sales Training   35
Brian Y. Gordon, CFA, CFP, CIM, MBA, FCSI
Email: info@examsuccess.ca

Solutions:

a)   The analyst is looking to test whether or not the mean difference between the Korean mutual fund and the Canadian mutual fund is zero.

   μ_d is the underlying mean difference between the two mutual funds. The null (H_0) says this difference is zero, that is the mean returns are the same, and the alternative (H_a) says the mean returns are different (the difference is not equal to zero).

b)   We are not given actual data points to work with, (ie x_a and x_b) to calculate the sample mean difference, so we just use the data in Table 1 and calculate this as the difference in the mean returns:

   $\overline{d} = 15 - 13 = 2$

c)   Again we are not give actual observations, so we just use the data in Table 1 and calculate this as the standard error of the mean difference:

   $S_{\overline{d}} = \dfrac{25 - 20}{\sqrt{24}} = 1.0206$

d)   We need calculate our t-statistic and to use the data in the t-table to find our rejection points.

   $t = \dfrac{\overline{d} - \mu_{d0}}{S_{\overline{d}}} = \dfrac{2 - 0}{1.0206} = 1.959$ $df = n - 1 = 24 - 1 = 23$

   Now we need to interpret the "p-values" from the t-table:

   $p = 0.05$ means there is 5% of the area in each tail, so 10% in both tails. This is the 10% significance level.

   $p = 0.025$ means there is 2.5% of the area in each tail, so 5% in both tails. This is the 5% significance level.

   i) Because $1.959 < 2.069$ we cannot reject H_0 and the analyst would conclude that the difference in mean returns for the two mutual funds is not statistically significant.

   i) Because $1.959 > 1.714$ we can reject H_0 and the analyst would conclude that the difference in mean returns for the two mutual funds is statistically significant.

Exam tip:
If the calculations are blowing you away…focus on understanding how to answer questions like part d). I think may be given all the numbers and just have to put them together to draw the correct conclusion.

Technical Analysis

Basic Principles

Technical analysts focus on interpreting *market data* to determine price targets for securities. Technicians believe that supply and demand determine security prices and that changes in supply and demand cause changes in security prices. Technicians believe that prices can be projected with charts and other tools and that trends and prices reflect the irrational behavior of investors.

Types of Charts

Charts are a graphic display of price and volume that provide information about past price behavior and infer likely future behavior.

Line Chart –a plot of data points with one data point per time interval, provides a broad overview of investor sentiment.

Bar Chart – four bits of data per entry, high, low, open and closing prices, provides a sense of the day's trading activity.

Candlestick Chart – four bits of data per entry, high, low, open and closing prices, the body of the candle is shaded if the opening price was higher than the closing price, allows for quicker analysis.

Point and Figure Chart – focus on entry and exit levels, horizontal axis represents number of changes in price, vertical axis represents discrete increment of price, X for up movement, O for down movement, box size and reversal size must be known.

Scale – can be linear or logarithmic

Volume – can be used to confirm signals

Time intervals – daily, weekly, monthly, hourly, etc.

Relative strength analysis – used to compare the performance of an asset to the benchmark, ratio of prices to the benchmark is used to show under- outperformance or neutral performance

Trend lines – market participants act together over periods of time, major breakdowns/minor breakdowns, link to supply and demand imbalances.

Support lines – low price range in which buying activity is sufficient to stop the decline in price.

Resistance lines – high price range in which selling activity is sufficient to stop the increase in price.

Change in polarity principle – once a support level is breached it becomes a resistance level and once a resistance level is breached it becomes a support level.

Common Chart Patterns

Formations in price charts that create some type of recognizable shape, represent human trading behavior, lead to subsequent price movements "without a clear trend in place prior to the pattern, the pattern has no predictive value".

Reversal Patterns – signals the end of a trend, eg. "head and shoulders", "double tops" and "double bottoms", once the head and shoulders pattern has formed, technicians expect the price to decline through the neckline.

Head and Shoulders Price Target:

Normal: Price Target = Neck – (Head – Neck)
Inverted: Price Target = Neck + (Head – Neck)

---

1. An analyst identifies the following information regarding a "head and shoulders" pattern:
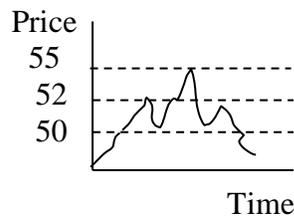
    Based on this information, the price target is closest to:

    A. $53
    B. $47
    C. $45

Continuation Patterns – used to predict the resumption of a market trend that was in place prior to the pattern forming, eg. "triangles and rectangles"

Common Technical Analysis Indicators

Price based – includes moving averages and Bollinger bands

Bollinger bands consist of a moving average surrounded by two lines:

1) the ***upper limit*** representing the moving average plus a set number of standard deviations from the average price.
2) the ***lower limit*** representing the moving average minus a set number of standard deviations from the average price.

Momentum based oscillators – signal over bought and over sold conditions, includes relative strength index and the MACD.

Sentiment indicators – include opinion polls, put/call ratio, CBOE volatility index (VIX) and short interest ratio.

Flow of funds indicators – includes margin debt, mutual fund cash positions and the "TRIN".

TRIN = [Advancing/Declining] ÷ [Advancing Volume/Declining Volume]

TRIN < 1 is a bullish signal, meaning more volume is rising than falling

Cycles – provide evidence of recurring patterns in the capital markets.

Elliott Wave Theory
Based on the concept that the market moves in waves, with "impulse waves" followed by a "corrective waves".

Fibonacci Numbers – common percentage retracements are Fibonacci ratios, may alert the technician to possible changes in the trend.

Intermarket Analysis – can be applied to technical analysis to identify market trends and connections between different categories of securities, eg. Bond yields vs. equity market movements, technicians can then identify asset classes that look relatively attractive for investment.

2. A low price range in which buying activity is sufficient to stop a price decline is *best* described as:

    A. support.
    B. resistance.
    C. change in polarity.

3. A technical analyst has detected a price chart pattern with three segments. The left segment shows a decline followed by a reversal to the starting price level. The middle segment shows a more pronounced decline than in the first segment and again a reversal to near the starting price level. The third segment is roughly a mirror image of the first segment. This chart pattern is *most* accurately described as:

    A. a triple bottom.
    B. a head and shoulders.
    C. an inverse head and shoulders.

Solutions:

1.C   Price Target = Neck – (Head – Neck) = 50 – (55 – 50) = $45

2.A   Support is defined as a low price range in which buying activity is sufficient to stop the decline in price.

3.C   An inverse head and shoulders pattern consists of a left segment that shows a decline followed by a reversal to the starting price level, a middle segment that shows a more pronounced decline than in the first segment and again a reversal to near the starting price level, and a third segment that is roughly a mirror image of the first segment.