

Methods Comprehensive Exam
May 2025

Instructions: Read all questions before answering any of them. You have 6 hours to complete the exam. When you use substantive examples in your answers, we prefer to see examples from political science. Answer all questions in part I. Answer 2 questions in part II. Answer 2 questions in part III. If you are completing the exam at GW, feel free to hand-write answers or parts of answers in a blue book, but carefully label them and note that you are using the blue book in your typed document. If you are completing the exam at home, you can similarly include photos/scans of hand-written material. Good luck!

This exam is open book (including articles from class) and notes. You may use a calculator or R for certain calculations. But you are not allowed to consult any other resource on the Internet, including chatGPT or other AI-based chat-bots.


Part I

In this section, please answer all questions.

1. Consider the contingency table (cross-tabulation) below using country-level data on the existence of judicial independence (coded as “yes” or “no”) and the history of coups in a country since 1950 (coded as: 0 coups, 1 or 2 coups, or 3 or more).
 - (a) How many “joint probabilities” can one calculate from this table? How many “marginal probabilities?” How many “conditional probabilities?” For each answer, briefly explain how you got your answer.
 - (b) Calculate $P(\text{JI}=\text{yes}, \text{no coups})$. Show your work.
 - (c) Calculate all of the marginal probabilities. Use proper notation and show your work.
 - (d) Calculate the following three probabilities: (1) $P(\text{JI}=\text{yes} \mid \text{coups}=0)$, (2) $P(\text{JI}=\text{yes} \mid \text{coups}=1 \text{ or } 2)$, and (3) $P(\text{JI}=\text{yes} \mid \text{coups}=3 \text{ or more})$. First, what types of probabilities are these, and what is the general equation for calculating these probabilities? Use proper notation. Second, show your work for each calculation. Third, use those probabilities to evaluate the extent to which judicial independence and coups are related. Briefly discuss how this exercise (in part d) is connected to interpreting crosstabs more generally.

<u>Number of Coups Since 1950</u>				
Judicial Independence	0	1 or 2	3 or more	<i>Total</i>
No	38	26	32	96
Yes	52	7	6	65
<i>Total</i>	90	33	38	161

2. You estimate a multiple regression model and suspect that the effect of x_1 on y has a nonlinear functional form. The following partial residual plot confirms your suspicion.



- What functional form does this plot suggest? Write out the equation for the OLS model that would estimate this functional form. From the figure above, what signs (negative or positive) would you expect the regression coefficients (central to this functional form specification) to have in your OLS model?
- Write out the equation for the marginal effect of x_1 on y . Show your work for how you derive this equation. Describe the nature of marginal effects in a nonlinear functional form (like this) compared to a linear functional form.
- Derive the equation for estimating the “tipping point,” that is, the value of x_1 at which the marginal effect of x_1 “flips” from being negative to being positive. Show your work.
- Briefly describe a substantive example from political science research that might capture such a relationship.

3. Consider the regression models below, which come from an article by Powers et al. in *Foreign Policy Analysis*. The article seeks to test the effects of nativism and free market attitudes on support for free trade in the U.S. and U.K. The analysis examines nationally representative public opinion survey data (observational data) in each country (a second study uses an experimental design). The results below are from the U.S. model.

- The **dependent variable is support for free trade globally** – specifically, whether respondents see the global benefits of free trade. It comes from a multi-item scale ranging from 0 to 1 (higher values are more supportive of free trade).

- **Nativism** comes from a multi-item scale measuring antipathy toward immigrants and preference for a homogeneous homeland. It ranges from 0 to 1, where higher values are more nativist attitudes.
- **Free market** taps agreement or disagreement with the sentiment that “people are better off in a free market economy.” The variable is measured using a 5-category ordinal scale and is recoded to range from 0 to 1, where higher values are more supportive of a free market.
- The authors also control for additional variables listed in model 3: perceptions of economic performance, support for government’s role in regulating the economy, left-right ideology, universalism values recognizing the interdependence of int’l relations, and conformity values stressing social solidarity for one’s in-group. Demographic controls (race, gender, age, education, and employment status) are included but not shown in model 3.

	(1)	(2)	(3)
Nativism	-0.42** (0.03)	-0.11 (0.06)	-0.01 (0.06)
Free market	0.07** (0.02)	0.28** (0.04)	0.30** (0.04)
Nativism × free market	—	-0.47** (0.08)	-0.43** (0.08)
Economic perceptions	—	—	0.30** (0.03)
Government regulation good	—	—	0.04 (0.02)
Right wing	—	—	0.09** (0.03)
Universalism	—	—	-0.02 (0.03)
Conformity	—	—	0.01 (0.02)
Controls	—	—	✓
Constant	0.71** (0.02)	0.58** (0.03)	0.38** (0.05)
<i>N</i>	2,186	2,186	2,186
<i>R</i> ²	0.10	0.12	0.19

* $p < 0.05$; ** $p < 0.01$.

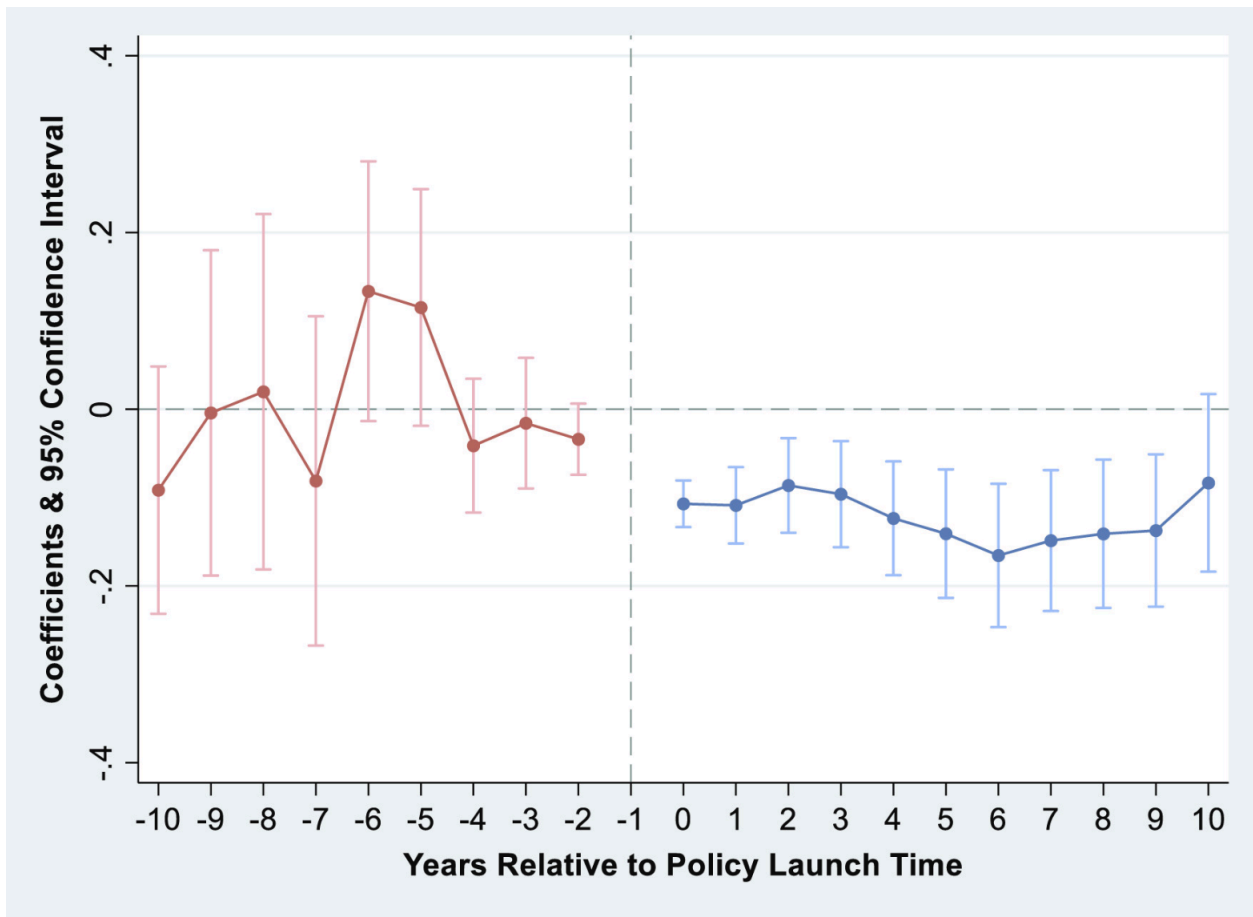
- (a) Support for free trade (the DV) follows a normal distribution (approximately) with mean=0.54, s.d.=0.26. (1) Approximately what percentage of respondents have extremely high support, i.e., values greater than 0.90? (2) What’s a sufficiently low score for free trade support such that just 10% of respondents are at or below that score? (3) Under these conditions and assumptions, what’s the median value of free trade support?

- (b) Interpret the *unconditional effects* of nativism and free market attitudes in terms of *direction*, *statistical significance*, and *size*. Which model do those effects come from?
- (c) From model 1, what is the approximate 95% confidence interval for the effect of nativism? Show your work. How do you use that interval to determine statistical significance? And how does it compare to the p-value approach?
- (d) Your colleague claims that “since the constant (intercept) is bigger in model 1 relative to model 3, model 1 must be ‘underspecified.’” Is such a conclusion warranted? What do the constants literally mean in each model, and how big of a deal should we make of the constant?
- (e) In model 2, interpret specifically what the “constituent terms” mean. Briefly describe why those effects are different from the effects in model 1.
- (f) From model 3, use the coefficient estimates to write out equations for: (1) the marginal effect of “nativism” on support for free trade, and (2) the marginal effect of “free market” attitudes on support for free trade. Also, (3) calculate the marginal effect of nativism when free market=0, 0.5, and 1; and (4) calculate the marginal effect of free market when nativism=0, 0.5, and 1. You do not need to calculate associated standard errors or 95% CIs. What do these calculations tell you substantively?
- (g) Describe and interpret in sufficient detail what the interaction term in model 3 means.
- (h) What are the best ways to visualize (via graphs) interaction model results? Illustrate one way of doing this using the calculations from part f.
- (i) How would you evaluate the causal validity of the results in model 3? Which OLS regression assumption pertaining to this issue does this model potentially violate? In your answer, provide a brief discussion of which independent variables are plausibly exogenous and which are potentially endogenous. For the latter, discuss the nature of the endogeneity you suspect and how it might be affecting the model results.
- (j) From a “robustness check” perspective, how does model 3 speak to potential concerns related to “omitted variable bias” that might plague model 2?

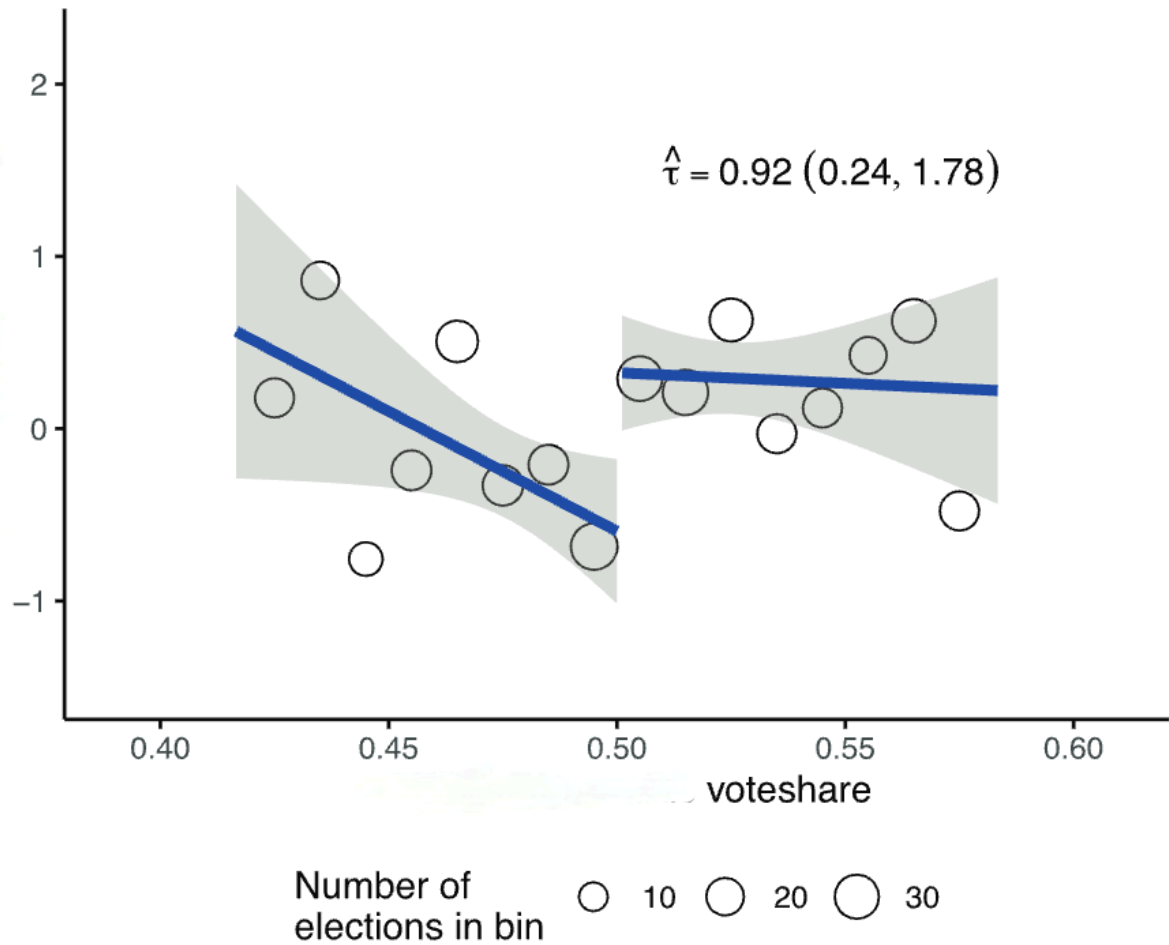
Part II

In this section, please answer two of the following questions.

1. Imagine that someone shows you the following graph with the coefficients from a difference-in-differences model. What should be some initial conclusions about the likelihood that the model's key identification assumptions are satisfied? What tests could you run to further check on the validity of these assumptions?



2. Imagine someone shows you the following plot from a study using a regression discontinuity (RD) design of the impact of a certain type of candidate winning on a policy outcome. The x-axis shows a running variable that represents that type of candidates' vote share among the top-two candidates (so 50% is the discontinuity). The y-axis shows an outcome of interest. Discuss the key causal identification assumptions and what tests you would run to verify these assumptions. Are there any specific concerns you might have based on the information on the graph?



- Imagine that you are a fan of the Alpine soccer (football) league. In this league, there are five divisions (the 1st-5th division, with the 5th division the lowest). Each year, the two teams with the most points in the 2nd-5th divisions are “promoted” to the next highest division, and the bottom two teams in the 1st-4th divisions are “relegated” to the next lowest division. As you’re watching a game one day, you wonder whether promotion and relegation in soccer impacts elections (e.g., maybe incumbents do better when teams from their city are promoted and/or worse when they’re relegated). Discuss three research designs you could use to examine this question. What are the key assumptions of each design? Which one do you think would yield the most reliable and credible causal inferences?

Part III

In this section, please answer two of the following questions.

1. Imagine that someone gives you a datafile with responses in 2024 from the American National Election Survey. The sampling frame is based on a probability sample of the national population but there is substantial non-response bias among the sample that actually completes the survey. You want to compare voting behavior (e.g., probability of voting for Donald Trump vs. Kamala Harris) for Americans with and without a college degree. The sampling weights included in the data are adjusted using raking for national targets for sex, race, Hispanic-ancestry, and age.
 - a. Are these weights sufficient for your project? Why or why not?
 - b. If not, how could you construct weights that adjust, as best as possible, for the probability of non-response? What criteria should you use as you're considering what variables to include in the construction of your weights?
 - c. Does your answer change if you decided you want to focus on comparing the voting behavior of Americans with and without a college degree in California?
2. Imagine that someone tells you that they want to use a Wordfish text analysis model to estimate the ideology of mayors in cities across Brazil based on their campaign platforms.
 - a. Briefly discuss the application of the Wordfish model for this substantive application in a manner similar to what we might see in an academic paper. What is a Wordfish model? What data would you need?
 - b. What assumptions would have to be necessary for this project to succeed?
 - c. How could we validate the outputs of the model?
 - d. Do you think this project is likely to succeed? Why or why not?
3. A team of researchers runs a field experiment to test the effect of a new civic education curriculum on political tolerance. The program is implemented in high school history classes. After running the study, the researchers discover that some of the teachers in the treatment group refused to implement the new curriculum. Why and under what conditions could this pose a challenge for estimating the causal effect of the curriculum? How should the researchers analyze the data? How should the results be interpreted?
4. In the multilevel models class, we spent a lot of time discussing multilevel regression and poststratification (MRP) models to estimate geographic variation in public opinion.
 - (a) Briefly discuss some of the strengths and weaknesses of MRP models for estimating geographic variation in public opinion.
 - (b) What data needs to be available in order to implement an MRP model?
 - (c) How could we evaluate the performance (validity) of an MRP model? Give concrete examples from the readings.
 - (d) Discuss some of the factors that influence the performance of MRP models.
 - (e) Give concrete examples of instances where we would expect MRP models to perform well and poorly.
 - (f) Should applied scholars interested in subnational public opinion switch to newer machine learning models?