Chapter 5
Hypothesis Test of the Mean

Section 5.4
Power Analysis

School of Business Administration
Portland State University

5.4a
Managerial vs Statistical Decisions

Managerial Decision versus Statistical Decision

Purpose of statistical analysis is to guide managerial decisions

- **Managerial decision**: Implement, or not, a specific action
- **Key Concept**: A technical decision regarding statistical inference provides information about reality that guides a subsequent managerial decision
- Statistical inference, confidence interval or hypothesis test, often focuses on a specific value in terms of the null hypothesis
- **Statistical decision** regarding a mean: Does a difference exist from the hypothesized value?
- The statistical decision is to either reject or not reject this reference value of interest, $\mu = \mu_0$
- Each outcome of the statistical decision, reject or not, may lead to a different managerial decision

Statistical Decision Guides the Managerial Decision

The statistical analysis informs the decision maker

- Consider the goal to assess a potential improved process with a trial version, implemented on a small scale, which management is considering to implement company wide
- To assess, compare a sample mean from the new process against the current mean, which is the null hypothesis
- One outcome is that the statistical analysis detects a difference from the null in the desired direction
  - Statistical decision → Reject the null hypothesis, so ...
  - typical Managerial decision → Implement the change
- Another outcome is that the statistical analysis does *not* detect a difference from the null
  - Statistical decision → Do not reject the null, so ...
  - typical Managerial decision → Retain the current process

Slippage Between Statistical and Managerial Decisions

Issue: These two decisions are not so tightly coupled

- However, this relation between statistical and managerial decisions must be qualified because the statistical and managerial decisions are distinct concepts
- **Key Concept**: The managerial decision does *not* necessarily follow directly from the statistical decision
- The statistical decision guides the managerial decision, but the statistical decision does not rigidly imply what action the manager should take
- In terms of deciding to retain or change to a new process
  - Rejection of the null hypothesis does *not necessarily* imply that the new process should be implemented
  - Failure to reject the null hypothesis does *not necessarily* imply that the current process should be retained
- This potential slippage is explored next

5.4b
Type I and Type II Errors

Binary Decisions

Definition and Examples of Binary Decisions

Choose one alternative or the other

- **Binary decision**: Choose between two alternatives
- Professional and personal life presents a series of binary decisions, from the trivial to the profound
  - Cross the road now, do not cross now
  - Hire the applicant, do not hire the applicant
  - Invest in increased manufacturing capacity, do not invest
  - Marry the person, do not marry
  - Reject the null hypothesis, do not reject

Consequences of Binary Decisions

Four possible outcomes from one binary decision

- Each binary decision provides two ways to be right (opportunity) and two ways to be wrong (error)
- Each outcome implies different consequences

☺ Correct: Cross road now and do not get run over

☺ Correct: Do not cross road now and watch the car go by

☹ Error: Do not cross now and no car goes by – consequence is to lose a little bit of time

☹ Error: Cross now and get run over – consequence is death or severe injury

- These two types of errors have names, rather unimaginatively called Type I and Type II errors

Four Possible Outcomes for an Hypothesis Test

Four possible outcomes from one binary decision

- The manager's decision regards an unknown reality, which at some later time reveals itself
- **Statistical Decision Rule** for analysis of the mean: If m is in the rejection region, decide H_0 is false, otherwise the data is consistent with H_0

☺ Correct: Decide the null is true, and it really is

☺ Correct: Decide a false null is false, a real difference detected

☹ Type I Error, False Positive: An unusually deviant event leads to decide the null is false, but it really is true

☹ Type II Error, False Negative: Decide the null is true, but fail to detect a difference that actually exists

Four Possible Outcomes for an Hypothesis Test

Reality \ Decision	Do Not Reject H_0	Reject H_0
H_0 is true	CORRECT: "accept" a true null	TYPE I ERROR: reject a true null
H_0 is false	TYPE II ERROR: accept a false null	CORRECT: reject a false null

Type I Error

Null Hypothesis is True: Probability of a Type I Error

Consider the null hypothesis to be true

- Distribution of sample mean, *m*, centered over hypothesized value, μ_0, assumed true *for purposes of the test*
- Presume the hypothesized value, μ_0, *really is true*, not just for purposes of the test
- IF μ_0 is true, then . . .
  - The applicable error is a Type I error, rejecting a true μ_0
  - A Type I Error occurs when *m* falls in the rejection region
  - What is the (conditional) probability that *m* falls in the rejection region given that the null hypothesis is true?

Null Hypothesis is True: Probability of a Type I Error

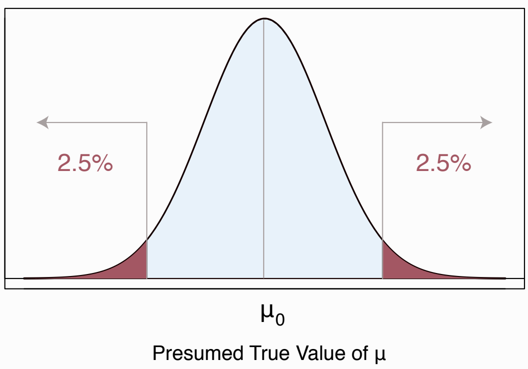

- Probability of a Type I Error for any one sample is not computed, but instead set at the probability of the complete rejection region, α, usually $\alpha = .05$ as shown here

Type II Error

More of the What-If Game

Assume $\mu = \mu_0$, but what if this assumption is false?

- To understand and estimate the value of the unknown value of μ, probe the consequences of assuming different values of μ
- The basis of analyzing Type II error is that the assumed null hypothesis, $\mu = \mu_0$ is wrong
- **Alternative value** of the mean: To presume $\mu \neq \mu_0$ means that some alternative value of μ, μ_{ALT}, is true, $\mu = \mu_{ALT}$
- Type II Error Analysis:
  - First, set up the hypothesis test by assuming
    WHAT IF $\mu = \mu_0$?
  - Second, add a second assumption beyond the first,
    WHAT IF we first assume $\mu = \mu_0$,
    but then ask actually WHAT IF, instead, $\mu = \mu_{ALT}$?
- Will we properly reject our first assumption of $\mu = \mu_0$ if, instead, the truth is that $\mu = \mu_{ALT}$?

Concepts and Notation

Many presumed values of μ when μ_0 is false

- Decision rule: Fail to reject μ_0 if m is reasonably close to μ_0, as assessed by the corresponding hypothesis test
- The error here, a Type II error, is to "accept" $\mu = \mu_0$ when some alternative value is actually the true mean, $\mu = \mu_{ALT}$
- The probability of a Type II error is defined in terms of one corresponding specific value of μ_{ALT} relative to the specified value of μ_0
- There are many probabilities of a Type II error, each probability calculated relative to a specific value of the presumed true mean, $\mu = \mu_{ALT}$
- There are many possibilities for μ_{ALT}, which each can be systematically investigated, as shown next, beginning with a specific example where $\mu_{ALT} = 45$

Probability of a Type II Error

Presume true mean is actually $\mu_{ALT} = 45$

- Hypothetical distribution of m is centered over the presumed true mean of 45, with most values between 40 and 50
- For example, if a sample were taken, values might be obtained such as m=48.8, m=47.1 or m=41.5

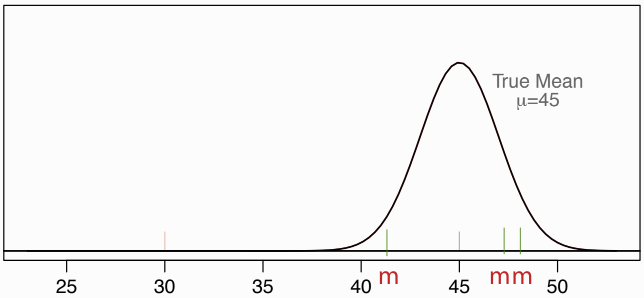

Probability of a Type II Error

Consider an extremely deviant sample mean, $m = 31.0$

- A value of $m = 31.0$ would be highly unlikely to ever occur for this distribution with $\mu_{ALT} = 45$
- Although theoretically possible, a value of m this far below the mean of 45 would not happen in practice

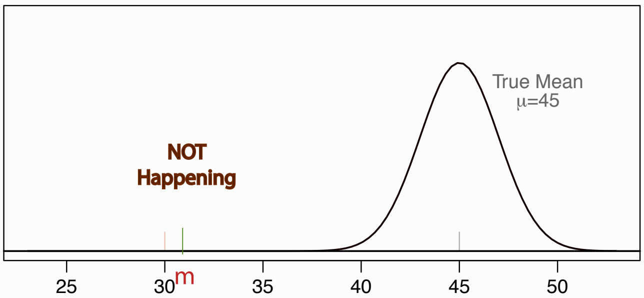

Probability of a Type II Error

Hypothesized mean is $\mu_0 = 30$

- Now consider a hypothesis test with μ_0 set at 30, which, in this scenario is *false* as the true mean is presumed to be 45
- Around the value of the hypothesized mean of 30 is the "acceptance" region of the hypothesis test

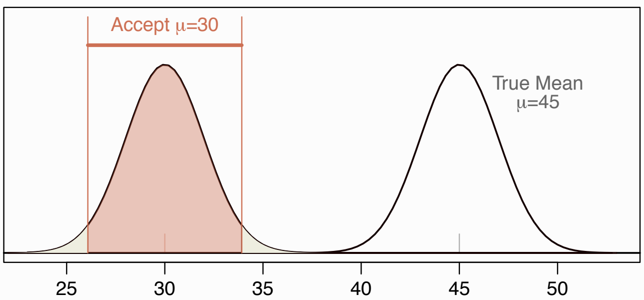

Probability of a Type II Error

Presume true mean is actually $\mu_{ALT} = 45$

- $\mu = 30$ is false in this scenario, so a Type II error occurs when *m* falls in the "acceptance" region, "accepting" the false null
- But *m* will virtually never fall in the "acceptance" region
- The false $\mu_0 = 30$ would almost certainly be correctly rejected for any one sample because the sample mean, *m*, would be much higher than the hypothesized mean, μ_0
- The value of the one observed *m* is not known until the data are collected, but whatever value it assumes, a difference from the null value will be properly detected in this situation

Basis of a Type II Error

Focus on a specific alternative to μ

- The hypothesis test "accepts" the null hypothesis when *m* falls in the "acceptance" region, close to μ_0
- If μ_{ALT} is the true mean, then the sample mean *m* is sampled from the distribution centered about μ_{ALT}
- **Key Concept**: If μ_{ALT} is the true mean, *m* can still fall within the "acceptance" region defined by the null hypothesis
- A Type II Error occurs when the null is false, but *m* is in the "acceptance" region:

  The difference from the null value μ_0 to the presumed reality μ_{ALT} exists, but the hypothesis test does not detect the difference
- What is the probability that a value of *m* is in the "acceptance" region *given* the value of μ_{ALT}?

5.4c
The Concept of Power

Power Defined and Computed

Probability of the Correct Choice: Power

Focus on a specific alternative to the hypothesized value of μ

- Typically prefer to state the probability of making the *correct* choice, of rejecting the null hypothesis when it is false
- **Power** of an Hypothesis Test: Probability of *correctly* detecting a real difference from a false null hypothesis
- When the null is false, so that μ is different from μ_0,
  - Either make the right choice, reject the false null, or
  - Make the wrong choice, do not reject the false null
- **Notation**: Probability of a Type II Error is β
- The probabilities of making the correct or incorrect choice, here for $\mu \neq \mu_0$, sum to 1,

  Power $+ \beta = 1$

  or

  Power $= 1 - \beta$

Calculate Power and Probability of a Type II Error

Use R as an electronic table

- To use R to calculate Power is to use R as an electronic table of probabilities – no data file, no data analysis
- Provide a specific sample size, *n*, and standard deviation, *s*, such as from an initial data analysis
- To calculate Power and β, the actual values of the alternative and hypothesized values of μ per se are not of interest
- Rather, the focus is on the *difference* or "delta" between the hypothesized value of μ and the specific proposed alternative

  delta $= \mu_{ALT} - \mu_0$
- For the same *n* and *s*, the power of detecting a real μ of 35 from a hypothesized value of 30 is exactly the same as detecting a real μ of 105 from a hypothesized value of 100

R Input: Power for a one-sample t-test

R function `power.t.test` calculates power

- Suppose for a specific data analysis:

  $n = 25$, $s = 10$, $\mu_0 = 30$ and $\mu_{ALT} = 35$
- R input:

```
> power.t.test(n=25, sd=10, delta=5,
    type="one.sample")
```

- Required values:
  - `n`: sample size
  - `sd`: standard deviation of data
  - `delta`: change from μ_0 to the alternative value of μ
- Defaults:
  - `alternative="two.sided"`, i.e., two-tailed test
  - `type="two.sample"`, i.e., for a mean difference

R Output: Power for a one-sample t-test

One-sample t-test power analysis

```
> power.t.test(n=25, sd=10, delta=5,
    type="one.sample")
```

```
     One-sample t test power calculation

              n = 25
          delta = 5
             sd = 10
      sig.level = 0.05
          power = 0.6697014
    alternative = two.sided
```

- For a sample of size 25 with a standard deviation of 10, and a hypothesized value of 30, the probability of correctly detecting a difference when $\mu = 35$ is 0.67, and so $\beta = 1 - 0.67 = 0.33$

Probability Distribution that Underlies a Type II Error

Technical Note

- The hypothesis test of the mean follows from the distribution of t_m based on the sample mean, assuming the null hypothesized value, μ_0, is true
- The probability of a Type II Error follows from a non-central t-distribution, the distribution of the t-statistic, t_m, when μ_0 is assumed true but is actually false
- For pedagogical simplicity, the following illustrations are based on the sample mean, m, and the normal distribution of m instead of t_m and the non-central t-distribution
- The basic concepts remain the same, and numerical computations only suffer when the sample size, n, is small
- Fortunately, Type II Error calculations with R are based the correct non-central t-distribution, so actual applications of these probabilities from R are accurate for all values of n

Power: An Example Continued

Illustrate Probability of Type II Error with Graphics

Consider a specific example

- Suppose the hypothesized value is $\mu_0 = 30$
- Collecting a sample of size $n = 25$ yielded a standard deviation of $s = 10$, for a standard error of the mean of $s_m = 2$
- Suppose the sample mean m was larger than 30, but close enough to $\mu_0 = 30$ that the null hypothesis was *not* rejected
- Although no difference from the null was detected, the analyst wonders if the test might have failed to detect a real difference with the true value of μ actually larger than 30
- What is the probability of detecting this difference if it actually exists?
- Power analysis provides the answer

Probability of a Type II Error

Null hypothesis is $\mu = 30$

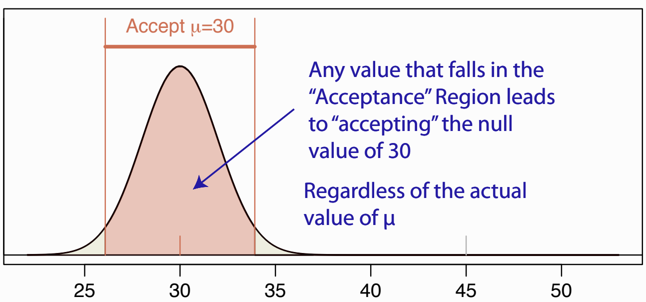

- Here the null hypothesized value, μ_0, is 30
- The hypothesis test sets up an "acceptance" region around 30
- Whatever the actual value of μ, a m that falls in the "acceptance" region leads to "accepting" the null

Probability of a Type II Error

Probability depends on the true value of μ

- So the computation of the probability of a Type II Error, β, depends on the true value of μ
- Presume the value of the null hypothesis, μ_0, is false, so what is the alternative value of μ, μ_{ALT} to use for the probability calculation?
- We do not know the true value of μ, which is the whole point of statistical inference, to estimate μ
- In absence of the actual value of μ, explore what happens for a range of *potential* different alternative values of μ, from values far away from μ_0 to values that are close to μ_0
- What are Power and β . . . IF the actual $\mu = \mu_{ALT}$ was equal to a value just one unit above or below μ_0, or two units, or . . .

Probability of a Type II Error

Hypothesize 30; Presume true mean is actually $\mu_{ALT} = 40$

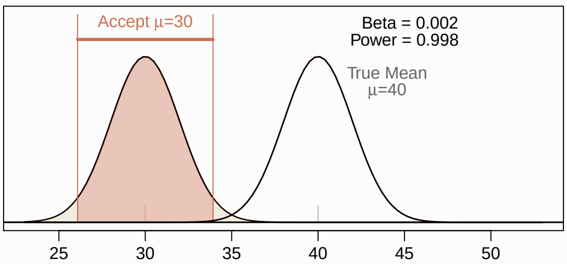


```
> power.t.test(n=25, sd=10, delta=10, type="one.sample")
```

- For $\mu_{ALT} = 40$, *m* likely varies somewhere between 34 and 46
- Rarely a *m* may just cross the line into the "acceptance" region, resulting in a Type II Error

Probability of a Type II Error

Hypothesize 30; Presume true mean is actually $\mu_{ALT} = 38$

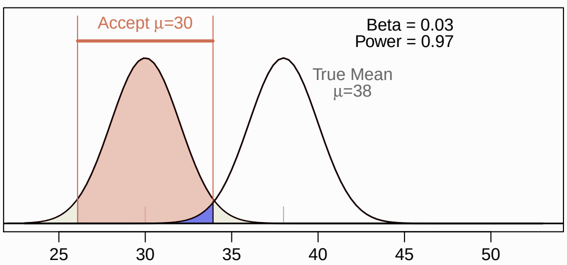


```
> power.t.test(n=25, sd=10, delta=8, type="one.sample")
```

- For $\mu_{ALT} = 38$, *m* likely varies somewhere between 32 and 44
- Here *m* occasionally lands in the "acceptance" region, as a Type II Error has a probability of only 0.03 (blue area)

Probability of a Type II Error

Hypothesize 30; Presume true mean is actually $\mu_{ALT} = 35$

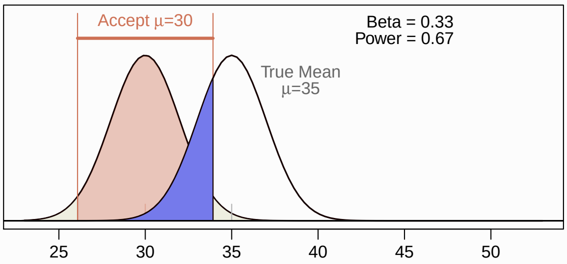


```
> power.t.test(n=25, sd=10, delta=5, type="one.sample")
```

- For $\mu_{ALT} = 35$, *m* likely varies somewhere between 29 and 41
- Here *m* lands in the "acceptance" region with reasonable frequency, as a Type II Error has probability of 0.330

Probability of a Type II Error

Hypothesize 30; Presume true mean is actually $\mu_{ALT} = 33$

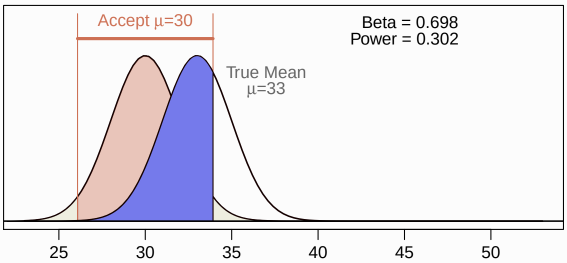


```
> power.t.test(n=25, sd=10, delta=3, type="one.sample")
```

- For $\mu_{ALT} = 33$, *m* likely varies somewhere between 27 and 39
- Now *m* lands in the "acceptance" region more often than not, as the probability of a Type II Error is 0.698

Probability of a Type II Error

Hypothesize 30; Presume true mean is actually $\mu = 31$

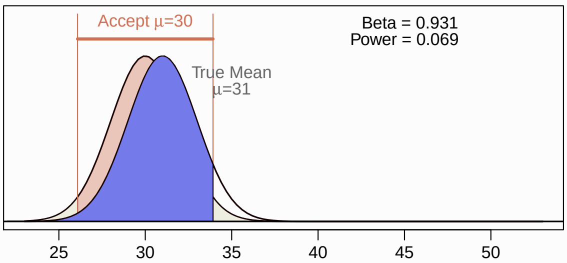


```
> power.t.test(n=25, sd=10, delta=1, type="one.sample")
```

- For $\mu_{ALT} = 31$, *m* likely varies somewhere between 25 and 37
- Now *m* lands in the "acceptance" region most of the time, with a probability of a Type II Error at 0.931

5.4d
The Power Curve

Power Over a Range of Alternative Values of μ

Simultaneously Consider Many Alternative Values of μ

Only one μ_0, but many possible alternative values of μ

- Organize this information regarding power for multiple alternative values of μ into a single graph
- **Power curve**: Plot the power relative to a specified null value, μ_0, across a range of alternative values of μ, μ_{ALT}
- By analyzing power over a range of possible values of μ, an interesting pattern emerges that can provide the analyst much information for the interpretation of a specific hypothesis test
  - Values of μ far from μ_0 are almost assuredly detected
  - Values of μ close to μ_0 are almost assuredly *not* detected
- Discover "far from" and "close to" for a specific analysis
- One tradition is to establish a minimum threshold of desirable power of 0.8 in general, which may be adjusted up or down for a specific situation

Power Curve for Previous Example with Sample Size of 25

- lessR `ttestPower` function provides the power curve

```
> ttestPower(n=25, s=10, mu=30)
```

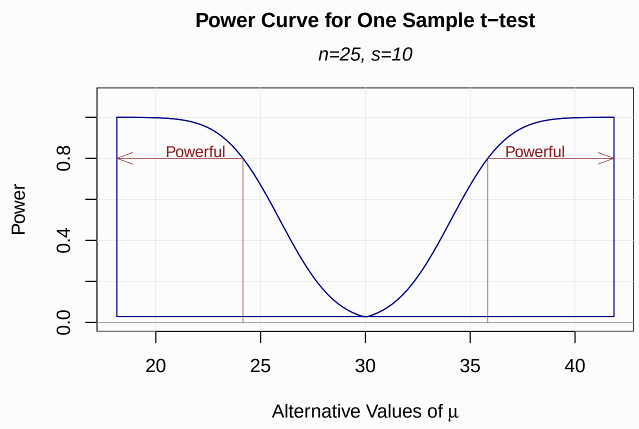

Power Curve Interpretation

What ranges of μ have high and low power, respectively?

- Small changes more difficult to detect, so always have low power against values close to the hypothesized value
- Minimum desired power is usually set at 0.8, here not obtained until real $\mu > 36$ or $\mu < 24$
- Probability of detecting a real change from 30 is less than 50% for values of μ all the way up to 34 and down to 26
- So what if the true mean is $\mu = 33$, which management decides is a sufficiently large increase from $\mu_0 = 30$ to indicate a meaningful, interesting difference from 30?
- If $\mu = 33$, then such a value is too close to the hypothesized value of 30 to be detected with any reasonable probability in this analysis, as the power for this value is only 0.302
- The next section shows how to address this issue

Needed Sample Size for Sufficient Power

Effect of Sample Size on Power

The more data, the more ability to detect a real difference

- The solution to the problem of trying to detect a smaller change from the null hypothesized value is the usual solution in statistical analysis, get more information, that is, more data
- The larger the sample size, n, the smaller the standard error, s_m, so the narrower the curves for the hypothesized and actual distributions of m
- Narrower curves for the two distributions means less overlap, which means fewer Type II Errors
- **Key Concept**: As n gets larger, Power increases, β decreases
- To visualize, compare the Power Curve for $n = 100$, on the next figure, with the previous Power Curve for $n = 25$
- For the larger sample size of $n = 100$, obtain power of at least 0.8 for a real μ of 33 or larger, or for a real μ of 27 or smaller

Power Curve for a Sample Size of 100

- Consider a larger sample size, $n = 100$, a power of 0.8 is obtained for a real $\mu \approx 33$ or $\mu \approx 27$, at $\mu_0 = 30$

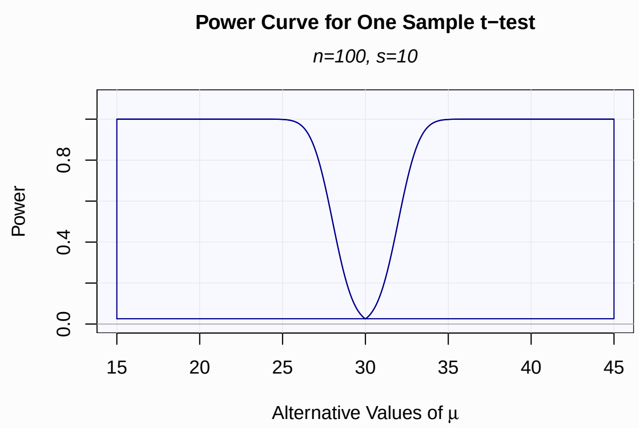

Power Curve for a Sample Size of 1000

- Consider a dramatically larger sample size, $n = 1000$, a power of 0.8 is obtained for a real $\mu \approx 31$ or $\mu \approx 29$, at $\mu_0 = 30$

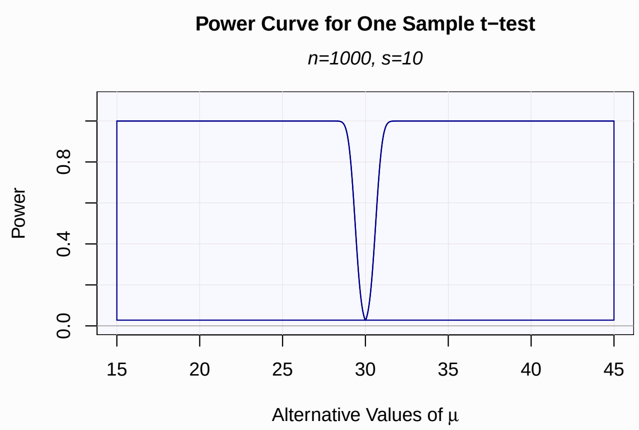

Power Curve for a Sample Size of 10000

- Consider a huge sample size, $n = 10000$, a power of 0.8 is obtained for a real $\mu \approx 29\frac{3}{4}$ or $\mu \approx 30\frac{1}{4}$, at $\mu_0 = 30$

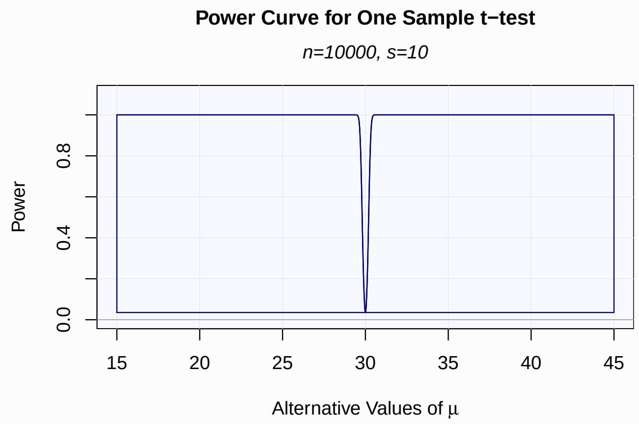

Obtain the Sample Size to Achieve the Needed Power

Set all but one, and then calculate the remaining value

- With the R `power.t.test` function, had previously set *n*, *s* and *delta* and then solved for power
- Now consider a related problem in which the information entered into the `power.t.test` function is the desired power of 0.8 and also *s* and *delta*, and then solve for *n*

```
> power.t.test(power=0.8, sd=10, delta=5,
          type="one.sample")
```

```
              n = 33.36720
                   ⋮
```

- For the original sample with a size of $n = 25$, a power of 0.67 is obtained for a delta of 5
- For a delta of 5, to achieve a power of 0.80, a sample of size $n = 34$ is needed

5.4e
Practical Importance

Practical Importance

Distinguish the meaningful from the trivial

- **Practical importance**: The importance of the extent or size of a deviation from the true value of the population mean, μ, from the null hypothesized value
- **Minimal mean difference** (mmd): Smallest deviation from the null hypothesized value that is of practical importance
- For a hypothesized value of 30, is a difference from
  - 30 to 32 sufficiently large to achieve practical importance?
  - 30 to 30.01 too small to be of practical importance?
- **Key Concept**: The size of mmd follows from the return on investment given implementation of the policy that resulted in the change of the observed magnitude
- Ex: Determine the size of the needed improvement following the re-engineering of a process to justify its cost

Managerial vs Statistical Decisions

Two levels of decisions

- **Key Concept**: The ultimate purpose of statistical analysis for the manager is to facilitate decision making
- Distinguish between the ...
  - Statistical decision: A difference from the null detected?
  - Managerial decision: Follows from a meaningful difference
- The statistical decision may corroborate the managerial decision, but not necessarily
  - Statistically significant but meaningless: A trivial difference is rendered statistically significant by a large sample that yields an extremely powerful test
  - Statistically insignificant but meaningful: A perhaps small but meaningful difference is not detected by a statistical test with insufficient power → *the possibility examined next*

Power Curve with `mmd=2`

```
> ttestPower(n=25, s=10, mu=30, mmd=2)
```

`mmd` option: Range of meaningful changes from μ_0

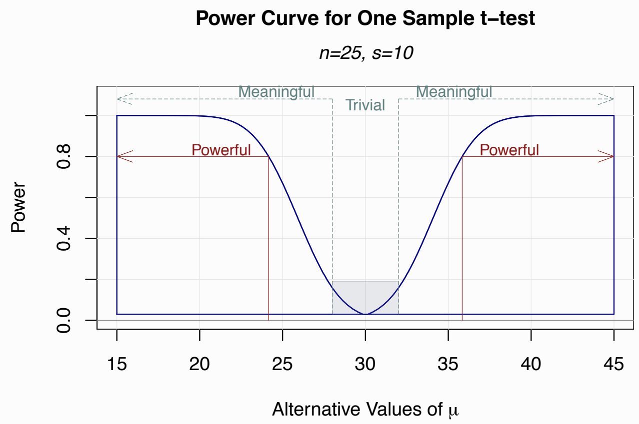

Analysis of Power Curve with Meaningful Changes

What changes are detected and which are meaningful?

- Remember, the power analysis is only of interest if the null is not rejected, a result *not* statistically significant
- To include the minimal mean difference in the analysis of the power curve, add the `mmd` option to the `lessR` `ttestPower` function, which then annotates the power curve accordingly
- The resulting power curve contains two different annotations that delineate two sets of values of μ_{ALT}, those that are ...
  - Powerful: Values likely to be detected, with power `> 0.8`
  - Meaningful: Values sufficiently far from μ_0
- The purpose of both annotations is to compare meaningfulness with power

Analysis of Power Curve with Meaningful Changes

What changes are likely detected and which are meaningful?

- Text output of previous `ttestPower` analysis includes

```
Given n1, n2 and s, Power for mmd of 2 is 0.159
Warning: Meaningful differences from 32 to 35.84
         have Power < 0.8
```

- **Key Concept**: Failure to detect a difference from μ_0 does not imply there is no meaningful difference if some meaningful differences have low power
- If a result is *not* significant, then either
  - There is no difference from the null hypothesis
  - Or, a meaningful difference exists but was not detected
- To help distinguish between these alternatives, discover if if all meaningful values of μ_{ALT} would have been likely detected with the test

Follow-up Analysis to the Hypothesis Test

If the hypothesis test *fails* to reject the null, then ...

- Look at the meaningful alternatives, particularly the smallest
- Do a power analysis against these meaningful alternatives ...
  - If power is low, then maybe a real, meaningful effect exists despite the failure to achieve significance
  - If power is high, then *not* finding a difference from the null means "no meaningful effect likely exists"

If the hypothesis test rejects the null, then ...

- Conclude that a difference from the null value has been detected
- Obtain a confidence interval to estimate the likely value of the true population value μ to understand the magnitude of the effect

5.4f
Comprehensive Strategy for Statistical Inference

Strategy for Managerial Decisions from Stat Inference

An integrated strategy

- Statistical inference is always in pursuit of unknown population values, such as the population mean, μ
- Two complementary forms of statistical inference: Confidence interval and hypothesis test, and related techniques such as power analysis
- Need an integrated strategy of statistical inference that provides a guide to form conclusions about μ or other population values of interest that involves
  - confidence interval
  - hypothesis test
  - power analysis
  - practical importance

Strategy for Managerial Decisions from Stat Inference

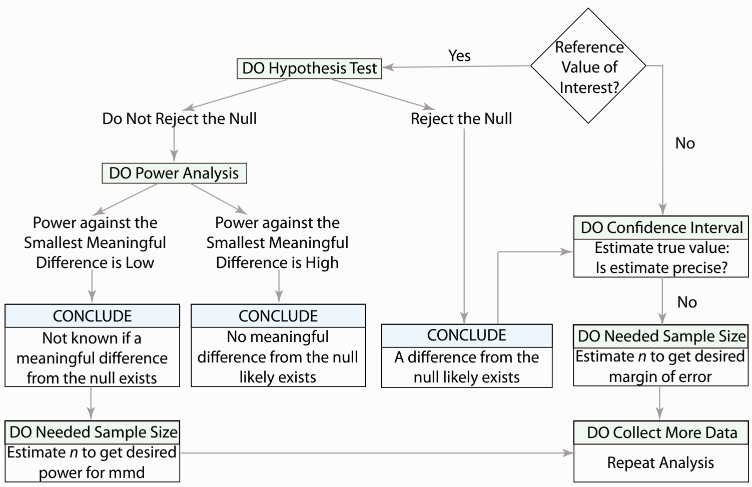

5.4g
Application

Follow-up Analysis from Hypothesis Test with No Rejection

Consider the previous application of filling cereal boxes

- A previous application provided a hypothesis test of the average weight of packaged cereal to be 350g
- Traditional Notation: $n = 25$ $m = 350.86$ $s = 2.66$
- Statistical Decision: p-value $= 0.119$ $>$ $\alpha = .05$
  so do *not* reject the null hypothesized value of 350g, the sample mean of 350.86 is close to 350
- No difference in average weight of cereal from 350g detected
- Failing to reject the null hypothesis, the question becomes:
  What is the power against meaningful alternatives that may exist but were not detected?
- Suppose management decides that any change from the target of 350g is not meaningful if it is 0.5g or less
- That is, `mmd=0.5`

Power Curve with `mmd=.5`

```
> ttestPower(n=25, s=2.66, mu=350, mmd=.5)
```

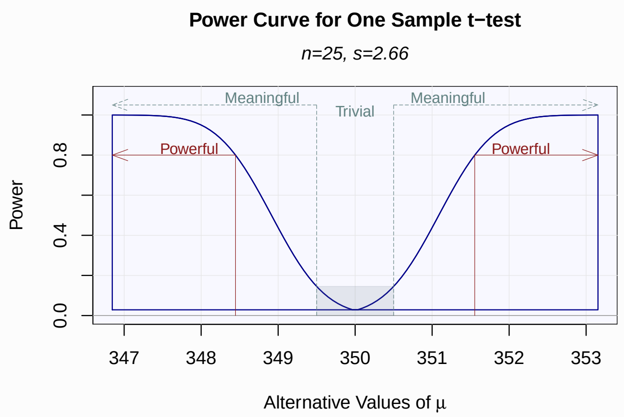

Analysis of Power Curve with Meaningful Changes

What changes are detected and which are meaningful?

- $m = 350.86$, but not so much larger than $\mu_0 = 350g$ to reject μ_0 as unreasonable
- Text output of previous `ttestPower` analysis includes

```
Given n1, n2 and s,  Power for mmd of 0.5 is 0.145
Warning: Meaningful differences from 350.5 to 351.554
         have Power < 0.8
```

- So maybe a meaningful difference exists from the null value of 350g that this hypothesis test failed to detect
- If a more definitive conclusion is desired, then the only alternative is to gather more data, of which `ttestPower` provides the relevant information

```
Needed n to achieve power=0.8 for mmd of 0.5: n=225
```


Managerial Summary

What is important from this analysis for the manager?

- Deviations from the target weight of cereal in the cereal box of 350g in either direction need to be corrected
- Deviations less than 0.5g are deemed to small to provide any reliable or meaningful corrected action that includes shutting down the production line to make the adjustment
- So management wishes to detect any deviation of the average weight from the current production process that is smaller than 349.5g and larger than 350.5g
- With an initial sample size of $n = 25$, the null hypothesis of a population mean of 350g was not rejected

Managerial Conclusion

What is important from this analysis for the manager?

- The failure to reject lead to an analysis of power to determine if the detection of some meaningful differences from the null were unlikely
- The primary conclusions are that
  - Differences less than 351.55g have low power
  - The minimal value of interest larger than the target of 350g is 350.5g, which has a power of only 0.145
- To pursue this analysis further would require a considerable increase of sample size from $n = 25$ to $n = 225$ to achieve high power against against the value of $\mu_{ALT} = 350.5$g

Index Subtract 2 from each listed value to get the Slide

- The End