Experimetrics: Power Analysis
實驗計量：統計檢定力分析

Joseph Tao-yi Wang (王道一)
Special Lecture, EE2-Replication

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: $\Pr(\text{confirm} \mid \text{infected patient})$
  - 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: $\alpha = \Pr(\text{reject null} \mid \text{null is true})$

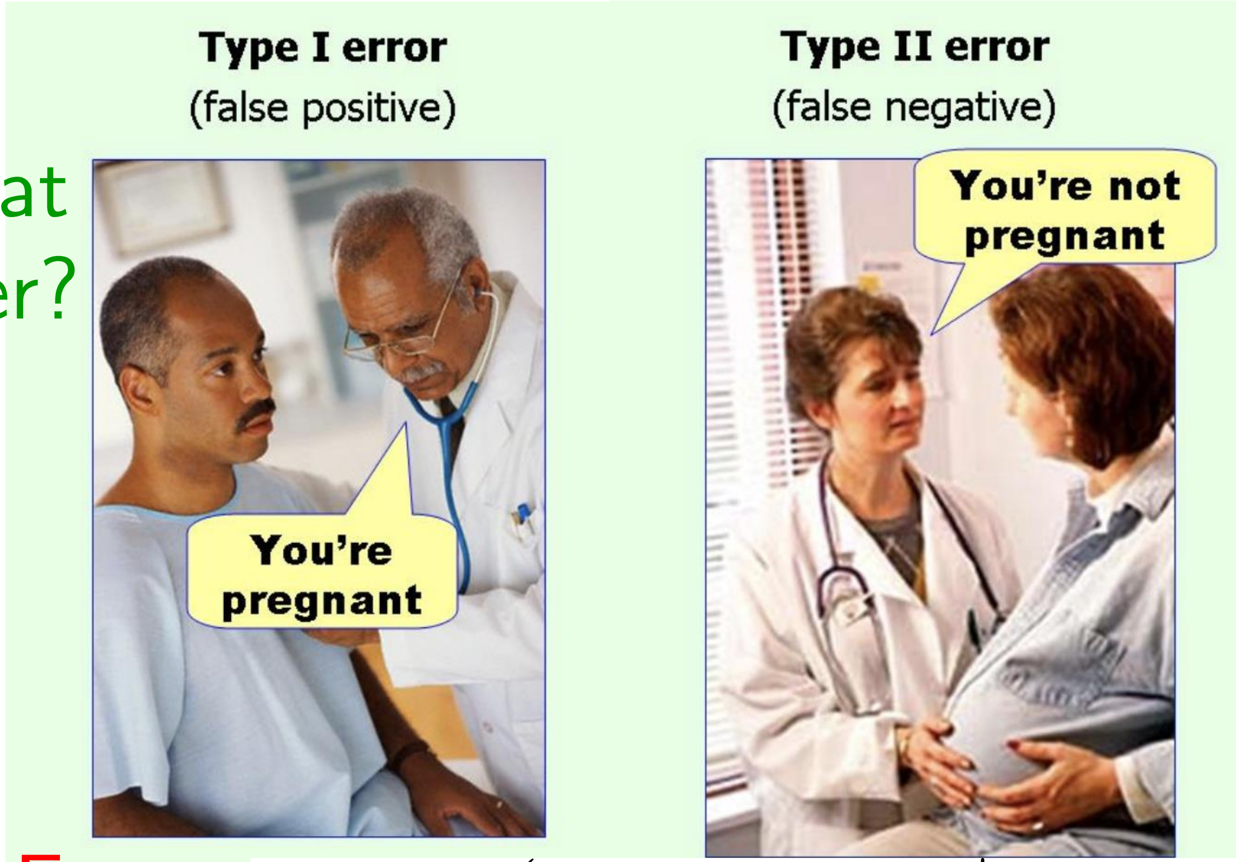

- But what is Power?

- 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

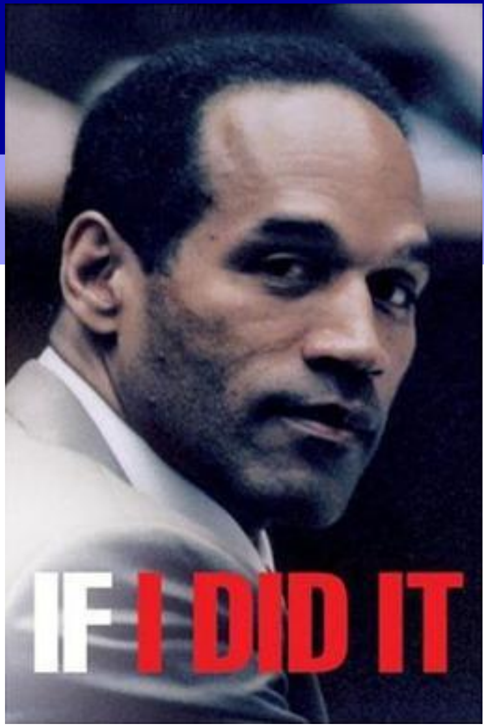

- How big can we allow Type 1 Error
- To convict a crime suspect,
  - Null Hypothesis: Not Guilty
  - Alternative Hypothesis: Guilty
  - Type 1: $\alpha = \Pr(\text{convict} \mid \text{innocent suspect})$
  - Type 2: $\beta = \Pr(\text{acquit} \mid \text{guilty suspect})$
- Type 1 Error more serious 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:

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

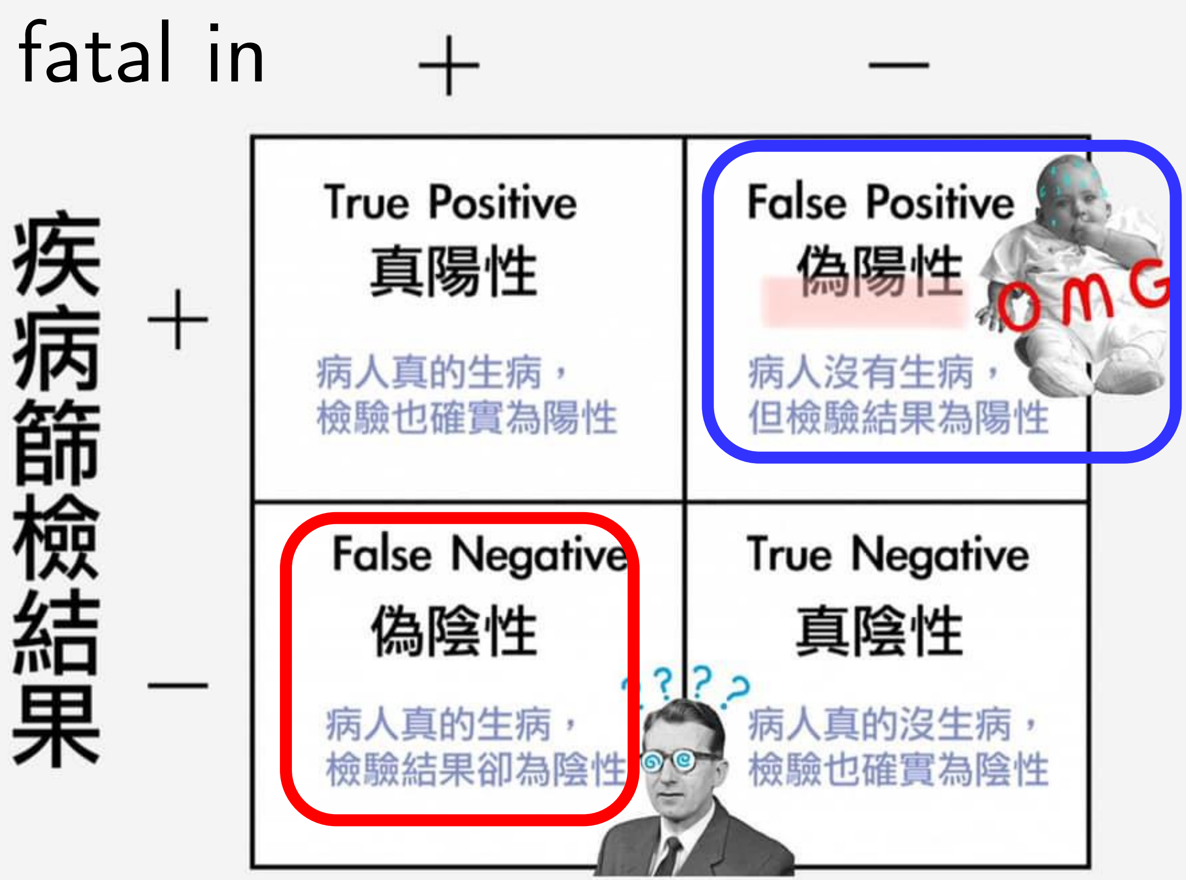

Treatment Testing Toolkit

- 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
  - 2^{nd} 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

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

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

$\mu_0 = 10$

$\mu_1 = 12$

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

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

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

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

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

Power Analysis: How Big Should n Be?

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

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

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

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

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

$s = 5$

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

$\mu_0 = 10$

$\mu_1 = 12$

$$= \underline{38.66}$$

- So we need $n \geq 39$

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

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

- sample std; 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:

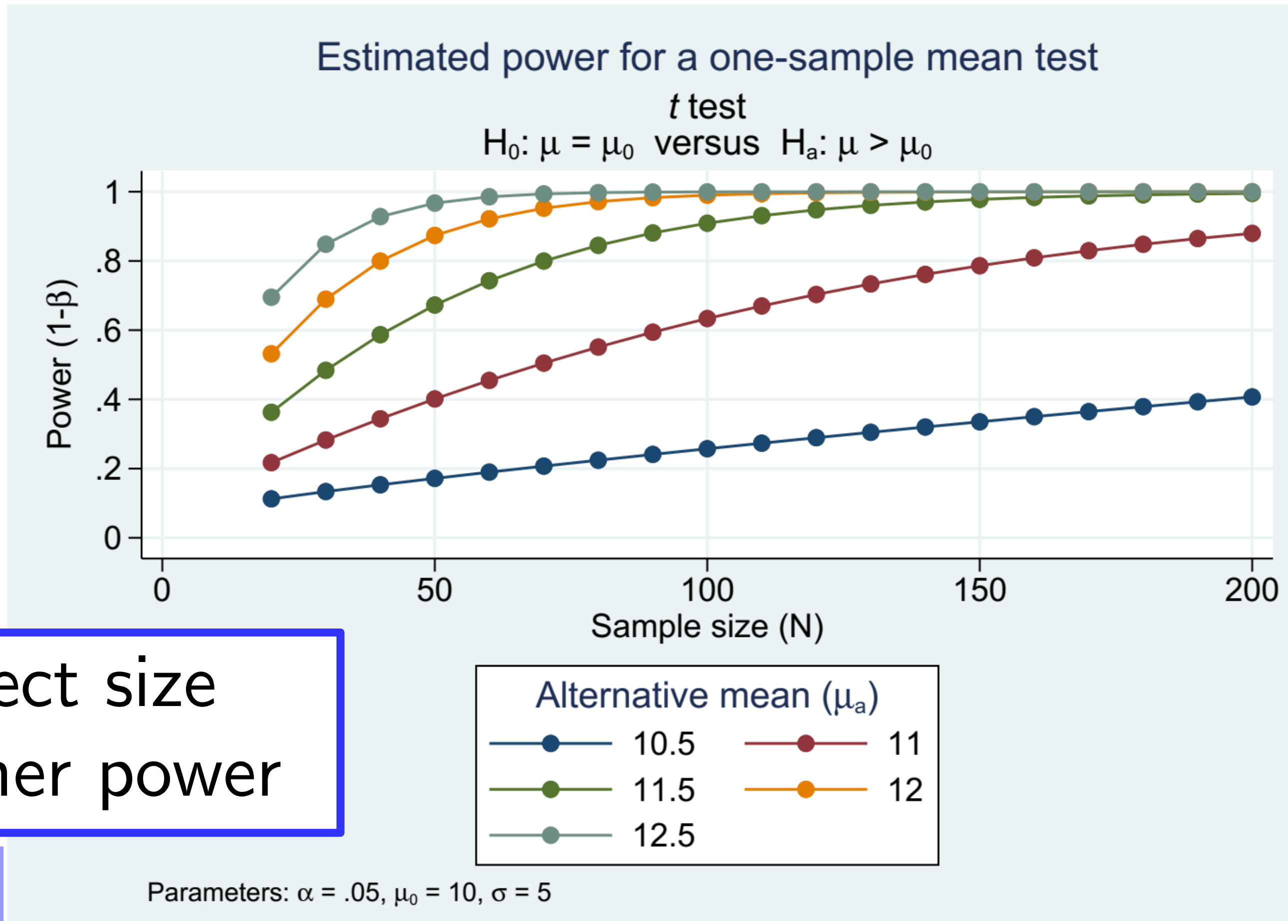

Larger effect size yields higher power

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: