Instructions: *Read all questions before answering any of them. When you use substantive examples in your answers, we strongly prefer to see examples from political science. Answer all questions in part I. Answer 3 questions in part II (but note the game theory question counts for two questions). 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!*

Part I

1. A casino tracks the losses and winnings of 1,000 gamblers who stayed at the casino in May 2019 and in May 2021. One casino employee finds that the biggest losers from 2019 improved the most when comparing across trips, speculating that their losses convinced them to try less risky bets. Another casino employee disputes this because the losers and winners from 2019 performed about the same in 2021. How would you account for these findings? Is there evidence that 2019 outcomes changed behavior in 2021?

2. In your dissertation research, you want to predict a continuous variable Y from a continuous variable X. Worried about endogeneity, you go to your advisors with a list of potential control variables. One advisor cautions that if these controls also predict X, you may get some multicollinearity and your standard errors on X will increase. Is it true that controlling for a variable always increases the standard errors on X? Should your decision on what controls to include be based on this?

3. In a study published in the *Journal of Politics* in 2007, David Karol and Edward Miguel present evidence that the Iraq War casualties negatively affected President Bush's reelection vote share. Karol and Miguel estimate the localized effect of U.S. casualties in the Iraq War on the change in President Bush's vote share between 2000 and 2004 at the state level. The main ordinary least squares (OLS) regression estimated is:

$$(\text{Bush Vote 2004})_i - (\text{Bush Vote 2000})_i = \alpha + X_i'\beta$$
$$+ \gamma(\text{Irap War Casualties})_i + \varepsilon_i \qquad (1)$$

Table 1 reports the main findings (standard errors in parentheses).

| | Dependent variable | | | | |
| | Bush vote share, 2004 | Change Bush vote share, 2000 to 2004 | | | |
	(1)	(2)	(3)	(4)	(5)
Bush vote share, 2000	1.01**		.015		
	(.03)		(.034)		
Total Iraq deaths and wounded per	−.0065**	−.0060**	−.0065**	−.0055*	
100,000 pop.	(.0022)	(.0021)	(.0022)	(.0023)	
Total Iraq deaths per 100,000 pop.					−.025*
					(.011)
Total Iraq wounded in action per					−.0045[†]
100,000 pop.					(.0024)
Army, Air Force reserve call-up per					−.00005
100,000 pop.					(.00003)
Select reserve per 100,000 pop., 8/2002					.00004[†]
					(.00002)
Proportion active armed forces (aged				.43	.25
18–64), 2000				(.26)	(.36)
Proportion veterans (aged 18+), 2000				−.29	−.20
				(.20)	(.18)
Change in unemployment, 9/2003 to				−.05	−.17
8/2004				(.65)	(.63)
Change Black pop. proportion,				2.15**	3.05**
2000–2003				(.66)	(.99)
Change White (non-Hispanic) pop.				−.35	−.87
proportion, 2000–2003				(.60)	(.62)
Proportional change in total pop.,				−.12	−.15
2000–2003				(.16)	(.16)
Observations	51	51	51	51	51
R^2	.97	.18	.19	.41	.48

(a) Note that the authors use two outcomes of interest: President Bush's vote share in 2004 (column 1) and the change in his share between 2000 and 2004 (columns 2-5). Which outcome would you say is preferable? Why?

(b) Compare and provide a substantive interpretation of the R-squared in columns (1) and (2).

(c) For columns (1) and (2), compute the 95% confidence intervals of the estimated coefficients on "Total Iraq deaths and wounded per 100,000 pop."

(d) In columns (4) and (5), the authors control for the sizes of the active armed forces population and veterans. Why might they have done this?

A variety of other robustness checks to the basic model are presented in Table 2.

	Dependent variable: Change Bush vote share, 2000 to 2004			
	(1)	(2)	(3)	(4)
Total Iraq deaths and wounded per 100,000 pop.	−.0060*	−.0087**	−.0058*	−.0072**
	(.0024)	(.0027)	(.0026)	(.0026)
Post-election Iraq deaths and wounded in action per 100,000 pop.			.0008	
			(.0014)	
Bush vote share, 2000				−.039
				(.042)
Change Republican Presidential vote share, 1996 to 2000				.059
				(.117)
Same-Sex marriage ban ballot initiative in 2004				.0038
				(.0066)
(Indicator Democratic President or VP candidate home state 2004)− (Indicator Democratic President or VP candidate home state 2000)				−.019**
				(.005)
Total Presidential or VP nominee visits to the state after the GOP convention				−.00035**
				(.00011)
Weight by state population	Yes	No	No	No
Two party vote share	No	Yes	No	No
State economic, demographic controls	No	No	Yes	Yes
Observations	51	51	51	51
R^2	.14	.26	.41	.55

(e) In column (1), the authors weight states by their 2003 populations. Why is this important? How does this result differ relative to the one reported in Table 1, column (2)?

(f) In column (3), the authors control for Iraq casualties *after* the election. Why would they do this and what do the results suggest?

(g) In column (4), the results show that visits by the GOP ticket to a state had a negative relationship with Bush vote share. Does this imply that voters reacted negatively to these visits?

(h) One reviewer points out that different regions in the U.S., or groups of states, may exhibit varying sensitivity to the Iraq casualties in terms of their support for President Bush. How would you address and test for this?

(i) Another reviewer says the models are incomplete because any effect of casualties should run through voters' opinions about the war. Therefore, the researchers should control for state-level support for the war. What would your response be to this? Is there a better way to test this idea?

(j) The final reviewer says they ran the same models for voting during World War II and did not find such an effect, speculating that the negative effect of casualties could depend on the public's support for the war as just or not. What would your response be to this and is there a way to test this idea?

Part II

1. You are analyzing some survey data from Pakistan to determine how exposure to varying levels of violence (X) affects people's preferences for militant groups (Y). During data cleaning, you discover that a large proportion of observations of the dependent variable are missing. What could cause a high number of missing observations of Y? Is it acceptable to omit the missing data from your sample and, if not, what could you do in response? Next suppose you instead had missing observations for X and answer the same question.

2. Researchers are increasingly interested in estimating causal mediation models. Give an example of a causal relationship that would be theoretically enriched by estimating the role of a mediator. When doing causal mediation work, what are the key principles that you think lead to credible and valuable findings? What are some of the inferential problems that one needs to account for in causal mediation analysis?

3. Policymakers have implemented a wide range of interventions to fight the spread of COVID-19. Suppose you want to evaluate the effectiveness of lockdown policies on COVID cases by comparing locality A to locality B. The former imposed a strict lockdown policy in mid-March, while the latter imposed a similar policy in early May. You will employ a difference-in-differences research design to estimate causal effects of such policy on covid cases. How would you do it? What are the required assumptions for this design to be valid? Discuss potential threats to the validity of this design in this particular example, as well as the broader challenges of this approach to estimating causal effects of counter-COVID measures.

4. Regression discontinuity (RD) designs are an increasingly common research design in political science. What is an RD design? Discuss the key assumption(s) that must be necessary for an RD design to generate a causal estimate (be precise). How would you evaluate the validity of the RD design? Provide at least one example of how an RD design has been used to evaluate an important research question in political science.

5. One of the most common criticisms of causal inference and experimental research is "external validity." First, explain what this means and contrast it with "internal validity." Second, why is external validity a special concern for causal and experimental work? Third, choose at least two causal inference techniques and describe some of the ways that researchers can address external validity concerns for each technique. Examples may be helpful.

6. Suppose you are interested in assessing whether incarceration causally affects political participation in the United States. To minimize concerns about selection bias and measurement error, you decide to use administrative records. By merging sentencing data to voter records, you expect to compare formerly incarcerated individuals to individuals convicted of crimes but not incarcerated. In other words, your empirical strategy is based on comparing the participatory patterns of these two groups of convicts. This would allow you to hold constant the fact that everyone in this sample is convicted of a crime. How would you estimate the effect of serving time in prison on voting for multiple elections? What are the required assumptions for this design to be valid? Discuss potential limitations and threats to the validity of this design.

7. Suppose we are trying to assess how well an LDA model predicts "fake" news articles from real ones on Facebook. For the purpose of this test, we consider a positive result to be a "fake" article and a negative result to be a "real" news article. After fitting the model in R, we compare predicted values from the actual values, as shown below:

```
> lda.fit = lda(fakenews ~ x1 + x2, data=train)
> lda.pred = predict(lda.fit, train)
> table(lda.pred, train$health)
              train
lda.pred  fake  real
fake       16    23
real       35    89
```

a) What is the misclassification rate of the LDA model?
b) How could we decrease the rate of false positives to false negatives and how is this likely to affect the misclassification rate?
c) Suppose you also ran a logit model for comparison and found the logit model had a smaller misclassification rate than the LDA on the same data. Why might logit perform better than LDA?

Game Theory: Counts as two questions

Consider an activist who wants a dictator to implement a political reform. The activist comes in three types: Radical, Moderate, and Quiet. The dictator's prior beliefs over these types are given by q_R, q_M, and $q_Q = 1 - q_R - q_M$. The order of the game is as follows:

1. The activist chooses to protest or not at cost c.
2. The dictator implements the reform or not.
3. The activist chooses to launch a revolution or not, at cost d to both players and with likelihood of success p.

The payoffs are such that Radical types will revolt no matter what. Quiet types will never revolt, but prefer getting the reform. Moderate types will revolt if and only if the reform is *not* granted (i.e., the reform satisfies them). Implementing the reform costs the dictator 1, with $d > 1$. The dictator also gets benefit W from ruling and 0 otherwise. If a revolution is attempted, the activist's payoff does not depend on whether the reform was granted (since they'll either be in charge or in jail), but assume the dictator still loses 1 by granting the reform.

(a) What is the total payoff to the dictator if they do not reform and face revolt? What is the total payoff to the dictator if they reform and avoid revolt?

(b) After seeing step 1, the dictator will update their beliefs on the type they are facing. Call the updated beliefs in step 2 q'_R, q'_M, and q'_Q. For what set of updated beliefs will the dictator implement the reform in step 2?

(c) What are the conditions for each type of activist to protest in step 1?

(d) Using (b) and (c), under what conditions is there a separating equilibrium? (This includes cases where two of the three types overlap, but the third does something different.)

(e) In the separating equilibrium, what is the probability that reform occurs (assuming the dictator's initial beliefs give the correct probabilities of each type)? What is the probability of revolt?

(f) How does the structure of signaling in step 1 and/or payoffs for the activist types need to change to get an equilibrium that is maximally beneficial for the dictator?