Day 4, Morning: Measures of Association with Nominal Variables

Instructor: Marshall A. Taylor
844 Flanner Hall
mtaylo15@nd.edu
marshalltaylor.net

Recap

- Up to this point we have been talking about **statistical significance.**

- A result is statistically significant to the extent that, over repeated observations, it is not likely due to random chance at some α level.

On to the substantive

- However, statistical significance should never be mistaken for **substantive significance**.

- Even if the result is not likely due to random chance, does the result amount to much in the real world?

- How strongly related are the variables, given that they are statistically dependent?

On to the substantive

- What does this independent samples difference-of-proportions test tell us about the relationship between the chances that a candidate will engage in negative campaigning and their being a Democrat or Republican?

```
. prtest GONEG, by(DEMOCRAT2)

Two-sample test of proportions               Republican: Number of obs =      377
                                               Democrat: Number of obs =      352

------------------------------------------------------------------------------
    Variable |       Mean   Std. Err.       z     P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  Republican |   .4217507    .025434                         .3719009    .4716004
    Democrat |   .5397727   .0265656                          .487705    .5918404
-------------+----------------------------------------------------------------
        diff |  -.1180221    .036778                        -.1901056   -.0459385
             |  under Ho:   .0370253    -3.19   0.001
------------------------------------------------------------------------------
        diff = prop(Republican) - prop(Democrat)                  z =  -3.1876
    Ho: diff = 0

    Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
 Pr(Z < z) = 0.0007         Pr(|Z| > |z|) = 0.0014          Pr(Z > z) = 0.9993
```

On to the substantive

- We can say that there is less than a 1% chance (α = .01) that we would observe a difference in proportions of at least .11 if the real difference was zero. These proportions for Republicans and Democrats are far enough apart—and our sample size is large enough—to say that this result is not likely due to random chance.

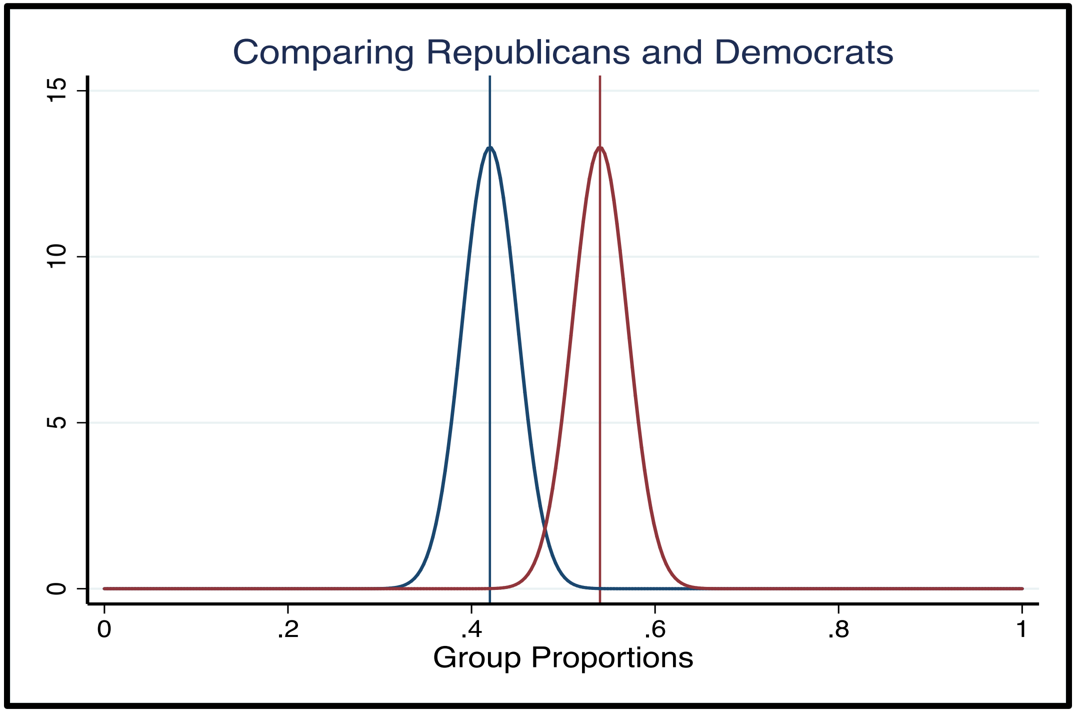

But…

- Is the difference *really* that big?

```
. prtest GONEG, by(DEMOCRAT2)

Two-sample test of proportions                  Republican: Number of obs =      377
                                                  Democrat: Number of obs =      352

    Variable |       Mean   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  Republican |   .4217507    .025434                      .3719009    .4716004
    Democrat |   .5397727   .0265656                       .487705    .5918404
-------------+----------------------------------------------------------------
        diff |  -.1180221    .036778                     -.1901056   -.0459385
             |  under Ho:   .0370253    -3.19   0.001
------------------------------------------------------------------------------
        diff = prop(Republican) - prop(Democrat)                  z =  -3.1876
    Ho: diff = 0

    Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
 Pr(Z < z) = 0.0007         Pr(|Z| > |z|) = 0.0014          Pr(Z > z) = 0.9993
```

Important Difference!

- We can say that Republicans and Democrats likely differ in the population of congressional candidates in their proclivity to “go negative,” but we can’t just assume that this means the difference is a big one!

- In some contexts, a difference in 12 percentage points may be “big” and in others “small.”

Measures of Association

- So how do we quantify **substantive significance**?

- How do we know how strongly related two variables might be in the population, given that there is a statistically significant relationship between the two?

- This is the purpose of **measures of association**.

Measures of Association

- A measure of association is a numerical representation of the strength of the relationship between variables.

- Like inferential tests (t-tests, chi-square tests, etc.), the appropriate measure of association is entirely dependent on how the variables are measured.
  - Nominal?
  - Ordinal?
  - Continuous?

Types of Association

- Since a measure of association necessarily implies more than one variable, there are different statistics for different level-of-measurement pairings.

	Nominal	Ordinal	Continuous
Nominal	Relative risks, phi, Cramér's V, odds ratios, tetrachoric correlation		
Ordinal	Gamma, Kendall's tau-b	Gamma, Kendall's tau-b, Spearman's rho, polychoric correlation	
Continuous	Pearson's correlation (when dichotomous), Cohen's *d*, regression slopes, R^2	Pearson's correlation (if ordinal variable has many categories), regression slopes, R^2	Pearson's correlation

Types of Association

- Today we are going to focus specifically on measures of association for nominal-nominal relationships.

	Nominal	**Ordinal**	**Continuous**
Nominal	**Relative risks, phi, Cramér's V, odds ratios, tetrachoric correlation**		
Ordinal	Gamma, Kendall's tau-b	Gamma, Kendall's tau-b, Spearman's rho, polychoric correlation	
Continuous	Pearson's correlation (when dichotomous), Cohen's *d*, regression slopes, R^2	Pearson's correlation (if ordinal variable has many categories), regression slopes, R^2	Pearson's correlation

A Disclaimer

- Unlike the inferential tests we discussed, there is a lot more gray area when it comes to “picking” an “appropriate” measure of association.

- Ideally, measures of association for the same types of relationships will yield similar results—albeit perhaps scaled differently.

Nominal-Nominal Measures

- We will cover a sample of four nominal-nominal measures of association (though there are many more):
  - Relative risks
  - Odds ratios
  - Phi
  - Cramér's V

Relative Risks

- Relative risks are exactly what they sound like: the risk of person A experiencing an event relative to person B experiencing that event.

- More formally, the relative risk is the ratio of two probabilities: the probability of person A experiencing event E and the probability of person B experiencing event E:

$$\mathrm{RR}(E) = \frac{P(A)}{P(B)}$$

An Example: Relative Risks

- So if the $\mathrm{RR}(E) = P(A)/P(B)$, then what is the relative risk of being comfortable with allowing homosexuals to teach at a college/university between those who are also comfortable with racists teaching at a college/university and those who are not?

Key
frequency *column percentage*

allow homosexual to teach	allow racist to teach: allowed	allow racist to teach: not allow	Total
allowed	683 94.08	637 83.71	1,320 88.77
not allowed	43 5.92	124 16.29	167 11.23
Total	726 100.00	761 100.00	1,487 100.00

Pearson chi2(1) = 40.0889 Pr = 0.000

An Example: Relative Risks

$$\mathrm{RR}(E) = \frac{P(A)}{P(B)} = \frac{.941}{.837} = 1.124$$

Key
frequency
column percentage

allow homosexual to teach	allow racist to teach: allowed	allow racist to teach: not allow	Total
allowed	683 94.08	637 83.71	1,320 88.77
not allowed	43 5.92	124 16.29	167 11.23
Total	726 100.00	761 100.00	1,487 100.00

Pearson chi2(1) = 40.0889 Pr = 0.000

An Example: Relative Risks

- Those who are comfortable with a racist teacher are about 12% more likely to be comfortable with a homosexual teacher.

Key
frequency
column percentage

allow homosexual to teach	allow racist to teach: allowed	allow racist to teach: not allow	Total
allowed	683 94.08	637 83.71	1,320 88.77
not allowed	43 5.92	124 16.29	167 11.23
Total	726 100.00	761 100.00	1,487 100.00

Pearson chi2(1) = 40.0889 Pr = 0.000

That pesky ratio!

- "Wait… why 12% more likely and not 112% more likely? After all, the relative risk was 1.124!"
- Note that the relative risk is a *ratio* of proportions —**so the number reflecting no difference is 1**.
- If your relative risk is **greater than 1**, then the group in your **numerator is more likely** to experience the event. If you relative risk is **less than 1**, then the group in your denominator **is more likely** to experience the event.

That pesky ratio!

- We can take our same example from before to illustrate this point.

Key
frequency
column percentage

allow homosexual to teach	allow racist to teach: allowed	allow racist to teach: not allow	Total
allowed	683 94.08	637 83.71	1,320 88.77
not allowed	43 5.92	124 16.29	167 11.23
Total	726 100.00	761 100.00	1,487 100.00

Pearson chi2(1) = 40.0889 Pr = 0.000

That pesky ratio!

$$\mathrm{RR}(E) = \frac{P(B)}{P(A)} = \frac{.837}{.941} = .88$$

Key: *frequency*, *column percentage*

allow homosexual to teach	allow racist to teach: allowed	allow racist to teach: not allow	Total
allowed	683 94.08	637 83.71	1,320 88.77
not allowed	43 5.92	124 16.29	167 11.23
Total	726 100.00	761 100.00	1,487 100.00

Pearson chi2(1) = 40.0889 Pr = 0.000

Just switch the numerator and denominator.

That pesky ratio!

- Those who are comfortable with racists teaching are 12% more likely to be comfortable with homosexuals teaching than those who are not comfortable with racists teaching ($\mathrm{RR}(E) = 1.12$).
- Those who are not comfortable with racists teaching are 12% less likely (or 88% *as* likely) to be comfortable with homosexuals teaching than those who are comfortable with racists teaching ($\mathrm{RR}(E) = .88$).
- **They mean the same thing—just the dividend and divisor have flipped positions in the equation!**

Another Example: Relative Risks

- Let's try one more. What is the relative risk of giving the correct answer on a probability test (correct answer is "no") between those who took at least one college-level science course and those who did not?

Key
frequency
column percentage

test of knowledge about probablity 1	r has taken any college-level sci course yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

Another Example: Relative Risks

$$\mathrm{RR}(E) = \frac{P(A)}{P(B)} = \frac{.936}{.856} = 1.093$$

Key
frequency *column percentage*

test of knowledge about probablity 1	r has taken any college-level sci course yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

Another Example: Relative Risks

- Those who took a college-level science course are about 9% more likely to give the correct answer.

Key
frequency
column percentage

test of knowledge about probablity 1	r has taken any college-level sci course yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

Another Example: Relative Risks

- Alternative (but analogous) interpretation: Those that did not take a college-level science course are about 91% as likely (1 – .09) as those that id to give the correct answer.

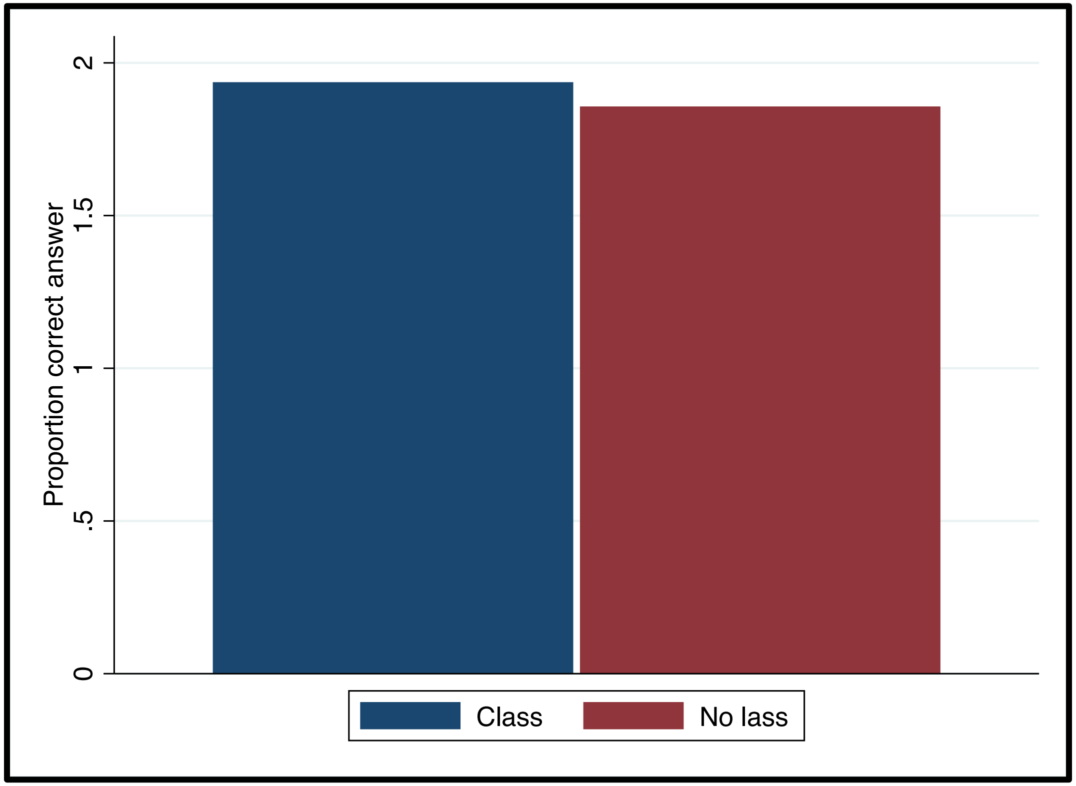

Odds Ratios

- Odds ratios are very similar to relative risks in that they focus on the substantive significance between two particular values in the table.
- Unlike relative risks, which are ratios of two proportions, the odds ratio is the **ratio of (you guessed it) two odds**.

Odds Ratios

- The odds ratio between Y_1 and Y_2 on X_1 can be found with the following:

$$\mathrm{OR}(Y \, \mathrm{on} \, X_1) = (\frac{P(A)}{P(C)})/(\frac{P(B)}{P(D)})$$

	Y_1	Y_2	**Margins**
X_1	a	b	a+b
X_2	c	d	c+d
Margins	a+c	b+d	n

An Example: Odds Ratios

- Let's go back to our scientific knowledge and probability test question from a moment ago. Recall that the χ^2-statistic is statistically significant (p < .001), so examining the substantive significance is warranted.

Key: *frequency* / *column percentage*

test of knowledge about probablity 1	r has taken any college-level sci course: yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

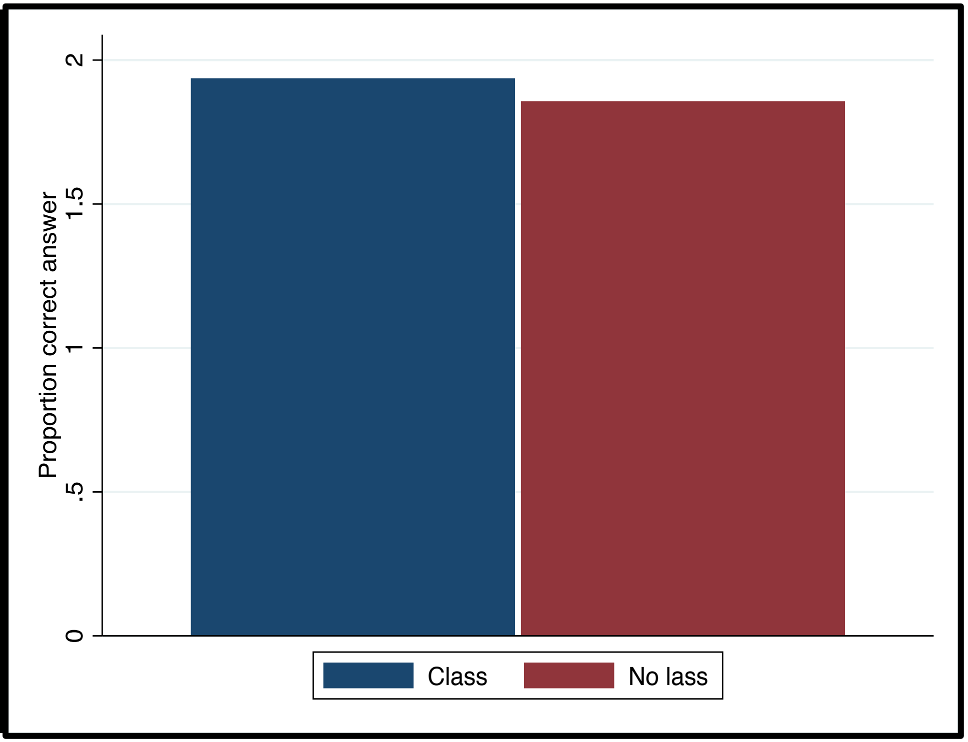

An Example: Odds Ratios

- Looking at those who answered the probability question correctly (“no”), what is the odds ratio between those who have taken a college-level science course and those who have not? What does it mean?

Key
frequency   *column percentage*

test of knowledge about probablity 1	r has taken any college-level sci course yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

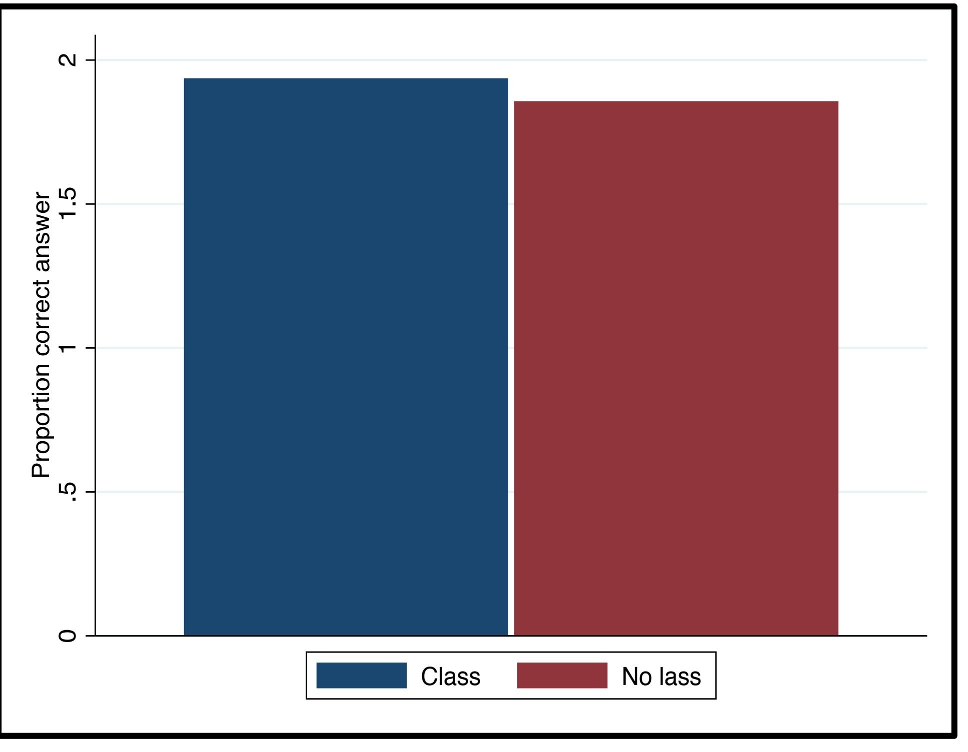

An Example: Odds Ratios

$$\text{OR} = \left(\frac{P(C)}{P(A)}\right) / \left(\frac{P(D)}{P(B)}\right) = \left(\frac{.94}{.64}\right) / \left(\frac{.86}{.14}\right) = 2.55$$

Key

frequency
column percentage

test of knowledge about probablity 1	r has taken any college-level sci course: yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

An Example: Odds Ratios

- The odds of a person who took a college-level science course getting the correct answer is about 2.5 times greater than the odds of a person who did not take a college-level science course.

Key
frequency *column percentage*

test of knowledge about probablity 1	r has taken any college-level sci course yes	no	Total
yes	31 6.40	95 14.37	126 11.00
no	453 93.60	566 85.63	1,019 89.00
Total	484 100.00	661 100.00	1,145 100.00

Pearson chi2(1) = 18.1100 Pr = 0.000

That pesky ratio again…

- A value of 1 takes on the same property in an odds ratio as it does in a relative risk.

- If your relative risk is 1:
  - Both groups are just as likely to experience X.

- If your odds ratio is 1:
  - The odds for both groups are the same on X.

That pesky ratio again…

- For example, remember that our odds ratio of those who took the course to those who did not was 2.55. We can flip the numbers to get an odds ratio of those who did **not** take the course to those who did. It turns out to be .39.

- Interpretation. The odds of a person who did not take the science course is only .39 times the odds of a person who did take the course.

- Odds ratios, like relative risks, have a lower bound (0), but not an upper bound!

Phi

- A measure of association between two dichotomous variables (a 2X2 table).

- Unlike the relative risk and odds ratio, phi is a single number that represents association in the table **as a whole** (i.e., not just between two values in the table).

Phi

- Can be found by (1) finding the difference between the products of cells in the inner-diagonal cells, then (2) dividing this by the square root of the product of the four marginal cells*:

$$\phi = \frac{(a \times d) - (b \times c)}{\sqrt{(a+b)(c+d)(a+c)(b+d)}}$$

	Y_1	Y_2	**Margins**
X_1	a	b	a+b
X_2	c	d	c+d
Margins	a+c	b+d	n

Formula and graph presentation from Richard Liu, "A Note on Phi-Coefficient Comparison." *Research in Higher Education* 13(1):3-8, 1980.

Phi

- Multiply these…

	Y_1	Y_2	**Margins**
X_1	a		
X_2		d	
Margins			

Phi

- Then these…

	Y_1	Y_2	**Margins**
X_1		**b**	
X_2	**c**		
Margins			

Phi

- Add up the two products **ad** and **bc**…

	Y_1	Y_2	**Margins**
X_1			
X_2			
Margins			

Phi

- Then divide this number by the square root of the product of the margins to get ϕ!

	Y_1	Y_2	**Margins**
X_1			**a+b**
X_2			**c+d**
Margins	**a+c**	**b+d**	

Phi

- ϕ ranges between –1 and +1, where –1 indicates a **strongly negative relationship** between the variables and +1 indicates a **strongly positive relationship**.
- Some yardsticks*:
  - –1 to –.7: strong weak association
  - –.7 to –.3: weak negative association
  - –.3 to .3: little/no association
  - .3 to .7: weak positive association
  - .7 to 1: strong positive association

Value range recommended on pmean.com (http://www.pmean.com/definitions/phi.htm).

An Example: Phi

- The relationship between sex and the belief that it should be possible for a woman to get a legal abortion if she is raped is statistically significant. How strong is the relationship?

```
. tab sex abrape, chi2

           | pregnant as result of
respondent |         rape
     s sex |       yes         no |     Total
-----------+----------------------+----------
      male |       595        131 |       726
    female |       681        222 |       903
-----------+----------------------+----------
     Total |     1,276        353 |     1,629

     Pearson chi2(1) =  10.1429   Pr = 0.001
```

An Example: Phi

$$\phi = \frac{(a \times d) - (b \times c)}{\sqrt{(a+b)(c+d)(a+c)(b+d)}} = \frac{(595 \times 222) - (131 \times 681)}{\sqrt{726 \times 903 \times 1,276 \times 353}} = .079$$

```
. tab sex abrape, chi2

           | pregnant as result of
respondent |         rape
     s sex |       yes         no |     Total
-----------+----------------------+----------
      male |       595        131 |       726
    female |       681        222 |       903
-----------+----------------------+----------
     Total |     1,276        353 |     1,629

          Pearson chi2(1) =  10.1429   Pr = 0.001
```

An Example: Phi

- So our result is statistically significant at the .01 level, but the strength of the association between sex and opinions on abortion after rape is only .079. This is a good example of how and why statistical significance **is not** substantive significance!

```
. tab sex abrape, chi2
```

respondents sex	pregnant as result of rape: yes	no	Total
male	595	131	726
female	681	222	903
Total	1,276	353	1,629

Pearson chi2(1) = 10.1429 Pr = 0.001

An Example: Phi

Confirm with Stata*:

```
. tab sex abrape, chi2 V

           | pregnant as result of
respondent |         rape
     s sex |       yes         no |     Total
-----------+----------------------+----------
      male |       595        131 |       726
    female |       681        222 |       903
-----------+----------------------+----------
     Total |     1,276        353 |     1,629

          Pearson chi2(1) =  10.1429   Pr = 0.001
                   Cramér's V =   0.0789
```

*In a 2X2 table, Cramér's V becomes phi.

Cramér's V

- Also a measure of association between nominal variables—but mostly used for >2X2 tables.

- Phi is the same value as Cramér's V when computed on a 2X2 table—that's why Stata just reports the latter.

Cramér’s V

- The formula for Cramér’s V is as follows:

$$V = \sqrt{\frac{\phi^2}{\min(r-1, c-1)}}$$

- It is simply the square root of phi-squared (ϕ^2), divided by the the smallest number of categories from the row or column variable (minus 1).
- Result is bound between 0 (total independence) and 1 (total dependence—meaning variation in one variable is completely determined by the other).

An Example: Cramér's V

- In 2004, there was a statistically significant relationship between opinions on which U.S. party was better equipped to handle the War on Terrorism and which party was better equipped to handle the nation's economy. How strong was this relationship?

```
. tab prtyterr prtyecon if prtyterr!=9 & prtyecon!=9, chi2 V
```

Which party do you think would do a better job of HANDLING THE WAR ON TERRORISM.	Which party do you think would do a better job of HANDLING THE NATION'S ECONOMY. Democrats	Republica	Wouldn't	Don't kno	Total
Democrats	230	5	57	6	298
Republicans	50	252	171	7	480
About the same by bot	148	26	203	9	386
Neither party	5	2	3	0	10
Don't know	5	3	12	14	34
Total	438	288	446	36	1,208

```
Pearson chi2(12) = 732.3495   Pr = 0.000
      Cramér's V =   0.4495
```

An Example: Cramér’s V

- Our Cramér’s V is .45, which indicates a moderate relationship. These means that, in 2004, opinions on which party was better equipped to handle issues on the War on Terror and the nation’s economy were moderately dependent on one another.

```
. tab prtyterr prtyecon if prtyterr!=9 & prtyecon!=9, chi2 V

   Which party do you |
     think would do a |
        better job of | Which party do you think would do a better
   HANDLING THE WAR ON|    job of HANDLING THE NATION'S ECONOMY.
           TERRORISM. | Democrats  Republica  Wouldn't   Don't kno |     Total
----------------------+--------------------------------------------+----------
            Democrats |       230          5         57          6 |       298
          Republicans |        50        252        171          7 |       480
  About the same by bot|      148         26        203          9 |       386
        Neither party |         5          2          3          0 |        10
           Don't know |         5          3         12         14 |        34
----------------------+--------------------------------------------+----------
                Total |       438        288        446         36 |     1,208

          Pearson chi2(12) = 732.2495   Pr = 0.000
                   Cramér's V =   0.4495
```

Conclusion

- Measures of association are used to get a grasp on the substantive significance of a relationship after the relationship has been identified as statistically significant.

- The appropriate choice of measure—like the appropriate choice of inferential test—is dependent on the level of measurement of the variables.

Conclusion

- This morning we talked about measures of association between nominal variables.

- Relative risks and odds ratios are useful for comparing the differences between individual cells in a table. As such, they are useful for making more nuanced claims and can be used regardless of the size of the table.

Conclusion

- Phi (ϕ) is a measure of association between two dichotomous variables. It summarizes the strength of the relationship in the whole table, and ranges from –1 (strong negative association) to 1 (strong positive association).

- Cramér's V is an extension of phi that works with non-dichotomous nominal variables (i.e., tables that are bigger than 2X2). It ranges from 0 (no association) to 1 (perfect association).

Conclusion

- Tomorrow morning we will talk about measures of association between ordinal and continuous variables.

Datasets Used

- ANES 2004 (http://thearda.com/Archive/Files/Descriptions/NES2004.asp).
- Druckman, James, Michael Parkin, and Martin Kifer. 2013. *Congressional Candidate Websites*. ICPSR-34895-v1. Ann Arbor, MI: Inter-University Consortium for Political and Social Research. Retrieved February 10, 2015 (http://doi.org/10.3886/ICPSR34895.v1).
- GSS 2014 (http://gss.norc.org/get-the-data/stata).