

Research Methods in International Relations and Political Science

Focus Questions for Unit 3

(Last updated: 11/2/2024; Any material that is changed after I initially post this guide will be bolded in revised editions)

Some ideas to think about before read through this guide

- This study guide is different than the other ones you have reviewed this term because it is very much focused on understanding and interpreting many of the statistical techniques most commonly used in economics, international relations, and political science. Because this is completely unfamiliar material for many students, this guide is longer because it often explains concepts rather than just giving you focus questions to answer on your own.

As you read through the guide the first time, you will *not* be familiar with many of the italicized concepts at the start of the unit. Everything below will be covered *before* your test, and you will notice that many of the questions I want you to focus on also include an answer.

- ***Don't let all of the specifics in this guide intimidate you.*** The sections below on “linear” and “logistic” regression mostly involve application examples, so they are precise and ask you to answer questions with hypothetical regression output. You will recall from your last unit test, that some of the multiple-choice items asked you to either identify what statistical method was appropriate for a given type of situation or to interpret a statistical finding. That is what some of the items below do.

These types of study guide questions that require you to do some calculations are meant to help you see if you are ready to interpret the types of statistical findings that you will find in PSC upper-division courses and most outlets for original research in economics, international relations, and political science. I also want you to know what statistical tests and other analytical options are available to learn more about why some people (or countries, groups, politicians, etc.) do or think certain things, while others do not. This type of work is what most specialists in international relations and political science get paid well to do.

You also will practice answering these types of specific questions in low-stake BlackBoard assignments before you ever see them on an exam. If you find yourself struggling with a block of questions that cover a topic *already reviewed* in class and you *already have carefully* reviewed the relevant supporting materials in the course schedule and your textbook, see your instructor for guidance.

- In the big picture: What I want you to understand best at the end of this unit is when, why, and how to use and interpret results that involve these six statistical methods: (1) t-tests, (2) correlation, (3) linear regression, (4) statistical controls, (5) the interpretation of interval, dummy, and interaction variables, and (6) logistic regression
- What *must* you be able to do for the final test? You should assume that you will have required essay items that 1) give you SPSS output for a linear regression model and ask you to interpret the appropriate R-square statistic, 2) ask you to do the same thing, but with a logistic regression model, and 3) give you the summary of a study and ask you to identify the study's theory, hypotheses, dependent variable/s, independent variable/s, control variable/s, and an intervening variable if there is one.

Statistical significance

- Most of the material covered in this unit involves “statistical tests” that tell us whether any finding we are seeing in a sample (in most cases, the results of a survey that had just a few thousand participants) could be expected to be found in repeated analyses of similar groups of individuals.
- Critical to know: Having a statistical test reveal a *statistically* significant finding is not the same as having a test uncover a *substantively* important finding. When are we most likely to find that a seemingly large and substantively important result in our sample is not statistically significant? Answer: When you are looking at a group for which there are few observations.
- And when are you most likely to find a statistically significant finding that has little substantive value

(answer: if you have a very large sample—the Cooperative Congressional Electoral Survey, for example, has tens of thousands of people in it—even a very slight difference in means or a tiny regression coefficient will generally be statistically significant)?

- T-test statistics, association measures like correlation, and regression coefficients are all reported with a corresponding *probability statistic* (i.e. a “p-value,” which is often listed in SPSS output in the “sig.” column). P values range from .000 to almost one. What exactly does the P-value measure for each type of statistic just listed? Hint: a one-sample T-test compares whether the means for two of a variable’s subgroups are the same for another variable (what’s an example of that from your SPSS work?). A one-sample T-test can be used to compare a group or subgroup’s mean to any value (what are some examples of test values against which we might test the mean of a group?).
- We interpret p-values to determine whether a finding we see for a sample is statistically significant. Sometimes SPSS labels p-values as “p-values” and sometimes, it labels them as “sig.” Here are the kind of p-values we might see: .012, .453, .000, <.001, .002, .054, and .050.
- Hint: the easiest way to interpret a p-value is to think of it as a percentage. A p-value of .054 means there is over a five percent chance that a second, comparable sample would reach a finding that runs in the opposite direction from the one seen in the sample. So, let’s say, we have a .101 p-value on a t-test that has found that women’s average GPA is higher than men’s. We would reject the hypothesis that women’s grades are higher because over 10% of the time, we would expect that repeated sampling would find that men’s grades are higher. A p-value of .001 means that there is less than .1% = 1/1000th of a chance that a second, comparable sample would have a different finding.
- To make tables easier to read, it is a common practice to place significance stars next to statistics (e.g., a coefficient of .326^{***}) to denote those that attain statistical significance instead of putting this information into an additional table common. Be able to interpret starred statistics assuming that * if $p < .05$, ** if $p < .01$ and ***if $p < .001$.

Using t-tests to see if two means or proportions differ

- What is a “T-test”? What is the difference between a “one-sample T-test” and an “independent samples T-test”?

Which test would be appropriate if you want to know if the typical age of a Trump supporter is different from non-Trump voters, and you have a variable that is coded 1=Voted for Trump, 0=Voted for someone else or not Trump)?

Which T-Test test would let you know whether the share of Latinos in your particular survey sample is statistically different than the 18% reported by the US Census?

Which T-test would you use if you wanted to know if the mean age of a Trump supporter is statistically different than the mean age of an American male?

- Let’s make sure you know what you can do with both kinds of T-tests (you actually can do everything with a one-sample T-test if you want). What exactly is being tested with an “independent samples” T-test? What is being tested with a one-sample test (hint: you specify the value, so you could run a T-test on a dummy variable where males are coded one and non-males are coded 0, to see if the sample share of men is different from 50%)? And, if you split your data on a given variable (as an example gender), you then run a “one-sample” T-test to determine if the average score for men on another variable is different than the mean for everyone, for non-males, or for any other group for whom you have calculated an average score.
- A T-test is the first time you will be conducting a *statistical test* in this class. What is a statistical test doing in a T-test, and how do we interpret the statistics for whether a test is statistically reliable (or, put in a more general way, what does it mean for a result to be “statistically significant”)?
- What does it mean if an independent-samples T-test looking at the average income for the variable Male (where group 1 = males and group 0=non-males) says that the difference in means has a “p-value” of <.001, or .001, or .002? What does it mean if it says .012 or .010? What about .05, .025, or .089, or .321?

How do you compare the confidence intervals for means for a given variable people belonging to different

groups that involve separate variables? (Hint: they are reported in T-tests, which by default in SPSS will tell us what the mean is for a group and the range of means we would expect to see 95% of the time in repeated sampling).

For example, we might want to compare the averages for Republicans, Catholics, and women who have rated a US president on a 1-100 scale or see if truly different proportions of these groups approve of a president (i.e., coded 0=no; 1=yes). What role do “confidence intervals” around the means for different types of people signify in this context? Assuming the standard setting for confidence intervals (CIs), which is 95%, how would we interpret the following data:

We have a dummy variable called Support_Capital_Punishment_Yes.

-The mean for Catholics is .234 with an upper 95% CI of .255 and a lower of .194

-The mean for men is .534 with an upper 95% CI of .555 and a lower of .455

-The mean for Republicans is .544 with an upper 95% CI of .575 and a lower of .525

Testing for an association between two variables

- Most of this unit of the course examines how researchers use *dummy* (sometimes called “binary” or “dichotomous”) and *interval* (sometimes called “continuous”) variables to examine and test the statistical significance of relationships between variables. While there are statistical methods designed specifically for *categorical* (often called “nominal”) and *ordinal* variables, your textbook focuses exclusively on statistical methods involving dummy and interval variables.

The textbook’s focus on dummy and interval variables reflects the fact that most researchers choose to convert categorical variables into dummy variables before using analyses that involve statistical tests.

And researchers typically convert ordinal variables into dummy variables or treat the ordinal measures as interval variables, depending on their number of categories and frequency distributions.

Be able to define and give an example of each of the four types of variables *italicized* above. Also, make sure you understand the key ideas communicated in this paragraph (i.e., this block of the study guide). This paragraph explains why this course focuses on a lot fewer statistics and methods than the typical social science methods course.

- Before moving on to look at several types of bivariate analyses, here are a couple of more questions about variable structure: What about the distribution of an ordinal variable—and even sometimes what a measure that seems to look like an interval variable—would make a researcher consider collapsing that variable into a dummy variable or perhaps two of them?

Thinking about how household income and education are measured in the datasets we have worked with in class, why do researchers typically group and order the response categories for these types of interval variables so that each incremental gain/group doesn’t have the same number of units (i.e., the number of years of education or range of income often is not the same for each unit)? Are these new variables still interval variables? (Hint: Yes, they better interval variables because have been modified so that each one-unit increase on the measure as a substantively similar effect).

- There are numerous association measures that social scientists can use to analyze pairs of variables where *one or both measures are categorical (AKA nominal) or ordinal variables*. However, as explained above, political scientists and economists most often focus their statistical analyses on variables that have been transformed into dummy and interval variables (an exception is regression analysis designed to work with ordinal dependent variables, which is quite common, but that method is outside of the scope of this course).

The only association tests you learn anything about in this course are *correlations* and a *chi-squared (χ^2) tests*. What kind of variable pairs are examined by a chi-squared (χ^2) test? Hint, this is the most commonly used statistical test of whether a *pair of nominal (aka categorical) variables* are associated with one another, such that knowing something about a subject’s category for one variable tells us something about what category they will belong to for the other variable.

As an example of when we would use a chi-squared test, we might want to know whether a person’s preferred

political party has a statistically significant association with that same individual's preferred color when given five options. Using crosstab results (or splitting the dataset by party and running a frequency on the color question), we could make a bar chart or frequency table to visually compare if there are differences in the share of Democrats, Republicans, independents, and third-party adherents who prefer each color. However, it is the chi-squared test results that will tell us whether there is a statistically meaningful association between the categorical variables that might not be found in repeated sampling (i.e., the probability of there being no consistent association is higher than .05).

To calculate a *chi-squared test* in SPSS, we use Analyze -> Descriptive Statistics -> Crosstab (DV on the row) -> Statistics and check the appropriate test.

- What if you want to see whether belonging to a variable's category is statistically associated with a higher or lower average score on a dummy or interval variable when compared to the average score of another category of the same variable? An example of this would be if you want to know whether the typical level of education among Democrats is consistently higher/lower than Republicans.

For this situation, you can use an *independent sample t-test* to compare whether belonging to one or another of the groups measured by the same categorical variable is associated with the average value of an interval or dummy dependent variable. If you hypothesize that gender is a statistically significant predictor of differences in *income* (an interval variable) or identifying as *Democrat* (a dummy variable), you can answer this question with an independent samples t-test that compares two gender groups and tests for differences on the variable income and then Democrat.

- What method do you use if you want to know whether belonging to one group or another is associated with a different typical score, and the two groups you care about are measured by two different categorical variables? For example, we might want to know if the proportion of whites who identify as Republican is statistically different than the proportion of men who are Republican. Or we might want to know if identifying as racially white corresponds to a higher average household income than being a male.

We can answer both questions if we split our dataset by gender and run separate **one-sample t-tests** on income and identifying a Republican, using white American's mean value for *Republican* and average income (that we have separately calculated) as the test values.

- When we want to examine the relationship between any pair of interval/dummy variables, *correlation* is the option most political scientists will choose, and it is the method discussed in your textbook. What does *Pearson's correlation* (typically referred to as just "the correlation" and denoted by the letter r) measure, and roughly what levels of a correlation statistic suggest a weak, moderate, or strong correlation?
- What else does a correlation statistic tell us? In addition to telling us (1) how much of an association there is, a correlation statistic is (2) signed to tell us which direction the association is (i.e., does an increase in one of the variables typically correspond to an increase or decrease (-.xxx) in the other. And (3) the squared value of a correlation can be interpreted as how much changes in one variable tell us about the value of the other variable. Finally, (4) a correlation statistic typically is reported with a p-value or significance stars that tell us how probable it is that repeated sampling would find that the association is zero or runs in the other direction. If the p-value is higher than .05, there is not a statistically significant correlation.
- What does correlation NOT measure despite how the term is used in everyday language? A correlation statistic does *not* measure how much a change in the value of x will change y, on average. You need bivariate linear regression to do that. It also does *not* tell you whether x causes y, or if y causes x, or even if there is a causal relationship between the two. In short, only your theory or the temporal/logical ordering—rather than a statistical test—can help you to sort out if there is a causal relationship between two variables (for example, going to religious services more can't cause your gender).
- Sometimes, two variables have a correlation when one doesn't cause the other at all. What is a *spurious* correlation (hint: it is due to "omitted variable bias")? What, for example, explains the correlation between how much ice cream Americans eat in a month and crime?: <https://www.foxweather.com/lifestyle/summer-crime-wave>).
- What key assumptions about the nature of two variables' structure are made when one tests for a correlation between them? As a reminder, correlation works best with two continuous interval variables; however, social

scientists widely use correlation analysis to analyze bounded variables (like dummy variables, percentages, Likert scale results, GPAs, >3-point indexes, 10-point scales, and so on). The only categorical variables suitable for correlation measures (including regression models) are dummy variables, which is why we code multi-category variables into a series of dummy variables before most statistical analyses. If we use correlation to analyze an ordinal variable, we are treating that variable as though it were an interval measure.

- What does correlation's (and regression's) *assumption of linearity* mean? What are some examples where two interval variables could have a logical, very consistent relationship and yet still have a very low or even correlation? You may find it helpful to think about the effect of (1) a higher income (the IV) on a person's happiness (DV), (2) more time on the value of a person's investments, and (3) age on a person's level of physical independence. None of these relationships is linear because the effect of each additional unit of the IV doesn't consistently have the same influence on the dependent variable. However, all three situations involve a causal relationship from which we can make predictions (i.e., up to a certain point, making more money will provide you with more necessary resources and security, but at a certain point, earning more money has very little influence on happiness).
- How can you deal with variables that have a non-linear relationship in statistical modeling (Hint: There are ways to quickly identify and address whether the relationship between an independent variable has a non-linear relationship with a dependent variable, but we will only consider—very briefly—how multiple dummy variables would allow you to analyze the three situations above).

For example, we could code three dummies for age—1-20, 21-75, and older than 75. If we used a regression model to predict the level of a person's physical independence and used 21-75 as the reference category, we would expect to find that our dummy variables for young and old would both be significant and negatively signed.

For what it's worth, these details won't be on any test you take, but I do want you to be aware of the fact that there are a variety of statistical techniques that we won't learn in this class that allow linear or logistic regression to work well when independent variables do not have a linear relationship with the dependent variable.

Regression with one independent and dependent variable

- What is the line of best fit (aka the regression line) measure for two variables, and how does that line relate to their correlation? Thinking about a scatterplot with a plotted regression line, which axis of the plot is the Y axis; which is the X axis? Specifically, by convention, which axis reports the value of the dependent variable? Answer: By convention, the dependent variable goes on the Y-Axis and the X-axis is where the independent variable is listed.
- A *bivariate linear regression* model calculates the *correlation* between a single independent and dependent variable AND provides other valuable information that is not available from just a correlation statistic. A regression model's results provide information about the correlation (both an *R* and a *R-square statistic*), the *Y-intercept* (aka, the model's "*constant*"), and the independent variable's *slope*. If we look at a regression table that includes the constant, how do we predict the value of a dependent variable for different values of the independent variable? The slope is reported as an *unstandardized regression coefficient*. What do all of these *italicized* terms mean? If you aren't sure, review textbook Chp. 9, the summary for linear regression in the schedule, and the handout of annotated SPSS output for regression.
- Assume we have a bivariate regression looking at the relationship between years of education (1-16 years) and income in \$10k increments (1-15, with 15 representing \geq \$150k). The two variables are called Edu16 and Income15. The model includes years of education as the independent variable and income as the dependent variable. The (adjusted) R-square statistic for the model is .54. The model's constant is -2.10, and the variable Edu16 has an unstandardized coefficient of 0.821. Its standardized coefficient (i.e., its Beta) is .425. The P-value for the Edu16 coefficient is .005. Be able to clearly explain these hypothetical findings for each of these underlined elements.
- What is a *scenario*, when talking about regression results? Selecting the relevant statistics from the hypothetical example just given, what would the predicted income be for someone with 6 years of education? (Hint: you can write an equation to calculate the first one and plug it into Google, which will

solve it for you:
 $-2.10 + (6 \times .821)$

- If your regression model has only one dummy variable, what is the “reference” category? For example, if the variable Republican is in a regression model and there are no other partisan dummy variables, is the reference group Democrats or non-Republicans? Why are those not the same thing? **Important to remember: the reference group is any group NOT in the regression model; if a Republican dummy variable is in your model, and it is the only partisan dummy variable, the reference group is non-Republicans, which includes Democrats and independents.**

Multivariate (i.e., more than one independent variable) linear regression

- Assume that you also want to add multiple dummy categories for a single variable to the model you just examined. Specifically, you add two: the first indicates whether a person lives in the American South, and the second identifies respondents living in the West. What is the reference category for these two new variables? If the results returned statistically significant coefficients for West (say .492) and South (say -.103), how would you interpret these coefficients? **Answer: The reference group for both of those new variables would be “people who don’t live in the West or South.”**
- Staying with the same example used above (i.e., predicting a person’s income bracket from their education level), assume that we decide to run a multivariate linear regression model by adding a second independent variable. The new variable, *Male*, is a dummy indicator. If the r-square statistic for the new model is .10 higher than our previous model’s statistic, what does that tell us? It is possible that the coefficient for Edu16 would be smaller than before with the addition of a new variable? **Answer, if a new variable adds valuable information, a regression model’s adjusted r-square statistic will increase. When new variables are added, the value of the other variables’ regression coefficients may increase or decrease as the model more accurately predicts the value of the outcome variable.**
- What is an interaction term? Let’s say our model of a person’s income score (in the example, that’s a 1-15 measure called *Income15*) includes *Edu16 Male* and *Edu16xMale*. Why would we put that last variable in the model (i.e., what is an example of a hypothesis that this interaction term can test)? What would it mean if our results showed it was statistically significant and positive? What if it was statistically significant and negative? What if the interaction term wasn’t statistically significant? **Answer: If the interaction term is not significant, the effect of education on income is the same for men and everyone else. If the interaction term is positive and significant, having more education increases the income of men more than everyone else. If the interaction term is negative, additional income has less of an effect for men than everyone else.**
- Assume that the unstandardized coefficient for *Male* is 2.13 and the standardized coefficient is .214. In the revised model, the corresponding values for *Edu16* have been reduced to .623 (the unstandardized coefficient) and .340 (for the beta). The *constant* for the new model is: -1.015. Interpret the meaning of the unstandardized and standardized coefficients for the male dummy variable. What explains the reduction in the value of the standardized and beta coefficients for *Edu16* in the new model? Looking at the standardized coefficients, which of the independent variables has more of an influence on *Income15*?
- Let’s use those results to create *a predicted scenario for Income15 that includes a control*. If we know that the mean for the variable *Male* in the sample is .521 (in other words, 52.5% of our sample is male), how would you calculate the estimated income bracket (i.e., a value somewhere between 1 and 10) for an individual with 3 years of education, controlling for the effect of gender by setting the value of that independent variable at its mean value? Answer: you would use this formula:
Constant + (x years of education times its regression coefficient) + (then mean for Male times its regression coefficient), which is:
 $-1.015 + (3 \times .623) + (.521 \times 2.13)$
If you plug that formula into Google, you will see that at those values, the expected income bracket (i.e. the value of *Income15*) is 1.96, or about \$20-30K (remembering that a 2 *Income15* corresponds to that income range. If you do nothing but change the 3 years of education to 16, you see that a college-educated individual has an estimated income bracket score of 10.06, which corresponds mostly closely with \$100-100K bracket for *Income15*.

- What is *multicollinearity*, and why is it sometimes a problem when using multivariate regression? Why would you, perhaps, need to be worried about multicollinearity if you had a regression model that tried to predict an outcome with the 7-point variable of party identification (typically coded so that 1=Strong Democrat thru 7=Strong Republican) and a 10-point measure of how conservative someone is? A very common use of correlation in social science research involves a “correlation matrix” measuring the correlation between each of the independent and control variables in multivariate regression models. How does a correlation matrix with all of the independent and control variables help us to identify and deal with collinearity that is so high that we may have a difficult time telling whether or how much a given independent variable is correlated with our dependent variable?

Logistic Regression

- What kind of regression do you use when we are dealing with a dependent variable that is a dummy variable? (*logistic regression* is the most common technique)
- Why is it often preferable to use logistic regression to analyze ordinal variables (especially Likert- scale items) that have been recoded into binary outcomes rather than using linear regression with the original variables? If you had a variable that originally asked respondents to rate the current US president’s performance as: Very good, good, typical, bad, or very bad, what would be a logical way to code a dummy variable for his performance?
- Assume that you are using logistic regression to predict whether a person voted. Your model includes *Male* and *Education16* (as defined above). In your results, you have multiple pseudo R-square measures. Hint: You can report and interpret either the Nagelkerke’s R-Square or the Cox & Snell R-Square (but not both), since these measures are both widely used. What does this statistic tell you and how do you interpret it?
- Continuing with our logistic example, why will you want to report data from the SPSS output column that lists odds ratios (i.e., the statistics in the column exp(B)) rather than the unstandardized coefficients that are so useful with linear regression? (Note, you will see tables in published articles that report b instead of exp(b), but this is usually the case only if the body of the paper analyzes odds ratios or predicted probabilities similar to what you see in the sample Setzler and Yanus paper on characteristics that predict attitudes about gender equality.
- Assume that our logistic model predicting voting finds that *Male* has an exp(B) of .875 and *Education16* has an exp(B) of 1.25. The Wald score for *Male* is 4.2 and for the education measure, it is 7.2. Interpret these odds ratios and explain why we know education is the more powerful predictor.
- We probably will learn how to predict scenarios from logistic regression models. Although it requires a little more work, scenarios are by far the best way to explain logistic regression results to an audience that does not regularly work with statistics. Be able to describe in very basic terms how we predict a logistic regression scenario and provide examples.