Introduction to Experimetrics and Power Analysis (實驗計量與統計檢定力分析)

Joseph Tao-yi Wang (王道一)
Experimetrics Lecture 1 (實驗計量第一講)

Outline: The Replication Size Trinity

1. Sample Size n : # of observations/subjects
2. Effect Size: How big is the true result
3. Power (1-β): How likely will your test show significance if there is truly an effect

Why Do We Care About This?

- Editor's Preface (JEEA 2015):
  - A necessary (but not sufficient) condition for publishing a replication study or null result
  - will be the presentation of power calculations.
- Test Resolution: $\Pr(\text{confirm} \mid \text{infected patient})$
  - In 2020, Taiwan requires 3 consecutive negatives to discharge for COVID-19, since even PCR has insufficient power (around 70%)...
- But what about structural estimation?

Key Concepts and Definitions

- Treatment Test:
  - Null ($H_0 : \theta = \theta_0$) Hypothesis - No Effect!
  - Alternative ($H_1 : \theta = \theta_1$) Hypothesis - Effective!
- Effect Size ($\theta_1 - \theta_0$): True size of effect
- Alternative Hypothesis can be Directional:
  - One-sided Alternative - One-tailed test
    - Usually comes from prior beliefs based on theory
  - Two-sided Alternative - Two-tailed test

Key Concepts and Definitions

- Two Stages of the Treatment Test:
  1. Compute Test Statistic of sample size n
  2. Compare Test Statistic with null distribution
- Rejection Region = Tail of null distribution
  - of a Size $\alpha = \Pr(\text{reject null} \mid \text{null is true})$
  - Critical Value: Rejection region starting point
- p-value $= \Pr\left(|T| \geq T_{CV} \middle| \text{ null is true}\right)$
  - $p < 0.05$ vs. $p < 0.01/0.001$ (strength of evidence)
  - Evidence vs. Strong/Overwhelming Evidence

Key Concepts and Definitions

- Type 1 Error: $\alpha = \Pr(\text{reject null} \mid \text{null is true})$

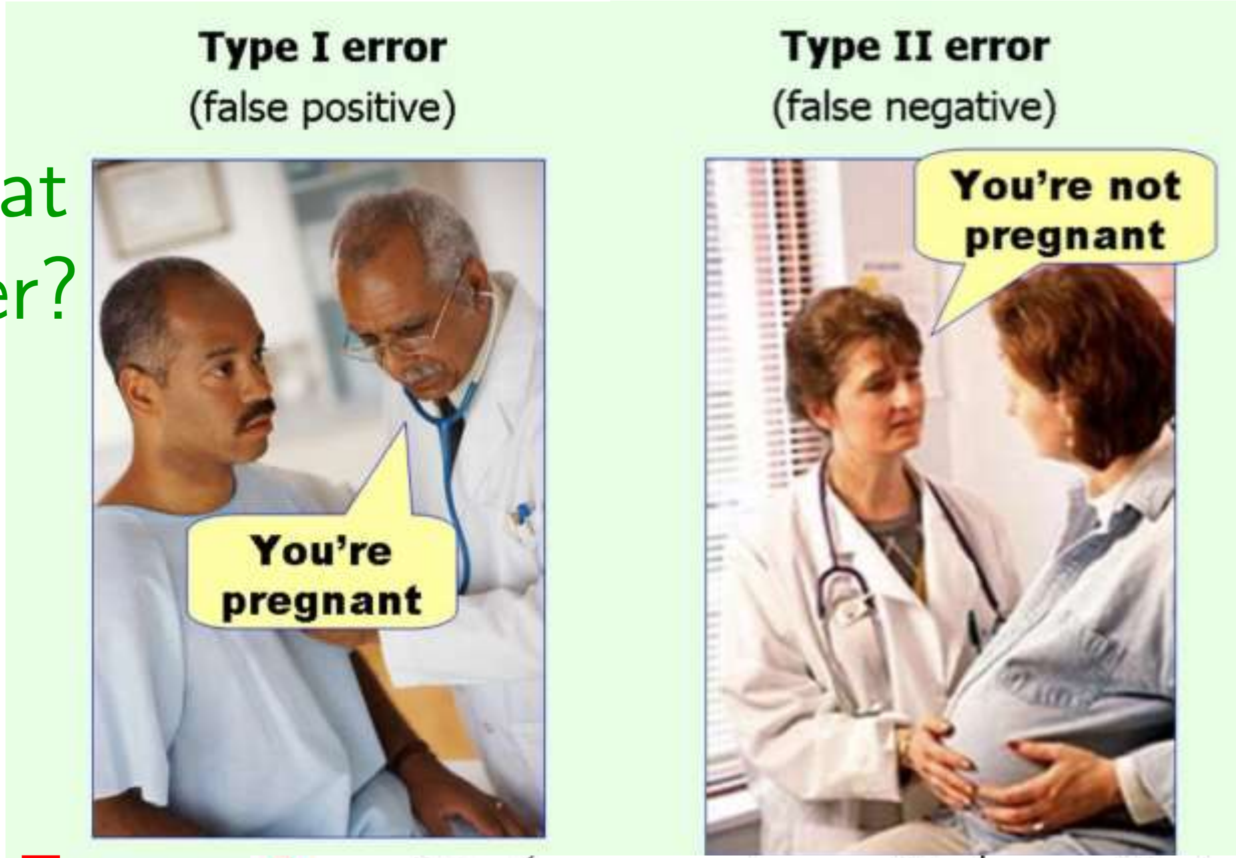

- But what is Power?

- Type 2 Error: $\beta = \Pr(\text{accept null} \mid \text{null is false})$

Key Concepts and Definitions

- Type 1 Error: $\alpha = \Pr(\text{reject null} \mid \text{null is true})$
- Type 2 Error: $\beta = \Pr(\text{accept null} \mid \text{null is false})$
- Power(π): $1 - \beta = \Pr(\text{reject null} \mid \text{null is false})$
  1. True effect size $\theta_1 - \theta_0$ (and one/two-tailed)
  2. Sample size n
  3. Size of the test α
- Trade-off: The higher α/n, the higher is π
  1. Power Analysis: Compute power $\pi = 1 - \beta$, or
  2. Find n to meet power requirement $\pi(n) \geq \overline{\pi}$

Choosing the Value of α

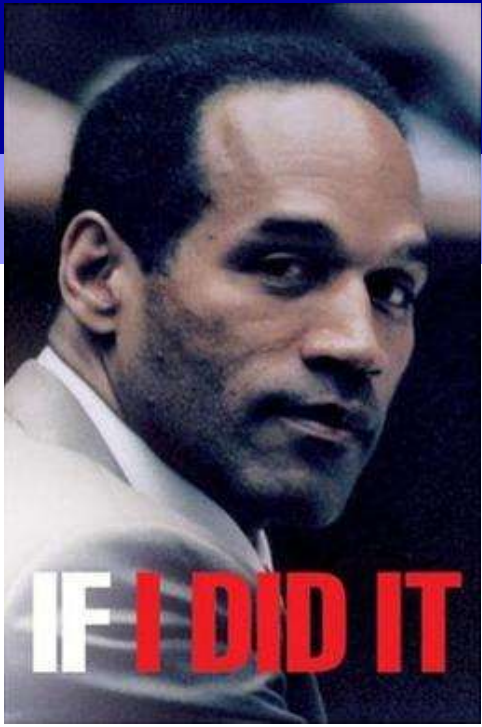

- How big can we allow Type 1 Error
- To convict a crime suspect,
  - Null Hypothesis: Not Guilty
  - Alternative Hypothesis: Guilty
  - Type 1: $\alpha = \Pr(\text{convict} \mid \text{innocent suspect})$
  - Type 2: $\beta = \Pr(\text{acquit} \mid \text{guilty suspect})$
- Type 1 Error more serious than Type 2 Error
  - Choose very low α at the expense of power:

$$1 - \beta = \Pr(\text{convict} \mid \text{guilty suspect})$$

Choosing the Value of α

- How big can we allow Type 1 Error?
- To test for COVID-19,
  - Null Hypothesis: Healthy
  - Alternative Hypothesis: Infected by COVID-19
  - Type 1: $\alpha = \Pr(\text{confirm} \mid \text{healthy patient})$
  - Type 2: $\beta = \Pr(\text{discharge} \mid \text{infected patient})$
- Type 2 Error more serious than Type 1 Error
  - Choose a higher α so get higher power:

$$1 - \beta = \Pr(\text{confirm} \mid \text{infected patient})$$

Choosing the Value of α

- Type 1 $\alpha = \Pr(\text{confirm} \mid \text{healthy patient})$
- Type 2 $\beta = \Pr(\text{discharge} \mid \text{infected patient})$
- Both errors not fatal in Experimental Economics,
- Convention is:

$$\alpha = 0.05$$
$$\pi = 1 - \beta = 0.80$$
$$\beta = 0.20$$

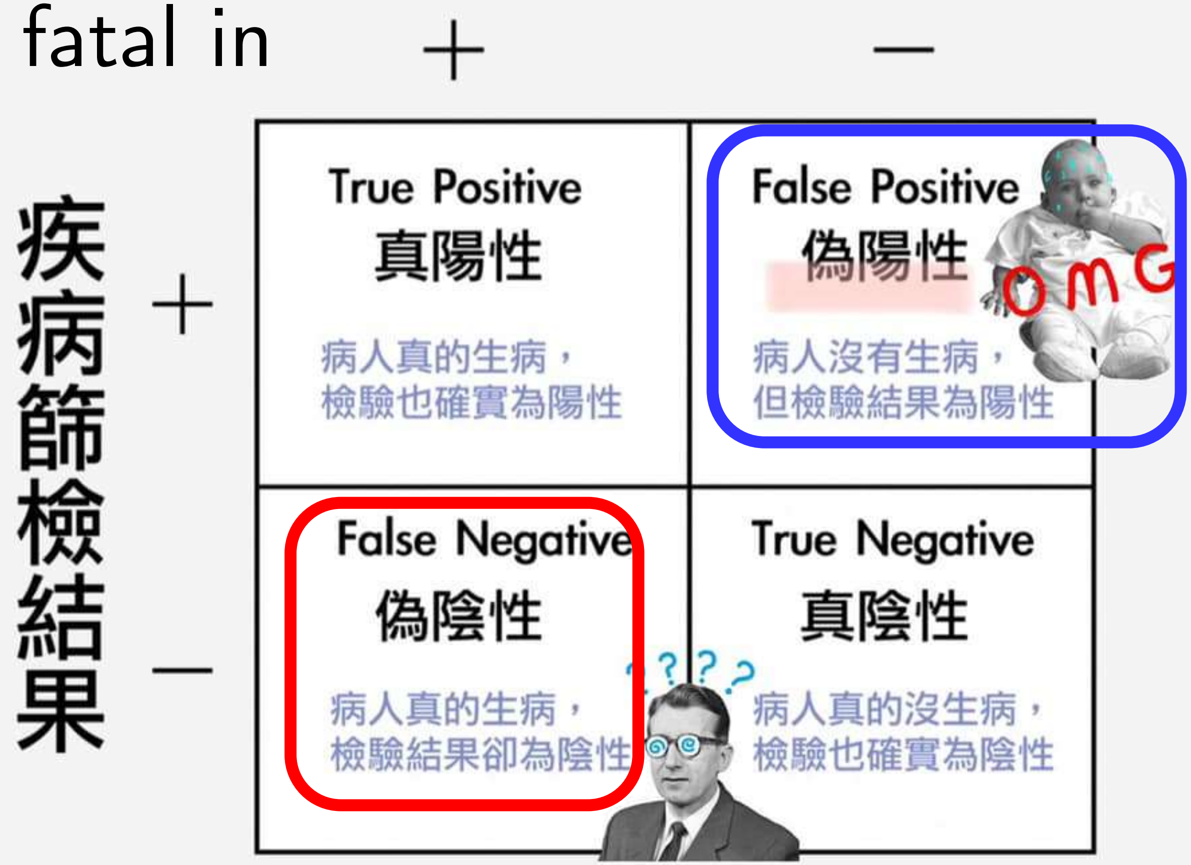

Treatment Testing Toolkit

- One-sample t-test
  - Does WTP $= £3$ ($=$ retail value of coffee mug)?
- Two-sample t-test (with equal variance)
  - If passes variance ratio test
  - Can be done using OLS!
- Two-sample t-test (with unequal variance)
  - If fails variance ratio test
  - Skewness-kurtosis test
- Need CLT: Okay if sufficiently large n $(\geq 30?)$

Treatment Testing Toolkit

- What if we do not have CLT/large n ?
  - Use non-parametric tests instead!
- Mann-Whitney Test (aka ranksum test)
  - Between-subject non-parametric treatment test
- Kolmogorov-Smirnov (KS) Test
- Epps-Singleton Test (discrete KS test)
  - Tests comparing entire distributions

Treatment Testing: WTP - WTA Gap

- What if we have within-subject data?
- Can use within-subject tests!
  - But, watch out for order effect...
- Paired t-test (assume CLT)
- Wilcoxon Signed Rank Test
  - Within-subject non-parametric treatment test
  - Assume symmetric distribution around median
  - (regarding paired difference). Without it, use:
- Paired-sample sign test

Treatment Testing: WTP - WTA Gap

- Isoni et al. (AER 2011)
  - Replicate Plott and Zeiler (AER 2007), which
  - Replicate Kahneman et al. (JPE 1990) (KKT)
- Measure WTP and/or WTA
  - Becker–DeGroot–Marschak (BDM) mechanism
  - 2nd price auction against (randomizing) computer
- Treatment Test:
  - Does WTP or WTA $= £3$ ($=$ retail value of the coffee mug)?

Power Analysis: Theory

1. Power Analysis: Find test power $\pi = 1 - \beta$, or
2. Find n to meet power requirement $\pi(n) \geq \overline{\pi}$

- One-sample t-test
  - Rarely used in experimental economics, but...
  - Isoni et al. (2011) test WTP of coffee mug = £3
- Y: Continuous outcome measure with mean μ
  - Null Hypothesis: $H_0 : \mu = \mu_0$
  - Alternative Hypothesis: $H_1 : \mu = \mu_1 > \mu_0$
- Collect data of sample size n

Power Analysis: Theory

1. What is the power of this test?
2. How big should sample size n be?

- Test Size $\alpha = 0.05 = \Pr(\text{reject null} \mid \text{null is true})$
- Type 2 $\beta = 0.20 = \Pr(\text{accept null} \mid \text{null is false})$
- Power $\pi = 1 - \beta = 0.80$
- One-sample t-test
  - Test Statistic: $t = \dfrac{\bar{y} - \mu_0}{s/\sqrt{n}} \sim t(n-1)$
  - Reject if $t > t_{n-1,\alpha}$ ($t > z_\alpha$ for large n)

$\bar{y} = \text{sample mean}$

$s^2 = \text{sample variance}$

Power Analysis: Power of the Test

$$\pi = \Pr(t > z_\alpha | \mu = \mu_1) = \Pr\left(\frac{\bar{y} - \mu_0}{s/\sqrt{n}} > z_\alpha \middle| \mu = \mu_1\right)$$

$$= \Pr\left(\bar{y} > \mu_0 + z_\alpha (s/\sqrt{n}) \middle| \mu = \mu_1\right)$$

$\mu_0 = 10$

$\mu_1 = 12$

$$= \Pr\left(\frac{\bar{y} - \mu_1}{s/\sqrt{n}} > \frac{\mu_0 + z_\alpha (s/\sqrt{n}) - \mu_1}{s/\sqrt{n}} \middle| \mu = \mu_1\right)$$

$z_\alpha = 1.645, \quad s = 5$

$$= \Phi\left(\frac{12 - 10 - 1.645(5/\sqrt{30})}{5/\sqrt{30}}\right) = \underline{0.71}$$

$n = 30, \quad \alpha = 0.05$

- What n is required to get $\pi = 0.80$?

Power Analysis: How Big Should n Be?

- Power $\pi = 1 - \beta = \Phi\left(\dfrac{\mu_1 - \mu_0 - z_\alpha (s/\sqrt{n})}{s/\sqrt{n}}\right)$

$$\Rightarrow z_\beta = \frac{\mu_1 - \mu_0 - z_\alpha (s/\sqrt{n})}{s/\sqrt{n}}$$

$\alpha = 0.05, \quad \beta = 0.20$

$$\Rightarrow z_\beta + z_\alpha = \frac{\mu_1 - \mu_0}{s/\sqrt{n}}$$

$z_\alpha = 1.645, \quad z_\beta = 0.842$

$s = 5$

$$\Rightarrow n = \frac{s^2 (z_\alpha + z_\beta)^2}{(\mu_1 - \mu_0)^2} = \frac{5^2 (1.645 + 0.842)^2}{(12 - 10)^2}$$

$\mu_0 = 10$

$\mu_1 = 12$

- So we need $n \geq 39$ $\qquad = \underline{38.66}$

Power Analysis: Power in STATA

- What is the power for sample size $n = 30$?
  - STATA command for power calculation

μ_0 / μ_1

```
power onemean 10 12 , sd(5) n(30) oneside
```

sample std; sample size

  - 1-sample t-test

one-tailed test

Power Analysis: Power Results in STATA

- What is the power for sample size $n = 30$?

```
power onemean 10 12 , sd(5) n(30) oneside
```

- STATA Results:

```
Estimated power for a one-sample mean test
t test
Ho: m = m0  versus  Ha: m > m0

Study parameters:

        alpha =    0.0500
            N =        30
        delta =    0.4000
           m0 =   10.0000
           ma =   12.0000
           sd =    5.0000

Estimated power:

        power =    0.6895
```

Slightly different since STATA did not use normal approximation...

Power Analysis: Sample Size in STATA

- What is the sample size to get power $\pi = 0.80$?
  - STATA command for power calculation

μ_0 / μ_1 — sample std

```
power onemean 10 12 , sd(5) oneside p(0.8)
```

required power

- 1-sample t-test

one-tailed test

Power Analysis: Sample Size Result/Stata

- What is the sample size to get power $\pi = 0.80$?

```
power onemean 10 12 , sd(5) oneside p(0.8)
```

  - STATA Results:

```
Performing iteration ...

Estimated sample size for a one-sample mean test
t test
Ho: m = m0  versus  Ha: m > m0

Study parameters:

        alpha =    0.0500
        power =    0.8000
        delta =    0.4000
           m0 =   10.0000
           ma =   12.0000
           sd =    5.0000

Estimated sample size:

            N =        41
```

Slightly larger n since STATA did not use normal approximation...

Power Analysis: Graph Power in STATA

- Plot power against sample size with `graph`
  - STATA command for power calculation

$\mu_0/\mu_1 \in [10.5, 12.5]$

```
power onemean 10 (10.5(0.5)12.5), sd(5) n(20(10)200) oneside graph
```

sample std; n=20-200

  - 1-sample t-test

one-tailed test

Power Analysis: Graph Power in STATA

- Plot power against sample size with `graph`

```
power onemean 10 (10.5(0.5)12.5), sd(5) n(20(10)200) oneside graph
```

- STATA Results:

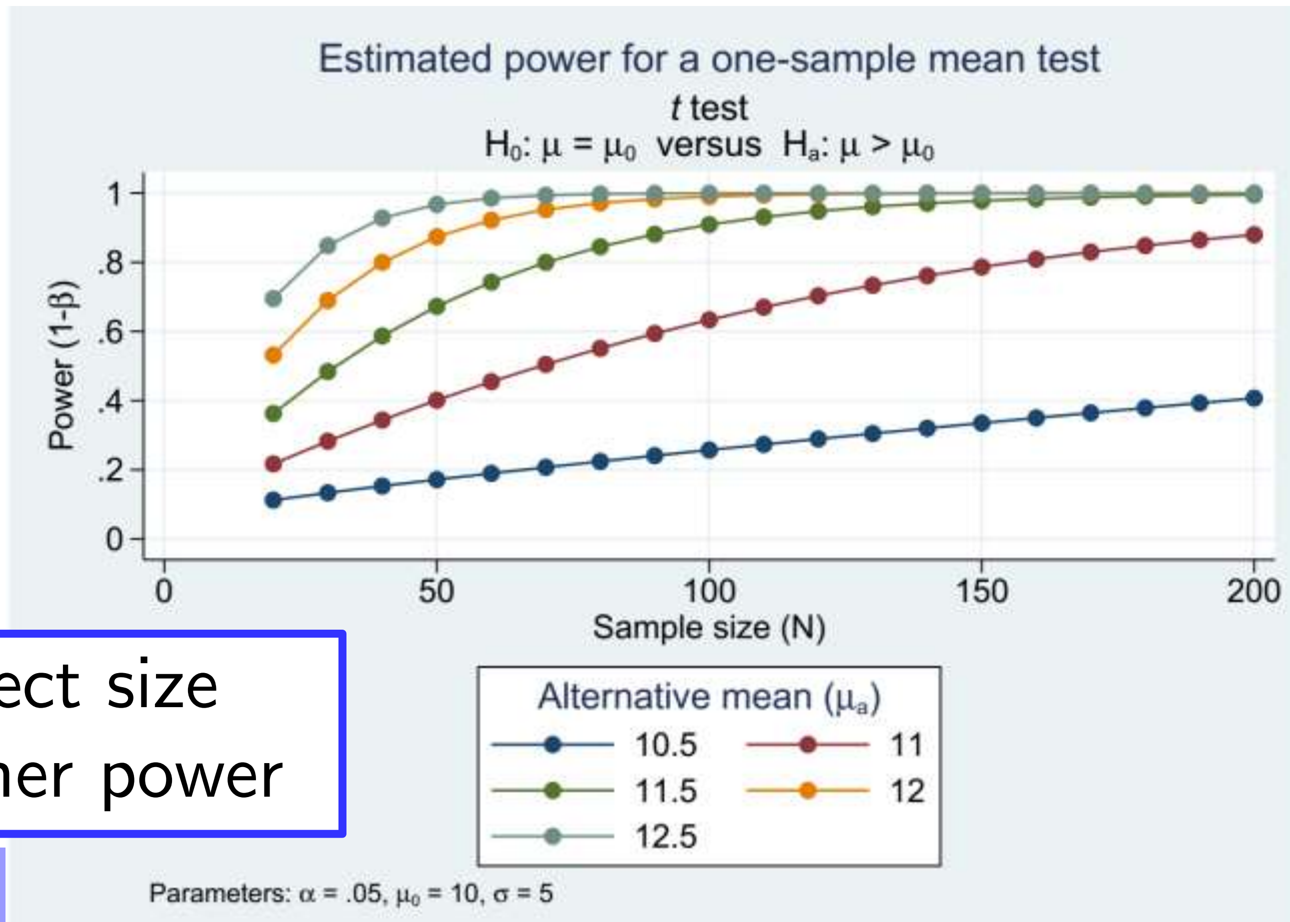

Larger effect size yields higher power

Power Analysis: Graph Sample Size/Stata

- Plot sample size against effect size
  - STATA command for power calculation

$$\mu_0/\mu_1 \in [10.5, 12.5]$$

```
power onemean 10 (10.5(0.25)12.5), sd(5) p(0.6(0.1)0.9) oneside graph
```

sample std; power=0.6-0.9

  - 1-sample t-test

one-tailed test

Power Analysis: Graph Sample Size/Stata

- Plot sample size against effect size

```
power onemean 10 (10.5(0.25)12.5), sd(5) p(0.6(0.1)0.9) oneside graph
```

- STATA Results:

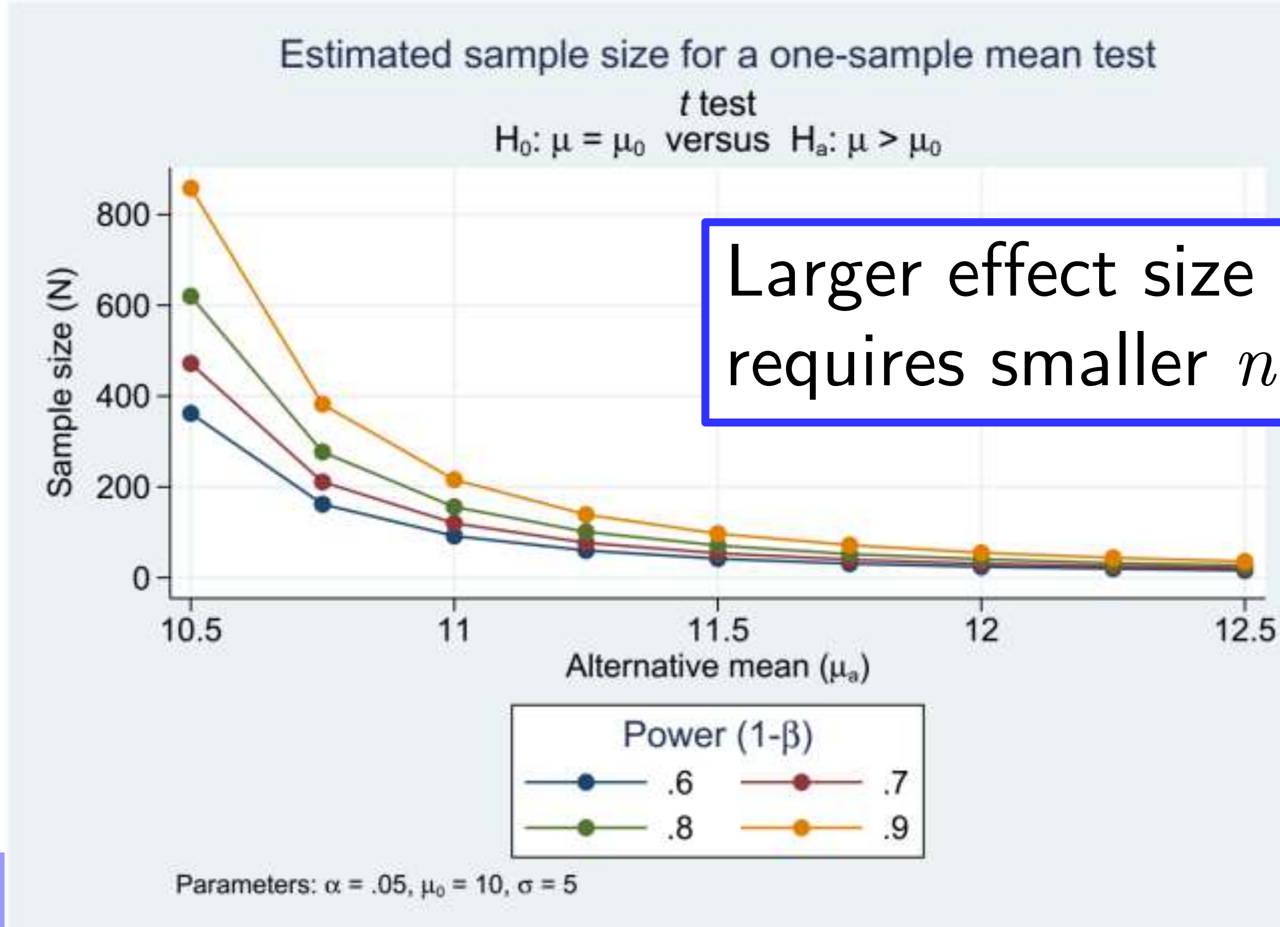

Larger effect size requires smaller n

Power Analysis: Two-sample t-test

1. Power Analysis: Find test power $\pi = 1 - \beta$, or
2. Find n to meet power requirement $\pi(n) \geq \overline{\pi}$

- Two-sample t-test
  - More common in experimental economics...
- μ_1: Population mean of control group
- μ_2: Population mean of treatment group
  - Null Hypothesis: $H_0 : \mu_2 - \mu_1 = 0$
  - Alternative Hypothesis: $H_1 : \mu_2 - \mu_1 = d$
- Collect data of sample size n_1 and n_2

Effect Size from prior

Power Analysis: Two-sample t-test

- Test Size $\alpha = 0.05$
- Type 2 $\beta = 0.20 =$
- Power $\pi = 1 - \beta = 0.80$

$\bar{y}_1, \bar{y}_2 = \text{sample means}$

$s_1^2, s_2^2 = \text{sample variances}$

$$s_p = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}}$$ (pooled sample s.d. for equal variance)

- Test Statistic: $$t = \frac{\bar{y}_2 - \bar{y}_1}{s_p\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}} \sim t(n_1 + n_2 - 2)$$
- Reject if $t > t_{n_1+n_2-2,\alpha}$ ($t > z_\alpha$ for large n)

Power Analysis: Two-sample t-test

- Equal sample size

  $n_1 = n_2 = n$

$\bar{y}_1, \bar{y}_2 = \text{sample means}$

$s_1^2, s_2^2 = \text{sample variances}$

- Pooled sample s.d. for equal variance is:

$$s_p = \sqrt{\frac{s_1^2 + s_2^2}{2}}$$

  - Test
  - Statistic:

$$t = \frac{\bar{y}_2 - \bar{y}_1}{s_p \sqrt{\frac{2}{n}}} \sim t(2n-2)$$

- Reject if $t > t_{2n-2,\alpha}$ ($t > z_\alpha$ for large n)

Power Analysis: Power of the Test

$$
\begin{aligned}
\pi &= \Pr(t > z_\alpha | \mu_2 - \mu_1 = d) \\
&= \Pr\left(\frac{\overline{y}_2 - \overline{y}_1}{s_p\sqrt{2/n}} > z_\alpha \middle| \mu_2 - \mu_1 = d \right) \\
&= \Pr\left(\overline{y}_2 - \overline{y}_1 > z_\alpha s_p \sqrt{2/n} \middle| \mu_2 - \mu_1 = d \right) \\
&= \Pr\left(\frac{\overline{y}_2 - \overline{y}_1 - d}{s_p\sqrt{2/n}} > \frac{z_\alpha s_p \sqrt{2/n} - d}{s_p\sqrt{2/n}} \middle| \mu_2 - \mu_1 = d \right) \\
&= \Phi\left(\frac{d - z_\alpha s_p \sqrt{2/n}}{s_p\sqrt{2/n}} \right)
\end{aligned}
$$

Power Analysis: How Big Should n Be?

- Power $\pi = 1 - \beta = \Phi\left(\frac{d - z_\alpha s_p \sqrt{2/n}}{s_p \sqrt{2/n}}\right)$

$$\Rightarrow z_\beta = \frac{d - z_\alpha s_p \sqrt{2/n}}{s_p \sqrt{2/n}}$$

$\alpha = 0.05, \quad \beta = 0.20$

$z_\alpha = 1.645, \quad z_\beta = 0.842$

$s_1 = 4.0, s_2 = 5.84$

$$\Rightarrow z_\beta + z_\alpha = \frac{d}{s_p \sqrt{2/n}}$$

$s_p^2 = \frac{s_1^2 + s_2^2}{2} = 5.0^2$

$$\Rightarrow n = \frac{2 s_p^2 (z_\alpha + z_\beta)^2}{d^2} = \frac{2(5^2)(1.645 + 0.842)^2}{2^2}$$

$d = 2$

- So we need $n \geq 78$

$$= \underline{77.32}$$

Power Analysis: Sample Size in Stata

- What is the sample size to get power $\pi = 0.80$?
  - STATA command for power calculation

2 sample std's

μ_0 / μ_1

```
power twomeans 10 12 , sd1(4.0) sd2(5.84) oneside p(0.8)
```

required power

- 2-sample t-test

one-tailed test

Power Analysis: Sample Size Result/Stata

- What is the sample size to get power $\pi = 0.80$?

```
power twomeans 10 12 , sd1(4.0) sd2(5.84) oneside p(0.8)
```

- STATA Results:

```
Performing iteration ...

Estimated sample sizes for a two-sample means test
Satterthwaite's t test assuming unequal variances
Ho: m2 = m1  versus  Ha: m2 > m1

Study parameters:

        alpha =    0.0500
        power =    0.8000
        delta =    2.0000
           m1 =   10.0000
           m2 =   12.0000
          sd1 =    4.0000
          sd2 =    5.8400

Estimated sample sizes:

            N =       158
  N per group =        79
```

Slightly larger n since STATA did not use normal approximation...

Power Analysis: Graph Power in STATA

- Plot power against sample size with `graph`
  - STATA command for power calculation

μ_0 / μ_1

```
power twomeans 10 12, sd1(4.0) sd2(5.84) n(20(10)200) oneside graph
```

sample std; n=20-200

  - 2-sample t-test

one-tailed test

Power Analysis: Graph Power in STATA

- Plot power against sample size with `graph`

```
power twomeans 10 12, sd1(4.0) sd2(5.84) n(20(10)200) oneside graph
```

- STATA Results:

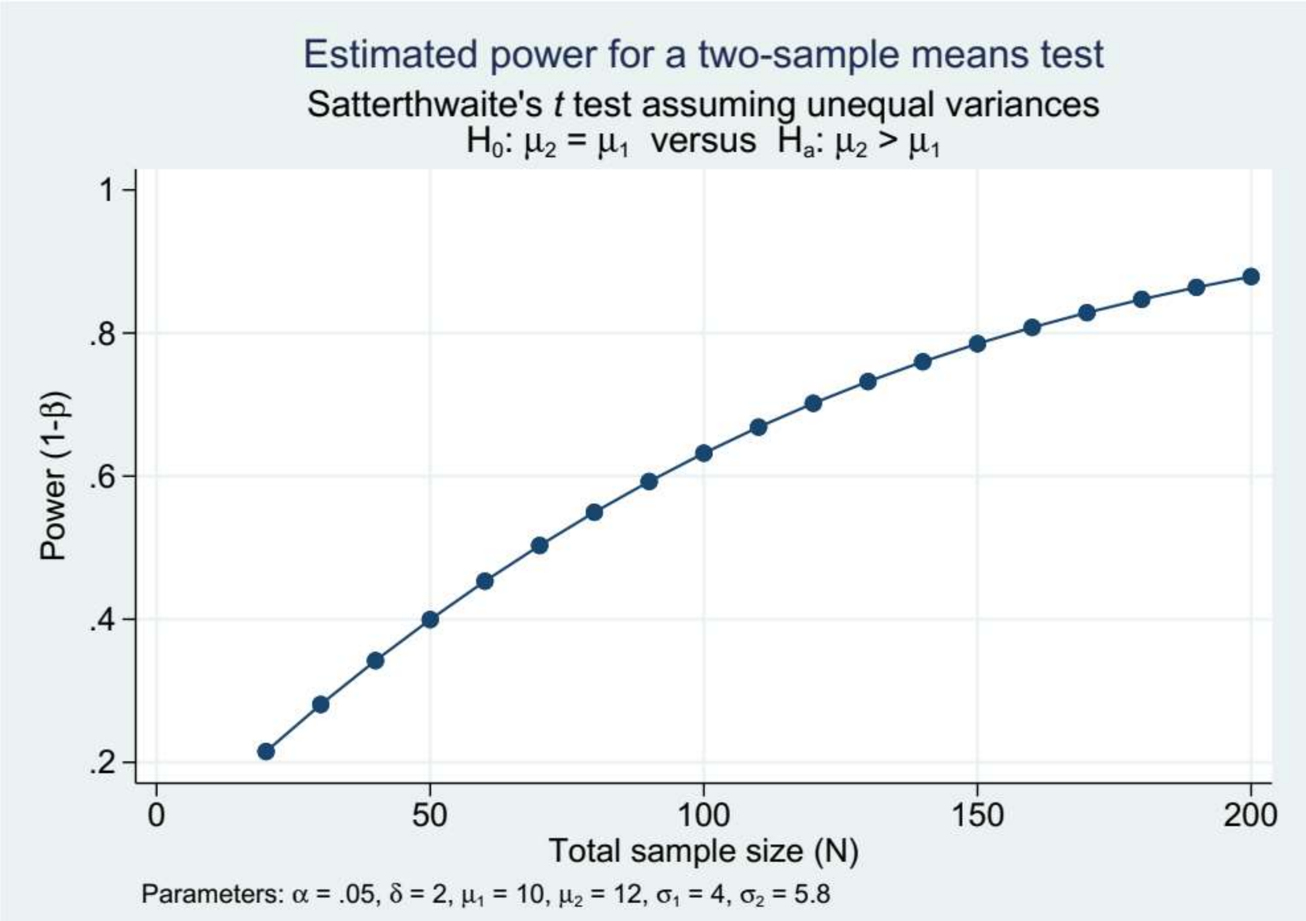

Larger total sample size yields higher power

Power Analysis: Paired Sample t-test

1. Power Analysis: Find test power $\pi = 1 - \beta$, or
2. Find n to meet power requirement $\pi(n) \geq \overline{\pi}$

- Paired Sample t-test
  - Observe subjects twice (with vs. w/o treatment)...
- μ_1: Population mean without treatment
- μ_2: Population mean with treatment
  - Null Hypothesis: $H_0 : \mu_2 - \mu_1 = 0$
  - Alternative Hypothesis: $H_1 : \mu_2 - \mu_1 = d$
- Collect data of sample size n (HW: Theory?!)

Effect Size from prior

Power Analysis: Paired Sample t-test

- What is the sample size to get power $\pi = 0.80$?
  - STATA command for power calculation

μ_0 / μ_1

```
power pairedmeans 10 12 , sd1(4.0) sd2(5.84) corr(0) oneside
```

2 sample std's & its correlation

- Paired Sample t-test

one-tailed test

Power Analysis: Paired Sample t-test

- What is the sample size to get power $\pi = 0.80$?

```
power pairedmeans 10 12 , sd1(4.0) sd2(5.84) corr(0) oneside
```

- STATA Results:

```
Performing iteration ...

Estimated sample size for a two-sample paired-means test
Paired t test
Ho: d = d0  versus  Ha: d > d0

Study parameters:

        alpha =    0.0500           ma1 =   10.0000
        power =    0.8000           ma2 =   12.0000
        delta =    0.2825           sd1 =    4.0000
           d0 =    0.0000           sd2 =    5.8400
           da =    2.0000          corr =    0.0000
         sd_d =    7.0785

Estimated sample size:

            N =        79
```

Same n per group as in two-sample t-test!

Power Analysis: Paired Sample t-test

- What if subjects' two responses are correlated?
  - STATA command for power calculation

2 sample std's

μ_0 / μ_1

```
power pairedmeans 10 12 , sd1(4.0) sd2(5.84) corr(0.5) oneside
```

2 sample's correlation

  - Paired Sample t-test

one-tailed test

Power Analysis: Paired Sample t-test

- What if subjects' two responses are correlated?

```
power pairedmeans 10 12 , sd1(4.0) sd2(5.84) corr(0.5) oneside
```

- STATA Results:

```
Performing iteration ...

Estimated sample size for a two-sample paired-means test
Paired t test
Ho: d = d0   versus  Ha: d > d0

Study parameters:

        alpha =     0.0500             ma1 =    10.0000
        power =     0.8000             ma2 =    12.0000
        delta =     0.3867
           d0 =     0.0000
           da =     2.0000
         sd_d =     5.1716

Estimated sample size:

            N =         43
```

Paired data positively correlated (from same subjects) require less n

- What if corr. $= 1$?
  - n=1 ok!

Power Analysis: Paired Sample t-test

- Paired Tests saves number of subjects!
  - Since the two responses are highly correlated
- But can cause Order Effect
  - AB and BA yield different results
- Crossover Design
  - Half AB vs. Half BA
  - Measures Order Effects (if any)
    - Can later control with order dummies, etc.

Real Example: WTP - WTA Gap

- Isoni et al. (AER 2011)
  - Replicate Plott and Zeiler (AER 2007), which
  - Replicate Kahneman et al. (JPE 1990) (KKT)
- Measure WTP and/or WTA
  - Becker–DeGroot–Marschak (BDM) mechanism
  - 2nd price auction against (randomizing) computer
- Treatment Test:
  - Does WTP or WTA = £3 (= retail value of the coffee mug)?

WTP - WTA Gap: Summary Statistics

- Summary Statistics of N=100

```
summ v_mug
```

  - STATA Results:

Variable	Obs	Mean	Std. Dev.	Min	Max
v_mug	100	2.0415	1.571287	0	7.5

Is the population mean value of a mug = 3.0?

WTP - WTA Gap: One-sample t-test

- Test $H_0 : \mu = 3.0$

```
ttest v_mug=3.0
```

- STATA Results:

```
One-sample t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
   v_mug |     100      2.0415    .1571287    1.571287    1.729723    2.353277
------------------------------------------------------------------------------
    mean = mean(v_mug)                                            t =  -6.1001
Ho: mean = 3.0                                   degrees of freedom =       99

   Ha: mean < 3.0              Ha: mean != 3.0               Ha: mean > 3.0
 Pr(T < t) = 0.0000        Pr(|T| > |t|) = 0.0000          Pr(T > t) = 1.0000
```

Power Analysis: One-sample t-test

- What is the power of this test?

```
power onemean 3.0 2.04, n(100) sd(1.571)
```

- STATA Results:

```
Estimated power for a one-sample mean test
t test
Ho: m = m0  versus  Ha: m != m0

Study parameters:

        alpha =    0.0500
            N =       100
        delta =   -0.6111
           m0 =    3.0000
           ma =    2.0400
           sd =    1.5710

Estimated power:

        power =    1.0000
```

Power = 1?
Not surprising
if mean=2.04
& s.d.=1.571

Power Analysis: One-sample t-test

- What is the sample size to get power $\pi = 0.80$?

```
power onemean 3.0 2.04, p(0.80) sd(1.571)
```

- STATA Results:

```
Estimated sample size for a one-sample mean test
t test
Ho: m = m0  versus  Ha: m != m0

Study parameters:

        alpha =    0.0500
        power =    0.8000
        delta =   -0.6111
           m0 =    3.0000
           ma =    2.0400
           sd =    1.5710

Estimated sample size:

            N =        24
```

n=24 is enough!

Power Analysis: Two-sample t-test

- Is there sufficient power to test WTA vs. WTP?

```
power twomeans 1.86 2.21 , n1(49) n2(51) sd1(1.29) sd2(1.80) oneside
```

- STATA Results:

```
Estimated power for a two-sample means test
Satterthwaite's t test assuming unequal variances
Ho: m2 = m1  versus  Ha: m2 > m1

Study parameters:

        alpha =    0.0500
            N =       100
           N1 =        49
           N2 =        51
        N2/N1 =    1.0408
        delta =    0.3500
           m1 =    1.8600
           m2 =    2.2100
          sd1 =    1.2900
          sd2 =    1.8000

Estimated power:

        power =    0.2973
```

- Power < 0.3!!

Power = Probability of rejecting null hypothesis

Power Analysis: Two-sample t-test

- What is the sample size to get power $\pi = 0.80$?

```
power twomeans 1.86 2.21 , sd1(1.29) sd2(1.80) oneside p(0.8)
```

- STATA Results:

```
Performing iteration ...

Estimated sample sizes for a two-sample means test
Satterthwaite's t test assuming unequal variances
Ho: m2 = m1  versus  Ha: m2 > m1

Study parameters:

        alpha =    0.0500
        power =    0.8000
        delta =    0.3500
           m1 =    1.8600
           m2 =    2.2100
          sd1 =    1.2900
          sd2 =    1.8000

Estimated sample sizes:

            N =       498
  N per group =       249
```

Need really large n due to tiny effect size assumed

Power Analysis: Two-sample t-test

- Plot power against sample size with `graph`

```
power twomeans 1.86 2.21 , sd1(1.29) sd2(1.80) oneside
                              power(0.1(0.1)0.9) graph
```

- STATA Results:

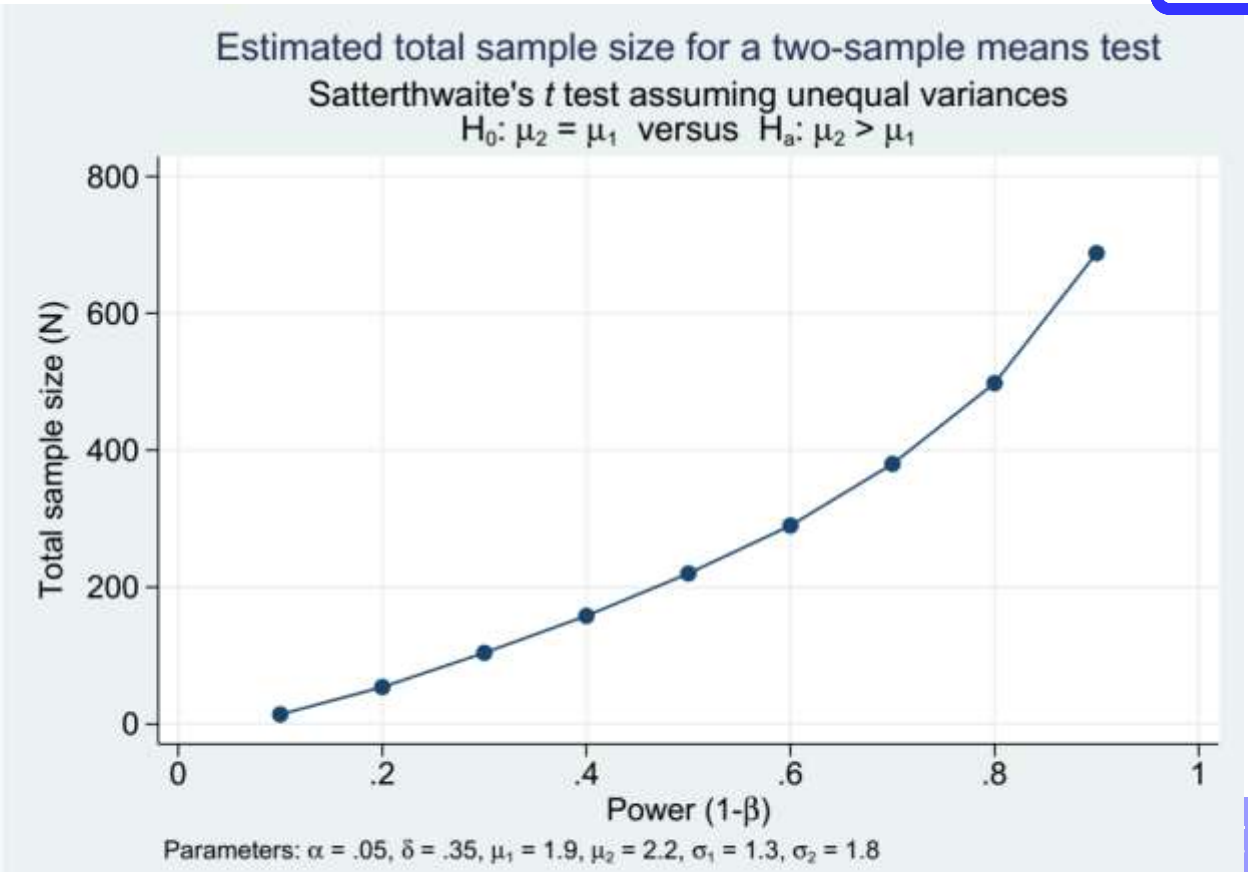

Larger total sample size yields higher power

Power Analysis: Equality of Variance Test

- Test $H_0 : \sigma_1^2 = \sigma_2^2$

```
sdtest v_mug, by(v_type)
```

- STATA Results:

Significant difference b/w var(WTP) and var(WTA)!

```
Variance ratio test
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
     WTP |      49    1.862245     .184379    1.290653    1.491526    2.232964
     WTA |      51    2.213725    .2515679    1.796554    1.708436    2.719015
---------+--------------------------------------------------------------------
combined |     100      2.0415    .1571287    1.571287    1.729723    2.353277
------------------------------------------------------------------------------
    ratio = sd(WTP) / sd(WTA)                                     f =   0.5161
Ho: ratio = 1                                    degrees of freedom =   48, 50

    Ha: ratio < 1               Ha: ratio != 1                 Ha: ratio > 1
  Pr(F < f) = 0.0114         2*Pr(F < f) = 0.0229          Pr(F > f) = 0.9886
```

Power Analysis: Equality of Variance Test

- What is the power of this test?

```
power twovariances 1.66 3.24, n1(49) n2(51)
```

$1.29^2 = 1.66,$
$1.80^2 = 3.24$

- STATA Results:

```
Estimated power for a two-sample variances test
F test
Ho: v2 = v1  versus  Ha: v2 != v1

Study parameters:

        alpha =     0.0500
            N =        100
           N1 =         49
           N2 =         51
        N2/N1 =     1.0408
        delta =     1.9518
           v1 =     1.6600
           v2 =     3.2400

Estimated power:

        power =     0.6402
```

Larger sample needed to get desired power of 0.8...

Power Analysis: Equality of Variance Test

- What is the sample size to get power $\pi = 0.80$?

```
power twovariances 1.66 3.24, p(0.8)
```

- STATA Results:

$1.29^2 = 1.66,$
$80^2 = 3.24$

```
Performing iteration ...

Estimated sample sizes for a two-sample variances test
F test
Ho: v2 = v1  versus  Ha: v2 != v1

Study parameters:

        alpha =    0.0500
        power =    0.8000
        delta =    1.9518
           v1 =    1.6600
           v2 =    3.2400

Estimated sample sizes:

            N =       146
  N per group =        73
```

Need $n = 73$!

WTP - WTA Gap: Paired Sample t-test

- $H_0 : WTA = WTP$ vs. $H_1 : WTA > WTP$

```
ttest WTA=WTP
```

- STATA Results:

1-sided p-value = 0.0096!

```
Paired t test
------------------------------------------------------------------------------
Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
     WTA |     100      2.5745    .0961546    .9615459    2.383708    2.765292
     WTP |     100      2.2415    .1115095    1.115095    2.020241    2.462759
---------+--------------------------------------------------------------------
    diff |     100        .333    .1398705    1.398705    .0554666    .6105334
------------------------------------------------------------------------------
     mean(diff) = mean(WTA - WTP)                                 t =   2.3808
 Ho: mean(diff) = 0                              degrees of freedom =       99

 Ha: mean(diff) < 0           Ha: mean(diff) != 0           Ha: mean(diff) > 0
 Pr(T < t) = 0.9904         Pr(|T| > |t|) = 0.0192          Pr(T > t) = 0.0096
```

Power Analysis: Paired Sample t-test

- Is there sufficient power to test WTA vs. WTP?

```
power pairedmeans 2.57 2.24 , sd1(0.96) sd2(1.11) corr(0)
                                        n(100) oneside
```

- STATA Results:

```
Estimated power for a two-sample paired-means test
Paired t test
Ho: d = d0  versus  Ha: d < d0

Study parameters:

        alpha =    0.0500             ma1 =    2.5700
            N =       100             ma2 =    2.2400
        delta =   -0.2249             sd1 =    0.9600
           d0 =    0.0000             sd2 =    1.1100
           da =   -0.3300            corr =    0.0000
         sd_d =    1.4675

Estimated power:

        power =    0.7219
```

Almost 0.8, but not yet!

Power Analysis: Paired Sample t-test

- What is the sample size to get power $\pi = 0.80$?

```
power pairedmeans 2.57 2.24 , sd1(0.96) sd2(1.11)
                              corr(0) p(0.80) oneside
```

- STATA Results:

```
Performing iteration ...

Estimated sample size for a two-sample paired-means test
Paired t test
Ho: d = d0  versus  Ha: d < d0
Study parameters:

        alpha =    0.0500             ma1 =    2.5700
        power =    0.8000             ma2 =    2.2400
        delta =   -0.2249             sd1 =    0.9600
           d0 =    0.0000             sd2 =    1.1100
           da =   -0.3300            corr =    0.0000
         sd_d =    1.4675

Estimated sample size:

            N =       124
```

$n = 124$ is slightly above 100

Power Analysis: Paired Sample t-test

- Is correlation truly 0?

```
corr WTA WTP
```

  - STATA Results:

```
(obs=100)

             |      WTA      WTP
-------------+------------------
         WTA |   1.0000
         WTP |   0.0987   1.0000
```

Should set correlation = +0.1...

Power Analysis: Paired Sample t-test

- What is the sample size to get power $\pi = 0.80$?

```
power pairedmeans 2.57 2.24 , sd1(0.96) sd2(1.11)
                         corr(0.10) p(0.80) oneside
```

- STATA Results:

```
Performing iteration ...

Estimated sample size for a two-sample paired-means test
Paired t test
Ho: d = d0  versus  Ha: d < d0

Study parameters:

        alpha =    0.0500            ma1 =    2.5700
        power =    0.8000            ma2 =    2.2400
        delta =   -0.2369            sd1 =    0.9600
           d0 =    0.0000            sd2 =    1.1100
           da =   -0.3300           corr =    0.1000
         sd_d =    1.3930

Estimated sample size:

            N =       112
```

Only need $n = 112$!

Acknowledgment

- This presentation is based on
  - Section 1.1-1.4 of the lecture notes on Experimetrics,
- prepared for a mini-course taught by Peter G. Moffatt (UEA) at National Taiwan University in Spring 2019