The Politics of Potholes:
Service Quality and Retrospective Voting in Local Elections[*]

Craig M. Burnett
Department of Public and International Affairs
University of North Carolina at Wilmington
burnettc@uncw.edu

Vladimir Kogan
Department of Political Science
Ohio State University
kogan.18@osu.edu

August 12, 2015

Abstract

By conditioning their support for political incumbents on observed government outputs, voters can motivate policymakers to faithfully represent their interests. Whether elections serve this function in sub-national U.S. jurisdictions remains unclear, however, because the existing research on retrospective voting in these contexts focuses largely on outcomes over which government officials have little, if any, influence. In this study, we examine one outcome — the quality of local roads — that is unquestionably under the control of city government. In doing so, we leverage variation in service quality within one of America's largest cities and show that such variation can explain neighborhood-level differences in support for incumbents in two political offices — mayor and city council — across two different electoral cycles.

Introduction

Representative democracy works best when voters are capable of using the ballot to motivate and incentivize policymakers to faithfully represent their interests.[1] Given strong evidence that voters remain woefully ignorant about the operation of American political institutions and the policy positions of competing candidates on Election Day (e.g., Campbell et al. 1960, Delli Carpini and Keeter 1997), the question of whether ordinary citizens indeed possess such capacity remains at the center of the study of political behavior.[2]

Retrospective voting provides one potential mechanism to resolve this democratic dilemma, allowing for effective representation with minimal voter effort. By conditioning their support for political incumbents on one piece of easily accessible information — the observed quality of government outputs — voters can potentially use electoral incentives to create the political equivalent of Darwinian "natural selection," ensuring that only faithful agents retain their offices over the long haul (e.g., Fearon 1999, Ferejohn 2014). A large body of literature examines whether voters behave in this way in national elections.

But not all policymaking takes place in Washington. A large number of political issues — including the policies that arguably have the most effect on voters' quality of life on

[1]Our statement is general enough to encompass both the "delegate" and "trustee" conceptions of representation (Pitkin 1967), as one can interpret "constituent interests" to reflect voters' stated preferences or more paternalistic understanding of their interests.

[2]Some scholars assert that voters do not need full information to make reasonably informed decisions, arguing that simple information shortcuts can substitute for more detailed knowledge (e.g., Downs 1957, Lau and Redlawsk 1997, Lupia and McCubbins 1998, Popkin 1994). Others, however, argue that cues and other shortcuts do not necessarily lead to the same quality of decisions and are subject to serious behavioral biases (e.g., Bartels 1996, Healy and Lenz 2014, Healy, Malhotra and Mo 2010).

a day-to-day basis — remain the purview of state and local governments. Only recently, however, have scholars shifted their attention to examine the incidence of retrospective voting in sub-national elections (e.g., Arnold and Carnes 2012, Berry and Howell 2007, Miller 2013, Oliver and Ha 2007, Rogers 2013).

In this paper, we build on this emerging literature to make three significant contributions to our understanding of democratic accountability in American local government. Our first goal is to provide a more nuanced theoretical account of voter behavior in these elections. In contrast to many prominent scholarly claims that the electoral dynamics at the city level are inherently different from national contests, we show that voters exhibit some of the same behavioral patterns. In particular, we provide convincing evidence that voters hold local officials accountable for observed quality of government outputs.

Second, we highlight an important theoretical distinction between two types of behavior often combined under the banner of "retrospective voting." Although constituents may condition their votes on perceived service quality, we argue it is important to consider whether they correctly attribute credit or blame for observed outcomes to the proper agents in government. When they fail to do so — and hold the wrong office accountable or blame public officials for outcomes beyond their control — retrospective voting might lead to *worse* democratic outcomes, a possibility that the existing literature on local accountability largely overlooks.

Third, our approach brings two important improvements over the extant empirical literature by using better measures and a more credible identification strategy. By examining the political impact of street and road quality on elections in one of America's largest cities, we analyze a policy area that is unquestionably under the control of city government. Doing so has a significant advantage over existing studies, which instead focus almost exclusively on outcomes such as unemployment rates and crime levels over which city officials have only

4

limited direct control. Our analysis focuses on variation in service quality *within* a political jurisdiction, recognizing that government services and outputs likely vary significantly within cities, and show that such variation can explain neighborhood-level differences in support for incumbents in two political offices — mayor and city council — across two different electoral cycles.

In the remainder of the paper, we proceed by providing a brief overview of the existing theoretical and empirical literature on voter behavior in city elections. In doing so, we situate the available research on local retrospective voting within the broader debate about the dynamics of local democracy. Next, we introduce our case — the City of San Diego — and describe how we combine geocoded municipal records on street repair complaints with a panel of precinct-level data on election outcomes, allowing us to examine how road quality influences voter behavior in San Diego city elections. After describing the unusual political context during the years included in our analysis — a context that provides a particularly close fit between the theoretical ideal of retrospective voting and our data — we report the results of our analysis and conclude with a broader discussion of the findings, focusing in particular on the general questions of voter competence and democratic accountability. The takeaway message for the reader is necessarily nuanced, neither wholly laudatory nor completely gloomy regarding voters' capacity and ability to use the ballot box to motivate government officials to pursue socially desirable policy.

Retrospective Voting in Local Elections

One key question motivating our analysis is how similar local political dynamics are to federal elections. To be sure, there are some obvious differences: While all federal elections are partisan, approximately 80 percent of local government positions are filled through non-partisan contests (Kaufmann 2004). But are these differences meaningful enough to

5

fundamentally affect voting behavior and, in turn, the nature and quality of democratic representation?

A large literature in urban politics argues that cities are indeed unique. Much of this work begins with the recognition that local governments face a number of constraints absent from the national scene — including competitive pressures from nearby jurisdictions (e.g., Peterson 1981, Tiebout 1956), interference from higher-level governments (e.g., Desmond 1955, Gamma and Kousser 2013), and binding tax-and-expenditure limitations (e.g., Brooks and Phillips 2010, Mullins and Wallin 2004) — that greatly limit local discretion and the menu of available policy options.[3] As a result of these constraints, political competition at the city level often focuses on allocational questions of how to divide a fixed amount of resources, pitting broad groups of citizens against one another. As Kaufmann (2004) summarizes:

> Local politics — and the kinds of issues that dominate local elections — are often more proximate and more discrete than are the larger symbolic issues of national elections. Local governments, while sometimes lawmakers, are principally service providers. ... To most Americans, the role of their local government is to maintain or enhance their immediate quality of life, to provide necessary services, and at times to ameliorate intercity conflict. Thus the citizen's expectations of local government are inherently more connected to daily life than are his or her expectations of other governmental bodies. And as such, how people view their local leadership and why they vote for them will likely reflect the proximity of their interests (pp. 18-19).

---

[3]A growing empirical literature challenges this view, showing that city governments have the capacity to tailor policies to respond to the ideological preferences of their voters (Einstein and Kogan 2015, Tausanovitch and Warshaw 2014).

One implication of this reality is that distributional issues are more likely to take central stage in local contests, which perhaps explains why voting falls along racial and ethnic lines to a much greater extent in these elections (Hajnal and Trounstine 2014, Kaufmann 2004). The second implication, however, is that constraints on local government power should shape how retrospective voting operates in city elections. If municipal government influence is limited to a small set of basic housekeeping services, it makes little sense for voters in local elections to hold city officials accountable for outcomes that are largely outside of their direct control.

Surprisingly, however, much of the empirical literature on retrospective voting in the local context focuses on just such issues. For example, Hopkins and Pettingill (2015) examine the importance of economic voting in big-city mayoral elections, although there is little evidence that variation in sub-national economic performance can be explained by local policy choices (e.g., Prillaman and Meier 2014), rather than the spatial distribution of industries and occupations and natural resource endowments, factors over which local governments have minimal control. Similarly, Berry and Howell (2007) analyze the extent to which student achievement affects the reelection prospects of school board members, although there is clear evidence that much of the variation in student achievement is explained by student-level factors, such as socioeconomic status (Chingos, Whitehurst and Gallaher 2015, Kogan, Lavertu and Peskowitz 2015), rather than local educational policies. Indeed, Miller (2013) finds the win-loss records of local professional sports teams — a statistic over which government officials clearly have no control — frequently determine local mayoral election outcomes. He interprets this result to mean that voters follow a "Prosperity Model" of voting, holding incumbents accountable for their overall well-being, regardless of whether government policies played any actual role in determining their level of happiness.

The existing literature rarely acknowledges the disconnect between what local governments actually do and where analysts search for electoral accountability. When scholars do note the apparent discrepancy, they nevertheless draw sanguine conclusions about the normative implications for local democracy. Miller's (2013) discussion of the likely welfare implications of voters relying on sports records when deciding whether to reelect their incumbent mayors offers a case-in-point:

> [T]he flip side to how voters choose their leaders is how this incentivizes politicians in office looking for reelection. Under the Prosperity Model, incumbent politicians are motivated to maximize voter happiness at the time of the election. This is hardly the gravest of threats, even if motivating politicians to maximize well-being in the long run is preferable. ...
>
> It is fair to say that voters are imperfect and occasionally irrational. It does not follow that democracy is seriously questioned. In fact, voting for incumbents when well-being is high is a sensible rule of thumb for voters who cannot reliably connect the dots between political actions and outcomes. Although irrelevant events may seep into this calculation, voting remains an effective method for selecting and disciplining leaders (pp. 16-17).

We contend that such conclusions are not always warranted. Consider the findings of Arnold and Carnes (2012), who use several decades of panel data on New York City mayoral approval to show that voters appear to blame local elected executives when crime levels rise and credit them when crime declines. In reality, however, there is little evidence that local governments have much control over crime rates. In a recent exhaustive analysis, Roeder, Eisen and Bowling (2015) show that much of the decline in crime between 1990 and 2013 was a result of broad national trends, such as falling alcohol consumption, rising income, an aging population, and the legalization of abortion. While local policies, such

as the introduction of data-driven law enforcement approaches (i.e., CompStat) and the growing size of local police forces, played a small role on the margin, forces over which local governments had no direct control dwarfed these effects. Yet voters — at least in New York City — appear to have credited local officials anyway.

If "retrospective voting" is merely a reflection of irrelevant considerations, the implications for local democracy are potentially far more troubling than Miller (2013) suggests.[4] If mayors know that voters will hold them accountable for crime, but also realize there are few tools at their disposal to influence crime levels directly, they may be tempted to shape the few things that *are* under their control — namely, how crime rates are measured and reported. Indeed, recent scandals in both Chicago (Bernstein and Isackson 2014) and Los Angeles (Poston and Rubin 2014) provide strong reason to worry. In both cities, mayoral administrations faced growing political pressure to reverse rising violent crime rates. In the absence of policy tools to do so quickly enough to reap the political payoff ahead of their elections, mayors in both cities put pressure on their police departments to misclassify serious crimes to make the statistics look less dire.

Thus, in determining whether local democracy "works," we argue it is insufficient to measure the extent to which observed policy outcomes shape incumbent electoral prospects. It is also necessary to consider the extent to which voters allocate credit and blame appropriately (see also Brown 2010, Lowery, Lyons and DeHoog 1990), correctly linking these outcomes to the government officials responsible. Our study moves in this direction by shifting the focus to precisely the kinds of services controlled by city government — in

---

[4]Arnold and Carnes (2012) take a similarly glass-half-full view, concluding, "The electoral connection provides powerful incentives for mayors to do their best in tackling whatever problems occur on their watches, just as it does for elected executives around the world" (p. 962).

our case, the quality of local roads — and by paying attention to important within-city heterogeneity in service quality. If voters indeed vote retrospectively, we should expect their behavior to reflect the quality of the services that *they* receive, which we can rarely infer from average citywide outcomes that are the typical focus of analysis in the existing literature.[5]

Fiscal Crisis and Government Services in 'Enron-by-the-Sea'

As is the case with most of the leading research on local political dynamics, our study examines a single city, San Diego. The primary disadvantage is that doing so potentially limits the generalizability of the findings. Fortunately, this is much less of a concern in our case, because the city we examine is similar in important respects to most other major and mid-size U.S. cities. In particular, San Diego is ethnically and racially diverse, overwhelmingly Democratic, utilizes a "strong mayor"[6] form of government, and elects city officials through nonpartisan ballots. Across these dimensions, and many others summarized in Table 1, America's eighth largest city is thus a representative case, so the political dynamics we document are likely to provide useful insights about city politics writ large.

They key advantage of a single-city case study, however, is that we can utilize detailed data on both service quality and electoral outcomes at a very low level of aggregation, an empirical strategy that is simply not practical to implement across multiple jurisdictions. Since our unit of analysis is the electoral precinct, and we focus on variation across neigh-

---

[5]Berry and Howell's (2007) finding that voters appear to place greater weight on the achievement levels in their neighborhood schools, rather than the average achievement in each school district, provides strong empirical support for this argument.

[6]Technically known as a mayor-council, as opposed to council-city manager, form of government

Table 1: San Diego Descriptive Statistics

Population (2010)	1.3 million
City Share of Metro Population	42.2%
Population Under 18	21.4%
Population Over 65	10.7%
White Population	45.1%
Latino Population	28.8%
Asian Population	15.9%
Black Population	6.7%
Median Household Income (2009-2013)	$64,058
Poverty Rate (2009-2013)	15.6%
Government Structure	Council-Mayor
Election Method	Nonpartisan
Obama Two-Party Vote-Share (2012)	62.7%

Sources: U.S. Census Bureau, California Secretary of State

borhoods, the effective N of our study is still significantly higher than in a typical analysis of local government election outcomes.[7] Most importantly for our purposes, San Diego is a useful context in which to examine voting behavior because the city faces — and, in some ways, exemplifies — many of the competitive and legal constraints confronted by many other local governments.[8]

Much of San Diego's population growth occurred in the three decades after World War II, an era that coincided with emerging fiscal pressures. In 1978, California voters passed Proposition 13, a constitutional initiative that permanently reduced and froze property taxes and limited the reassessment of property values. Although the measure produced

---

[7]More typically, scholars focus on aggregate, city-wide election outcomes for a small handful of cities, ranging from from fewer than a dozen to a few hundred (e.g., Hajnal and Trounstine 2014, Hopkins and Pettingill 2015).

[8]The political evolution of the city in the second half of the 20th century is examined in detail by Erie, Kogan and MacKenzie (2011), so we summarize only the most relevant facts in this section.

a sizable negative revenue shock for many California local governments, the effect was particularly stark in San Diego, which had unusually low property taxes prior to the change. As a result, Proposition 13 permanently locked in the city's low revenue (Erie, Kogan and MacKenzie 2011, Chapter 3).

As the city continued to grow, local officials found creative ways to supplement local tax revenues to balance the annual budget. Starting in the early 1980s, for example, the city began diverting "surplus" earnings from its municipal pension system to pay day-to-day operating expenses. In 1996, when police overtime for the Republican national convention and the construction of a new downtown baseball stadium for the San Diego Padres stretched the budget even thinner, the city went further and struck a secret deal with its independent pension board to systematically underfund the pension system. In exchange for their support to allow the city to pay less than its actuarially-required pension contribution, city employees received a significant increase in pension benefits. When city finances faced even greater pressures after the 2001 recession, this deal was expanded, with the city continuing to underfund the pension system in exchange for yet more benefit increases for city workers.

After journalists brought to light these underfunding schemes in the early 2000s, the city was left on the brink of municipal bankruptcy. It significantly increased its pension contributions, to help close a now-multi-billion-dollar unfunded pension liability. Because the city did not adequately disclose the pension underfunding on the its financial statements, San Diego's external auditor refused to sign off on the city's books, which effectively banished San Diego from the municipal bond markets. With the city poised at a political and fiscal precipice, San Diego attracted growing national attention, and *New York Times* headline writers called the city "Enron-by-the-Sea." Unable to borrow and facing

a ever-increasing pension bill,[9] the city made dramatic reductions to city services and allowed basic infrastructure repairs to go unaddressed. By 2010, the bill for unfunded capital maintenance exceeded \$1 billion, more than the city's annual day-to-day operating budget.

Road Quality and Citizen Satisfaction

Table 2: Selected San Diego Public Service Levels

Activity	Fiscal Year 2001	Fiscal Year 2010
No. of customer service centers	15	0
No. of aquatic program users	712,000	255,380
No. of youth program after-school sites	37	15
Average no. of weekly rec center hours	62.3	40.3
Average no. of weekly library hours	48	36
Firefighters (per 100,000)	80	71
Police officers (per 100,000)	166	158

**Sources:** Erie, Kogan and MacKenzie (2011, Table 4.4)

As Table 2 documents, San Diego's fiscal and pension crisis resulted in substantial reduction in city service levels. Library and recreation center hours were cut, the city's work force was downsized, and essential services were reduced. One of the most striking changes, however, was the continued deterioration of city streets and sidewalks. With no bonding capacity and little spare operating revenues with which to carry out even the most basic street maintenance, road quality declined rapidly. One of the city's major news organizations created a website on the photo sharing service Tumblr, which it named The Stumblr, for readers to submit pictures of the worst streets and sidewalks. Even in 2015,

---

[9]The 2008 recession made things even more difficult. As a consequence of the recession, city revenues declined and the pension system suffered investment losses equal to roughly 20 percent of its asset portfolio. This further increased the pension payment even as revenues declined, squeezing basic city services.

years after the city's fiscal crisis passed and revenues rebounded, a national assessment found that over half of San Diego's streets were in poor quality, ranking it the 8th worst in the country (*Bumpy Roads Ahead: America's Roughest Rides and Strategies to Make Our Roads Smoother* 2015).

When San Diego officials commissioned a representative survey of city residents in 2010, satisfaction with the quality of city streets was ranked second to last among the 25 services evaluated by respondents, just ahead of downtown parking availability. In a Priority Spending Index, which combined citizen satisfaction with existing services and their willingness to pay higher taxes or fees to fund maintenance or avoid further cuts, city streets led all other priorities by a large margin (see Figure 1).

Figure 1: Constituents' Top Spending Priorities

Priority Spending Index
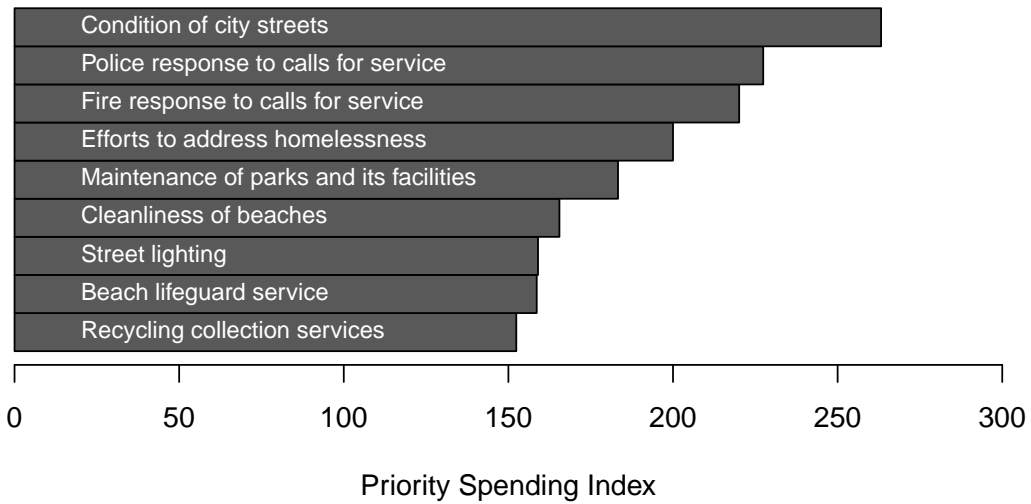
Source: *City of San Diego Residents' Opinions on City Services* (2010)

14

Empirical Strategy

Given the magnitude of service declines, and the salience of street quality in particular among voters, San Diego provides the perfect political context in which to examine whether voters hold local government officials accountable at the ballot box for observed performance. In this section, we provide an overview of our empirical strategy for carrying out just such an analysis. It describes the variables we use and the statistical models we estimate to measure the impact of road quality on voter support for local political incumbents.

Independent Variable: Citizen Pothole Complaints

To identify the effect of municipal service quality on incumbent political prospects, we leverage geographic variation in road conditions between neighborhoods in San Diego. In doing so, we assume that voters care most about the quality of government services that they personally receive, rather than the average service quality citywide. Although there are a number of potential justifications for this assumption, the most plausible one is that voters are likely to form their retrospective judgments based on casual observations from their own day-to-day experiences (Popkin 1994). It seems much more realistic that voters will tend to notice whether a pothole has developed on their city block than to assume that they calculate the average asphalt quality index for all city-maintained streets, including far-away areas that they never visit.

The key advantage in focusing on within-city variation in road quality is that such differences between neighborhoods are plausibly independent of incumbents' expected electoral performance. This is particularly true in our case, because we condition on each incumbent candidates' lagged level of support in the previous election. Conditional on previous electoral performance, whether a pothole develops on one city block but not a few blocks

away should be as-if randomly assigned, allowing us to identify the effect of service quality on voting behavior.

Our key treatment variable is the number of citizen pothole complaints, aggregated at the level of the electoral precinct. The data, covering the years 2008 through 2011, was first obtained by the *Voice of San Diego*, an online investigative journalism organization, through a public records request.[10] The database contains more than 52,000 complaints and service requests recorded by the San Diego's street maintenance department during these years.[11] Focusing on citizen complaints, rather than aggregate measures of road condition, also ensures that we measure service quality using the information that is clearly visible and salient to local residents.[12]

Table 3: Summary of Service Calls

Problem Code	% of Total Complaints
Major Asphalt Repair	22.7%
Minor Asphalt Repair	8.9%
Pothole	68.4%

Table 3 provides a breakdown of service calls recorded by the streets department. Overall, more than two-thirds of the recorded complaints focused on potholes, although the data include a significant number of other, more serious road repair requests. For convenience, we refer to all of these service requests as pothole complaints in the text below.

---

[10] *Voice of San Diego* generously agreed to share this data with the authors.

[11] Constituents have a number of ways to transmit these complaints to the city. First, they can contact their local city council member, whose staff pass the information on to the streets department. Alternatively, residents can send the information directly to the department using an online form or a phone-based application.

[12] If no one noticed or reported a pothole to the city, there is little reason to expect it to affect electoral outcomes.

Both theoretical and empirical literature on voting decisions illustrate that voters' use of performance information is subject to behavioral biases, such as myopia. Recent performance information, such as macroeconomic conditions in the six months before an election, has a disproportionate effect on electoral outcomes (e.g., Healy and Lenz 2014, Huber, Hill and Lenz 2012, Kogan, Lavertu and Peskowitz Forthcoming). We expect voters to exhibit similar myopia in our local context. As a result, we calculate a count of total complaints filed in each precinct in the six month window immediately preceding each election using the notification dates recorded in the city's database.[13]

Dependent Variable: Incumbent Electoral Performance

The dependent variable used in our analysis is the share of votes won by incumbent candidates in each election. In particular, we focus on two offices most responsible for making policy and setting funding levels for city services, the mayor and the city council. During our period of study, San Diego held two citywide elections, in 2008 and 2010. In each year, local candidates competed in a primary in June. If no candidate received more than 50 percent of all votes cast, the top two vote-getters proceeded to a runoff in November.

Although San Diego had eight city council districts during this period, term limits mean that we observed incumbent council members running for reelection in just two districts in these years. Thus, our analysis focuses on three office holders: Mayor Jerry Sanders and Councilmen Kevin Faulconer and Tony Young.[14] Sanders stood for re-election in 2008,

---

[13]The choice of a six-month window also reflects, to a certain extent, data availability. The earliest complaints we observe are from January 2008, and the first election during the period covered in our data took place in early June 2008.

[14]We exclude open city council races for two reasons. First, it is less plausible that responsibility follows the incumbents' co-partisans in these races, because partisan labels do not appear on the ballot, making partisanship less salient. Second, most of these races

while the two councilmen both ran in 2010. All three candidates won more than 50 percent of the vote in the June primaries, securing their reelection without a November runoff. Both Sanders and Faulconer are Republicans, while Young is a Democrat. Supplemental Appendix A describes how we handle changing precinct boundaries over time, to make the election data comparable across electoral cycles.

Statistical Models

Our baseline specification of each incumbents' vote share in precinct i at the time of their reelection, $Incumbent_{it}$, is:

$$Incumbent_{it} = \alpha + \tau \text{Potholes}_{it} + \beta \text{Incumbent}_{it-1} + \epsilon_i$$

We include each candidates' lagged vote share from the previous election, $Incumbent_{it-1}$, to capture several possible dynamics. First, San Diego exhibits significant racial — and as a result, partisan — residential segregation. The heavily Democratic areas are concentrated in the oldest, poorest, and most industrial parts of the city south of Interstate 8. This creates the potential for a spurious correlation between local road quality and partisan political dynamics. Second, it is also possible that that local officials strategically target city spending to certain constituents based in part on their expected level of support in each neighborhood (e.g., Cox and McCubbins 1986, Stokes 2005). If policymakers direct road resurfacing funds to neighborhoods containing their strongest supporters, we would see a correlation between voting behavior and road quality, but the direction of the causal arrow would be reversed. Conditioning on lagged vote share addresses both of these potential

---

feature multiple candidates from the same party, and there are few theoretical expectations about which co-partisans voters should hold responsible for the incumbents' performance.

problems.

Since service levels are likely correlated between nearby heighborhoods — for example, voters on two adjacent streets are likely to patronize the same local library branch or recreation center — the simple OLS specification does not account for the spatial relationships between precincts. To account for the fact that service quality is likely to be more similar among nearby precincts than between distant locations, resulting in spatial clustering in the error terms, we also estimates models that add a spatial error component to our baseline specification:

$$Incumbent_{it} = \alpha + \tau \text{Potholes}_{it} + \beta \text{Incumbent}_{it-1} + + \lambda \mathbf{w}_i \xi_i + \epsilon_i$$

where $\mathbf{w}_i$ is the spatial weights matrix specifying the relationship between precincts and ξ_i is a vector of error terms.

As a final robustness check, we also estimate a spatial lag model, which allows voting behavior in one precinct to be affected by its proximate neighbors. The motivation for the spatial lag is that voters likely notice and care about road quality not only in their neighborhood but also the condition of the other streets they use regularly, perhaps during their daily commutes. This specification thus allows vote share in precinct i to be affected by potholes in the same precinct but also in other nearby areas. The precise specification is:

$$Incumbent_{it} = \alpha + \tau \text{Potholes}_{it} + \beta \text{Incumbent}_{it-1} + + \rho \mathbf{w}_i Incumbent_{it} + \epsilon_i$$

where $\mathbf{w}_i$ is again the spatial weights matrix and $Incumbent_{it}$ is a vector of election outcomes.[15]

---

[15]As an additional robustness check, we have also estimated a spatial Durbin model.

19

Theoretically, it is most plausible that the interrelationships between precincts are inversely proportional to their distance from one another, following Tobler's first law of geography: "Everything is related to everything else, but near things are more related than distant things" (**?**). To reflect this, our preferred specifications use spatial weights that are an inverse of the quadratic distance between precincts. This corresponds to spatial effects that decay exponentially as the distance between precincts grows.[16]

For the city council elections, we pool the results from the two districts, but modify the models above by adding the subscript d to the intercept term, α_d, representing the district-specific fixed effects. Because we observe city council elections with both Democratic and Republican incumbents, this ensures that the effects we document are not the result of partisan attribution of responsibility.[17]

Results

In order to map each pothole complaint to its corresponding electoral precinct, we rely on two data fields included in the street complaint database. The first is *StreetAddress* and lists either the street block containing the pothole (e.g., "12000-12199 LOMICA DR") or the intersection (e.g., "WESTMORE CR & WESTMORE RD"). The second field is *ProblemDescription* and contains a more precise description of the location. In some case *ProblemDescription* includes a exact address (e.g., "ALLEY Behind 2315 31ST -

Since the results were identical to those below, we do not present to conserve space.

[16]Our results are not sensitive to this choice. We find the same substantive effects if we instead use an inverse of the linear distance or apply equal weights to all observations within a mile of each precinct.

[17]This might occur if voters blame the mayor's co-partisans for potholes, regardless of whether they are the incumbents or challengers in the race.

POTHOLES") while in others it reports only a vague description of the location and problem (e.g., "SEMILLON&KINGSPINE-P/H NEXT TO PEDESTRIA").

As a first step, we used an automated text parsing algorithm to extract an exact numerical address from the *ProblemDescription* field when one was reported. If no numerical address could be found, we used the midpoint of the block or the intersection reported in *StreetAddress*.[18] Using this procedure, we successfully obtained latitude and longitude coordinates for 50,642 of the 52,489 potholes in the full dataset, a match rate of 96.5 percent.

Of course, the validity of our analysis depends on the accuracy of the coordinates we obtained. It is worth noting that geocoding errors would result in measurement error that would attenuate the results, so such errors cannot explain the strong effects we report below. Nevertheless, it would be helpful to benchmark the accuracy of our procedure against some measurable baseline. We can do so by mapping the coordinates in our data to city council districts and comparing these against the actual district locations, which are reported in the original city pothole database.[19]

Table 4 provides just such a validity check. Overall, it shows that our geocoding procedure was quite effective, reaching or surpassing 98 percent accuracy in all but one city

---

[18]As a first cut, we ran the resulting location through the Texas A&M web-based geocoding service. This produced an exact hit for most observations, although some could not be matched to specific address (usually when the block mid-point address did not correspond to a valid address, and for all street intersections). These remaining addresses were then geocoded with Google Maps, which interpolates coordinates for nonexisting addresses based on the locations of valid nearby addresses and can also parse intersections.

[19]Of course, when the two differ, one should not necessarily assume that our coordinates are more likely to be wrong than those used by city staff.

council district. In District 3, where the procedure produces the lowest level of accuracy, we still correctly geocoded more than 95 percent of the observations. These results strengthen our confidence in the accuracy of our key treatment variable.

Table 4: Assessing Accuracy of Geocoded Coordinates

Actual Location	Correct (Count)	Total Geocoded	Correct (Percent)
District 1	4,666	4,708	99.1%
District 2	9,237	9,424	98.0%
District 3	7,572	7,922	95.6%
District 4	6,660	6,697	99.4%
District 5	4,803	4,861	98.8%
District 6	6,073	6,151	98.7%
District 7	5,742	5,774	99.4%
District 8	4,790	4,845	98.9%

Tables 5 and 6 report the regression results for the mayoral and city council elections, respectively. The first column in both tables corresponds to the baseline OLS model, the second reports the spatial error specification, and the third contains the spatial-lag model. The primary quantity of interest is the coefficient on the pothole count variable.

Across all specifications and both office types, we find a significant negative electoral effect for pothole complaints. When constituents submit more road work requests in the six months before the election, incumbent office holders suffer at the polls. The coefficients are also similar in size for both the mayoral and council elections, with each additional pothole complaint reducing incumbent vote share by roughly 0.2 percentage points. Increasing the number of potholes from the 25th percentile in 2008 (0 potholes) to the 75th percentile (6 potholes) is thus expected to reduce the incumbents' share of the vote by about 1.2 percentage points. Although this effect is not particularly large, it is meaningful in substantive terms. In a close election, local road quality could determine whether the incumbent wins another term or loses the election. Note that the coefficients on the spatial-lag specifications are somewhat smaller because they include only the short-run rather than

the full equilibrium effects that account for the dynamics of the spatial spillovers between precincts, and are thus not directly comparable to the other two columns (see Ward and Gleditsch 2008, pp. 44-49).

Table 5: Mayoral Election Results

	(Model 1)	(Model 2)	(Model 3)
Potholes Count	-0.178*	-0.20**	-0.13*
	(0.07)	(0.06)	(0.06)
Lagged Vote Share	0.36**	0.31**	0.19**
	(0.02)	(0.03)	*0.03)
Precincts	742	742	742
Specification	OLS	Spatial Error	Spatial Lag

** p<0.01, * p<0.05, using two-sided tests

Table 6: City Council Election Results

	(Model 1)	(Model 2)	(Model 3)
Potholes Count	-0.26*	-0.22*	-0.22*
	(0.12)	(0.11)	(0.10)
Lagged Vote Share	0.77**	0.69**	0.65**
	(0.09)	(0.09)	(0.09)
Precincts	164	164	164
Council District FEs	Yes	Yes	Yes
Specification	OLS	Spatial Error	Spatial Lag

** p<0.01, * p<0.05, using two-sided tests

It is useful to put these effects in context by comparing them to the findings on economic voting in U.S. presidential elections. In these contests, incumbent presidents gain roughly 4 percentage points of vote share for each percentage point of election-year income growth (Healy and Lenz 2014). Thus, the electoral impact of a single pothole on city contests is roughly one-twentieth the magnitude of a single percentage point of income growth in national presidential elections. When considered in such comparative terms, the impact of road quality on local democracy is surprisingly large.

We conclude by noting that, across all elections and specifications, we find a strong

relationship between current electoral performance and lagged vote share. That should not be surprising — incumbents who perform well in a neighborhood in the previous election are likely to do well there again next time they run for office. The coefficient in the mayoral election, however, is not particularly large, with Mayor Sanders gaining roughly 0.3 votes in 2008 for each vote he won during the 2005 special election that first brought him to office. This relatively modest relationship likely reflects important differences in dynamics between these two elections. In the 2005 mayoral runoff, Sanders' opponent was Democratic City Councilwoman Donna Frye, a self-styled progressive and environmental activist. In 2008, by contrast, local Democrats failed to recruit a viable opponent. Instead, Sanders' leading challenger was Republican businessman Steve Francis. Although Francis tried to recast himself as an independent populist in 2008, he was widely considered to be more conservative than Sanders, a moderate Republican who publicly embraced same-sex marriage long before it was fashionable to do so. Given these important differences, it is not surprising that Sanders' winning coalition would change significantly between these two elections.

Discussion and Conclusion

Overall, we find strong evidence that voters' retrospective evaluations of government performance have a meaningful impact on local elections and the political fortunes of incumbents. That the results are consistent across three different modeling approaches, each of which makes different statistical assumptions about voter behavior, gives further credence to their validity. These findings challenge prominent claims in the literature that local nonpartisan elections are systematically different from partisan national contests (e.g., Kaufmann 2004). In addition, our empirical approach makes several important theoretical and methodological contributions to the literature on retrospective voting, particularly

in sub-national contexts. First, our analysis offers a much closer fit between theory and empirics by focusing our analysis on outcomes that are actually under the control of local government. Second, our analysis relies on a more credible identification strategy by exploiting plausibly exogenous variation in service equality between neighborhoods, rather than differences in average outcomes between cities. Perhaps most importantly, our study shows the significant scholarly potential for applying rigorous social science methods to the analysis of fine-grained government administrative data. In the modern era of "big data," governments everywhere are dramatically increasing the types and availability of the metrics they collect, which will provide important opportunities for bringing evidence to bear on outstanding theoretical debates.

While our findings are hopeful in normative and welfare terms, they do not offer convincing answers for the debate about the quality of American democracy. Whether the type of retrospective voting we document improves democratic accountability or perversely leads to less optimal policies depends on a variety of additional considerations. First, our analysis focuses on the one government output that voters themselves report caring most about. It remains to be seen, however, the extent to which voters also incorporate performance information about less salient or visible but arguably equally important services (e.g., storm and water drains) into their electoral calculation. Second, it is important to consider whether voters accurately attribute credit and blame for outcomes to the right individuals. Indeed, all of the incumbents we examine in this study were elected to office in the aftermath of San Diego's pension crisis, and none of them are arguably to blame for the dire service declines the city has experienced over the past two decades. It is unclear how holding Mayor Sanders and Councilman Faulconer and Young accountable for the sins of their predecessors improves democratic responsiveness. Finally, scholars must also pay attention to how politicians respond to the electoral incentives created by retrospective

voting. For example, if voters indicate that they value well-maintained streets, local politicians might feel political pressure to reallocate funding to road maintenance from other spending priorities that benefit primarily well-connected special interest groups, such as publicly financed sports stadiums (e.g., Erie, Kogan and MacKenzie 2010) and convention center expansions (e.g., Sanders 2014). Alternatively, elected officials might leverage the biases in voter behavior and systematically cut spending on essential street maintenance and use the savings to ramp up spending on repairs only in the months leading up to each election, thereby increasing their credit claiming opportunities but also arguably the long-term costs. All of these possibilities offer important questions that should be the subject of future research.

In sum, we argue the available evidence is insufficient to draw clear conclusions about the health of American democracy, particularly at local levels of government. The analogy to biology we refer to in the introductory section of this paper is worth returning to: Natural selection has helped humans evolve essential survival traits, but it has given rise to pathogens, such as antibiotic resistant bacteria, that pose significant long-term risks to our species. Similarly, the logic of political survival may lead politicians to pursue welfare improving policies, but it could also lead them to prioritize expedient short-term kludges over long-term solutions. Much depends on the types of retrospective evaluations voters bring to the ballot box, and how elected officials respond to the resulting political incentives.

References

Arnold, Douglas R. 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. 1996. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40(1):194–230.

Bernstein, David and Noah Isackson. 2014. "The Truth About Chicago's Crime Rate: Part 2." *Chicago Magazine* .

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

Brooks, Leah and Justin H. Phillips. 2010. "An Institutional Explanation for the Stickiness of Federal Grants." *Journal of Law, Economics, and Organization* 26:243–264.

Brown, Adam R. 2010. "Are Governors Responsible for the State Economy? Partisanship, Blame, and Divided Federalism." *Journal of Politics* 72(3):605–615.

Bumpy Roads Ahead: America's Roughest Rides and Strategies to Make Our Roads Smoother. 2015. Technical report TRIP Washington, DC: .
URL: *http://www. tripnet. org/ docs/ Urban_ Roads_ TRIP_ Report_ July_ 2015. pdf*

Campbell, Angus, Philip E. Converse, Warren E. Miller and Donald E. Stokes. 1960. *The American Voter.* New York: Wiley.

Chingos, Matthew M., Grover J. Whitehurst and Michael R. Gallaher. 2015. "School Districts and Student Achievement." *Education Finance and Policy* Forthcoming.

City of San Diego Residents' Opinions on City Services. 2010. Technical report Behavior Research Center Phoenix, AZ: .

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

Delli Carpini, Michael X. and Scott Keeter. 1997. *What Americans Know About Politics and Why It Matters*. New Haven, CT: Yale University Press.

Desmond, T.C. 1955. "States Eclipse the Cities." *National Municipal Review* 44:296–300.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York, NY: Harper.

Einstein, Katherine Levine and Vladimir Kogan. 2015. "Pushing the City Limits: Policy Responsiveness in Municipal Government." *Urban Affairs Review* .

Erie, Steven P., Vladimir Kogan and Scott A. MacKenzie. 2010. "Redevelopment, San Diego Style: The Limits of Public-Private Partnerships." *Urban Affairs Review* 45(5):644–678.

Erie, Steven P., Vladimir Kogan and Scott A. MacKenzie. 2011. *Paradise Plundered: Fiscal Crisis and Governance Failures in San Diego*. Palo Alto, CA: Stanford University Press.

Fearon, James. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, ed. Adam Przeworski, Susan C. Stokes and Bernard Manin. Cambridge University Press.

Ferejohn, John. 2014. "Incumbent Performance and Electoral Control." *Public Choice* 50(1-3):5–25.

Gamma, Gerald and Thad Kousser. 2013. "No Strength in Numbers: The Failure of Big-City Bills in American State Legislatures, 1880-2000." *American Political Science Review* 107(4):663–678.

Hajnal, Zoltan and Jessica Trounstine. 2014. "What Underlies Urban Politics? Race, Class, Ideology, Partisanship, and the Urban Vote." *Urban Affairs Review* 50(1):63–99.

Healy, Andrew and Gabriel S. Lenz. 2014. "Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy." *American Journal of Political Science* 58(1):31–47.

Healy, Andrew J., Neil Malhotra and Cecilia Hyunjung Mo. 2010. "Irrelevant Events Affect Voters' Evaluations of Government Performance." *Proceedings of the National Academy of Sciences* 107.

Hopkins, Daniel J. and Lindsay M. Pettingill. 2015. "Economic Voting in Big-City U.S. Mayoral Elections." Available at SSRN: http://ssrn.com/abstract=2567294.

Huber, Gregory A., Seth J. Hill and Gabriel S. Lenz. 2012. "Sources of Bias in Retrospective Decision Making: Experimental Evidence on Voters' Limitations in Controlling Incumbents." *American Political Science Review* 106:720–741.

Kaufmann, Karen M. 2004. *The Urban Voter: Group Conflict and Mayoral Voting Behavior in American Cities.* Ann Arbor: University of Michigan Press.

Kogan, Vladimir, Stéphane Lavertu and Zachary Peskowitz. 2015. "Performance Federalism and Local Democracy: Theory and Evidence from School Levy Elections." *American Journal of Political Science* .

Kogan, Vladimir, Stéphane Lavertu and Zachary Peskowitz. Forthcoming. "Performance

Federalism and Local Democracy: Theory and Evidence from School Levy Elections." *American Journal of Political Science* .

Lau, Richard R. and David P. Redlawsk. 1997. "Voting Correctly." *American Political Science Review* 91(3):585–598.

Lowery, David, William E. Lyons and Ruth Hoogland DeHoog. 1990. "Institutionally-Induced Attribution Errors: Their Composition and Impact on Citizen Satisfaction with Local Government Services." *American Politics Research* 18:169–196.

Lupia, Arthur and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge University Press.

Miller, Michael K. 2013. "For the Win! The Effect of Professional Sports Records on Mayoral Elections." *Social Science Quarterly* 94(1):59–78.

Mullins, Daniel R. and Bruce A. Wallin. 2004. "Tax and Expenditure Limitations: Introduction and Overview." *Public Budgeting & Finance* 24:2–15.

Oliver, Eric J. and Shang E. Ha. 2007. "Vote Choice in Suburban Elections." *American Political Science Review* 101(3):393–408.

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

Pitkin, Hanna Fenichel. 1967. *The Concept of Representation.* Berkeley: University of California Press.

Popkin, Samuel L. 1994. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns.* Chicago: University of Chicago Press.

Poston, Ben and Joel Rubin. 2014. "LAPD Misclassified Nearly 1,200 Violent Crimes as Minor Offenses." *Los Angeles Times* .

Prillaman, Soledad Artiz and Kenneth J. Meier. 2014. "Taxes, Incentives, and Economic Growth: Assessing the Impact of Pro-business Taxes on U.S. State Economices." *Journal of Politics* 76(2):364–379.

Roeder, Oliver, Lauren-Brooke Eisen and Julia Bowling. 2015. *What Caused the Crime Decline.* Brennan Center for Justice at New York University School of Law.

Rogers, Steven Michael. 2013. "Accountability in a Federal System." Ph.D. dissertation, Princeton University.

Sanders, Heywood T. 2014. *Convention Center Follies: Politics, Power, and Public Investment in American Cities.* Philadelphia, PA: University of Pennsylvania Press.

Stokes, Susan C. 2005. "Perverse Accountability: A Formal Model of Machine Politics with Evidence from Argentina." *American Political Science Review* 99:315–325.

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

Tiebout, Charles. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 54:416–424.

Ward, Michael D. and Kristian Skrede Gleditsch. 2008. *Spatial Regression Models.* Thousand Oaks, CA: Sage Publications Inc.

Supplemental Appendix

A. Adjusting Precinct Boundaries Between Elections

In California, election administration is the responsibility of local county registrars of voters, and the basic procedures are government by state law adopted by the state legislature and regulations promulgated by the Secretary of State. Each county registrar is responsible for creating two types of precincts, registration and consolidated or voting precincts. When a voter registers to vote, she is assigned to a home registration precinct based on her address. In San Diego County, new home precincts are drawn once a decade, following the census, and do not change over time. Registration precincts are generally very small, and all voter registration data is reported at this level.

In preparation for Election Day, home registration precincts are combined (or "consolidated") into voting precincts, which correspond to polling locations. These consolidation may not cross the boundary "of any supervisorial district, congressional district, senatorial district, Assembly district, board of equalization district, judicial district, incorporated city, ward, or city council district" (Election Code §12222). This ensures that all voters using the same polling location vote with an identical ballot, and election results are reported only at the consolidation level. State law allows voting precincts to be redrawn between elections, subject to the constraint that they cannot contain more than 1,000 voters, as measured by voter registration counts 88 days before the election (Election Code §12223). As a result, voting precinct boundaries sometimes change between elections. This might be done because new voters result in more than 1,000 registered voters living in a consolidated precinct, requiring it to be split, or to provide for smooth administration of the election. When turnout is expected to be high — for example, during during presidential elections — registration precincts are generally combined into a large number of geographically smaller

voting precincts. When turnout is expected to be low — e.g., during primary elections — registration precincts are consolidated into a small number of geographically larger voting precincts.

Because the level at which election results are reported changes across elections, we must standardize these boundaries to calculate both contemporaneous and lagged candidate vote share at the same level of aggregation. To do so, we follow the general approach used in the redistricting literature. Suppose we are interested in comparing candidate vote shares between an elections in year $t-1$ and year t. First, we calculate the total number of votes cast for each candidate at the election precinct level in year $t-1$. We then allocate these votes to the registration precincts that make up each consolidation in proportion to the registration precincts' shares of registered voters.[20] This step assumes that the registration precincts making up a single consolidation are political homogenous, an assumption that does not stray too far from the truth because consolidation boundaries do not cross census tracts, and these tracts are drawn by design to be homogenous in racial and economic terms. We then re-aggregate the registration precinct vote counts into the consolidations used in election year t.

_____

[20]Using the actual number of registered voters is preferable to using the citizen voting-age population, which is usually how redistricting projections are made. We obtain the registered voter counts for each election from the California Statewide Database: `http://statewidedatabase.org/`.

A.2