Title
Essays on Public Policy, Representation, Accountability, and Elections

Author
Hamel, Brian

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Essays on Public Policy, Representation, Accountability, and Elections

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Political Science

by

Brian Thomas Hamel

2021

ABSTRACT OF THE DISSERTATION

Essays on Public Policy, Representation, Accountability, and Elections

by

Brian Thomas Hamel
Doctor of Philosophy in Political Science
University of California, Los Angeles, 2021
Professor Lynn Vavreck Lewis, Chair

The three papers in this dissertation examine public policy, representation, accountability, and elections. The first paper studies the effects of two New Deal-era economic relief programs on voting behavior. I find that voters rewarded the Democratic Party for providing direct economic benefits, but not for indirect beliefs delivered through the private sector. The paper suggests that policy, and policy design, influences voting behavior. The second paper tests whether city governments are more responsive to certain demographic or economic groups at the expense of others. To study this, I collect 25 million citizen-initiated 311 calls for city goods and services, and find that cities respond marginally faster to calls from high-income and white neighborhoods. The third paper tests the effects of a popular electoral policy reform — independent redistricting commissions — on electoral competition in U.S. House and state legislative elections. I compare competition under the enacted districts to competition under simulated districts as well as competition under a set of of district maps that were considered but not enacted. I find that districts drawn by partisan legislators are just as competitive as those drawn by nonpartisan commissions.

The dissertation of Brian Thomas Hamel is approved.

Christopher N. Tausanovitch

Jeffrey B. Lewis

Martin Isaac Gilens

Lynn Vavreck Lewis, Committee Chair

University of California, Los Angeles

2021

To my parents, Neil and Linda Hamel, whose unending support turned my dream into a reality

Table of Contents

List of Figures

List of Tables

Acknowledgments

Though only my name is on this dissertation, so many others are just as responsible for it. I've been lucky to be surrounded by incredible mentors and colleagues, and the most supportive family and friends. Without them, none of this would have ever happened. I'll never truly be able to thank everyone enough, but it's worth a shot.

I came to American University thinking I wanted a career in "real" politics — jumping from campaign to campaign, perhaps with a pitstop on the Hill in between. But it turns out that being unemployed every other first Wednesday of November isn't particularly stable. Jan Leighley and Antoine Yoshinaka took me under their wing, showed me what life as a political scientist is like, and put me on the path to a Ph.D. In many ways, they are more responsible for this dissertation than anyone else. Jan and Antoine started as mentors, turned into co-authors, and are now dear friends. I cannot imagine where I'd be — professionally and personally — without them.

The decision to pursue my Ph.D. at UCLA was made easy because of one person: Lynn Vavreck. I met Lynn and instantly knew that I wanted to work with her. Lynn is a dream advisor. She encouraged me to pursue my own ideas and interests, no matter what those were or how "out of the box" they may have been, and always devoted the resources needed to pursue them. She was critical, too — in a good way. She regularly sat with me for hours at a time working through a new research idea. I always left with an *actual* research idea. Her enthusiasm is endless. She makes coming to campus every day and doing political science fun. She is exactly the kind of advisor I hope to be one day. I apologize to my students in advance for inevitably falling short of that standard.

The other members of my committee have made an indelible impact on me, as well. Jeff Lewis is one of the smartest and kindest person I've ever met. He has a knack for taking complex statistical issues and making it all seem so simple — with great humor, too. Chris Tausanovitch did more than most to shape how I think about representation and accountability. His class on representation was the best and most memorable class I took at UCLA. Marty Gilens returned to UCLA about midway through my studies. He is a master at the "big picture" and made me think critically about why my research matters. My work is stronger because of it. Several other faculty at UCLA deserve thanks. Miriam Golden hired me as a research assistant, and taught me so much about the research process — and Italian politics. I am grateful for her mentorship and friendship. Dan Thompson and Graeme Blair spent countless hours reading my job market paper and helping me prepare my job talk. They made a bad talk decent enough.

I also want to thank several friends and collaborators in the discipline. Lonna Atkeson and Bob Stein are the best co-authors, and I've learned so much about how to be a political scientist from them. Lonna also deserves recognition as the world's best karaoke partner. I couldn't have survived the job market without their support, either. Tyler Reny, Josh McCrain, and Eunji Kim provided much needed camaraderie while on the job market. They are three of the kindest and brightest young scholars out there, and I am lucky to know them.

Graduate school takes a non-academic village, too. My family and friends always supported in this crazy journey, and most importantly, they made sure that I took a break from my computer and had some fun in the process. Sara Curtis has been in my life since her knee gave out while standing for prayer in our freshman year of high school. I saved her then, but she's saved me many times since. Grant Conway and Jenn Geuther have been my best friends since "Welcome Week" at American. They were the first to know that I wanted to go to graduate school, and they even tolerated me talking about political science nonstop (for the most part). They have supported me every step of the way, and encouraged me to keep pushing even when I wanted to give up. I won the lottery of life when I met Grant and Jenn. Meet you at the "Smoker's Quad" at 7:55.

I've also been fortunate to meet some amazing people in Los Angeles. James Lewis deserves a special shoutout, though he'll (pretend to) hate the attention. He's been there for the good and the bad. From Vegas to Disney World to Alaska, we've had great adventures together — though none better than our visits to America's finest dining establishment, Chili's. He's also singlehandedly responsible for my "exquisite" taste in beer. My other friends don't thank him for that. Robin Wivell is the most positive person I know. He can put a smile on anyone, and he has surely put one on me —-except when he makes me dance. Robin was always there for me, and he has always brought out the best in me. For that, I'm grateful. My WeHo Dodgeball friends — Lesley McGovern Kupiec, Liz Turner, Jamell Dorton, Devon Stone Mills, and Ryan McLaughlin — are, simply put, the best people in the world. I've never had more fun in my life than when I've been with them. They make getting hit in the face with a ball — or breaking my pinky finger — worth it. If I could turn back time, you'd all have been a part of my life much earlier. I've also been so lucky to have the best roommates ever, Alex Simons and Catherine Saiz. I am going to miss our roommate nights, and especially Jon and Vinny's spicy fusilli, so much. At UCLA, I've been lucky to make some lifelong friends, too: Derek Holliday, Carolyn Steinle, and Ciara Sterbenz. I'll miss our late night Zoom venting sessions, and especially wrangling prospective students at preview weekend for a disastrous trip to WeHo. I'm most sad, though, that I won't be able to turn up in Vegas in Carolyn's minivan with you all.

I owe the biggest debt to my parents, Neil and Linda Hamel. You never once wavered in your support of me. No matter what I needed, you were there for me. More than you know, you are responsible for this dissertation. You instilled in me a curiosity about the world and a love for learning and discovery, and encouraged me to pursue my interests no matter what they may be. After all these years, you can finally say that your son has a job. I am so proud to be your son. My sister, Lisa — my "big sis" — is my biggest cheerleader, and I hope I'm hers, too. We got off to a rocky start — seriously, rug burn hurts, and I just don't look good in a tutu — but I really couldn't ask for a better sister.

Some of the material in this dissertation is co-authored, or has been published in an academic journal. Chapter 2 is co-authored with Derek Holliday. Chapter 3 was published in 2017 in the *Journal of Politics* as "Gerrymandering Incumbency: Does Nonpartisan Redistricting Increase Electoral Competition?" and is co-authored with John Henderson and Aaron Goldzimer.

Vita

2019 M.A. Political Science, University of California, Los Angeles

2016 B.A., Political Science, American University

Publications

Hamel, Brian T., and Bryan Wilcox-Archuleta. Conditionally Accepted. "Black Workers in White Places: Daytime Racial Diversity and White Public Opinion." *Journal of Politics*

Asquer, Raffaele, Miriam A. Golden, and Brian T. Hamel. 2020. "Corruption, Party Leaders, and Candidate Selection: Evidence from Italy." *Legislative Studies Quarterly* 45 (2): 291-395

Atkeson, Lonna Rae, and Brian T. Hamel. 2020. "Fit for the Job: Candidate Qualifications and Vote Choice in Low Information Elections." *Political Behavior* 42 (1): 59-82

Hamel, Brian T., and Antoine Yoshinaka. 2020. "Who Donates to Party Switchers?" *American Politics Research* 48 (2): 295-307

Hamel, Brian T., and Michael G. Miller. 2019. "How Voters Punish and Donors Protect Legislators Embroiled in Scandal." *Political Research Quarterly* 72 (1): 117-131

Henderson, John A., Brian T. Hamel, and Aaron M. Goldzimer. 2018. "Gerrymandering Incumbency: Does Nonpartisan Redistricting Increase Electoral Competition?" *Journal of Politics* 80 (3): 1011-1016

Panagopoulos, Costas, Jan E. Leighley, and Brian T. Hamel. 2017. "Are Voters Mobilized by a 'Friend-and-Neighbor' on the Ballot? Evidence from a Field Experiment." *Political Behavior* 39 (4): 865-882

CHAPTER 1

Direct vs. Indirect Benefits: Policy Design and the Electoral Effects of New Deal Spending

Elections offer citizens a mechanism for democratic accountability. That is, elections allow citizens the chance to reward and punish incumbent politicians for their decisions while in office. Theory and evidence suggests that voters care a great deal about government performance (e.g., Ferejohn 1986; Key 1966), and the way voters tend to evaluate performance is through observing the conditions around them instead of evaluating politicians through the policies and programs enacted under their watch. Fiorina (1978), for instance, argues that "the citizen looks at *results* rather than the policies and events which produce them" (430). In other words, voters do not evaluate the particulars of economic policies, but simply whether times are good or bad. Along these lines, there is ample empirical evidence that voters punish politicians for poor economic conditions (Kramer 1971; Nadeau and Lewis-Beck 2001; Healy and Lenz 2017), local war casualties (Grose and Oppenheimer 2007; Karol and Miguel 2007), high crime rates (Arnold and Carnes 2012), and poor school performance (Berry and Howell 2007; Kogan, Lavertu, and Peskowitz 2016).

Is public policy largely irrelevant to voters? My paper addresses this question. To do so, I study the electoral effects of the New Deal. The New Deal included some of the largest and most significant public spending programs in history. The initial appropriation for the the Works Progress Administration (WPA) alone amounted to almost 7% of GDP in 1935. Yet, recent work suggests that voters did not even respond to massive and consequential policy programs like these. In studying partisan behavior in the New Deal era, Achen and Bartels (2016) argue that "voters rendered no meaningful verdict on the substance of New Deal policies . . . specific aspects of the New Deal programs were largely irrelevant . . .

voters were focused on concrete economic conditions that they could see and feel when they went to the polls. " (136). My paper reevaluates this claim, and offers fresh evidence on the effects of policy, and policy design, on voting.

I study the effects of the two largest New Deal-era programs designed to promote economic relief and recovery: the Works Progress Administration (WPA) and the Public Works Administration (PWA). Combined, these two programs spent $12 billion between 1933 and 1939, or about 13% of nominal GDP in 1939.[1] These programs put more than 9 million people to work, built or repaired more than 650,000 miles of road, and were responsible for about 70% of new schools and one third of new hospitals built between 1933 and 1939.

I link county-level expenditure data on WPA and PWA expenditures between 1933 to 1939 to county-level presidential election returns both before and after the Depression. Using a generalized difference-in-differences design leveraging within-county changes in voting patterns before and after spending, I find that counties that received more WPA money became more Democratic than counties that received less WPA money. However, counties that received more PWA money became no more or less Democratic than counties that received less PWA money. These findings are consistent with theory on the importance of policy traceability in voting (Arnold 1990; Pierson 1993; Mettler 2011). Though both programs promoted economic relief, the WPA provided *direct*, government-to-citizen benefits, giving jobs to those without work. Workers on WPA projects were employees of the government. In contrast, the PWA operated *indirectly*, providing money to private sector firms who went out on the open market to hire workers to execute infrastructure projects. Simply put, the WPA's design made much more clear that the government was responsible for the policy and its benefits. I also show that these effects appear to persist through at least 1988, suggesting that the WPA permanently altered the partisan tendencies of counties.

Critically, using new data on the county economy measured both before and after spending began, I show that the effects of the WPA on Democratic voting are not a simple

[1] As a point of comparison, President Obama's 2009 economic stimulus package represented about 5% of nominal GDP in 2009. President Trump's first COVID-19 stimulus package was just over 10% of nominal GDP in 2019.

reflection of voters responding to economic growth and rewarding the party in power because of it. In fact, the WPA had no positive impact on county economies, suggesting instead that voters responded to the policy itself. In contrast to Achen and Bartels (2016), I argue that voters in the New Deal era did not just behave as they always do by responding to the state of the economy. Rather, voters responded to policy.

My results have a number of important consequences for the study of representation and accountability. For starters, the results suggest that policy — and policy design — matters. With a *traceable* policy design, voters will not just assess government performance by considering the conditions around them, but instead will evaluate performance on the basis of actual policy actions and deliverables. These results therefore offer an important antidote to a significant body of research suggesting that voters prioritize outcomes and conditions over policy. They also suggest that politicians may be constrained by voter preferences in the policies that they can pursue. As a consequence, leaders may prioritize delegate forms of representation that emphasize implementing the preferences of the public (Canes-Wrone, Herron, and Shotts 2001; Canes-Wrone and Shotts 2007; Fox and Shotts 2009).

Yet, these results also point to some challenging normative implications. On one hand, voters are not blind, and are responding to the particular actions of those in office. On the other hand, voters are not all-seeing, either. While the hand of government was more obscured in the PWA relative to the WPA, it was still a large and significant government-funded program that provided millions of jobs and stimulated economic activity across the country. In this way, voters' capacity to see and evaluate government action is still limited. At the elite level, these results suggest that politicians may have electoral incentives to design policies that deliver goods straight to citizens. Doing so may improve waning citizen trust in political elites (Hetherington 1998; Lerman 2019), but it also may open the door to the manipulation of public policy solely for political purposes. The "magic" of the WPA was in focus on directing benefits to those in the most need — those without work — in a way that also connected citizens directly to government. Policies, though, can just as easily be crafted to be direct without targeting aid to those who need it most, but rather to those whose votes are most important to a reelection-seeking politician or party (Cox and McCubbins 1986;

Dixit and Londregan 1996).

1.1 Policy and Voting: Theory and Evidence

While most work emphasizes the effects of broader national and local conditions on voting, scholars have theorized about *when* voters might use policy in their vote choice calculus. Two conditions must be met. First, policies must be *visible*. That is, voters must be able to readily observe a particular policy action, and its benefits. A change in the minimum wage, for instance, would constitute a visible policy because the policy will directly impact a sizable constituency. Even those who are unaffected by such a change would likely observe and understand the policy change, in part through interactions with those who are affected by the change. Other policies may easily go unnoticed. For instance, most citizens are unlikely to think twice about a one cent increase in the price of a gallon of milk, let alone wonder why such a change happened in the first place and whether government is the culprit.

Second, a policy and the benefits it accrues must be *traceable* to government action (see also Soss and Schram 2007). That is, they must hold government wholly responsible for the policy and the benefits its delivers. Voters can only hold an incumbent accountable if they perceive the incumbent as deserving of credit or blame. Policy traceability is thus akin the "clarity of responsibility" in economic voting (Powell Jr. and Whitten 1993; Rudolph 2003; Samuels 2004) For instance, in times of divided party government, voters are less likely to vote on the basis of the state of the economy because they are less certain which party is most responsible for the economy, good or bad. With policy, voters must be certain that government — not some other actor, like the private sector — is at the root of a particular change or benefit. Indeed, an emerging literature suggests that traceability in government programs is critically important for citizen awareness of government action (Mettler 2011) in (see Kettl 1988 and Morgan and Campbell 2011 for similar arguments).

There is very little evidence, however, that *certain* kinds of policies are more discernible to voters, and as such, are more likely to impact voting behavior. Kogan (2021) finds that voters rewarded the Democratic Party for creating the food stamps program. Kone and

Winters (1993) and Niemi, Stanley, and Vogel (1995) both find evidence that voters punish incumbent governors for raising state taxes, while Margalit (2011) finds that while voters punish incumbents for trade-related job losses, providing direct job loss benefits in response to those losses — such as job retraining opportunities — attenuates the negative effect of job loss on incumbent support. Each of these papers, however, *selects* on visibility and traceability in the policies under study. As a result, we do not know whether these two particular features of public policy matter for translating policies into votes.[2] Two papers have considered the effects of benefits that are delivered straight from government to citizens, as opposed to those delivered more generally to a geographic area. The results are mixed. Though Healy and Malhotra (2009) find that only direct disaster relief payments affected voters, Levitt and Snyder (1997) find that only direct payments had *no* effects on voting behavior.

My paper seeks to fill this gap and add additional evidence by studying the effects of two government programs — the Works Progress Administration (WPA) and the Public Works Administration (PWA) — that were similar in their goals and objectives, but differed substantially in their levels of traceability to government action. The goal of the analyses is to both address whether public policy impacts voting, and also offer a critical test of how important features of policy design influence whether new policies are translated into votes.

1.2 The WPA and the PWA

The seismic economic collapse of the early 1930s is well-documented: industrial production declined by 47% between 1929 and 1933. Similarly, GDP declined by 30%. At its maximum, unemployment reached more than 20%.[3] In response, newly-elected FDR and the Democrats

[2] There is much larger body of work on the effects of policy on political participation. Past work finds positive turnout effects of Social Security (Campbell 2002), the G.I. Bill (Mettler 2002), welfare (Soss 1999), FEMA disaster aid (Chen 2011), and Medicaid (Clinton and Sances 2018; Michener 2018). The argument here is that government benefits increases voter resources and in turn the capacity to participate in politics. Recipients are thought to participate primarily to protect their stake in the program and ensure the continuation of benefits. These studies, too, select, on visibility and traceability.

[3] These figures are from `https://www.britannica.com/story/causes-of-the-great-depression`.

enacted a “New Deal,” a series of policies, programs, and agencies designed to provide relief, recovery, and reform. New Deal policies dramatically and permanently expanded the size, scope, and power of the federal government. Indeed, some of these programs and agencies — such as the Social Security Administration —- remain cornerstones of U.S. public policy today.

Though some of these some programs focused on structural reforms (e.g., the FDIC), most of these programs were focused on alleviating economic stress and stimulating the economy in the short-run. Two of the largest of these programs were the Works Progress Administration (WPA) and the Public Works Administration (PWA). The goal of the WPA was to provide jobs to those without work. Most WPA projects were small-scale, “make work” construction projects where those without work would be taken off the “dole,” hired to work on the project, and paid through the federal government. The WPA paved 650,000 miles of road, and built 16,000 miles of water mains and 6,000 parks. WPA workers also cooked school lunches, sewed 300 million pieces of clothing for the poor, and taught the illiterate how to read and write. In the end, the WPA employed more than 8 unique million workers, meaning that over its period of operation, nearly 25% of U.S. families relied on WPA jobs and wages (Federal Works Agency 1947).

The PWA’s main goal was to stimulate economic production. The PWA operated in the same way that most economic stimulus packages and spending on public infrastructure does today. In practice, the PWA contracted with local private sector construction firms who then hired workers on the open market to execute large-scale, public works projects. The PWA was established in 1933 with a $3.3 billion authorization, and in the end, $6 billion was spent on PWA projects between 1933 and 1939,. PWA projects included 11,428 road projects and 7,488 educational buildings, and by March 1939, PWA workers had constructed 70% of all school buildings constructed and and 35% of hospitals and health facilities constructed since 1933 (Public Works Administration 1939).

These two programs were quite similar in their goals, and in their output. Both were designed to promote economic recovery, and both provided significant benefits for individuals — work — and for local communities. In this way, both the WPA and the PWA were very visible programs in the 1930s. Where the two differ, however, is in their traceability to government action. The WPA offered *direct* benefits. Those who received WPA jobs became employees of the federal government. Their paychecks came from the federal government. Thus, we can conceive of the WPA as a traceable program. In contrast, in operating through the private sector, the PWA was far less traceable. Those who worked on a PWA project were not paid by the federal government. Rather, these workers were hired and paid by the private construction firms responsible for the execution of the project. Similarly, firms providing supplies and materials for PWA infrastructure projects worked not with the federal government, but with the firm executing the project. As a result, government responsibility is much more blurred relative to the WPA. PWA benefits are therefore *indirect*.

Below I bring to bear county-level data on *where* the federal government allocated WPA and PWA money between 1939. Though all three were national programs active across the country at the same time, states and counties differed in how much money they received from each program. This feature allow me to estimate the effect of variation in these federal policies on voting.

1.3 Data

County-level data on New Deal expenditures were published in a 1940 report of the U.S. Office of Government Reports, and digitized by Fishback, Kantor, and Wallis (2003). These data give the total amount of money spent in each county by program aggregated across the years 1933 to 1939. Year-by-year data are not available. The data do not include expenditure information about every New Deal program — data on Civilian Conservation Corps (CCC) expenditures, for instance, were not reported — but does include data on many of the most prominent programs. Figure 1.1 gives an example of a full county report.[4]

[4]In 1939, the Works Progress Administration was renamed the Work Projects Administration.

Figure 1.1: Example of a County Expenditure Report

COUNTY REPORT OF FEDERAL EXPENDITURES

COOK County, ILLINOIS (16)

March 4, 1933 through June 30, 1939

LOANS		Number		Amount
Current Programs				
1. Reconstruction Finance Corp. (From Feb.2,1932)			1.	$256,015,152
2. Farm Credit Admin. - - Land Bank Commissioner		87	2.	250,100 2/
3. Farm Credit Admin. - - Emergency Crop and Feed		148	3.	13,780 2/
4. Farm Security Admin. - - Rural Rehabilitation			4.	27,925
5. Farm Security Admin. - - Farm Tenant Purchase			5.	---
6. Public Works Admin. - - Non-Federal projects			6.	52,674,442 3/
7. U.S. Housing Authority - Loan Contracts Signed		1,708 Units	7.	8,674,000
8. Rural Electrification Admin. - (Total project cost divided by No. of counties participating)			8.	---
9. Disaster Loan Corporation		4	9.	1,400
Completed Programs				
10. Farm Credit Admin. - 1934-1935 Drought Relief		1	10.	50
11. Home Owners' Loan Corporation - - - 1933-1936		45,984	11.	212,716,567
Total Repayable				$ 530,373,416
EXPENDITURES				
Current Programs				
1. A.A.A. Conservation Programs				
1936	$ 62,501			
1937	43,900		1.	$ 106,401
2. Farm Security Admin. - - Rural Rehabilitation			2.	10,655
3. Public Works Admin. - - Non-Federal projects			3.	75,130,850
4. Public Works Admin. - - - - Federal projects			4.	5,758,160
5. U.S.H.A. - Housing (Former PWA Housing only)(4)		2414 Units	5.	15,425,656
6. Public Buildings Admin. - - Federal buildings			6.	6,352,777
7. Public Roads Admin. - - - Completed projects			7.	13,656,974
8. Work Projects Administration			8.	273,705,121
9. Other Projects under Works Program		7	9.	660,674
10. Social Security Board				
Old Age Assistance	$13,163,944	44,415		
Aid to the Blind	---			
Aid to Dependent Children	---		10.	13,163,944
Completed Programs				
11. Federal Emergency Relief Administration			11.	156,770,075
12. Civil Works Administration			12.	33,235,984
13. A.A.A. Rental and Benefit Payments			13.	186,159
Total Non-Repayable				$594,163,430
GRAND TOTAL REPAYABLE AND NON-REPAYABLE				$1,124,536,846

As discussed, I focus on the WPA and PWA. Combined, these three programs make up about 45% of spending across the 15 New Deal programs included in the available report. Throughout the paper, I use spending per capita, taking each county's expenditure level and dividing by its population in 1930. Per capita measures of spending ensures that we can compare allocations and the effects of spending across counties of different population sizes. In addition, I adjust for inflation and use spending in 1944 dollars throughout.

The median county received about $34 per person in WPA spending. The median county received about $10 per person in PWA spending. Still, there is tremendous variation across counties and programs. The IRQ value for the WPA is about $38 per person — larger than the median. The IQR value for the PWA is $17. In addition, spending is clearly right-skewed, as the average county received much more than the median county. In other words, some counties received much more than the average county. These descriptive statistics suggest significant variation in who benefitted from these two New Deal programs, and my analyses below exploit this county-level variation to estimate the effects of policy on voting.

Figure 1.2 places each county into a spending quintile for each of the PWA and WPA. As is clear, counties in the West, upper Midwest, and through parts of the Mid-Atlantic tended the receive the most money per capita. These patterns are consistent across programs, though regional differences in allocations were noticeably sharper for the PWA. Here, most of the high-spending counties are completely concentrated in the West.

Figure 1.2: Geographic Distribution of New Deal Spending

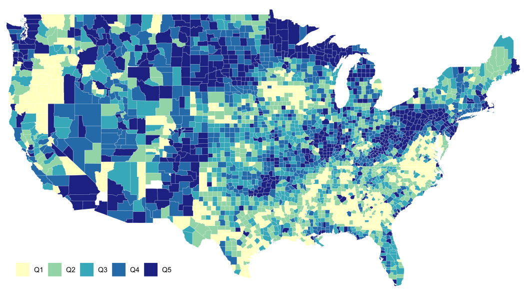

(a) WPA

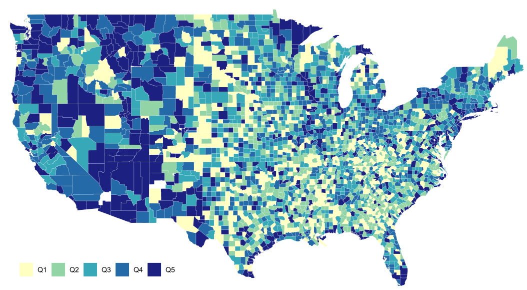

(b) PWA

1.4 Empirical Strategy: Spending and Voting

I merge my data on New Deal spending with county-level presidential elections voting returns from 1924 to 1944. Note that in the New Deal data, some counties are reported as one. For instance, spending in each of New York City's five borough's — each of which is in a different county — were collapsed into one geographic unit. The data are reported similarly for many counties and incorporated cities in Virginia (e.g., Alexandria and Arlington County). In these cases, I also aggregated votes into one "county." In sum, my dataset includes data on spending, economic conditions, and voting behavior for 3,067 "counties." To estimate the effects of spending and the economy on voting, I use a generalized difference-in-differences design. I estimate the following linear model:

$$\%\ Democrat_{cst} = \beta_1 log(Spending)_{cs} + \theta_{cs} + \alpha_{st} + \epsilon_{cst}$$

where $\%\ Democrat_{cst}$ is the Democratic share of the two-party vote in county c and state s at time t. $log(Spending)_{cs}$ is logged New Deal spending per capita in county c and state s. For all t prior to 1936, the value of $log(Spending)_{cs}$ is 0. θ_{cs} are county fixed effects, and α_{st} are state-year fixed effects. In all specifications, I also cluster the standard errors by county. County fixed effects mean that my estimates reflect within-county changes in voting patterns before and after New Deal spending. State-year fixed effects mean that my model uses only other counties in the same state as the counterfactual comparison group for each county. Combined, the inclusion of these fixed effects means that I am comparing the within-county change in Democratic support in counties with large New Deal allocations to the within-county change in Democratic support in same-state counties with smaller New Deal allocations.

For β_1 to be causal, we must assume parallel trends. Parallel trends means that the vote share in counties with smaller New Deal allocations must provide valid counterfactuals for the trends *we would have observed* in same state counties with larger New Deal allocations had these counties instead had smaller New Deal allocations. It cannot be the case that counties that received larger New Deal allocations were already trending in the direction of the Democratic Party faster than counties received smaller allocations. If this is the case, we cannot attribute the estimated effect to the effect of spending itself.

Is parallel trends a reasonable assumption? It is impossible to rigorously test. But I take two approaches to try to make the most apt county counterfactual comparisons. The first, mentioned above, is to make comparisons only between same-state counties. The second involves making comparisons between same-state counties that have similar characteristics. Or, in other words, removing cross-county demographic, economic, and political differences. This is particularly important in this setting given the dramatic changes in voting behavior that occurred from the 1920s to the late 1930s. We may be concerned, for instance, that certain county characteristics may both predict spending, as well as changes in voting behavior that may have occurred *even in the absence of the spending* simply because of changing partisan voting coalitions over-time. To estimate just the pure effect of spending, we need to remove the influence of these other characteristics, and one way to do so is to make counties "look" similar on those very characteristics. To do this, I use covariate balancing propensity score (CBPS) weights (Fong, Hazlett, and Imai 2019; Imai and Ratkovic 2014). CBPS weights minimize the correlation between New Deal program spending and pre-treatment county characteristics, meaning that counties that received large vs. small allocations will now "look" the same on other characteristics, such as their voting behavior prior to spending occurred. I use 1930 (pre-treatment) county Census data to create these weights.

Table 1.1 and Table 1.2 give the correlation coefficients before and after weighting for the WPA and PWA, respectively. As is clear, urban counties, counties with fewer Blacks, and those with higher unemployment received more WPA and PWA money. Counties that were less supportive of the Democratic Party from 1924-1932 also received more money. However, after weighting, all of these differences disappear. The weights will allow me to compare counties that are similar, removing the possibility that these covariates may be time-varying confounders that could bias the effect of spending in Democratic support.

Table 1.1: CBPS Weights for the WPA

	Unweighted	Weighted
% Urban	0.28	-0.00
% Black	-0.38	0.00
% Unemployed	0.39	-0.00
% Foreign Born	0.26	0.00
% Democrat, 1924	-0.38	0.00
% Democrat, 1928	-0.27	0.00
% Democrat, 1932	-0.35	0.00

Table 1.2: CBPS Weights for the PWA

	Unweighted	Weighted
% Urban	0.32	-0.00
% Black	-0.12	0.00
% Unemployed	0.22	-0.00
% Foreign Born	0.23	-0.00
% Democrat, 1924	-0.16	0.00
% Democrat, 1928	-0.11	0.00
% Democrat, 1932	-0.14	0.00

1.5 Results

Table 1.3 and Table 1.4 report the main results. The first column in both reports the effects without the CBPS weights, while the second column includes the CBPS weights. Focusing first on the WPA, we find a positive and statistically significant effect of WPA spending on voting. Counties that received more WPA money became more Democratic than those that received less money. Notably, while the effect attenuates when the weights are included, the effect remains positive and significant. Using the estimates in the second column, I find that, on average, moving from the 25th percentile ($18.98 per capita) to the 75th percentile ($57.34 per capita) in WPA spending would increase Democratic vote share by 1.27 percentage points. The average Democratic vote share is 57.90%, meaning that a 1 IQR change in represents about a 2% increase in Democratic vote share. In the third column, I estimate the effects in the each election. I also "treat" 1932, allowing me to test for pre-trending. If there is a positive and significant effect of spending in 1932, it would suggest that counties that received more WPA spending were in fact already trending in the direction of the Democratic Party prior to spending. I find no evidence of pre-trending in 1932, offering additional credibility to the research design and the parallel trends assumption.

Table 1.4 repeats the same models for the PWA. In the first column, I find a positive and significant effect of the PWA on voting, albeit one smaller in magnitude than that of the WPA. However, once the weights are applied — making counties that received large vs. small PWA allocations similar across other dimensions — the effect of the PWA attenuates

Table 1.3: Effects of the WPA on Democratic Voting

	DV: Democratic Vote Share (0-100)		
	(1)	(2)	(3)
ln(WPA)	1.88***	1.15***	
	(0.24)	(0.32)	
ln(WPA) × 1932			0.39
			(0.37)
ln(WPA) × 1936			0.70*
			(0.36)
ln(WPA) × 1940			1.47**
			(0.46)
ln(WPA) × 1944			1.57***
			(0.47)
County FEs	✓	✓	✓
State-Year FEs	✓	✓	✓
Weights	✗	✓	✓
Observations	18,388	18,330	18,330
R^2	0.94	0.94	0.94
Notes:	*p<0.05; **p<0.01; ***p<0.001		

significantly and is no longer statistically distinguishable from zero. In other words, the effects of the PWA recovered in the first column largely reflect demographic and economic differences between counties receiving more and less money. There does not appear to be a direct effect of the PWA on voting. The third column shows that there are no PWA effects in any of the individual post-spending elections.

Table 1.4: Effects of the PWA on Democratic Voting

	DV: Democratic Vote Share (0-100)		
	(1)	(2)	(3)
ln(PWA)	0.73***	0.16	
	(0.11)	(0.17)	
ln(PWA) × 1932			-.06
			(0.16)
ln(PWA) × 1936			0.18
			(0.21)
ln(PWA) × 1940			0.19
			(0.22)
ln(PWA) × 1944			0.05
			(0.23)
County FEs	✓	✓	✓
State-Year FEs	✓	✓	✓
Weights	✗	✓	✓
Observations	18,388	18,330	18,330
R^2	0.94	0.94	0.94
Notes:			***p<0.001

These results provide strong evidence for the importance of policy and especially policy traceability. As noted, both the WPA and PWA were large programs, employing millions of people and spending billions of dollars on infrastructure and economic production. Nevertheless, my results suggest that voters only responded to the WPA. Voters responded to the program that directly linked citizens to government, but not the program that placed a middle-man — private enterprise — between government and citizens. In an even broader sense, the results suggest that policy matters. With a traceable policy design, voters will not just assess government performance by considering the conditions around them, but instead will evaluate performance on the basis of actual policy actions and deliverables.

1.6 Counterclaim: It's the Economy, Stupid

One possibility, however, is that voters are *still* only responding to the state of the economy. Perhaps the WPA improved local economies, and the PWA did not. In this case, the observed positive effect of the WPA on voting may just be a reflection of voters observing positive economic growth and rewarding the Democratic Party for it. In other words, voters, as Achen and Bartels (2016) would argue, may have behaved as they always do by rewarding the incumbent party for the economic conditions that they can see and feel, without really knowing or caring much about policy at all.

To test for this, I examine the effect of the the WPA and PWA on the county economy. Doing so requires measures of the county economy taken both before and after spending began. Common contemporary measures of the economy such as income and GDP are not available at the county-level during this time period. Those that are available — such as unemployment — are not available at regular intervals throughout the 1930s. Rather, we only know county-level unemployment in 1930 and 1940 as recorded in the decennial Census.

As alternatives, I draw on measures of county-level retail sales and manufacturing output. Retail sales data were collected in 1929, 1933, 1935, and 1939 through the Census of Business, and digitized by Fishback, Horace, and Kantor (2005). These data report the total value of sales from retail establishments in the county for the year, giving a sense for county consumption before, during, and after the Depression. Businesses included in this measure range from grocery stores to motor vehicle dealers. Manufacturing output was collected in 1929, 1933, 1935, 1937, and 1939 by the Census of Manufactures. Manufacturing output is a value added measure; it i is the value of all manufacturing outputs minus the value of the inputs used in production. It represents manufacturing's contribution to overall GDP.

Higher values are indicative of greater production. The manufacturing data are incomplete, as the Census Bureau does not report data for counties with few manufacturers. I have complete data for 1,852 counties (60%). Analyses below using these should be viewed with the caveat that only a subset of (often larger) counties are included.

Figure 1.3: County Economic Conditions in the 1930s

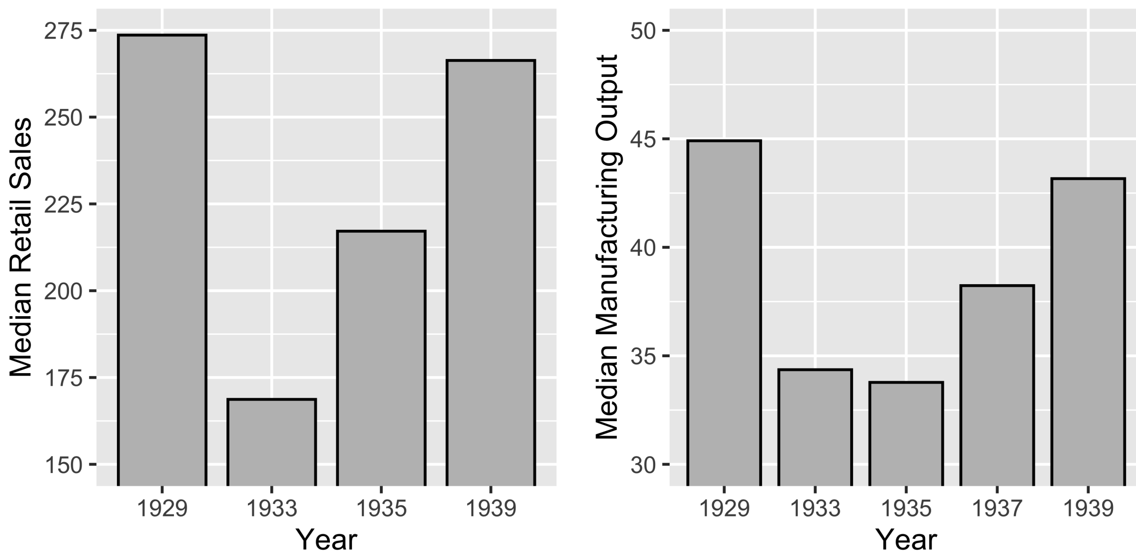

Figure 1.3 plots the median per capita value of both measures by year.[5] As expected, retail sales activity and manufacturing output dropped significantly from 1929 to 1933 as the Depression reached its trough. Consumption and production improved year-by-year between 1933 and 1939, and by 1939, conditions had recovered to 1929, pre-Depression levels. In fact, both retail sales and manufacturing output were marginally higher in 1939 than 1929. All told, while these measures are not standard in the literature on economic voting, both appear to be reasonable proxies for local economic conditions during this time period.

My empirical strategy is the same as in the analysis of spending and voting. Using the before-and-after spending economic data, I estimate equations of the following form:

$$County\ \ Economy_{cst} = \beta_1 log(Spending)_{cs} + \theta_{cs} + \alpha_{st} + \epsilon_{cst}$$

[5]As with the New Deal spending data, all values are in 1944 dollars.

where $County\ \ Economy_{cst}$ is logged retail sales per capita, or logged manufacturing output per capita, in county c and state s at time t. $log(Spending)_{cs}$ is logged WPA spending per capita, or logged PWA spending per capita, in county c and state s. For all t prior to 1936, the value of $log(Spending)_{cs}$ is 0. θ_{cs} are county fixed effects, and α_{st} are state-year fixed effects. With these fixed effects, I am again able to estimate the change in the county economy in counties with large WPA and PWA allocations as compared to the change in the county economy in counties with small WPA and PWA allocations. Again, I restrict comparisons to same-state counties. I focus on two "types" of economic effects. First, I consider the effects from pre-Depression (1929) to post-spending (1939). This comparison allows me to test whether these two programs put counties in an economically stronger position relative to before the Depression. The second considers the effects from the beginning of the spending period to the end of the spending period. For instance, the WPA began in 1935, and the spending data available ends in 1939. This analysis allows me to test whether the programs increased economic growth relative to where the county was at the start of spending.

Table 1.5 and Table 1.6 reports the results for the retail sales and manufacturing output, respectively. The first two columns give the effects of the WPA on each economic comparison described above: the first comparing the economy in 1929 to 1939, and the second comparing the economy in 1935 to 1939. The final two columns repeat these analyses for the PWA. As is clear across both economic outcomes and time comparisons, there are no positive effects of the WPA on the economy. Counties that received more WPA money saw no more or less economic growth than those that received less WPA money. If anything, the WPA had an economically negative effect on manufacturing output relative to before the Depression began. In contrast, the PWA appears to have had positive economic effects on both consumption and manufacturing production.

These results are not meant to suggest that the WPA had zero positive economic benefit or impact. Indeed, many would have otherwise gone without work if not for the WPA. These jobs surely put food on the table and kept a roof over the head of many families. However, at least along these two economic dimensions, it does not appear that the WPA improved *broader* county economic conditions. For the purposes of my analyses of spending and voting, it seems unlikely then that the effects of the WPA on Democratic voting are a simple reflection of an improved economy. Rather, voters were responding to the direct benefits of the policy itself.

Table 1.5: Effects of the WPA and PWA on County Retail Sales

	DV: ln(Retail Sales)			
	(1)	(2)	(3)	(4)
ln(WPA)	-0.02	0.01		
	(0.01)	(0.01)		
ln(PWA)			0.05^{***}	0.02^{***}
			(0.00)	(0.00)
Effect	1929 to 1939	1935 to 1939	1929 to 1939	1933 to 1939
County FEs	✓	✓	✓	✓
State-Year FEs	✓	✓	✓	✓
Observations	6,130	6,131	6,130	6,132
R^2	0.95	0.98	0.95	0.97
Notes:				$^{***}p<0.001$

1.7 Did the Effects of the WPA Persist?

Throughout this paper, I have focused on the immediate differences in voting patterns pre- and post-spending. For instance, the estimates in Table 1.3 and Table 1.4 tell us whether Democratic voting patterns changed — and by how much — from just prior to spending (1924-1932) to just after spending (1936-1944). The estimates from these models therefore give the effects of the WPA and PWA on Democratic support in the short-run.

Recent research looks for persistent effects of policies, historical institutions, and significant events on public opinion and elections over time. For instance, Acharya, Blackwell, and Sen (2016) find that counties that had large slave populations in 1860 are more Republican today than those counties that had smaller slave populations. Southerners living in these counties today are also more racially resentful than Southerners in Southern counties with a

Table 1.6: Effects of the WPA and PWA on County Manufacturing Output

	DV: ln(Manufacturing Output)			
	(1)	(2)	(3)	(4)
ln(WPA)	-0.08***	-0.02		
	(0.02)	(0.02)		
ln(PWA)			0.05***	0.02+
			(0.01)	(0.01)
Effect	1929 to 1939	1935 to 1939	1929 to 1939	1933 to 1939
County FEs	✓	✓	✓	✓
State-Year FEs	✓	✓	✓	✓
Observations	4,864	4,608	4,864	4,434
R^2	0.96	0.97	0.96	0.97

Notes: $^{+}$p<0.10; ***p<0.001

weaker slaveholding history. Along the same lines, Mazumder (2018) shows how civil rights protests in 1964 affect contemporary racial attitudes and county election outcomes.

One obvious question is whether the effects of New Deal spending were substantively large during FDR's re-election campaigns in 1940 and 1944, and weak or potentially even non-existent in subsequent elections and for candidates other than FDR. The answer to this question also addresses ongoing debates in work on explanations for the New Deal realignment (Achen and Bartels 2016; Caughey, Dougal, and Schickler 2020). To test for persistent effects, I reestimate the same model as before — inclusive of county-confounder fixed effects — but expand the dataset to 2016, and interact New Deal spending with each election year. Doing so gives me a separate coefficient for the effect of spending (relative to 1924-1932) for every election cycle. As before, because of the fixed effects, I make counterfactual comparisons only among similar counties (pre-treatment) in the same state.

Figure 1.4: Long-Run Effects of the WPA on Democratic Voting

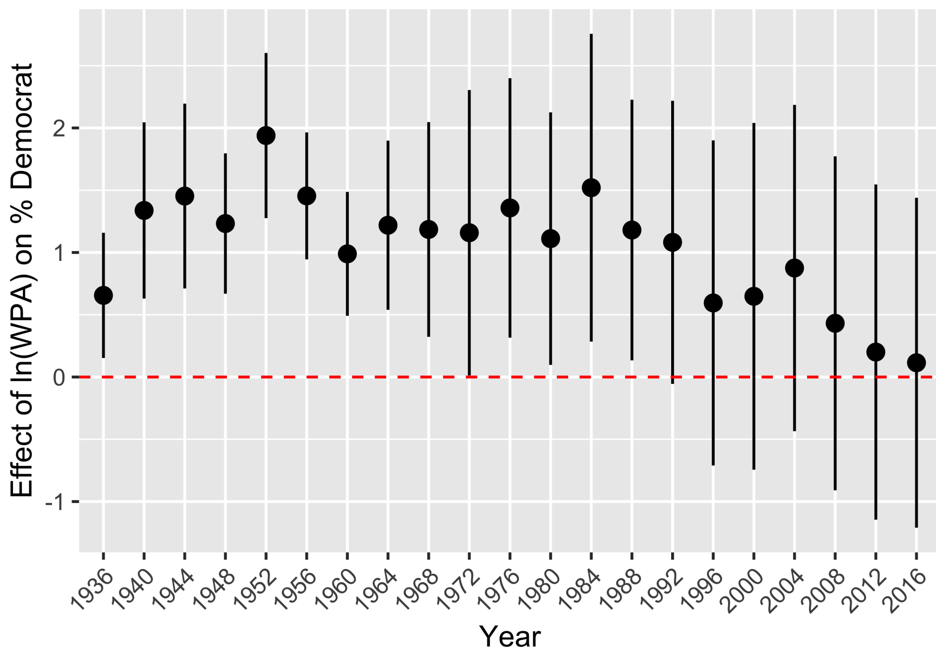

Figure 1.4 presents the results. I find positive and significant effects of New Deal spending in every presidential election cycle through 1988, suggesting that variation in New Deal allocations has permanently affected voting behavior. To be sure, estimating persistence across such a long time-series is not easy, and we should view the estimates as suggestive. It is possible that confounders have been introduced over time that bias these estimates. Beyond estimation, outlining a theory that can account for the persistence of effects, even effects that appear substantively reasonable, is challenging. For instance, one possibility is that the WPA socialized across generations of citizens. For instance, the children of adults who received work from the WPA may have been much more likely to be Democrats over the long-term, thereby permanently affecting the partisan proclivities of those counties. Despite these questions, what we can say with more certainty is that the effects of New Deal spending on county voting patterns did not quickly converge back to zero once the spending stopped. It appears that public policy can have such a profound affect that it can also can fundamentally change the politics of the places who feel the impact of government action

the most.

1.8 Conclusion

In 2001, George W. Bush and the Republican Congress passed a tax relief bill. Soon thereafter, almost 100 million taxpayers received a letter from the IRS stating the following: “We are pleased to inform you that the United States Congress passed and President George W. Bush signed into law the Economic Growth and Tax Relief Reconciliation Act of 2001, which provides for long-term tax relief for all Americans who pay income taxes.”[6]

After passage of the CARES Act in response to COVID-19, Americans received a similar letter alerting them to the arrival of their $1,200 stimulus check: “As we wage total war on this invisible enemy, we are also working around the clock to protect hardworking Americans like you from the consequences of the economic shutdown.”[7] The letter was signed by President Trump.

These are just two examples of incumbent politicians promoting their policy actions. Both policies were designed to put more money in the hands of Americans and boost the economy. Yet, though it is no secret that politicians often promote their accomplishments (Grimmer, Messing, and Westwood 2012; Mayhew 1974), political science theory would predict that these policy decisions and actions will not affect subsequent voting behavior. Rather, it is expected that voters would only evaluate the incumbent on the basis of the state of the conditions these policies sought to affect. Indeed, scholars even claim this to be the case in times of tremendous change in government policy, like in the New Deal era (Achen and Bartels 2016).

In this paper, I test this claim using data from two of the largest New Deal-era economic relief programs. These programs were designed to put millions of Americans to work and

---

[6] https://www.cnn.com/2001/ALLPOLITICS/06/19/tax.letter.text/?mod=article_inline

[7] https://www.usatoday.com/story/news/politics/2020/04/28/coronavirus-trump-letter-stimulus-check-rec 3040031001/

stimulate economy production and activity. Yet, some counties received much more of these benefits than others. Using a generalized difference-in-differences approach, I find robust evidence that voters rewarded the Democrats for WPA spending, but not for PWA spending. The WPA was a direct government assistance program, directing jobs and benefits straight from government to citizens. In contrast, the PWA was an indirect program, offering assistance to the private sector that trickled down to citizens. Voters, in other words, rewarded the Democratic Party for the program most easily attached to government action. These results provide important evidence for the importance of policy design, and in particular, policy traceability (Arnold 1990; Mettler 2011). Crucially, I find that these effects are not a reflection of an improved economy. The WPA, in fact, does not appear to have had positive effects on the local county economy. Voters rewarded the party for a policy that does not appear to have had its intended effects on the broader economy. I also offer some suggestive evidence that the positive effects of the WPA on voting lasted for decades. In short, while more work is needed to probe the persistence of these effects, it appears that policy can powerfully and perhaps even permanently affect the partisan proclivities of places.

My findings speak to how, and for what, voters hold politicians accountable, but in turn they also have significant implications for democratic representation. Fox and Shotts (2009) develop a theory of when elected officials will behave as trustees or delegates, and argue that whether officeholders pursue the interests of the public (delegate) or pursue their own interests (trustee) depends on how voters vote. When voters prioritize competence in their voting calculus, politicians will tend to act as trustees. But if voters are policy-oriented, politicians will emphasize ideological congruence with their constituents. In finding that voters are responsive to policies, my results suggest that politicians may be more constrained in their choices than theories of retrospective voting allow. Indeed, making the wrong policy choice, or failing to act at all, may hurt the incumbent just as much as failing to improve the economy itself.

My study also speaks to the challenge — and potential opportunity — of governing in times of crisis. Past work shows that voters punish and reward politicians for their responses to natural disasters (Bechtel and Hainmueller 2011; Gasper and Reeves 2011; Healy and

Malhotra 2009). My paper builds on this work, showing how policy action in the wake of economic catastrophe influences voters in a similar way — and potentially for many decades thereafter. These findings suggest that politicians may be able to mitigate the negative electoral effects of economic collapse by pursuing policies and offering provisions that may stem the tide (see also Margalit 2011). Simply put, voters care about whether and how leaders respond to crisis.

The purpose of this paper was to provide some much-needed baseline evidence on the direct effects of both policy and particular features of policy design like traceability. The next step is to extend my design to other policy areas (e.g., school performance and crime) and other contexts. As noted, the Depression and the subsequent New Deal was a unique event in American history, and we may not expect voters to be responsive to *all* policy actions, particularly those that may go completely unnoticed by voters or be significant part of an incumbents re-election story. Nevertheless, extending my framework across a variety of policy areas will allow for a more complete assessment of policy, policy design, and voting.

CHAPTER 2

Who Gets a Streetlight? Unequal Responsiveness in City Service Delivery

An essential ingredient of democratic government is the "responsiveness of the government to the preferences of its citizens, *considered as political equals [emphasis added]*" (Dahl 1971, 1). Past work suggests that federal and state policy reflects public preferences (Erikson, Wright, and McIver 1993; Lax and Phillips 2012; Simonovits, Guess, and Nagler 2019), and that *changes* in mass preferences correspond to similar movement in policy outputs (Caughey and Warshaw 2018; Page and Shapiro 1983; Stimson, Mackuen, and Erikson 1995; Soroka and Wlezien 2010).

There is increasing concern, however, that government tends to prioritize the preferences of the advantaged: there is a tighter link between the preferences of legislators and the preferences of white and black citizens than with those of non-white and poor citizens (Bartels 2008; Ellis 2012; Griffin and Newman 2007; 2008; Rigby and Wright 2011), and affluent opinion is a stronger predictor of policy changes than non-affluent opinion (Gilens 2005; 2012). These effects may also be exacerbated because candidates and thus elected officials tend to be members of advantaged groups (e.g., Juenke 2014; Shah 2014), and politicians' personal backgrounds often shape their decision-making while in office (Burden 2007; Butler 2014; Carnes 2013). In sum, despite a clear link between aggregate preferences and policy, citizens appear far from "political equals" in national and state politics and policymaking.

Though there is much less work on (unequal) responsiveness in local politics, findings thus far comport with expectations from the state and national politics literature: local policies reflect citizen policy preferences (Tausanovitch and Warshaw 2014) and cities spend more on municipal services (e.g., parks and recreation, public transportation, etc.) as support for Democratic presidential candidates increases (Einstein and Kogan 2016), even as the left-right ideological positions of white and affluent citizens appear better represented among city elected officials than the positions of non-whites and affluent citizens (Schaffner, Rhodes, and La Raja 2016). Responsiveness in the local context is unique because localities are often constrained by state and national mandates in policymaking (e.g., Gerber and Hopkins 2011; Peterson 1981). Still, they have tremendous latitude in providing and maintaining access to critical public goods and services — roads, bridges, streetlights, police, fire, water, sewers, waste collection, parks, recreation, libraries, public transportation, and much more. Indeed, these services are "at the heart of city government" (Yates 1977, 26). In this paper, we offer a test of whether cities are more responsiveness to the *service* demands and needs of the advantaged — white and affluent citizens — than those of the less advantaged.

Whether there are racial and economic inequalities in how governments allocate goods and services remains an open question. On one hand, diverse and segregated cities tend to invest less in collective goods and services (e.g., Alesina, Baqir, and Easterly 1999; Glaser 2002; Hopkins 2009; Trounstine 2016). Cities may even use public policy to ensure that goods primarily benefit predominantly white, wealthy, and politically powerful neighborhoods (Trounstine 2018). Further, non-white and poor respondents tend to evaluate city services much less favorably than do white and rich respondents, suggesting that cities may prioritize the advantaged in providing access to goods and services (e.g., DeHoog, Lowery, and Lyons 1990; Hajnal and Trounstine 2014). At the same time, other work argues that cities distribute services — and respond to citizen-initiated complaints and demands — on the basis of administrative procedures and professional bureaucratic values rather than on the basis of race or class (e.g., Jones et al. 1977; 1978; Lineberry 1977; Mladenka 1981; Nivola 1978). Understanding whether there are inequities in city services is a particularly important question given the strong link between constituent service and citizen evaluations of government and elected officials (e.g., Hajnal and Trounstine 2014; Harden 2016; Tucker 2019).

Our approach to this question mirrors contemporary work on (unequal) policy responsiveness: we measure group demands, government's response to those demands, and then assess whether government prioritizes the demands of white and wealthy citizens, relative to non-white and poor citizens. To do so, we introduce the largest and most comprehensive dataset of citizen-initiated demands for public goods and services — almost 25 million 311 requests from 15 of the largest 25 cities in the U.S. between 2008 and 2018. 311 is a government-sponsored phone number available in many U.S. cities, and is the primary way citizens can report issues in their community. For researchers, 311 data represent the main source of information on the demands citizens place on city governments (Minkoff 2016; White and Trump 2018). The data tell us the nature of demand (e.g., a broken streetlight or a missed trash collection), where it was requested, when it was requested, and when the government closed the request. We use these data to assess whether government responds *faster* to service demands from advantaged neighborhoods than to demands from less advantaged neighborhoods. In other words, we assess whether cities prioritize whiter and richer neighborhoods using an empirical approach akin to contemporary work on policy responsiveness (e.g., Gilens 2012; Tausanovitch and Warshaw 2014).

We find *some* evidence of two kinds of inequality in responsiveness: request-based and need-based. First, consistent with audit experimental studies of constituent communication (e.g., Butler and Broockman 2011; Einstein and Glick 2017; White, Nathan, and Faller 2015), we find that cities respond to requests from whiter and and more affluent neighborhoods faster than those *same* requests from less white and affluent neighborhoods. We also find evidence of need-based inequality. We show that, to some extent, neighborhoods of different racial and economic compositions tend to request different goods and services. That is, some service demands tend to come much more often from non-white areas than white areas (and vice-versa). On average, government is slower to respond to needs of less advantaged communities than to the needs of more advantaged neighborhoods. Cities appear to prioritize both requests from and the needs of the advantaged.

The much more encouraging news for democratic governance is that these differences in responsiveness are quite small in substantive magnitude. On average, we find that service demands from white neighborhoods are responded to about 1.9 *hours* faster than the same demands from less white neighborhoods. Likewise, we find that each 1 percentage point increase in relative white neighborhood need (relative relative non-white neighborhood need) for a given good or service corresponds to just a 36 *minute* decrease in average wait time. In practice, these kinds of differences appear rare, as relative needs are ultimately quite similar across advantaged and less advantaged communities. The magnitude of these effects suggests to us that, on average, cities tend to respond to requests with the same relative urgency regardless of neighborhood characteristics. Unlike with national policymaking, we see that citizens of all economic and demographic backgrounds are much closer to "political equals" than not in access to municipal public goods and services. In the conclusion, we discuss the broader implications of these findings for American democracy.

2.1 Inequality and City Services

One well-established finding in political economy is the negative relationship between local racial diversity/segregation and public goods provision (Alesina, Baqir, and Easterly 1999; Glaser 2002; Hopkins 2009; Trounstine 2016). In new work, Trounstine (2018) shows that government exacerbates racial segregation through housing and land policy. In four cities, she shows these policy decisions had corresponding effects on public goods provision, as segregated black neighborhoods had fewer sewer systems and extensions. These findings replicate across other public goods, too (Trounstine 2016). Cities may be particularly reticent to offer collective goods as white residents (and white political elites) come to see these services as primarily benefitting racial minorities (e.g., Gilens 1996; Luttmer 2001). All told, race appears to shape whether — and how — municipalities distribute goods.

White and the affluent residents also tend to view the quality of city services much more favorably than do non-white and poor residents (e.g., Aberbach and Walker 1970; DeHoog, Lowery, and Lyons 1990; Hajnal and Trounstine 2014; Schuman and Gruenberg 1972). For instance, Hajnal and Trounstine (2014) find, across 26 cities, that non-white and economically disadvantaged citizens are less satisfied with city police, fire, libraries, and schools, and that these evaluations shape general impressions of local government. These differences are substantively large, too: relative to white respondents, black respondents evaluate city services about 6 percentage points less favorably. Other work finds that black respondents' service evaluations depend on who holds political office. Marschall and Ruhil (2007) find that blacks report higher levels of satisfaction with their neighborhood conditions, police services, and public schools when black officials serve on the city council and on the school board.

Critically, subjective measures reflect objective measures of local conditions (e.g., Holbrook and Weinschenk 2019), and tend to capture citizens' interactions and experiences with city officials, bureaucrats, and services more generally (e.g., Kelly 2003; Kelly and Swindell 2002; Percy 1986). DeHoog, Lowery, and Lyons (1990) and Hajnal and Trounstine (2014) argue that differential satisfaction reflects differential experiences with city services and treatment by elites, as the effects of race and class on service satisfaction disappear as measures of neighborhood conditions, issues, and service quality are accounted for. Moreover, Marschall and Ruhil (2007) find feelings toward police are particularly high as black mayoral tenure increases, as police forces are more responsive to service requests, and as the police force sees increases in black representation, again suggesting a strong link between evaluations and actual government performance and action. These findings are also consistent with accounts of how minority descriptive representation in Congress positively affects the provision of services focused on minority constituents (Canon 1999; Grose 2011; Lowande, Ritchie, and Lauterbach 2019). In short, racial and economic differences in evaluations reflect differential responsiveness: cities — except those with black officials serving in elected office — prioritize the demands and needs of white and affluent residents.

On the other hand, a separate literature sees the distribution of city services and resources as driven by "bureaucratic decision-rules" and *not* by neighborhood economic or racial characteristics (e.g., Jones et al. 1977; 1978; Lineberry 1977; Mladenka 1980; 1981; Mladenka and Hill 1978; Nivola 1978). Mladenka (1980) studies fire protection, educational resources, and waste collection services in Chicago, and finds that service delivery is routinized and reflect technical-rational criteria. For instance, whether a neighborhood receives waste collection services depends not on political support or demographic characteristics, but on the distribution of home ownership, the distance to be travelled between pickups, and the amount of waste generated per neighborhood. These technical-rational criteria benefitted the advantaged and the disadvantaged; that is, in some neighborhoods, these criteria meant more and better waste collection services for the advantaged, and in others, these processes benefited to less advantaged. Others found similar patterns in Oakland (Levy, Meltsner, and Wildavsky 1974), San Antonio (Lineberry 1977), Detroit (Jones et al. 1977; 1978), and Houston (Mladenka and Hill 1978).

Some work in this tradition even test responsiveness directly using data on citizen contacts. Jones et al. (1977) and Mladenka (1981) study a small-set of citizen-initiated complaints — akin to the kinds of data we use below — and find no evidence in Chicago, Dallas, Detroit, and Houston of substantive or statistical relationship between processing time — the time it took a city to respond to a complaint — and neighborhood race or class. Likewise, Vedlitz and Dyer (1984) find no relationship between the "clear rate" — the number of calls responded to versus the number received — and race, class, or a neighborhood's political composition in Dallas. In all three studies, government's response followed administrative procedures and resource constraints (i.e., time and effort to meet a particular kind of demand).

Combined, it is unclear whether we should expect significant racial and economic inequalities in how cities respond to public goods requests. Regardless, we see past empirical tests — with the exception of some early studies noted above (e.g., Jones et al. 1977; Mladenka 1981) — as ill-equipped to detect unequal responsiveness. We argue that the strongest and most appropriate empirical approach mirrors contemporary work on unequal policy responsiveness (e.g., Gilens 2012; Griffin and Newman 2008): (1) measure the preferences — in this case, demands — of different racial and economic groups as well as how government responds to those demands; and (2) assess the statistical relationship between group demands and government responses. The results offer insight into the degree of inequality in responsiveness.

2.2 Two Pathways to Unequal Responsiveness

Disadvantaged groups may experience inequality through a variety of mechanisms. In particular, cities may prioritize *requests* from the advantaged over those from the disadvantaged, and they may prioritize the *needs* of the advantaged over the needs of the disadvantaged — or both. More specifically, we propose two (mutually non-exclusive) pathways to unequal responsiveness. First, cities may respond to service requests — i.e., a 311 call to remove graffiti from public property — from whiter and richer neighborhoods faster than they do the *same* kind of request from less white and affluent neighborhoods. Second, the needs of neighborhoods — i.e., whether a neighborhood needs better infrastructure or better waste collection — may differ along economic and demographic lines, and in responding to all needs, government may prioritize responding to the needs of the advantaged over the disadvantaged. We further articulate both channels below.

Cities may prioritize requests from the advantaged. A growing body of audit experimental research shows that political elites — elected legislators and executives, as well as bureaucrats and other government actors — use discretion in responding to constituent communication. Butler and Broockman (2011) show legislators are 5 percentage points more likely to respond to e-mails about registering to vote from a white alias than the same email from a black alias. Critically, there appears to be limited evidence of legislators responding to strategic or partisan, electoral incentives, as legislators of both parties discriminate against black aliases at about the same rate. Rather, the biases appear to be "taste-based" in that politicians simply prefer prioritizing communications from white aliases relative to black aliases. Subsequent work reports similar biases against immigrants and Latino/as in communications with legislators, housing agency officials, and local election officials (e.g., Einstein and Kogan 2016; Gell-Redman et al. 2018; Mendez and Grose 2018; White, Nathan, and Faller 2015).

These racial and ethnic biases tend to generalize. In a meta-analysis of audit studies in the U.S. and abroad, Costa (2017) finds that minority constituents are about 10 percentage points less likely to receive a response than a non-minority constituent. There is, however, more limited evidence of similar effects among other social and demographic groups. Though Lajevardi (2018) and Pfaff et al. (2019) find evidence of religious bias toward Muslims, others find no evidence discrimination against low-income constituents (Carnes and Holbein 2019) and same-sex couples seeking to obtain a marriage license (Lowande and Proctor 2019).

Like audit studies, we define request-based inequality as city's propensity to respond to requests from whiter and more affluent communities faster than the same request from less white and affluent neighborhoods. In short, holding demand-side considerations constant, cities prioritize the advantaged: when a predominantly white neighborhood and a predominantly non-white neighborhood both report an overflowing sewer, government tends to fix the sewer in the whiter neighborhood before the sewer in the less white neighborhood. In empirical terms, for a given service, we would observe a *faster* response time for requests from whiter and more affluent neighborhoods relative to requests from a less white and poorer neighborhoods.

In responding to all requests, cities may also discriminate against the *needs* of the less advantaged. Different neighborhoods are likely to need different kinds of services, and in turn place different demands on government. For instance, Thornton et al. (2016) find that predominently low-income and non-white neighborhoods tend to have more graffiti, broken windows, and litter than wealthier neighborhoods and areas with more white residents. Similarly, they show that more advantaged neighborhoods have poorer sidewalk conditions. Neckerman et al. (2009) report that poorer areas have fewer street trees and less clean streets. If demands follow needs, poorer and less white neighborhoods should ask government for different kinds of goods and service than wealthier and whiter neighborhoods. Minkoff (2016) examines 311 data — as we do below — in New York City, and finds broad support for the link between need and revealed demand via contacting These findings are also consistent with early work suggesting that the primary predictor of citizen-initiated contacting is need (e.g., Thomas 1982; Vedlitz, Dyer, and Durand 1980), even above and beyond the influence of socioeconomic predictors of civic and political participation more generally (e.g., Verba and Nie 1972).

Unlike request-based inequality, demand-side considerations drive need-based inequality; if predominantly white neighborhoods tend to need sidewalk repairs and predominantly non-white neighborhoods tend to ask for graffiti removal, government will respond to requests for sidewalk repairs faster than to requests for graffiti removal. Need-based inequality would mean observing a *faster* response time for the kinds of services most often requested in whiter and more affluent neighborhoods relative to those most often needed in less white and affluent areas.

Two different mechanisms may drive need-based inequality. Differential responsiveness to needs may reflect the less advantaged tending to ask for goods and services that simply take more government resources to respond to. If so, need-based inequality may reflect government capacity; it is less a deliberate expression of a pro-white or pro-rich bias on the part of government, and more coincidental. On the other hand, government may recognize what neighborhoods of different economic and racial makeups need, and choose to devise a bureaucratic resource and response system that ensures the advantaged receive better representation than the less advantaged. In this case, need-based inequality reflects the same kind of "taste-based" discrimination outlined above. Teasing these mechanisms apart is difficult with the data we bring to bear below; it would likely require a case-by-case study of 311 response systems across cities and the many agencies and departments of each city government. We leave this important work to future researchers, but note that through either mechanism, the needs of the advantaged are prioritized and tended to first.

2.3 Data

311 is a government-sponsored phone number that provides access to non-emergency municipal services. It offers citizens a way to communicate with their local government about issues in their community. In short, 311 makes citizens the "eyes and ears" of the city, allowing city officials to devote resources to fixing — rather than searching for — issues. In most cities, citizens can call 311 directly, and in some cities, can report issues online or via a smartphone application. Complaints can be filed 24 hours a day, seven days a week. In some cities (e.g., Boston and San Francisco), callers can include a photo of the problem along

with their written description. After filing the complaint, an estimated time to completion is often given, and the request can be tracked from there. In most cities, calls are routed to a central command center before city employees direct the call to the appropriate agency or department. City workers then investigate the report and (if needed) fix the issue. Once fixed, city workers either mark the request as "closed" in the 311 log, or report back to the central command center.

311 data provide a wealth of information about citizen-initiated complaints, including: the kind of good or service requested/nature of demand (e.g., waste collection, streetlight repair, etc.), where it was requested (the latitude and longitude of the request location), when it was requested, and when the government closed the request. Though several scholars have used 311 data as measures of civic and political participation and neighborhood conflict (e.g., Cohen et al. 2019; Feigenbaum and Hall 2016; Legewie and Schaeffer 2016; Lerman and Weaver 2013; Levine and Gershenson 2014), the primary — and perhaps most appropriate (e.g., White and Trump 2018) — use of the data is as a measure of demand for public services. Given these features, 311 data meet our own needs: it provides information on citizen demands, and because it captures when government responded to each demand, information about responsiveness. Examples of 311 calls in our data include the following (edited by the authors for clarity and grammar):

> "Streetlight is hanging by one wire over the road. Looks like it will fall at any moment." — Detroit, 5/13/2018

> "HUGE pothole on the freeway entrance ramp. This hole is very dangerous to motorcycles entering the freeway. Someone is going to get hurt if they hit it." — San Diego, 1/2/2017

> "Trash collection was missed last Friday and yesterday. Trash cans are full, and trash is overflowing in the whole alley. Please pick up ASAP." — Washington, DC, 1/14/2015

We sought to collect 311 data for each of the 25 largest cities in the U.S. for a many years as possible (up to the end of 2018). To be included in our data, the city 311 log must include the four features described above: data on the kind of request, where it was requested, when it was opened, and when it was closed. Of the top 25 cities, 15 met these criteria (in order by population): New York, Los Angeles, Chicago, Houston, Philadelphia, San Diego, Dallas, Austin, Jacksonville, San Francisco, Denver, Washington, Boston, Detroit, and Nashville.[1]

Cities receive calls about hundreds of different kinds of issues, some much more common than others. For instance, Austin's 311 call log includes 147 unique services. San Francisco's includes 101, many of which only received a handful of calls. One challenge then is determining which of these calls to analyze in an assessment of unequal responsiveness. We considered two options. The first was to create a broad list of goods and services that we anticipated would be found in each city call log (e.g., street repairs), work through the data city-by-city, and categorize as many as possible into these broad categories. The downside of this approach is that we may lose a significant amount of within-city variation in the kinds of requests made and the frequency with which they are made. In looking at the data, we found the most common requests were often not the same across cities. In Dallas, complaints about tall grass and weeds are quite common, but we do not see a similar pattern in any other city. Likewise, only in New York are calls about heat and hot water issues and plumbing common.

We choose instead to maintain the variation across cities, and emphasize the most common needs within cities. To do so, for each city, we first collapsed similar but related services into one category. For example, we combine requests for trash collection and recycling collection into one. We then selected for analysis only the top 10 service areas — or as many service areas as are in the data up to 10 — within each city, aggregated across the entire time period. Our approach ensures our analyses do not "miss" important features of each city. Further, it ensures that our investigation is focused on the most pressing challenges that each city community faces. Table 2.1 summarizes the 311 data for each city. It gives the date range of our data, the total number of calls, as well as the number and share of

[1] Indianapolis and Seattle do not currently have 311. Phoenix, San Antonio, Columbus, El Paso, and Memphis have 311, but their call logs do not include latitude and longitude. City officials in San Jose did not respond to our FOIA requests. Data from Fort Worth are not freely available to the public.

calls included in the city's top 10 service categories (and in turn, our analyses). In total, our dataset includes 38.5 million service requests. Over 25 million of these — about 65% — fall in the top 10. City-by-city, we see that our sampling strategy picks up most of the calls. In 13 of 15 cities, our analysis includes more than 60% of all calls; in 8 of 15 cities, it includes more than 70% of all calls.

Table 2.1: Summary of Requests by City

City	First Request	Last Request	# of Requests	# of Requests (Top 10)	% Top 10
New York, NY	01/01/2010	12/31/2018	17,655,254	8,644,127	48.96
Los Angeles, CA	08/05/2015	12/31/2018	3,083,707	2,784,236	90.29
Chicago, IL	02/11/2008	12/31/2018	4,231,886	4,231,886	100.00
Houston, TX	11/07/2011	12/31/2018	2,161,509	1,313,892	60.79
Philadelphia, PA	12/08/2014	12/31/2018	791,334	558,533	70.58
San Diego, CA	05/20/2016	12/31/2018	370,888	274,598	74.04
Dallas, TX	10/01/2016	09/27/2018	733,833	284,691	38.80
Austin, TX	12/31/2013	12/31/2018	622,294	395,950	63.63
Jacksonville, FL	01/01/2012	12/30/2016	1,195,220	718,006	60.07
San Francisco, CA	07/01/2008	12/31/2018	2,969,713	2,407,340	81.06
Denver, CO	01/01/2008	12/31/2014	1,054,924	698,240	66.19
Washington, DC	01/01/2012	12/31/2018	2,037,846	1,483,042	72.77
Boston, MA	07/01/2011	12/31/2018	1,417,930	1,161,932	81.95
Detroit, MI	07/21/2014	12/31/2018	156,810	137,799	87.88
Nashville, TN	07/17/2017	12/31/2018	76,896	60,705	78.94
			38,560,044	**25,154,977**	**65.24**

We merged our 311 call log with economic and demographic data from the American Community Survey (ACS). The ACS is sent to approximately 295,000 addresses monthly and 3.5 million per year, making it the largest household survey administered by the U.S. Census Bureau. Today, it is the primary source for economic and demographic information. We use the latitude and longitude given for each 311 call to place each call within a block-group, and we append to each call block-group estimates of per capita income and percent non-Hispanic white.[2] Before doing so, we split the 311 data into two groups: calls made between 2008 and 2012, and calls made after 2013. For 2008-2012 calls, we use estimates from the 2012 ACS, reflecting income and race estimates averaged across survey data collected between 2008 and 2012. Post-2013 calls are merged with the 2017 ACS (reflecting ACS data collected between 2013 and 2017).

[2]We exclude observations where the latitude and longitude given in the data do not fit within the bounding box of the city. We suspect these few cases reflect record-keeping errors.

For each call, we created a measure of responsiveness: wait time. Wait time measures how long it took the government to respond to the demand; it is simply the difference (in days) between the given open date and the given close date. Larger values indicate a longer wait time. The measure makes a key assumption: the close dates given in the data are precise to the day. That is, we assume that city officials tend to close requests in "real-time," and not in batches at the end of the week or month. There is no obvious reason to suspect that the given close dates are inaccurate. First, governments tend to use 311 data as a performance metric and as one of many ways to determine how to best allocate resources. Closing requests in batches would give the government imprecise information and almost no insight into actual performance. Still, Figures A1 and A2 looks for clustering in the data within each city in two ways. First, Figure A1 simply counts the number of requests closed on each day of the week within each city. If the closing dates given in the data reflect the internal process that a city uses to "clear their deck," then we may expect to see that calls tend to be closed on only one or two days of the week.

Figure A2 takes the same approach for each day of the month — e.g., the 1st of the month, the 15th, etc. Here, we would be looking to make sure that cities do not tend to batch close requests on the 1st of the month, the last of the month, etc. Either would add raise questions about our measure of responsiveness, but we find limited to no evidence of obvious clustering in the data. Instead, calls seem to closed throughout the work week, suggesting that cities at least do not tend to batch completed requests from each week on Friday (for example). There is some evidence of clustering around particular dates in particular cities — e.g., in Philadelphia, many more calls than average are closed on the 26th of the month. Likewise, there appears to be some clustering around the 22nd in Dallas, the 17th in Denver, and the 13th in San Francisco. In general, though, we see limited consistent patterns in the data to suggest that the closure dates in the data are wholly inauthentic. These diagnostics suggest to us that our measure of responsiveness comes with limited measurement error and record keeping-induced bias.

We bring to bear the largest and most comprehensive dataset of citizen-initiated service demands to date. The data offer both measures of citizen demands, as well as government responsiveness. Merged with neighborhood-level economic and demographic data, we are able to assess whether government is more responsive — in our setting, whether government responds faster — to requests from the advantaged and to the needs of the advantaged. We do so below, presenting our design and results in turn.

2.4 Pathway #1: Request-Based Inequality

We start with a test of request-based inequality. Recall that we define request-based inequality as the propensity to respond to requests from whiter and more affluent communities faster than the same request from less white and affluent neighborhoods. Using the dataset describe above, we first categorize each call as coming from a neighborhood either above or below the city median percent white and the city median logged per capita income. From here, we calculate the mean wait time within each city-service, month, and year combination separately for calls both above and below the city medians. Before doing so, we trim the data in two ways. First, we exclude outlier values in wait time that may not reflect the true data-generating process. Specifically, we exclude all calls above the 90th percentile in wait times to exclude the (very rare) wait times that exceeded months or even years. We suspect these few cases likely reflect errors in record-keeping or required a much more substantial mobilization of municipal resources. Second, to ensure that the average wait times within each group do not reflect only a small number of calls, we exclude city-service-month-year observations where the total number of calls falls below the 5th percentile. Taken together, we create a dataset that gives our measure of responsiveness (within each city-service, month, and year) for the advantaged — defined here as for those neighborhoods above the city median in percent white and logged per capita income — and for the less advantaged. We assess request-based inequality with the following equation estimated using OLS:

$$Average\ Wait\ Time_{ijk} = \beta_0 + \beta_1 \cdot Advantaged\ Request_{ijk} + \gamma_i + \delta_j + \theta_k + \epsilon_{ijk} \qquad (2.1)$$

where $Average\ Wait\ Time_{ijk}$ is the mean wait time for city-service i in month j and year k. $Advantaged\ Request_{ijk}$ is an indicator variable equal to 1 for requests from neighborhoods above the city median percent, or city median logged per capita income for city-service i in month j and year k. γ_i, δ_j, and θ_k are fixed effects for each city-service, month, and year in the data. ϵ_{ijk} is the error term. The city-service fixed effects are essential to our test of request-based inequality. The inclusion of these fixed effects mean that β_1 gives the average estimated difference in average wait times between advantaged and less advantaged communities *for a given service*. With city-service fixed effects, our model is akin to audit experimental studies that estimate the effect of a treatment — i.e., a white alias relative to a black alias — on responsiveness. Month and year fixed effects account for seasonal variation and over-time trends in response times common to both advantaged and less advantaged neighborhoods. Standard errors are calculated using a block bootstrapping method that takes into account the variation within our aggregate measure of wait times. For 1,000 iterations per model, we resample (with replacement) within the groups over which we aggregate — city-service, month, and year — replicate the same aggregation procedure and reestimate the model. We take the standard deviation of the distribution of regression coefficients as the standard error of the coefficient, taking into account the variation across individual request wait times masked via aggregation.

Table 2.2: Estimates of Request-Based Inequality

	(1)	(2)
White Request	−0.079	
	(0.009)	
Rich Request		−0.107
		(0.009)
City-Service FEs	✓	✓
Month FEs	✓	✓
Year FEs	✓	✓
Observations	17,438	17,449
Adjusted R^2	0.736	0.739

Table 2.2 presents our results. Column 1 assess the relationship between race and

responsiveness; column 2 considers income and responsiveness. We find a negative and statistically significant relationship between requests from the advantaged and wait time, suggesting that cities tend to prioritize requests from advantaged neighborhoods. On average, cities tend to respond faster to requests from advantaged neighborhoods than to the same requests from less advantaged neighborhoods. At the same time, the substantive magnitudes of these differences suggest weak evidence of unequal responsiveness. For instance, we find that, on average, cities respond about 1.9 *hours* faster to requests from whiter communities than to the same request from less white communities. Our estimate of differential responsiveness by class is similar, though a bit larger in magnitude; on average, cities respond to about 2.57 *hours* faster to requests from more affluent neighborhoods than to requests from less affluent neighborhoods. In short, our results suggest that — once the good or service itself is accounted for in them model — there are limited differences on the basis of neighborhood economic or demographic characteristics.

Consider these effects relative to the effects found in audit experimental work. As noted, Costa (2017)'s meta-analysis finds that non-minority constituents are about 10 percentage points more likely to receive a response than a minority constituent. While our estimates are not directly comparable to those in audit studies, it is clear that political elites exhibit much more "taste-based" bias in responding to constituent communications than to 311 complaints for goods and services. We conclude that while cities do appear to prioritize requests from the advantaged, the magnitude of these effects are small and for most community issues, trivial and unnoticeable to those making these requests.

2.5 Pathway #2: Need-Based Inequality

It remains possible that — even with limited request-based inequality — cities may prioritize the *needs* of the advantaged. Here, cities consider the different demands of the advantaged and the less advantaged, and respond to those needs particular to more advantaged communities faster than those particular to less advantaged communities. Need-based inequality requires first that there are differences in demand for particular goods and services. Needs must vary along economic and demographic lines. Consider 10 service categories within a given city. If, in the aggregate, neighborhoods above the city median percent white place 10% of their demands — their calls — in each category, and if neighborhoods below the city median do the same, prioritizing needs becomes much harder for cities because the needs and demands are the same across groups. If, however, advantaged and disadvantaged neighborhoods differ in *where* they place their demands, cities can observe these differences and much more easily prioritize the needs of some over others.

We aggregate our data across months and years, and for each city-service, calculate the percentage of requests from neighborhoods above the city median and the percentage from below the city median placed in each service category. Each tells us about the *relative* need of a given good or service for advantaged neighborhoods and less advantaged neighborhoods: it tells us which goods — among the set of possible goods and services to demand from government — the advantaged (and less advantaged) need and demand the most. We take the difference between these two values, such that higher, positive values reflect greater need among the advantaged. We calculate the standard errors for these differences using a similar block bootstrapping technique as in our analysis of request-based inequality.

Figure 2.1 plots these results. Aggregating across time, we find some areas of significant differences in need between more white and less white neighborhoods, and richer and poorer neighborhoods. For instance, in New York, advantaged neighborhoods tend to call about street conditions, street lights, and illegal parking violations. In contrast, less advantaged areas in New York tend to call about heat, hot water, and plumbing issues. In Houston, white communities and wealthy communities call about water issues, traffic sign and signal repairs, and potholes, whereas less white and less affluent areas call about much different concerns: dead animals, illegal dumping and waste collection. Many of the instances where we find substantial differences in need comport with conventional wisdom and expectations. In Washington, DC, needs are quite similar across neighborhoods, except in one case: parking meter issues. We should expect parking meters to be found in areas of higher density and in business districts, where residents to be more white and have higher incomes. While these differences are notable, we also note significant similarity in needs. In some cities, our

Figure 2.1: Relative Need of Goods and Services

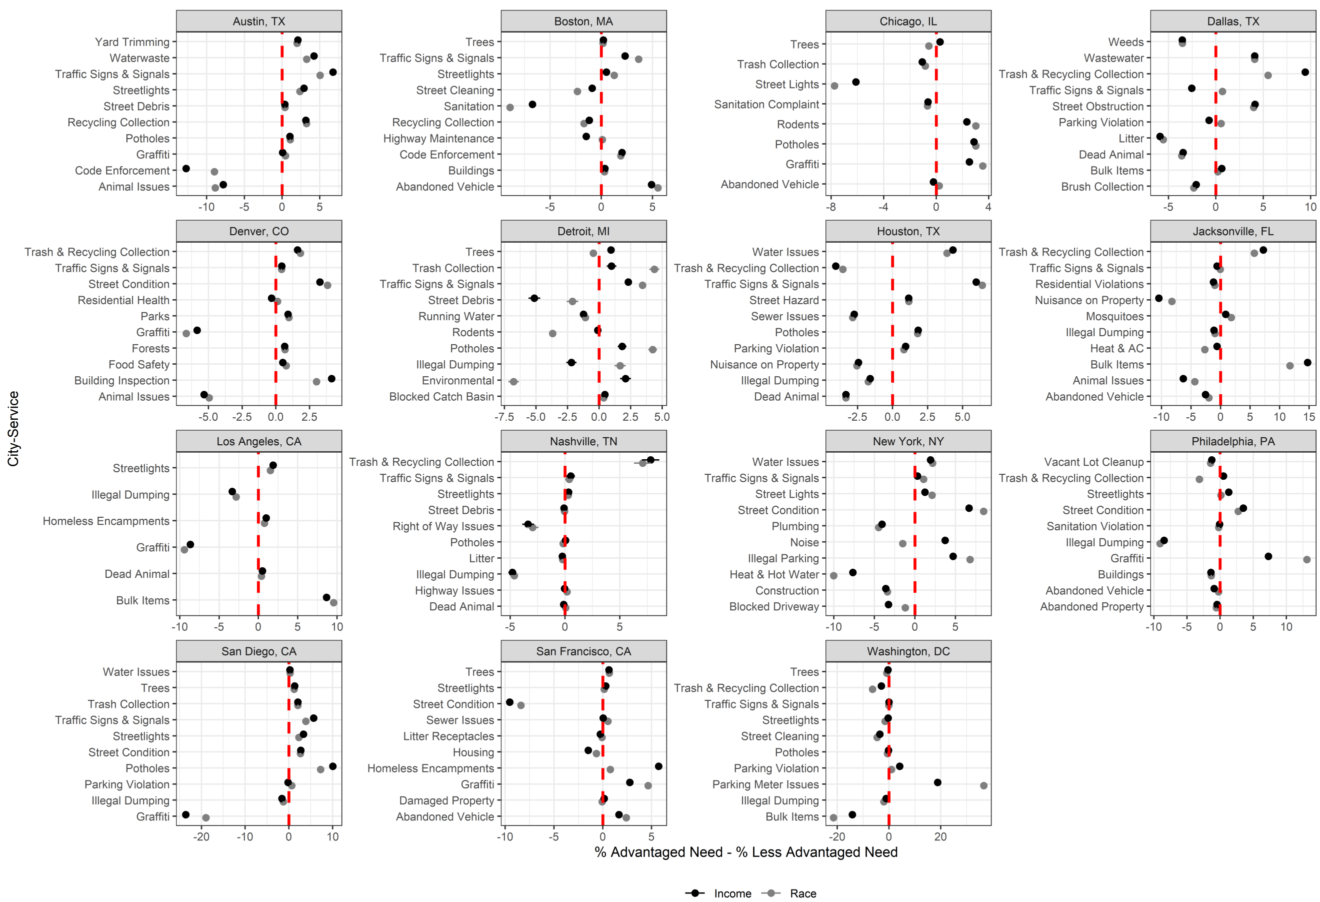

measure hovers around 0% for all services. In many others, the magnitude of the differences — to the extent they exist — are quite small. In short, while we find some evidence that needs often do vary across neighborhoods within cities — suggesting that need-based inequalities may manifest if cities prioritize the needs of the advantaged at the expense of the needs of the less advantaged — our expectations for significant differences may be tempered by the relative uniformity in needs.

As before, we categorize each call as coming from a neighborhood either above or below the city median percent white and the city median logged per capita income. We trim the data — excluding outliers in wait times, and observations with wait times generated with few calls — in the same way as in our test of request-based inequality. We calculate the mean wait time for each city-service, month, and year combination. As in our descriptive analysis of need, we calculate the percentage of requests from neighborhoods above the city median placed in each service category. We do this for each city-service, month, and year combination in the data, and we do the same for neighborhoods below the city median. We again take the difference between these two values, giving us a measure of whether the advantaged or the less advantaged demanded more of the given good or service in each month and year in the data. We test for need-based inequality using the following question estimated using OLS:

$$Average\ Wait\ Time_{ijk} = \beta_0 + \beta_1 \cdot Advantaged\ Need_{ijk} + \gamma_i + \delta_j + \theta_k + \epsilon_{ijk} \tag{2.2}$$

where the $Average\ Wait\ Time_{ijk}$ is the mean wait time for city-service i in month j and year k. $Advantaged\ Need_{ijk}$ is one of two variables. The first is the difference in relative need as described above, measured for each city-service i, month j, and year k. As noted, positive values indicate greater need among the advantaged. The second dichotomizes this variable, taking the value 1 if city-service i is a greater relative need among the advantaged (i.e., if the difference measure discussed above is positive) than the less advantaged in month j and year k. As before, γ_i, δ_j, and θ_k are fixed effects for each city-service, month, and year in the data. ϵ_{ijk} is the error term. The city-service fixed effects are important again, as it allows us to test how *changes* in relative need over-time corresponds to *changes* in wait times. That is, it allows us to assess whether wait times change as a service moves from a "need of the disadvantaged" to a "need of the advantaged" from month-to-month. As before, standard errors are calculated using the same block bootstrapping method.

Table 2.3 shows our results. We find consistent evidence of need-based inequality: cities respond faster to goods and services that white and wealthy neighborhoods tend to ask for, relative to the goods and services that non-white and poor areas to ask for. In particular, our estimates suggest that as advantaged demand for a given service increases over-time, wait time decreases. But, once again, we find that these changes in wait time are small. For example, averaging across city-services, we find that each 1 percentage point increase in white neighborhood demand relative to non-white neighborhood demand decreases average time about 36 *minutes*. Further, white neighborhood needs tend to be responded to about 3.8 hours faster than less white neighborhood needs. The point estimates are remarkably similar across models of race and income, with slighter larger effects for affluent needs. As mentioned, we are unable to tease apart the mechanisms for these differences. Indeed, these effects may reflect government discretion and discrimination, or government capacity and the time and resources needed to meet a demand. Nevertheless, we can conclude that the extent of need-based inequality is minimal — perhaps even more minimal than our findings of request-based inequality. At worst, government is 4 hours faster to respond to the services most needed and demanded by affluent neighborhoods. At best, because of the relative similarity in needs, the potential for disparities as large as 4 hours are quite unlikely.

2.6 Conclusion

Services are an important component of government work (Cain, Ferejohn, and Fiorina 1987; Fenno 1978) and are central to the way in which citizens evaluate government and their representatives in office. For example, Harden (2016) and Tucker (2019) both show that constituent service affects vote choice in elections, perhaps even more so than policy

Table 2.3: Estimates of Need-Based Inequality

	(1)	(2)	(3)	(4)
% White Need - % Non-White Need	−0.025			
	(0.001)			
White Need		−0.157		
		(0.018)		
% Rich Need - % Poor Need			−0.026	
			(0.001)	
Rich Need				−0.203
				(0.018)
City-Service FEs	✓	✓	✓	✓
Month FEs	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓
Observations	8,763	8,763	8,763	8,763
Adjusted R^2	0.754	0.754	0.754	0.754

considerations. In the aggregate and at the local level, goods and services affects outcomes, too. Burnett and Kogan (2017) show that pothole complaints correspond to decreases in incumbent vote share in local elections. Political elites tend to respond to these incentives, too, as cities respond faster to 311 calls in city council districts with incumbent politicians seeking reelection (Christensen and Ejdemyr 2018). More generally, elected officials may even prioritize service work over policymaking (Butler, Karpowitz, and Pope 2012; Parker and Goodman 2009).

Despite the centrality of service to citizens and elites, most work on representation focuses on policy. Research has documented strong relationships between public opinion and policy (e.g., Caughey and Warshaw 2018; Erikson, Wright, and McIver 1993). Nevertheless, the opinions of the advantaged — white and affluent citizens — appear to be better represented in government than those of the less advantaged (e.g., Gilens 2012; Griffin and Newman 2008; Schaffner, Rhodes, and La Raja 2016). In this paper, we build off of this work and provide one of the first contemporary tests of unequal responsiveness in city goods and services. To do so, we leverage a new dataset of 25 million 311 service requests, and an empirical approach that approximates new work on policy and representation.

Our results are consistent in direction with work on unequal policy responsiveness: cities prioritize requests from the advantaged (request-based inequality) and the needs of the advantaged (need-based inequality). Yet, at the same time, we caution against interpreting our results as a ringing indictment against city governments. The magnitude of the differential responsiveness we show is small. Though an extra hour of a downed streetlight in a business district would be quite inconvenient, in most cases, and for most goods and services, the average difference in response time would go unnoticed. In relation to past work, we find a middle-ground between research suggesting significant racial biases, and work suggesting that city service distribution follows technocratic administrative procedures. As noted, the magnitude of these findings stand in stark contrast to work on unequal policy responsiveness. For instance, Gilens (2005) shows when the rich and the poor disagree on policy, the probability of a proposed policy change being implemented increases about 30 percentage points as support among the rich increases. In comparison, low and median-income support has no effect on the probability of policy change. While these substantive findings do not map neatly onto our own empirics, the differences are nevertheless notable. We hesitate to speculate on why there lacks a substantive effect of race and class on goods provision, but note that bureaucrats and elected officials may be responding to different incentives. Elected officials may be responding to perceived electoral incentives to respond to particular groups. In contrast, as past qualitative and empirical work indicates (e.g., Mladenka 1980), city services follow bureaucratic decision-rules, and bureaucrats may prioritize upholding professional values and standards of performance. As a result, though elected officials clearly monitor service provision and see to it that demands are met (Christensen and Ejdemyr 2018), government action on services may still be divorced from much of the politics that drives policymaking.

We see our work as underscoring the importance of political participation — here via a government-run "complaint" line. A substantial body of work shows that political elites tend to respond to the preferences of those who vote (Griffin and Newman 2005; Martin 2003), and that patterns of unequal policy representation are attenuated as participation among the less privileged increases (e.g., Griffin and Newman 2007; Hajnal 2009; Hill and Leighley 1992; Hill, Leighley, and Hinton-Andersson 1995). Our findings corroborate this view, and suggest that non-voting participation — e.g., direct communication with government officials about needs and demands — may be particularly effective at achieving equal voice and representation.

Nevertheless, neighborhoods may not be fully equal in access to city goods and services. Our data only allow us to address unequal *responsiveness*. We cannot address government action unrelated to 311-delivered demands, unequal conditions, or inequities in 311 participation that may exacerbate unequal conditions. First, cities surely fix streetlights, repair roads and bridges, and clean graffiti without the prompt of a citizen call. City workers may be "out-and-about" improving advantaged neighborhoods without being asked to do so, but may not do the same in less advantaged neighborhoods. Related, less advantaged neighborhoods may also simply be in worse condition — relative to more advantaged areas — at baseline. de Benedictis-Kessner (2019) finds evidence for "inequality in conditions." He draws on a unique dataset of Boston sidewalk conditions, and finds that sidewalks in minority and poor neighborhoods are in worse condition. And while greater 311 use in these communities attenuates these differences in conditions, inequities in participation — the advantaged tend to use 311 more than the advantaged (see also Feigenbaum and Hall 2016) — mean that sidewalks in non-white and poor areas are still in worse shape on average. As such, though we find participation fosters relative equality in responsiveness, inequities may persist because of differing baseline conditions and because the costs to political participation remain too high for many less advantaged communities to overcome. Given the latter, less advantaged communities may not benefit as much from a system that emphasizes equality in response. In short, though our study offers a direct test of responsiveness to citizen demands, we do not rule out other ways in which access to collective goods and services may be unequal. Future work should focus on the possibilities identified above.

Finally, our work suggests a pathway for less bias in government. Our intent in this paper is not to make a global claim about discrimination in government. More simply, we do not argue that there is no discrimination in government. For one, audit experimental work suggests that important biases exist (e.g., Costa 2017). However, our work does suggest that bureaucratized service provision may be essential to limiting the discretion of individual political actors and in turn, achieving equal responsiveness. In our setting, citizen-contacts are routed from a central command center to particular agencies of government. Processes and procedures are in place to respond to these demands, and cities tend to follow these protocols. These processes stand in stark contrast to the audit experimental setting, where individual elites are given the choice over whether to respond or not. Normatively, our work suggests that government entities — cities and others — should adopt similar practices in handling all kinds of constituent service work, such as constituent communication.

2.7 Supplemental Appendix

Figure 2.2: Clustering of Close Dates — Day of Week

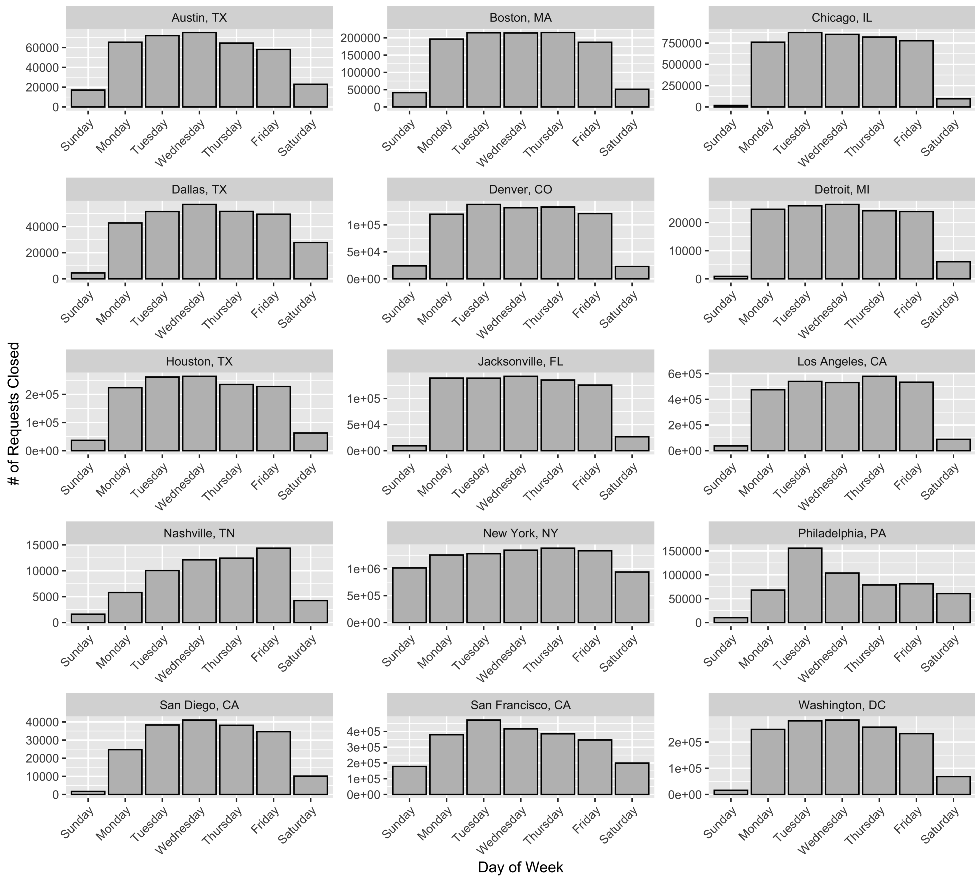

Figure 2.3: Clustering of Close Dates — Day of Month

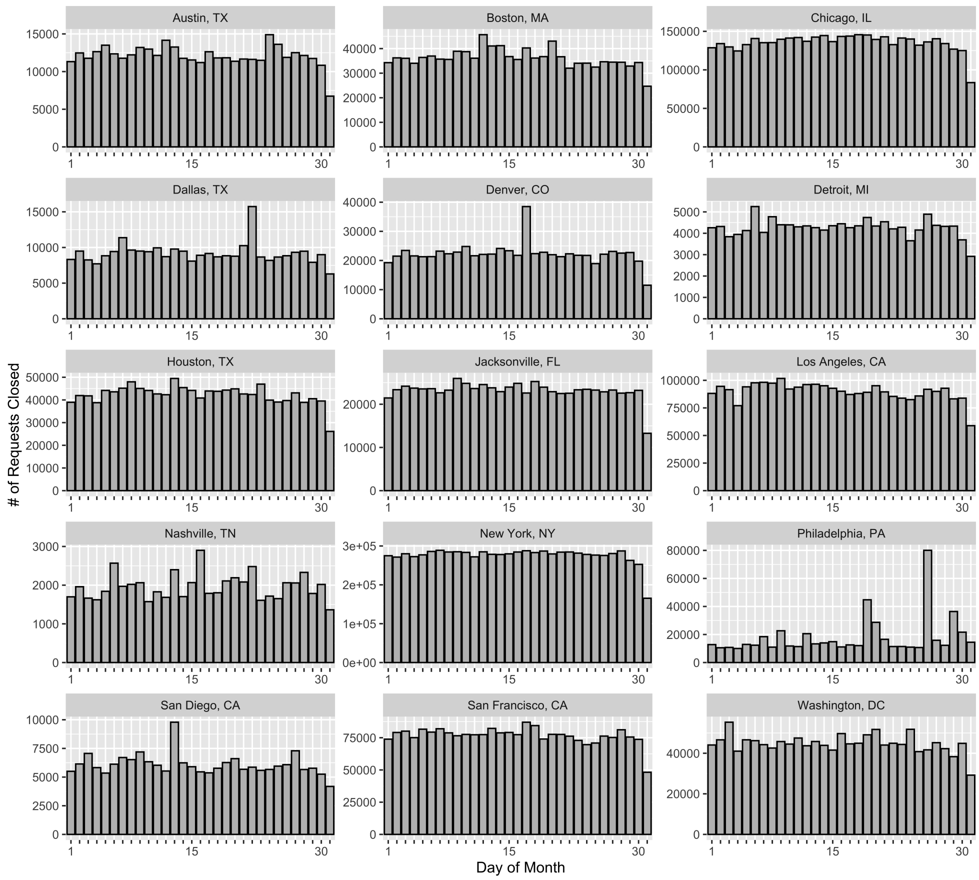

CHAPTER 3

Gerrymandering Incumbency: Does Nonpartisan Redistricting Increase Electoral Competition?

What can explain the electoral security of legislative incumbents in the U.S.? Popular accounts often point to the insulating effects of redistricting, with legislators drawing their own district boundaries to ensure reelection. Characteristic of this concern, a recent *New York Times* column argues that

> "political control of the map-drawing process ... aided by sophisticated computer programs that can micro-target political affiliation ... has stuffed Congress and state legislatures with increasingly safe seats, making lawmakers difficult to dislodge no matter what they do" (Hulse 2016).

This view is not uncommon, including among scholars who worry that such insulation may allow polarized legislators to enact policies out of step with public opinion, or to obstruct policymaking (e.g., Carson, Crespin, and Williamson 2014). Given this concern, reformers are pushing to wrest control of redistricting away from partisan legislatures, and to empower nonpartisan commissions as a way to increase electoral competition and improve representation. And these efforts are gaining momentum, in part due to their broad public support.[1] Voters in Ohio and South Dakota recently considered ballot initiatives to create independent commissions, with Ohio passing its measure. Petitioners in Illinois collected nearly half a million signatures to initiate a ballot question, and citizens in Colorado filed a similar proposal before both were rejected by state courts.

[1]Polls in a number of states (e.g., NC) show large majorities of voters support non-partisan redistricting.

When redistricting is controlled by elected officials, political considerations can loom large. Legislatures with unified party control may gerrymander to increase the majority's seat share, potentially insulating incumbents of their own party. Divided legislatures or bipartisan agencies may log roll maps to protect incumbents of *both* parties. In contrast, independent commissioners need not compete in elections or for political appointments, and thus promise a redistricting process divorced from political pressures and ambition. Consequently, commissions are expected to draw districts that more closely adhere to constitutional constraints (e.g., compactness, minority non-dilution), and in turn, expose incumbents to greater electoral risk (Edwards et al. 2017; Lindgren and Southwell 2013).

Yet, empirical evidence about the consequences of adopting independent commissions is decidedly mixed. Some studies find that districts drawn by legislatures are less competitive than districts drawn by 'politically-neutral' commissions (Carson, Crespin, and Williamson 2014; Grainger 2010; Lindgren and Southwell 2013; McDonald 2006). Yet, other work indicates that redistricting institutions have little effect on electoral competition, neither decreasing incumbent win margins nor increasing challenger entry (Abramowitz, Alexander, and Gunning 2006; Forgette, Garner, and Winkle 2009; Masket, Winburn, and Wright 2012). Some scholars uncover countervailing trends – incumbents may be *less likely* to face strong challengers when legislatures redistrict, but paradoxically get unseated at similar or higher rates (Cottrill 2012; Gelman and King 1994).

We offer new evidence on the link between redistricting and competition. Expanding on innovations by Chen and Rodden (2013) and Altman and McDonald (2015), we compare measures of district competitiveness between enacted redistricting plans and two sets of counterfactual maps that *could have been, but were not* enacted during the 2010 cycle. Specifically, we observe whether independent commissions adopt maps that are as competitive as the plausible set of alternatives *for the same state*, comparing this to the distortion found in states where legislatures conduct redistricting. We first analyze maps built from randomized redistricting simulations that only incorporate equal population and contiguity requirements (Chen and Cottrell 2016). We augment this with an original collection of all alternative maps officially debated by state redistricting institutions (Altman and Mc-

Donald 2015). Together, these data provide the strongest evidence to date about whether independent agencies increase electoral competitiveness.

We find that redistricting *sustains* the electoral security of incumbents in general. Maps that are eventually implemented show consistently larger incumbent vote advantages than would be expected given geographic and legal-political constraints. Most notably, we find that the incumbency bias persists regardless of whether politicians or citizens are in charge of redistricting. Independent redistrictors produce virtually the same degree of insulation as (bi)partisan plans devised in legislatures or by political appointees. Our findings suggest that replacing parties with independent redistrictors is unlikely to increase competition in legislative elections, offering little remedy to contemporary concerns about representation.

3.1 Data and Empirical Strategy

We draw on data collected by Levitt (2010), and our own analysis, to categorize each state's actual redistricting procedure for each legislative body. Our analyses focus on five categories: Democratic plans (legislature with unified Democratic control), Republican plans (legislature with unified Republican control), bipartisan plans (legislature under divided government or a political commission), court plans, and independent commissions. Though simplifying, this categorization captures the range of institutional devices of most interest to political actors, reformers, and scholars of redistricting (Carson, Crespin, and Williamson 2014; Chen and Cottrell 2016; Edwards et al. 2017; Gelman and King 1994; Levitt 2010).[2] In 2010, most U.S. House maps were drawn by Republican-controlled state legislatures. While most state-level jurisdictions were drawn by bipartisan legislatures or politician commissions, many more were created by Republican- rather than Democratic-controlled legislatures. Independent commissions redistricted only a handful of states.

Rather than compare final maps across states with and without independent commissions (Edwards et al. 2017; Forgette, Garner, and Winkle 2009), we adopt an approach

[2] For more detail on how states conduct redistricting, see http://redistricting.lls.edu.

similar to Chen and Rodden (2013), Chen and Cottrell (2016), and Altman and McDonald (2014; 2015) in their analyses of partisan gerrymandering.[3] These authors make comparisons between the final maps adopted by redistricting agents, and counterfactual sets of alternative maps that *could have been adopted within the same state*. Using 2000 cycle data, Chen and Rodden (2013) simulate randomly-generated redistricting plans that use only equal apportionment and contiguity as criteria. Political motivations, like partisanship or incumbency, are never considered. In being randomly drawn, these maps should capture whatever baseline electoral security we would expect absent any 'politics' in redistricting. For each simulated plan, Chen and Rodden (2013) analyze precinct-level vote data to determine which hypothetical U.S. House, Assembly, and State Senate districts would have received more Bush or Gore votes in the 2000 election. From this, the authors estimate the expected share of districts likely to elect Republicans for each simulation, and compare this baseline to the number of Republican-won districts in the final enacted plans. Chen and Cottrell (2016) extend this analysis for House districts in the 2010 cycle using 2008 presidential vote. Any surplus of party seats is interpreted as evidence of partisan gerrymandering.

We conduct a similar analysis of electoral competition. Initially, we analyze the simulated House maps produced by Chen and Cottrell (2016) for redistricting in 2010. Our main measure of competitiveness, Expected Margin of Victory, is the absolute value of the party win margin in a district, or roughly, the surplus of votes a winning incumbent might expect on average for a given plan. This measure incorporates the logic that incumbents with larger win margins would survive bigger electoral downturns year-to-year, and deter strong challengers from contesting their seats.[4] Following prior work, we use the 2008 two-party presidential vote (Chen and Cottrell 2016; Chen and Rodden 2013).[5]

---

[3] Cross-state comparisons are likely misleading. Independent redistricting is only found in Western states, with lopsided majorities that naturally dampen district competitiveness.

[4] We find identical results examining median win margins, counts of competitive seats at empirically-derived thresholds, and seat flip probabilities. See the Online Appendix for details.

[5] We use 2008 elections because these are not conditional on the redistricting process, while elections in 2012 and after would be 'post-treatment' to enacted plans. Many things could change across precincts in being assigned to a particular district and map, and we cannot observe precinct voting under any alterna-

Despite the usefulness of simulation approaches, we recognize their limitations. By design, simulations produce districts that adhere to minimalist constitutional requirements, approximated algorithmically. These exclude real legal and political constraints mapmakers face, including mandates to preserve majority-minority districts. A possible consequence is that simulations may overly concentrate Democrats in urban districts, thus overstating the baseline effect of geography on redistricting (Altman and McDonald 2014).[6]

Given these concerns, we introduce and examine a novel data source – the set of all alternative maps redistrictors publicly considered during the apportionment process. In many states, regardless of the redistricting procedure in place, redistrictors solicit map proposals from the public, including legislators, researchers, citizens, and interest groups.[7] In some states, these maps are made publicly available, alongside all the proposals introduced as legislation or considered during the meetings of the redistricting commissioners. We follow a similar approach innovated by Altman and McDonald (2014; 2015) in their study of party gerrymandering in Florida, Ohio and Virginia. We collected 1,627 maps across 15 states that meet equal population and contiguity constraints from the full set of data made publicly available by state legislatures or redistricting commissions. Importantly, the states in our sample vary in the way redistricting is conducted.[8]

To construct measures of 2008 district win-margins for each alternative map, we use both block-to-precinct data made available by McDonald and Altman (2011) and voting

tives redistrictors did not implement. We use 2008 party registration for Florida and Nevada as precinct presidential vote data were not available for those states.

[6]However, such an effect would likely *understate* the degree to which redistrictors adopt plans that insulate incumbents above what would emerge as a product of population or geography.

[7]In North Carolina, community meetings were held across the state to collect public input. Citizens were encouraged to propose maps online, and comment on official proposals.

[8]We collected virtually every map publicly considered in AK, AZ, CA, CO, FL, ID, MT, NC, NM, NV, OH, SC, TX, VA, and WA. Any maps restricted to the public would be excluded. Our convenience sample includes every non-partisan state, but no Democratic legislative plans, and only a handful of Republican, court and bipartisan maps. This may limit the generalizability of our findings if we expect that Democratic legislatures redistrict in fundamentally different ways than Republicans. We have no *ex ante* reason to suspect this as other research suggests that Democrats (e.g., Maryland) and Republicans (e.g., Texas) both gerrymander when possible. Future research should examine this, and other generalizability concerns.

district shapefile data collected by Ansolabehere, Palmer, and Lee (2015). In contrast to simulations, these maps incorporate the fuller set of legal and political considerations, like minority vote protections, redistrictors must balance. These counterfactuals allow us to assess the real tradeoffs redistrictors made in selecting final maps from among the reasonable alternatives, and especially how these choices differ depending on who is in charge of the process.

3.2 Results from 2010 U.S. House Redistricting

Turning to analysis of 2010 congressional maps, Figure 1 plots the distribution of (a) simulated House plans for the 42 states analyzed by Chen and Cottrell (2016), and (b) both simulated and alternative proposals for 12 (of 15) states which publicized alternative House plans. In both Figures 1a and 1b, grey dots represent the average district Expected Margin of Victory for all House seats in a state under one simulated plan. In Figure 1b, dark blue dots indicate the average seat Expected Margin of Victory for each alternative plan. Circles indicate the mean margin of victory under the final adopted map, with circumferences proportional to the standardized distance from the average of the distribution of simulated or alternative maps (interpretable as a z score).

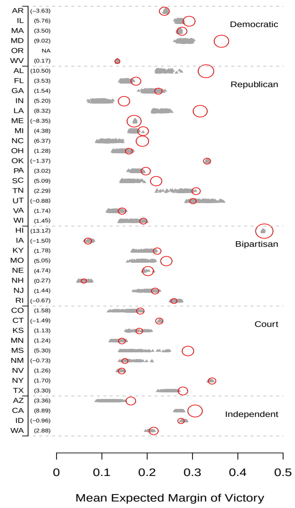

(a) House Simulations Only

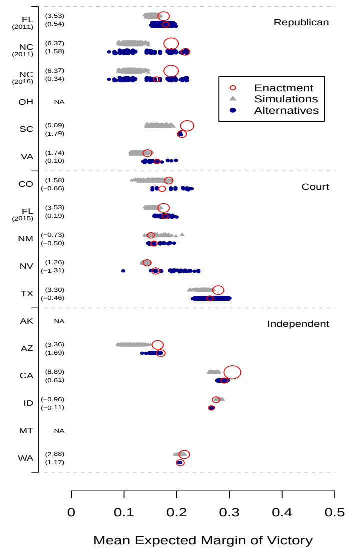

(b) House Simulations & Alternatives

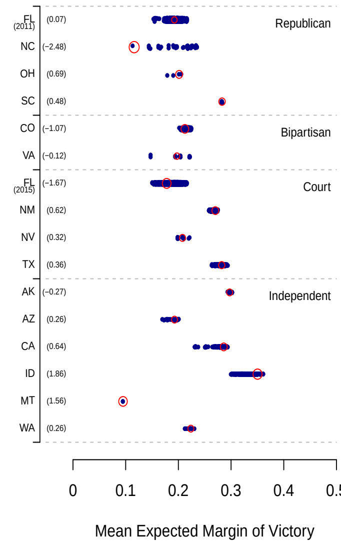

(c) State Senate Alternatives Only

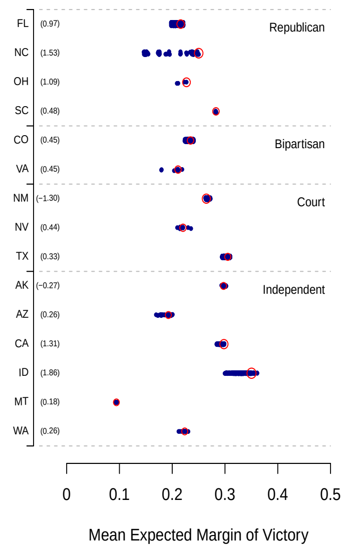

(d) State Assembly Alternatives Only

Figure 3.1: Competitiveness of U.S. House, State Senate and Assembly Districts for Enacted, Simulated and Alternative Plans in 2010

Shown in Figure 1a, we find that redistrictors typically adopt U.S. House maps in 2010 that are consistently less competitive than simulated plans. Across the states, adopted plans yield vote margins that are 2.4 percentage points safer than the average simulated map. Further, 43% of states have final maps that are less competitive than every single simulation, with the median plan being less competitive than 99% of simulations. Turning to Figure 1b,, we see again that redistrictors enact plans that are much less competitive than the alternatives debated during the mapmaking process. Of the non-court plans (i.e., by legislature or commission), expected win margins under the final maps were 1.2 percentage points less competitive than the alternative proposals, with 33% of these House plans being safer than nearly all the maps proposed to redistrictors. Perhaps most striking, the median final (non-court) House plan is more insulated than 76% of the remaining, unimplemented proposals. This distortion amounts to roughly 8,529 to 17,058 additional votes for each incumbent, and while modest, substantiates the general concern that the redistricting process may dampen electoral competitiveness in Congress.

Examining both simulated and proposed maps, we find this distortion persists *regardless* of whether redistricting is done by incumbents or independent commissioners. For instance, highly contested, Republican-drawn plans in Florida, North Carolina and South Carolina produced districts with expected win margins that were 2.1, 7.5, and 6.3 percentage points greater, respectively, than the average district under each state's simulated alternatives. A similar distortion is found (0.4, 7.0, and 0.4) when comparing enacted plans to those considered by the legislatures. Yet, in Arizona, California, and Washington – states with independent redistricting commissions designed to be insulated from political considerations – we see relatively little improvement. Incumbent vote margins under the final plans enacted in Arizona, California and Washington were 4.7, 3.5, and 0.7 points less competitive, respectively, than those recovered in the average simulated districts, and even 1.0, 0.3, and 0.2 points safer than the alternative plans considered by commissioners.

We next compare differences in the level of competitiveness found between each state's adopted and counterfactual plans, across the different redistricting institutions. Here, we compute the proportion of comparison plans that are less competitive than the adopted map overall, using the average of each plan's Expected Margin of Victory. We consider whether states with independent commissions adopt maps with relatively competitive districts compared to states where politicians conduct redistricting, using difference-in-means

t-tests. Initially, we find that politicians produce maps that are safer than 77.1% of the simulated alternatives for their states, in terms of average win margins. Yet, this insulation is virtually identical to that uncovered for states with independent commissions, with 74.9% of simulations ($p = 0.935$) being more competitive than adopted plans. An analysis of the publicized alternative maps indicates again that independent commissions choose maps that are as uncompetitive as those enacted by politicians. Legislative plans are less competitive than 77.3% of the alternative map submissions overall, compared to 76.2% ($p = 0.949$) for non-politician maps. In sum, independent commissions do not draw House maps that encourage any greater electoral competition than partisan legislatures.

3.3 Results from 2010 State Assembly and Senate Redistricting

We restrict analysis of State Assembly and Senate redistricting to the alternative proposals considered, but not enacted in 2010. Figures 1c and 1d compare expected competitiveness of final plans to the alternatives publicly considered for the Senate and Assembly, respectively. Results mirror the findings for the U.S. House. With few exceptions, competition in adopted plans is lower than that produced under the plausible alternatives redistrictors considered. Moreover, we find no systematic improvement in competition from non-partisan redistricting. Independent commissions enact Senate plans that are considerably *less competitive* relative to alternatives (0.659), than (bi)partisan legislatures (0.328; $p = 0.058$), though this somewhat weakly reverses for Assembly maps (0.559 - 0.785; $p = 0.111$). Counter to expectations, our results again confirm that independent commissions do not consistently increase electoral competition at the statehouse.

3.4 Discussion and Conclusion

Redistricting is an oft-cited source of American political dysfunction. Reformers posit that most incumbents, in facing minimal electoral competition, have leeway to discount constituent demands, encouraging polarization and poor representation. We show that simply changing how legislative districts are drawn, even in a process ostensibly divorced of political ambition, may not bring about competitive elections, at either the federal or state level. Redistricting marginally dampens electoral competitiveness as a whole, but these effects are similar regardless of whether maps are drawn by (bi)partisan legislatures or independent commissions. Independent redistricting may foster competition in other complex ways (e.g., incumbent retirement). Importantly, our data and design can only assess how redistrictors systematically put their thumb on the scale prior to adoption, and not what downstream consequences different choices may have in equilibrium.

Our work introduces novel data to the study of redistricting – the set of alternative maps considered by redistrictors. We believe these are representative of the set of maps that *could* have been adopted, but were not. Of course, it is possible that mapmakers (un)intentionally censor their deliberations over plans, such that our data overweight especially competitive maps relative to all feasible alternatives. If so, we would also overestimate the dampening 'effect' of redistricting on election competition overall. We doubt redistrictors censor maps strategically, and think it particularly unlikely that they publicize only those that make the final enacted plans look 'worse' than most available alternatives. Still, we suggest caution in concluding that redistrictors systematically insulate incumbents as this requires an untestable assumption that our simulated and collected maps approximate all feasible alternatives. Our study relies on a weaker parallel trends assumption to interpret the marginal impact of redistricting modes: roughly that any bias in collecting or simulating counterfactual maps is uncorrelated with how states conduct redistricting. We strongly doubt that non-partisan redistrictors censor the least competitive alternatives they consider, while legislatures obscure the most competitive ones. Simulations guard against this concern by using fixed constraints and geographies that ignore politics by construction.[9] Future research should assess this concern more extensively, however, with particular focus on how mapmakers determine which plans to consider or implement.

In closing, we suggest that the apparent differences in district competition across states

[9]California bolsters this interpretation, switching from partisan to independent redistricting without any appreciable change in relative competitiveness of enacted maps.

with and without independent redistricting may be an artifact of natural or demographic variation, rather than any politics in the process. Or, it may be that partisan and political forces play a larger role in citizen redistricting than anticipated by scholars or reformers (Pierce and Larson 2011). Notably, we also find when courts intervene, they produce very competitive maps, suggesting politics and not geography are to blame. Further research is needed to uncover precisely why independent redistricting falls short of expectations, especially as reform efforts gain momentum across the country. Greater understanding of the politics of non-partisan redistricting can help improve the match between institutional design and the intended effects of political reform. Nevertheless, our findings caution against replacing state legislatures with non-politician commissions wholesale, at least solely on the basis of increasing electoral competition.

3.5 Supplemental Appendix

3.5.1 Seat Flip Probabilities and Counts for 2010 Redistricting

We analyze a number of alternative measures of competition. One such measure, Seat Flip Probability, follows an approach similar to Chen and Rodden (2013), to estimate the likelihood a seat will change party control given an expected vote margin under various plans. This measure is built from a simple bivariate model of Congressional party seat switches. Denote a party seat switch indicator Y_d for congressional district d, which equals 1 if the seat changes party control between 2012 and 2014, and 0 otherwise. Given binary seat flips, we use a logit model to regress Y_d on 2012 (absolute) district vote win-margin V_d. Precisely, denote R_d to be the two-party Republican vote share in a district. Then $V_d = \text{abs}\{R_d - (1 - R_d)\}$. The resulting logit linear model is $Pr(Y_d = 1) = \Phi(-1.837 - 9.646 \times V_d)$, where Φ is the logit density function. From this model, we estimate predicted flip probabilities for each district using either observed or hypothetical vote margins across each redistricting plan. Data used for this measure originate from the Congressional Quarterly (2014). The results are virtually identical to those drawn from average win margins, and are presented in Figure 3.2.

Table 3.1: Minimum and Maximum Party Vote Shares for States, 2002 to 2010

State	Min. R	Max R.	Dif. R	Min. D	Max D.	Dif. D	Avg. τ
AK	50.20	74.51	24.31	17.00	45.00	28.00	26.16
AZ	44.48	55.28	10.81	40.98	51.62	10.64	10.72
CA	36.64	42.62	5.99	54.10	60.30	6.20	6.09
CO	42.44	52.95	10.51	43.04	56.01	12.98	11.74
FL	48.38	60.97	12.59	38.91	59.37	20.46	16.52
ID	55.95	70.10	14.15	28.65	39.80	11.15	12.65
MT	51.50	64.62	13.12	32.40	46.30	13.90	13.51
NV	40.83	57.70	16.87	36.98	52.13	15.16	16.01
NM	39.60	55.78	16.18	44.19	56.13	11.94	14.06
NC	44.85	53.88	9.03	45.38	54.71	9.32	9.18
OH	45.74	55.75	10.01	42.88	53.31	10.42	10.21
SC	49.98	57.27	7.28	40.88	49.22	8.33	7.81
TX	49.13	56.07	6.94	41.15	49.43	8.28	7.61
VA	50.77	55.24	4.47	44.13	47.83	3.71	4.09
WA	39.13	51.84	12.70	48.16	60.32	12.16	12.43

A nonparametric way to measure incumbent vulnerability is to estimate the widest possible partisan shift we observe across legislative districts in a state, and then calculate the number of seats that would switch party control under different redistricting plans if that maximal vote swing occurred. We estimate this maximal interval by calculating the maximum and minimum Democratic (or Republican) vote share across all districts for each state over the five elections from 2002 to 2010. These minimum and maximum party vote shares are displayed in Table 3.1. To illustrate, in Alaska's 2002 midterm, the best year for Republicans in that decade, Don Young won 74.51% of the vote, while his Democratic opponent got 17.28%. Six years later in 2008, Don Young's margin was reduced to 50.2%, with Ethan Berkowitz getting 45% of the vote. The remaining 2004, 2006 and 2010 elections in Alaska all fall somewhere in between. Thus, we consider a maximal shift of 26.16 percentage points to bound the most extreme volatility in electoral security Don Young, and other Alaska legislative incumbents, might experience in a worst case election.

As seen in Table 3.1, states differ substantially in the size of their maximal party vote swings. Consequently, we doubt that using some fixed competitiveness threshold is appropriate. For instance, researchers will often set some pre-determined value, like 5% or 10%, and then examine which seats are won by a margin smaller than that competitiveness threshold. The analogy here would be to add (or subtract) 2.5 or 5 percentage points to all incumbent's vote share and see how many districts would flip party control. Yet, the values in Table 3.1 illustrate that using a hypothetical swing (say of 10%) would simultaneously under- and over-estimate vulnerability across different states. Instead, we use this maximal shift to measure the state-specific vulnerability incumbents experience, when calculating how often legislative seats would switch party control across different redistricting maps. To do this, we add or subtract each states' maximal party shift to incumbents' Expected Margin of Victory under enacted, simulated and publicized plans, and observe how many seats would shift across the different plans. These results are presented in Figure 3.3, and mirror our main findings.

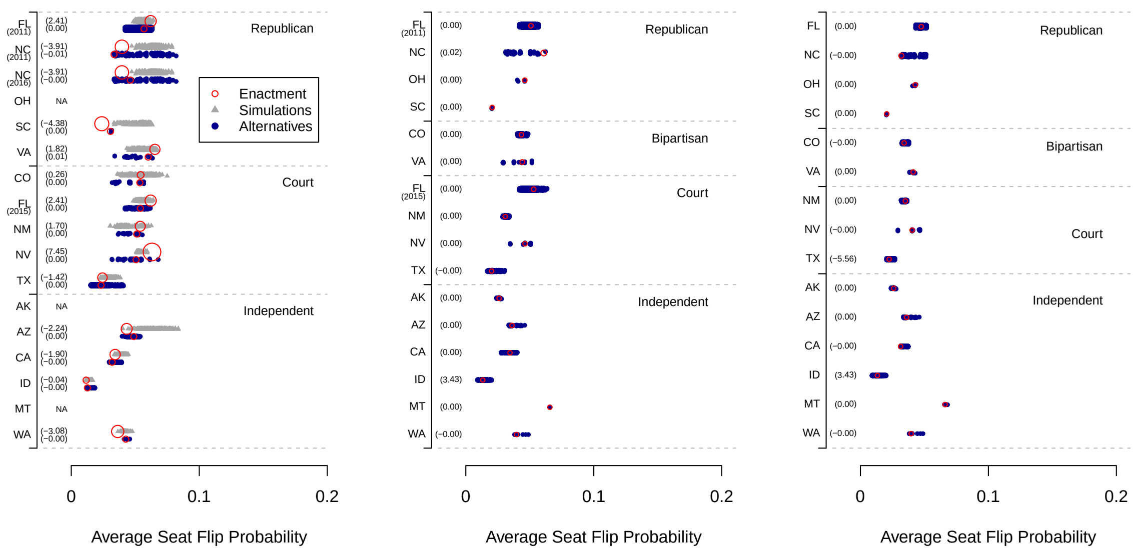

(a) House Simulations & Alternatives

(b) State Senate Alternatives Only

(c) State Assembly Alternatives Only

Figure 3.2: Incumbent Loss Probabilities of U.S. House, State Senate and Assembly Districts for Enacted, Simulated and Alternatives Plans in 2010

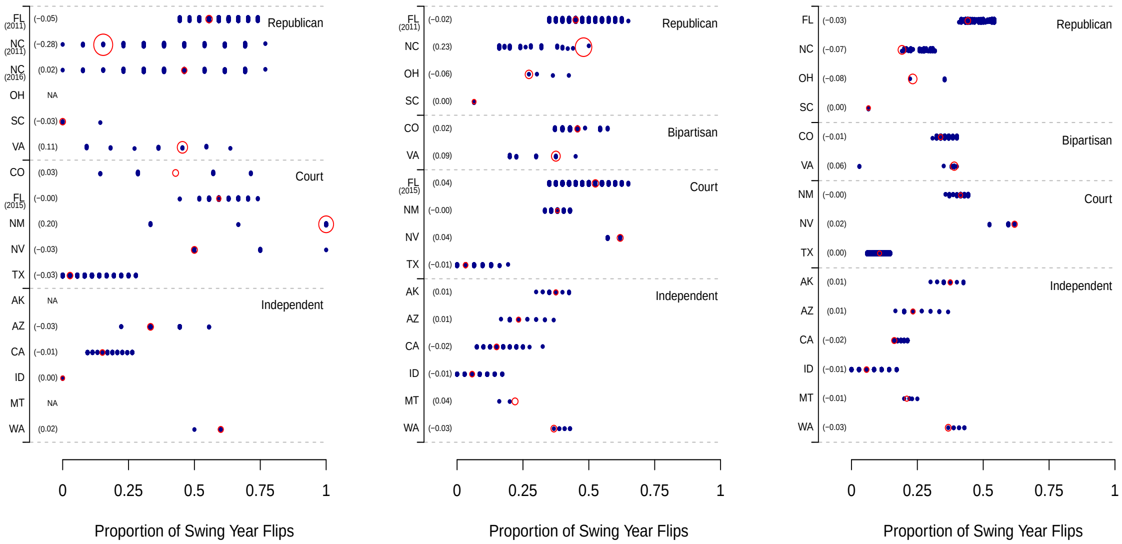

(a) House Alternatives Only (b) State Senate Alternatives Only (c) State Assembly Alternatives Only

Figure 3.3: Counts of Incumbent Seat Losses of U.S. House, State Senate and Assembly Districts for Enacted, Simulated and Alternatives Plans in 2010

3.5.2 Results from 2000 U.S. House, State Assembly and Senate Redistricting

We focus on the 2010 redistricting cycle in the main paper because we have simulated maps and alternative maps available to compare to the actual redistricting plans in a large number of states. Chen and Rodden (2013), however, also have made available simulated data for the 2000 redistricting cycle at both the federal and state level in a sample of states. These results are presented in Figure 3.4. Our 2000 results are consistent with the findings using simulations in 2010 for the U.S. House, and broadly mirror our other findings using publicized statehouse maps in 2010. A shortcoming here is that we only have data for one state (ID) with an independent commission, and so cannot say much about non-partisan redistricting in the 2000 cycle.

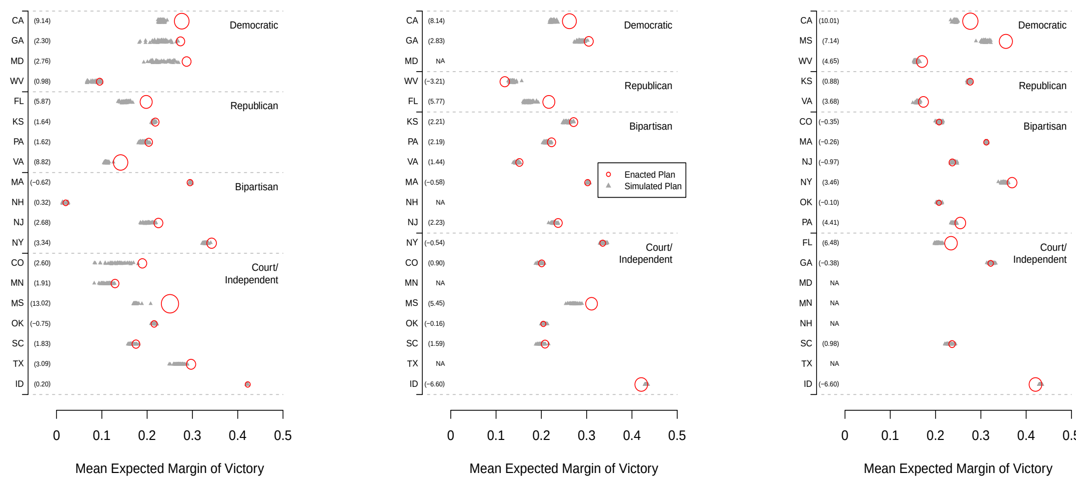

(a) House Simulations (b) State Senate Simulations (c) State Assembly Simulations

Figure 3.4: Competitiveness of U.S. House, State Senate and Assembly Districts for Enacted and Simulated Plans in 2000

3.5.3 California's Redistricting Experiment

Estimating the marginal impact of independent redistricting using our data and approach requires a *parallel trends* assumption. Specifically, this does not require that simulations or the publicized counterfactuals be an unbiased sample from some true distribution of all feasible alternatives, but only that any bias in our comparison maps be independent of how states conduct redistricting. For example, this would be violated if it were the case that independent redistrictors under-reported the non-competitive alternatives they debated, while partisan legislators over-reported these uncompetitive maps when making their deliberations public. A similar sort of bias would need to be present in the simulations as well for our findings to be faulty: that simulating using minimalist constitution conditions also under-samples uncompetitive maps in non-partisan states, but over-samples them in partisan places. We doubt either of these, much less both occur.

An admittedly imperfect test of this looks at simulations data for California, which switched from having a partisan legislative process in 2000 to an independent commission in 2010. Such a test mirrors that found in Grainger (2010), who analyzed multiple redistricting cycles in California. While many things are changing between 2000 and 2010, the basic geography of the state did not, including the distribution of census blocks and voting precincts, which are the fundamental units redistrictors use to draw maps. We would expect that whatever geographical factors bias simulations to over- or under-sample competitive maps are unlikely to change much for the same state within a decade. If this stability is the case, then we can look at the relative competitiveness of implemented plans for California in 2000 and 2010 against simulated alternatives drawn from the same simulation approach for both years.

In doing so, we observe that California's 2010 non-partisan map was 8.89 points less competitive than the average simulated map in terms of standardized win margins, and less competitive than all the simulations. For 2000, the party legislative map was 9.14 standardized points less competitive than the average simulation, and also outside the range of simulations. We can compare this estimate to the changes in the relative competitiveness

Table 3.2: Comparing the Improvement in Estimates of Relative Competitiveness from Adopting Non-Partisan Redistricting in California to Other Non-Changing States

State	2000 Difference	2010 Difference	Diff.-in-Diff.
MS	13.02	5.30	-7.72
VA	8.82	1.74	-7.08
FL	5.87	3.53	-2.34
NY	3.34	1.70	-1.64
NJ	2.68	1.44	-1.24
ID	0.20	-0.96	-1.16
CO	2.60	1.58	-1.02
WV	0.98	0.17	-0.81
GA	2.30	1.54	-0.76
MN	1.91	1.24	-0.67
OK	-0.75	-1.37	-0.62
KS	1.64	1.13	-0.51
CA	**9.14**	**8.89**	**-0.25**
NH	0.32	0.27	-0.05
TX	3.09	3.30	0.21
PA	1.62	3.02	1.40
SC	1.83	5.09	3.26
MA	-0.62	3.50	4.12
MD	2.76	9.02	6.26

we observe for 18 other states where we have simulated House maps in 2010 and 2000, and that did not alter their mode of redistricting. As shown in Table 3.2, we see that California's improvement in relative competitiveness (-0.25) over the decade is smaller than that observed in 12 of 18 (67%) states that did not change their redistricting, with the average change in relative competitiveness for all 18 states being -0.58. These coefficients estimate relative competitiveness taken from the 2010 simulation results in Figure 1(a) and the 2000 results presented in Figure 3.4(a) in the Appendix. Thus, in a place where we can be reasonably assured that any bias in simulations cannot be correlated with how redistricting is conducted, we again observe no substantial improvement in the relative competitiveness of maps produced by non-partisan redistrictors.

3.5.4 Generalizability Using Convenience Sample of Maps From 15 States

For 2010, we managed to collect 1,627 maps from the full set of data made publicly available by 15 state legislatures and redistricting commissions. These data contain maps from AK, AZ, CA, CO, FL, ID, MT, NC, NM, NV, OH, SC, TX, VA, and WA. This includes every state that has independent redistricting, but no Democratic plans, and only a handful of maps from Republican, court and bipartisan redistrictors. This incomplete coverage may limit the generalizability of our findings, especially if we expect that Democratic legislatures redistrict in fundamentally different ways than Republicans. Scant research suggests this possibility, typically finding that both Democrats (e.g., Maryland) and Republicans (e.g., Texas) gerrymander whenever they can. Nevertheless, there are clear differences in the types of states which collect and publicize alternative maps and those which do not. If there is heterogeneity in how redistricting is conducted in those states excluded from our counterfactuals sample, then this could alter our conclusions about redistricting generally, and about non-partisan redistricting in particular.

Our combined analysis of both simulations and counterfactual maps helps us address this generalizability issue. An important benefit in using both simulations and publicized counterfactuals for data is that these are generated in different ways and cover different, but overlapping sets of states and seats. Though imperfect, this overlap can give us greater assurance that our results are not idiosyncratic to the particular places where data are available. On this front, we were able to analyze U.S. House simulations produced by Chen and Cottrell (2016) for 42 (of all 43) multi-district states in 2010. Again, these results, presented in Figure 1(a) in the manuscript, affirm our main findings. Moreover, these results include the full range of states that redistrict House seats, including those with Democratic-controlled or bipartisan processes.

A limitation using these data, however, is the lack of simulations for 2010 State Assembly and Senate districts. In Figure 3.4, we present simulations for 19 states at the House, State Assembly and State Senate jurisdictions for 2000 (Chen and Rodden 2013). These mirror our findings for House districting in 2010. Note though that we coverage for all non-partisan states, save Idaho. Though limited in coverage of states, our counterfactual maps data allow us to better understand the consequences of non-partisan districting in state legislative districts and not just the U.S. House, which has been the major focus in work using simulations (e.g., Chen and Cottrell 2016).

In summary, though each source of data faces particular limits in coverage, these limits are different across each type of data. Consistent results across different sets of states and political jurisdictions at different redistricting cycles, using these very different data, strongly suggest that our findings are not idiosyncratic to any particular dataset or analysis. At the very least, these data can help bound our conclusions under different beliefs about how redistricting is conducted in left-out states. Namely, either (a) the simulations data significantly understate how uncompetitive redistricting really is in excluded partisan (i.e., Democratic) states, and independent redistricting does produce more competitiveness or (b) the simulations miss how much more competitive redistricting is in excluded partisan (i.e., Democratic) places, and independent redistricting is actually much worse at producing competitive districts.

3.5.5 Some Theoretical Expectations About Redistricting Institutions

Very little theoretical work has been done to explain the differences in how various institutions (i.e., legislatures, courts, political commissions, non-partisan commissions) draw maps. Most prior work relies on an explanation of the strategic behavior of political elites. This work starts with the theoretical prediction that partisan redistricting would lead to uncompetitive districts because incumbent politicians wish to stay in office, and have the legal means of facilitating this, either through party gerrymanders or incumbent protection plans. In contrast, independent commissions are thought to not have these same political incentives. Decisions are not made by (or influenced by) politicians seeking office, and so independent commissioners should produce maps freed from the desire to protect incumbents. The logic behind court plans is relatively similar. The courts usually draw maps only when a partisan legislature or politician commission fails to agree on a map, or when aggrieved plaintiffs sue (typically) on the basis of voting rights (VRA) violations. Given this, the job of the court is simply to draw maps as guided by the Constitution and prior court rulings. We should not expect that court plans introduce a pro-incumbent bias either.

Our data suggest that partisan redistricting marginally protects incumbents, and that independent redistricting does too. As a result, what we suggest is that the conventional wisdom surrounding the (non)-political incentives of independent redistricting does not comport with the evidence. Interestingly, we do find that courts produce the most relatively competitive plans, as these typically fall inside the range of plans (often near median) considered by redistrictors. We hesitate to offer a theoretical account of this finding given the many complexities associated with how courts make decisions about maps, including how and when they intervene, and on what grounds they base their decisions. This result certainly suggests that courts are the most politically-neutral remedy available, something that deserves greater scholarly attention in the future.

Clearly,the next step in this research area is a serious theoretical treatment, in light of our findings, on the incentives surrounding the actors involved in each redistricting method. Without such a theory, it will be difficult to explain the mechanisms driving our findings.

Nevertheless, we see our work as a critical first step in addressing whether claims put forth by pundits and the reform community pass empirical muster.

3.5.6 Additional Details on Data Collection and Construction

We clarify here how we constructed our dataset of publicized maps, briefly including our approach to acquiring the maps from redistrictors. On the latter, we manually searched for public maps from the webpages of every state's official redistricting authority, including legislatures. Most, but not all states have a public webpage describing how redistricting is conducted. Of those that do, we were only able to find any alternative maps for 15 states. Collecting maps for some of these states necessitated us to use the Wayback Machine to get archived links to public deliberations occurring in 2010 and after, but that were subsequently removed from state webpages. We searched each state until we were satisfied that no such maps were ever made public in electronic form online, or that if maps were made public, that no viable means existed to collect them during our online search. We then contacted the redistricting authorities for all 35 states where we could not find any maps to see if any such maps were available in any other form. Few agencies responded, and none that did provided us with additional maps.

The maps we were able to collect typically took one of two forms, either a matrix of block equivalency file linking census blocks to districts for each public map, or a database of shapefiles. We also collected a small number (≈ 0.005) of maps that were just compiled .pdf files visualizing district boundaries. We discarded these .pdf maps since we could not reliably link census blocks or precinct/voting districts (VTD) to the proposed jurisdictions. We collected election outcomes for VTDs for most states using the Ansolabehere, Palmer, and Lee (2015) dataset. These data include shapefiles of VTDs along with presidential voting and party registration in 2008. With these data, we used spatial geocoding packages in **R** to identify which VTDs/precincts were incorporated into proposed districts for the various alternative redistricting plans. When a VTD was entirely subsumed by the geographical boundary of a proposed district, that precinct/VTD's election data were added to that legislative jurisdiction's vote summary. Though relatively rare, VTDs were sometimes geographically split over multiple proposed legislative districts – in these cases we divided the vote summary data for that VTD equally across the connected legislative districts.

We used a different approach for maps that were provided as block equivalency files. Here, we required vote summaries, disaggregated from VTDs and linked to each census block, so we could re-aggregate block-level voting back to proposed legislative districts. Of course, voting data is collected at the VTD, and not block level. Thus, an initial step is required to split vote figures from VTDs into the various census blocks that are contained within each VTD. We did not do this, and rather relied on the data provided by McDonald and Altman (2011). These data connect 2008 election data to blocks, which we aggregated to districts across the various proposed maps.

Naturally, both approaches will contain random measurement error, and this error may vary somewhat when using shapefiles rather than block equivalencies. Importantly, any such measurement error will be constant within a state, since competitiveness across the alternative proposals are produced using the same basis in data within each state. Further, this error can only attenuate our relative comparisons, so that our comparisons of enacted and alternative plans may understate the level of uncompetitiveness we find following the redistricting process. Finally, we have a handful of maps that were provided as both block equivalencies and as shapefiles, and so we can evaluate how the results might vary due to differential measurement error from disaggregating and aggregating election data across different geographies. In inspecting these maps, we find differences are small and ignorable.

Once linked, our Expected Margin of Victory measure is produced for each district in a given map, by taking the absolute value of the head start provided a party incumbent. For example, assume a state with four districts had the following two-party Democratic vote share: $x = 0.72, 0.55, 0.51, 0.37$. The district expected margin of victory is the absolute winning margin, or $\text{abs}((1 - x_i) - x_i) = 0.44, 0.10, 0.02, 0.26$. This is then averaged over the plan to produce a map-level measure, $\text{E}\{\text{abs}((1 - x_i) - x_i)\}$, which is 0.205 in this example.

Notably, we do not discard any legislative districts from enacted or alternative maps based on whether or not these were contested or open seats in any elections, including 2008, 2010, and 2012. We do not discard any post-2010 districts either as this would likely bias our findings by conditioning on post-treatment effects of states enacting a particular plan. We also do not discard any districts depending on whether or not incumbents retired or were challenged in the pre-redistricting period. Any variation in voting patterns across precincts or VTDs due to pre-redistricting strategic choices would be absorbed into the entire distribution of counterfactual maps or simulations, and so would not affect the relative differences between enacted and alternative plans, or simulations. We also use 2008 presidential elections in our analyses since presidential voting is unaffected by down-ballot turnout and vote (e.g., Broockman 2009).

Bibliography

Aberbach, Joel D., and Jack L. Walker. 1970. "The Attitudes of Blacks and Whites toward City Services: Implications for Public Policy." In *Financing the Metropolis: Public Policy in Urban Economics*, ed. John P. Crecine. Beverly Hills: Sage.

Abramowitz, Alan, Brad Alexander, and Matthew Gunning. 2006. "Incumbency, Redistricting, and the Decline of Competition in U.S. House Elections." *Journal of Politics* 68(1): 75–88.

Acharya, Avidit, Matthew Blackwell, and Maya Sen. 2016. "The Political Legacy of American Slavery." *Journal of Politics* 78(3): 621–641.

Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press.

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics* 114(4): 1243–1284.

Altman, Micah, and Michael McDonald. 2014. "Paradoxes of Political Reform: Congressional Redistricting In Florida." In *Jigsaw Puzzle Politics in the Sunshine State*, ed. Seth McKee. Gainesville, FL: University Press of Florida , 163–184.

Altman, Micah, and Michael McDonald. 2015. "Redistricting and Polarization." In *American Gridlock: The Sources, Character, and Impact of Political Polarization*, eds. James A. Thurber, and Antoine Yoshinaka. Cambridge: Cambridge University Press , 45–67.

Ansolabehere, Stephen, Maxwell Palmer, and Amanda Lee. 2015. "Precinct-Level Election Data, 2002-2012." Unpublished paper, *Harvard Dataverse*, Available at: `goo.gl/OPUfWM`.

Arnold, R. Douglas. 1990. *The Logic of Congressional Action*. New Haven: Yale University Press.

Arnold, R. Douglas, and Nicholas Carnes. 2012. "Holding Mayors Accountable: New York?s Executives from Koch to Bloomberg." *American Journal of Political Science* 56(4): 949–963.

Bartels, Larry M. 2008. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton: Princeton University Press.

Bechtel, Michael M., and Jens Hainmueller. 2011. "How Lasting Is Voter Gratitude? An Analysis of the Short- and Long-Term Electoral Returns to Beneficial Policy." *American Journal of Political Science* 55(4): 852–868.

Berry, Christopher R., and William G. Howell. 2007. "Accountability and Local Elections: Rethinking Retrospective Voting." *Journal of Politics* 69(3): 844–858.

Broockman, David. 2009. "?Do Congressional Candidates Have Reverse Coattails? Evidence from a Regression Discontinuity Design." *Political Analysis* 17(4): 418–434.

Burden, Barry C. 2007. *Personal Roots of Representation*. Princeton: Princeton University.

Burnett, Craig M., and Vladimir Kogan. 2017. "The Politics of Potholes: Service Quality and Retrospective Voting in Local Elections." *Journal of Politics* 79(1): 302–314.

Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality*. New York: Cambridge University Press.

Butler, Daniel M., and David E. Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents? A Field Experiment on State Legislators." *American Journal of Political Science* 55(3): 463–477.

Butler, Daniel M., Christopher F. Karpowitz, and Jeremy C. Pope. 2012. "A Field Experiment on Legislators' Home Styles: Service versus Policy." *Journal of Politics* 74(2): 474–486.

Cain, Bruce E., John Ferejohn, and Morris P. Fiorina. 1987. *The Personal Vote: Constituency Service and Electoral Independence*. Cambridge: Harvard University Press.

Campbell, Andrea Louise. 2002. "Self-Interest, Social Security and the Distinctive Participation Patterns of Senior Citizens." *American Political Science Review* 95(3): 565–574.

Canes-Wrone, Brandice, and Kenneth W. Shotts. 2007. "When Do Elections Encourage Ideological Rigidity?" *American Political Science Review* 101(2): 273–288.

Canes-Wrone, Brandice, Michael C. Herron, and Kenneth W. Shotts. 2001. "Leadership and Pandering: A Theory of Executive Policymaking." *American Political Science Review* 45(3): 532–550.

Canon, David T. 1999. *Race, Redistricting, and Representation: The Unintended Consequences of Black Majority Districts*. Chicago: University of Chicago Press.

Carnes, Nicholas. 2013. *White-Collar Government: The Hidden Role of Class in Economic Policymaking*. Chicago: University of Chicago Press.

Carnes, Nicholas, and John Holbein. 2019. "Do Public Officials Exhibit Social Class Biases When They Handle Casework? Evidence from Multiple Correspondence Experiments." *PLOS One* 14(3): 1–9.

Carson, Jamie, Michael Crespin, and Ryan Williamson. 2014. "Reevaluating the Effects of Redistricting on Electoral Competition, 1972-2012." *State Politics & Policy Quarterly* 14(2): 165–177.

Caughey, Devin, and Christopher Warshaw. 2018. "Policy Preferences and Policy Change: Dynamic Responsiveness in the American States, 1936–2014." *American Political Science Review* 112(2): 249–266.

Caughey, Devin, Michael C. Dougal, and Eric Schickler. 2020. "Policy and Performance in the New Deal Realignment: Evidence from Old Data and New Methods." *Journal of Politics* 82(2): 494–508.

Chen, Jowei. 2011. "Voter Partisanship and the Effect of Distributive Spending on Political Participation." *American Journal of Political Science* 57(1): 200–217.

Chen, Jowei, and David Cottrell. 2016. "Evaluating Partisan Gains from Congressional Gerrymandering: Using Computer Simulations to Estimate the Effect of Gerrymandering in the U.S. House." *Electoral Studies* 44(2): 329–340.

Chen, Jowei, and Jonathan Rodden. 2013. "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures." *Quarterly Journal of Political Science* 8(3): 239–269.

Christensen, Darin, and Simon Ejdemyr. 2018. "Do Elections Improve Constituency Responsiveness? Evidence from U.S. Cities." *Political Science Research and Methods* .

Clinton, Joshua D., and Michael W. Sances. 2018. "The Politics of Policy: The Initial Mass Political Effects of Medicaid Expansion in the States." *American Political Science Review* 112(1): 167–185.

Cohen, Elisa, Anna Gunderson, Kaylyn Jackson, Paul Zachary, Tom S. Clark, Adam N. Glynn, and Michael Leo Owens. 2019. "Do Officer-Involved Shootings Reduce Citizen Contact with Government?" *Journal of Politics* 81(3): 1111–1123.

Congressional Quarterly. 2014. "CQ Congress Collection." Unpublished paper, Accessed Online at: `http://library.cqpress.com/congress/`.

Costa, Mia. 2017. "How Responsive are Political Elites? A Meta-Analysis of Experiments on Public Officials." *Journal of Experimental Political Science* 4(3): 241–254.

Cottrill, James. 2012. "The Effects of Non-Legislative Approaches to Redistricting on Competition in Congressional Elections." *Polity* 44(1): 32–50.

Cox, Gary W., and Mathew D. McCubbins. 1986. "Electoral Politics as a Redistributive Game." *Journal of Politics* 48(2): 370–389.

Dahl, Robert A. 1971. *Polyarchy: Participation and Opposition*. New Haven: Yale University Press.

de Benedictis-Kessner, Justin. 2019. "Where the Sidewalk Ends: How Participation Contributes to Inequity in Government Service Provision." Paper presented at the Annual Meeting of the American Political Science Association, Washington, DC.

DeHoog, Ruth Hoogland, David Lowery, and William E. Lyons. 1990. "Citizen Satisfaction with Local Governance: A Test of Individual, Jurisdictional, and City-Specific Explanations." *Journal of Politics* 52(3): 807–837.

Dixit, Avinash, and John Londregan. 1996. "The Determinants of Success of Special Interests in Redistributive Politics." *Journal of Politics* 58(4): 1132–1155.

Edwards, Barry, Michael Crespin, Ryan Williamson, and Maxwell Palmer. 2017. "Institutional Control of Redistricting and the Geography of Representation." *Journal of Politics* 79(2): 1–5.

Einstein, Katherine Levine, and David M. Glick. 2017. "Does Race Affect Access to Government Services? An Experiment Exploring Street-Level Bureaucrats and Access to Public Housing." *American Journal of Political Science* 61(1): 100–116.

Einstein, Katherine Levine, and Vladimir Kogan. 2016. "Pushing the City Limits: Policy Responsiveness in Municiapl Government." *Urban Affairs Review* 52(1): 3–32.

Ellis, Christopher. 2012. "Understanding Economic Biases in Representation: Income, Resources, and Policy Representation in the 110th House." *Political Research Quarterly* 65(4): 938–951.

Erikson, Robert S., Gerald C. Wright, and John P. McIver. 1993. *Statehouse Democracy: Public Opinion and Policy in the American States*. New York: Cambridge University Press.

Federal Works Agency. 1947. *Final Report on the WPA Program, 1935-1943*. Washington, DC: U.S. Government Printing Office.

Feigenbaum, James J., and Andrew B. Hall. 2016. "How High-Income Areas Receive More Service From Municipal Government: Evidence From City Administrative Data." Unpublished paper, .

Fenno, Richard F. 1978. *Home Style: House Members in Their Districts*. Boston: Little, Brown and Company.

Ferejohn, John. 1986. "Incumbent Performance and Electoral Control." *Public Choice* 50(1/3): 5–25.

Fiorina, Morris P. 1978. "Economic Retrospective Voting in American National Elections: A Micro-Analysis." *American Journal of Political Science* 22(2): 426–443.

Fishback, Price V., Shawn Kantor, and John Joseph Wallis. 2003. "Can the New Deal's Three R's be Rehabilitated? A Program-by-Program, County-by-County Analysis." *Explorations in Economic History* 40(3): 278–307.

Fishback, Price V., William C. Horace, and Shawn Kantor. 2005. "Did New Deal Grant Programs Stimulate Local Economies? A Study of Federal Grants and Retail Sales During the Great Depression." *Journal of Economic History* 65(1): 36–71.

Fong, Christian, Chad Hazlett, and Kosuke Imai. 2019. "Covariate Balancing Propensity Score for a Continuous Treatment: Application to the Efficacy of Political Advertisements." *Annals of Applied Statistics* 12(1): 156–177.

Forgette, Richard, Andrew Garner, and John Winkle. 2009. "Do Redistricting Principles and Practices Affect U.S. State Legislative Electoral Competition?" *State Politics & Policy Quarterly* 9(2): 151–175.

Fox, Justin, and Kenneth W. Shotts. 2009. "Delegates or Trustees? A Theory of Political Accountability." *Journal of Politics* 71(4): 1225–1237.

Gasper, John T., and Andrew Reeves. 2011. "Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters." *Journal of Politics* 55(2): 340–355.

Gell-Redman, Micah, Neil Visalvanich, Charles Crabtree, and Christopher J. Fariss. 2018. "It's All about Race: How State Legislators Respond to Immigrant Constituents." *Political Research Quarterly* 71(3): 517–531.

Gelman, Andrew, and Gary King. 1994. "Enhancing Democracy Through Legislative Redistricting." *American Political Science Review* 88(3): 541–559.

Gerber, Elisabeth R., and Daniel J. Hopkins. 2011. "When Mayors Matter: Estimating the Impact of Mayoral Partisanship on City Policy." *American Journal of Political Science* 55(2): 326–339.

Gilens, Martin. 1996. "'Race Coding' and White Opposition to Welfare." *American Political Science Review* 90(3): 593–604.

Gilens, Martin. 2005. "Inequality and Democratic Responsiveness." *Public Opinion Quarterly* 69(5): 778–796.

Gilens, Martin. 2012. *Affluence and Influence: Economic Inequality and Political Power in America*. Princeton: Princeton University Press.

Glaser, James M. 2002. "White Voters, Black Schools: Structuring Racial Choices with a Checklist Ballot." *American Journal of Political Science* 46(1): 35–46.

Grainger, Corbett. 2010. "Redistricting and Polarization: Who Draws the Lines in California?" *Journal of Law & Economics* 54: 545–567.

Griffin, John D., and Brian Newman. 2005. "Are Voters Better Represented?" *Journal of Politics* 67(4): 1206–1227.

Griffin, John D., and Brian Newman. 2007. "The Unequal Representation of Latinos and Whites." *Journal of Politics* 69(4): 1032–1046.

Griffin, John D., and Brian Newman. 2008. *Minority Report: Evaluating Political Equality in America*. Chicago: University of Chicago Press.

Grimmer, Justin, Solomon Messing, and Sean J. Westwood. 2012. "How Words and Money Cultivate a Personal Vote: The Effect of Legislator Credit Claiming on Constituent Credit Allocation." *American Political Science Review* 106(4): 703–719.

Grose, Christian R. 2011. *Congress in Black and White: Race and Representation in Washington and at Home*. New York: Cambridge University Press.

Grose, Christian R., and Bruce I. Oppenheimer. 2007. "The Iraq War, Partisanship, and Candidate Attributes: Variation in Partisan Swing in the 2006 U.S. House Elections." *Legislative Studies Quarterly* 32(4): 531–557.

Hajnal, Zoltan, and Jessica Trounstine. 2014. "Identifying and Understanding Perceived Inequities in Local Politics." *Political Research Quarterly* 67(1): 56–70.

Hajnal, Zoltan L. 2009. *America's Uneven Democracy: Race, Turnout, and Representation in City Politics*. New York: Cambridge University Press.

Harden, Jeffrey J. 2016. *Multidimensional Democracy: A Supply and Demand Theory of Representation in American Legislatures*. Cambridge: Cambridge University Press.

Healy, Andrew, and Gabriel S. Lenz. 2017. "Presidential Voting and the Local Economy: Evidence from Two Population-Based Data Sets." *Journal of Politics* 79(4): 1419–1432.

Healy, Andrew, and Neil Malhotra. 2009. "Myopic Voters and Natural Disaster Policy." *American Political Science Review* 103(3): 387–406.

Hetherington, Marc J. 1998. "The Political Relevance of Political Trust." *American Political Science Review* 92(4): 791–808.

Hill, Kim Quaile, and Jan E. Leighley. 1992. "The Policy Consequences of Class Bias in State Electorates." *American Journal of Political Science* 36(2): 351–365.

Hill, Kim Quaile, Jan E. Leighley, and Angela Hinton-Andersson. 1995. "Lower-Class Mobilization and Policy Linkages in the U.S. States." *American Journal of Political Science* 39(1): 75–86.

Holbrook, Thomas M., and Aaron C. Weinschenk. 2019. "Are Perceptions of Local Conditions Rooted in Reality? Evidence from Two Large-Scale Local Surveys." *American Politics Research* .

Hopkins, Daniel J. 2009. "The Diversity Discount: When Increasing Ethnic and Racial Diversity Prevents Tax Increases." *Journal of Politics* 71(1): 160–177.

Hulse, Carl. 2016. "Seeking to End Gerrymandering?s Enduring Legacy." Unpublished paper, *New York Times*, Available at: `goo.gl/qt59PI`.

Imai, Kosuke, and Marc Ratkovic. 2014. "Covariate Balancing Propensity Score." *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 76(1): 243–246.

Jones, Bryan D., Saadia R. Greenberg, Clifford Kaufman, and Joseph Drew. 1977. "Bureaucratic Response to Citizen-Initiated Contacts: Environmental Enforcement in Detroit." *American Political Science Review* 71(1): 148–165.

Jones, Bryan D., Saadia R. Greenberg, Clifford Kaufman, and Joseph Drew. 1978. "Service Delivery Rules and the Distribution of Local Government Services: Three Detroit Bureaucracies." *Journal of Politics* 40(2): 332–368.

Juenke, Eric Gonzalez. 2014. "Ignorance Is Bias: The Effect of Latino Losers on Models of Latino Representation." *American Journal of Political Science* 58(3): 593–603.

Karol, David, and Edward Miguel. 2007. "The Electoral Cost of War: Iraq Casualties and the 2004 U.S. Presidential Election." *Journal of Politics* 69(3): 633–648.

Kelly, Janet M. 2003. "Citizen Satisfaction and Administrative Performance Measures: Is There Really a Link?" *Urban Affairs Review* 38(6): 855–866.

Kelly, Janet M., and David Swindell. 2002. "Performance Monitoring and Citizen Satisfaction: Correlating Administrative Outcomes and Citizen Evaluation of Service Quality." *Public Administration Review* 62(5): 610–620.

Kettl, Donald F. 1988. *Government By Proxy: (Mis?)Managing Federal Programs*. Washington, DC: Congressional Quarterly Press.

Key, V.O. Jr. 1966. *The Responsible Electorate: Rationality in Presidential Voting, 1936–1960*. Cambridge: Harvard University Press.

Kogan, Vladimir. 2021. "Do Welfare Benefits Pay Electoral Dividends? Evidence from the National Food Stamp Program Rollout." *Journal of Politics* 83(1): 58–70.

Kogan, Vladimir, Stéphane Lavertu, and Zachary Peskowitz. 2016. "Performance Federalism and Local Democracy: Theory and Evidence from School Tax Referenda." *American Journal of Political Science* 60(2): 418–435.

Kone, Susan L., and Richard F. Winters. 1993. "Taxes and Voting: Electoral Retribution in the American States." *Journal of Politics* 55(1): 22–40.

Kramer, Gerald H. 1971. "Short-Term Fluctuations in U.S. Voting Behavior, 1896-1964." *American Political Science Review* 65(1): 131–143.

Lajevardi, Nazita. 2018. "Access Denied: Exploring Muslim American Representation and Exclusion by State Legislators." *Politics, Groups, and Identities* .

Lax, Jeffrey R., and Justin H. Phillips. 2012. "The Democratic Deficit in the States." *American Journal of Political Science* 56(1): 148–166.

Legewie, Joscha, and Merlin Schaeffer. 2016. "Contested Boundaries: Explaining Where Ethnoracial Diversity Provokes Neighborhood Conflict." *American Journal of Sociology* 122(1): 125–161.

Lerman, Amy, and Vesla Weaver. 2013. "Staying Out of Sight? Concentrated Policing and Local Political Action." *The Annals of the American Academy of Political and Social Science* 651(1): 202–219.

Lerman, Amy E. 2019. *Good Enough for Government Work: The Public Reputation Crisis in America (And What We Can Do to Fix It)*. Chicago: University of Chicago Press.

Levine, Jeremy R., and Carl Gershenson. 2014. "From Political to Material Inequality: Race, Immigration, and Requests for Public Goods." *Sociological Forum* 29(3): 607–627.

Levitt, Justin. 2010. *A Citizen's Guide to Redistricting*. New York: Brennan Center.

Levitt, Steven D., and James M. Snyder, Jr. 1997. "The Impact of Federal Spending on House Election Outcomes." *Journal of Political Economy* 105(1): 30–53.

Levy, Frank S., Arnold J. Meltsner, and Aaron Bernard Wildavsky. 1974. *Urban Outcomes: Schools, Streets, and Libraries*. Berkeley: University of California Press.

Lindgren, Eric, and Priscilla Southwell. 2013. "The Effect of Redistricting Commissions on Electoral Competitiveness in U.S. House Elections, 2002-2010." *Journal of Politics and Law* 6: 13–18.

Lineberry, Robert L. 1977. *Equality and Urban Policy: The Distribution of Municipal Public Services*. Beverly Hills: Sage.

Lowande, Kenneth, and Andrew Proctor. 2019. "Bureaucratic Responsiveness to LGBT Americans." *American Journal of Political Science* .

Lowande, Kenneth, Melinda Ritchie, and Erinn Lauterbach. 2019. "Descriptive and Substantive Representation in Congress: Evidence from 80,000 Congressional Inquiries." *American Journal of Political Science* 63(3): 644–659.

Luttmer, Erzo. 2001. "Group Loyalty and the Taste for Redistribution." *Journal of Political Economy* 109(3): 500–528.

Margalit, Yotam. 2011. "Costly Jobs: Trade-Related Layoffs, Government Compensation, and Voting in U.S. Elections." *American Political Science Review* 105(1): 166–188.

Marschall, Melissa J., and Anirudh V.S. Ruhil. 2007. "Substantive Symbols: The Attitudinal Dimension of Black Political Incorporation in Local Government." *American Journal of Political Science* 51(1): 17–33.

Martin, Paul S. 2003. "Voting's Rewards: Voter Turnout, Attentive Publics, and Congressional Allocation of Federal Money." *American Journal of Political Science* 47(1): 110–127.

Masket, Seth, Jonathan Winburn, and Gerald Wright. 2012. "The Gerrymanders Are Coming! Legislative Redistricting Won't Affect Competition or Polarization Much, No Matter Who Does It." *PS: Political Science and Politics* 45(1): 39–43.

Mayhew, David R. 1974. *Congress: The Electoral Connection*. New Haven: Yale University Press.

Mazumder, Soumyajit. 2018. "The Persistent Effect of U.S. Civil Rights Protests on Political Attitudes." *American Journal of Political Science* 62(4): 922–935.

McDonald, Michael. 2006. "Drawing the Line on District Competition." *PS: Political Science and Politics* 39(1): 91–94.

McDonald, Michael, and Micah Altman. 2011. "Public Mapping Project." Unpublished paper, Available at: `goo.gl/hTxUwx`.

Mendez, Matthew S., and Christian R. Grose. 2018. "Doubling Down: Inequality in Responsiveness and the Policy Preferences of Elected Officials." *Legislative Studies Quarterly* 43(3): 457–491.

Mettler, Suzanne. 2002. "Bringing the State Back in to Civic Engagement: Policy Feedback Effects of the G.I. Bill for World War II Veterans." *American Political Science Review* 96(2): 351–365.

Mettler, Suzanne. 2011. *The Submerged State: How Invisible Government Policies Undermine American Democracy.* Chicago: University of Chicago Press.

Michener, Jamila. 2018. *Fragmented Democracy: Medicaid, Federalism, and Unequal Politics.* Cambridge: Cambridge University Press.

Minkoff, Scott L. 2016. "NYC 311: A Tract-Level Analysis of Citizen-Government Contacting in New York City." *Urban Affairs Review* 52(2): 211–246.

Mladenka, Kenneth R. 1980. "The Urban Bureaucracy and the Chicago Political Machine: Who Gets What and the Limits to Political Control." *American Political Science Review* 74(4): 991–998.

Mladenka, Kenneth R. 1981. "Citizen Demands and Urban Services: The Distribution of Bureaucratic Response in Chicago." *American Journal of Political Science* 25(4): 693–714.

Mladenka, Kenneth R., and Kim Quaile Hill. 1978. "The Distribution of Urban Police Services." *Journal of Politics* 40(1): 112–133.

Morgan, Kimberly J., and Andrea Louise Campbell. 2011. *The Delegated Welfare State: Medicare, Markets, and the Governance of Social Policy.* New York: Oxford University Press.

Nadeau, Richard, and Michael S. Lewis-Beck. 2001. "National Economic Voting in U.S. Presidential Elections." *Journal of Politics* 63(1): 159–181.

Neckerman, Kathryn M., Gina S. Lovasi, Stephen Davies, Marnie Purciel, James Quinn, Eric Feder, Nakita Raghunath, Benjamin Wasserman, and Andrew Rundle. 2009. "Disparities in Urban Neighborhood Conditions: Evidence from GIS Measures and Field Observation in New York City." *Journal of Public Health Policy* 30(1): S264–S285.

Niemi, Richard G., Harold W. Stanley, and Ronald J. Vogel. 1995. "State Economies and State Taxes: Do Voters Hold Governors Accountable?" *American Journal of Political Science* 39(4): 936–957.

Nivola, Pietro S. 1978. "Distributing a Municipal Service: A Case Study of Housing Inspection." *Journal of Politics* 40(1): 59–81.

Page, Benjamin I., and Robert Y. Shapiro. 1983. "Effects of Public Opinion on Policy." *American Political Science Review* 77(1): 175–190.

Parker, David C.W., and Craig Goodman. 2009. "Making a Good Impression: Resource Allocation, Home Styles, and Washington Work." *Legislative Studies Quarterly* 34(4): 493–524.

Percy, Stephen L. 1986. "In Defense of Citizen Evaluations as Performance Measures." *Urban Affairs Quarterly* 22(1): 66–83.

Peterson, Paul E. 1981. *City Limits*. Chicago: University of Chicago Press.

Pfaff, Steven, Charles Crabtree, Holger L. Kern, and John B. Holbein. 2019. "Does Religious Bias Shape Access to Public Services? A Large-Scale Audit Experiment Among Street-Level Bureaucrats." Unpublished paper, .

Pierce, Olga, and Jeff Larson. 2011. "How Democrats Fooled California?s Redistricting Commission." Unpublished paper, *ProPublica*, Available at: `goo.gl/TuoIwa`.

Pierson, Paul. 1993. "When Effect Becomes Cause: Policy Feedback and Political Change." *World Politics* 45(4): 595–628.

Powell Jr., G. Bingham, and Guy D. Whitten. 1993. "A Cross-National Analysis of Economic Voting: Taking Account of the Political Context." *American Journal of Political Science* 37(2): 391–414.

Public Works Administration. 1939. *America Builds: The Record of the PWA*. Washington, DC: U.S. Government Printing Office.

Rigby, Elizabeth, and Gerald C. Wright. 2011. "Whose Statehouse Democracy? Policy Responsiveness to Poor Versus Rich Constituents in Poor Versus Rich States." In *Who Gets Represented?*, eds. Peter K. Enns, and Christopher Wlezien. New York: Russell Sage Foundation , 189–222.

Rudolph, Thomas J. 2003. "Who's Responsible for the Economy? The Formation and Consequences of Responsibility Attributions." *American Journal of Political Science* 47(4): 698–713.

Samuels, David. 2004. "Presidentialism and Accountability for the Economy in Comparative Perspective." *American Political Science Review* 98(3): 425–436.

Schaffner, Brian F., Jesse H. Rhodes, and Raymond J. La Raja. 2016. "Race- and Class-Based Inequality and Representation in Local Government." Paper presented at the Annual Meeting of the American Political Science Association, Philadelphia.

Schuman, Howard, and Barry Gruenberg. 1972. "Dissatisfaction with City Services: Is Race an Important Factor?" In *People and Politics in Urban Society*, ed. Harlan Hahn. Beverly Hills: Sage.

Shah, Paru. 2014. "It Takes a Black Candidate: A Supply-Side Theory of Minority Representation." *Political Research Quarterly* 67(2): 266–279.

Simonovits, Gabor, Andrew M. Guess, and Jonathan Nagler. 2019. "Responsiveness without Representation: Evidence from Minimum Wage Laws in U.S. States." *American Journal of Political Science* 63(2): 401–410.

Soroka, Stuart N., and Christopher Wlezien. 2010. *Degrees of Democracy: Politics, Public Opinion, and Policy.* New York: Cambridge University Press.

Soss, Joe. 1999. "Lessons of Welfare: Policy Design, Political Learning and Political Action." *American Political Science Review* 93(2): 363–380.

Soss, Joe, and Sanford F. Schram. 2007. "A Public Transformed? Welfare Reform as Policy Feedback." *American Political Science Review* 101(1): 111–127.

Stimson, James A., Michael B. Mackuen, and Robert S. Erikson. 1995. "Dynamic Representation." *American Journal Science Review* 89(3): 543–565.

Tausanovitch, Chris, and Christopher Warshaw. 2014. "Representation in Municipal Government." *American Political Science Review* 108(3): 605–641.

Thomas, John Clayton. 1982. "Citizen-initiated Contacts with Government Agencies: A Test of Three Theories." *American Journal of Political Science* 26(3): 504–522.

Thornton, Christina M., Terry L. Conway, Kelli L. Cain, Kavita A. Gavand, Brian E. Saelens, Lawrence D. Frank, Carrie M. Geremia, Karen Glanz, Abby C. King, and James F. Sallis. 2016. "Disparities in Pedestrian Streetscape Environments by Income and Race/Ethnicity." *SSM - Population Health* 2: 206–216.

Trounstine, Jessica. 2016. "Segregation and Inequality in Public Goods." *American Journal of Political Science* 60(3): 709–725.

Trounstine, Jessica. 2018. *Segregation by Design: Local Politics and Inequality in American Cities.* New York: Cambridge University Press.

Tucker, Patrick. 2019. "The Demand for Non-Ideological Representation in a Polarized World." Unpublished paper, .

Vedlitz, Arnold, and James A. Dyer. 1984. "Bureaucratic Response to Citizen Contacts: Neighborhood Demands and Administrative Reaction in Dallas." *Journal of Politics* 46(4): 1207–1216.

Vedlitz, Arnold, James A. Dyer, and Roger Durand. 1980. "Citizen Contacts with Local Governments: A Comparative View." *American Journal of Political Science* 24(1): 50–67.

Verba, Sidney, and Norman H. Nie. 1972. *Participation in America: Political Democracy and Social Equality.* New York: Harper & Row.

White, Ariel, and Kris-Stella Trump. 2018. "The Promises and Pitfalls of 311 Data." *Urban Affairs Review* 54(4): 794–823.

White, Ariel R., Noah L. Nathan, and Julie K. Faller. 2015. "What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials." *American Political Science Review* 109(1): 129–142.

Yates, Douglas T., Jr. 1977. *The Ungovernable City: The Politics of Urban Problems and Policy Making*. Cambridge: MIT Press.