CrossMark

Ideology and Vote Choice in U.S. Mayoral Elections: Evidence from Facebook Surveys

Michael W. Sances[1]

Abstract City elections in the U.S. are widely thought to be low-information contests decided by non-ideological factors. This consensus casts doubt on the possibility of electoral accountability in cities, and renders recent evidence of municipal responsiveness puzzling. However, our knowledge of how voters actually behave in local elections is severely limited by a lack of individual-level survey data collected from local contests. Using three such original surveys, I re-examine the role of ideology in mayoral elections, recruiting samples of local voters via geo-targeted Facebook advertisements. In two large cities, I find ideology is a powerful and independent predictor of vote choice. Using a panel design, I find voters learn the relative ideological positions of candidates over the course of a campaign, and that learning causally impacts vote choice. The effect of ideology also replicates in a conjoint experiment fielded to a sample of small-city voters in another region. Electoral accountability is thus a plausible explanation for ideological

✉ Michael W. Sances
  msances@memphis.edu

[1] University of Memphis, Memphis, TN, USA

🖄 Springer

responsiveness in U.S. cities, and the methodological tools introduced here can now be applied to a variety of questions about local voter behavior.

Do city elections present voters with meaningful policy choices, or do structural forces limit the scope of city conflict? Urban politics scholars have long debated the nature of local political conflict, yet they have mostly agreed on one point: city elections are not about ideology. Rather than providing voters with a choice between two ideologically distinct paths for the future, city elections are either secondary to the process of locational choice (Peterson 1981), extensions of racial or economic conflict (Hajnal and Trounstine 2014), or referendums on performance (Arnold and Carnes 2012; Hopkins and Pettingill 2017; Burnett and Kogan 2017).

The apparent absence of ideological voting in city elections is both normatively concerning and theoretically puzzling. Normatively, if voters do not choose candidates based on their policy stances, then city elections fail to serve their intended purpose of allowing voters to select leaders who represent their policy views. This might be of no consequence if cities did not actually hold much discretion over policy—but they do. As shown by several recent studies (Gerber and Hopkins 2011; Tausanovitch and Warshaw 2014; Einstein and Kogan 2016; de Benedictis-Kessner and Warshaw 2016), U.S. cities vary widely both in the generosity of their spending and the progressivity of their tax systems. If voters in these cities do not choose their leaders on the basis of these issues, then city officials may wield their policy discretion without fear of electoral consequences.

Surprisingly, however, city officials do behave *as if* they were constrained by electoral preferences. Using different measures and samples, both Tausanovitch and Warshaw (2014) and Einstein and Kogan (2016) find a robust, positive relationship between voter preferences and city policies. While normatively reassuring, this evidence of responsiveness raises a theoretical puzzle: if local elections are not about ideology, how are citizens able to achieve ideological responsiveness in U.S. cities?

In this paper, I conduct a novel examination of the role of ideology in city elections. Prior conclusions about the lack of ideological voting in cities tend to be drawn either from theory alone, or from limited data. This is understandable, given the relative dearth of existing individual-level survey data from city elections, and the cost of collecting original city survey data. Yet as I demonstrate in this paper, researchers now have the opportunity to field city surveys more easily and inexpensively than ever before. Using Facebook advertisements targeted to particular local audiences, I recruit samples of hundreds of voters from two large cities in the state of Tennessee, as well as an additional sample of over 1000 voters from smaller cities in Illinois.

These data allow me to estimate the relationship between voter ideology and candidate choice, adjusting for other observable and unobservable determinants of voting. In cross-sectional analyses, I find that ideology is a powerful predictor of vote choice in each election, controlling for demographic characteristics such as

race. To further disentangle ideology from other factors, I next conduct an analysis of voter learning in the Nashville mayoral election. Using a panel design where the same respondents are interviewed before and after the election, I show that voter knowledge of candidates' ideological positions increases significantly over the course of the election campaign. Further, this learning causally impacts voters' choices, with liberals (conservatives) who learn becoming more (less) likely to vote for the liberal candidate. Finally, I show that local voters in smaller Illinois cities, when presented with hypothetical candidates for village and city mayors, also weight ideology heavily in their voting decisions.

These results show that electoral accountability for policy is indeed possible in both large and small U.S. cities, and offer an explanation for recent evidence of municipal responsiveness. Yet they also showcase the potential of using geotargeted online surveys to study local political behavior—a subject which scholars know surprisingly little about compared to national-level behavior. A recent count of election-related articles published in major journals over 20 years shows less than 1% focus on local elections (Berry and Howell 2007). In part, this ratio is due to the relative difficulty of collecting original survey data from local contexts. By introducing a new, relatively low-cost method of collecting geographically targeted data, the current study breaks down barriers that have long kept the study of local elections separate from research conducted at other levels of government.

Ideology and the Study of City Elections

Ideology has long been recognized as an important predictor of vote choice in U.S. national elections (Campbell et al. 1960; Knight 1985; Alvarez and Nagler 1995; Jacoby 2009). Even if voters do not vote in a purely spatial manner (Downs 1957), the two major political parties do offer voters a clear choice, with the perceived ideological gap between parties growing over time (Hetherington 2001). In contrast, city elections are commonly characterized as non-ideological. A recent review by Hajnal and Trounstine (2014) casts the literature as divided between those who argue city politics is largely pre-determined by locational choice, on one hand (Peterson 1981; Oliver 2012), and those who argue that racial and class conflict define city elections, on the other hand (Barreto 2007; Hajnal 2006; Kaufmann 2004). In their extensive review, Hajnal and Trounstine cite only one study examining the role of ideology (Abrajano et al. 2005). Similarly, in their discussion of possible mechanisms for ideological responsiveness in cities, Tausanovitch and Warshaw (2014, p. 606) do not cite any existing evidence that city elections are ideological contests.

The empirical foundation for the conclusion that ideology plays little role in city elections is, however, limited. As Berry and Howell (2007) note, there are dramatically fewer empirical studies of city elections compared to national elections—about 1% of the literature on elections concerns local elections, by their count—despite the fact that 95% of all elected officials in the U.S. are chosen at the local level. Few of these local election studies, moreover, actually analyze the relationship between voter ideology and candidate choice at the individual level, due to the simple fact that such data is hard to obtain. In their review of the state of

local elections research, Marschall et al. (2011) note that while "survey data...-tend[s] to dominate the study of local political behavior...the high costs of survey research have historically limited studies to a few cities and a single point in time" (98). Thus existing surveys of voter behavior in local elections are rare, and those that do exist are limited in the conclusions they can draw.[1]

Using the admittedly limited available survey data, a handful of studies do estimate the relationship between ideology and city vote choice. Kaufmann (1998) finds ideology predicts vote choice in the 1993 mayoral elections in New York and Los Angeles, but not the 1989 New York election. Stein et al. (2005) also find ideology is a statistically significant predictor of vote choice in the 1997, 1999, and 2001 Houston city elections, though less important than race. Using exit poll data from five large cities, Hajnal and Trounstine (2014) similarly find ideology is a significant predictor of vote choice, though they conclude racial and economic conflict are more important factors. Abrajano et al. (2005) find ideology predicts votes for mayor and city attorney in the 2001 Los Angeles elections, while Arceneaux (2006) finds no effect of ideology in three separate mayoral races in 2002.

While the few existing studies that examine this question do find evidence of ideological voting, they generally conclude that other factors dominate. There are two notable exceptions. First, Boudreau et al. (2015) measure voter and candidate ideal points in the 2011 mayoral election in San Francisco. They find that ideological distance is the most important predictor of vote choice, even more so than race. Second, Oliver and Ha (2007) find that subjective perceptions of issue agreement strongly predict vote choice in suburban city council races. However, given these studies rely on cross-sectional data, alternative explanations are possible. In particular, voters in local elections may first decide who to support based on non-ideological factors, and then shift their ideological and issue positions accordingly. Indeed, Lenz (2012) finds evidence of such following on highly salient national policy issues, such as whether to privatize the Social Security program. In the case of local elections, such following may be even more likely given the lower salience of city elections, as well as voters' predicted tendency to first choose mayoral candidates based on group identities such as race.

Studying Local Voter Behavior Using Facebook

The advent of online surveys has transformed the political behavior literature, but this revolution has so far left the study of local behavior untouched. In part, this is because the new online data collection tools are best suited to national-level research questions. Omnibus online surveys such as the Cooperative Congressional Election Study (Vavreck and Rivers 2008) aim to collect nationally representative samples of voters, while services such as Amazon.com's Mechanical Turk

---

[1] For instance, Hajnal and Trounstine (2014) assemble perhaps the most comprehensive data set on local voter behavior by merging exit poll data from five cities, noting that their sample is simultaneously as "broad as possible" and yet not representative of "the entire urban arena" (69).

(Berinsky et al. 2012) do not allow requesters to restrict their convenience samples to geographies smaller than the country level.

Yet new methods for geographically targeting survey recruitment allow local politics researchers to collect data more easily than ever before. Social media web sites such as Facebook sell advertisers the opportunity to place advertisements in front of specific audiences, which advertisers may customize based on age, gender, interests, and location. While this feature is aimed at private companies wishing to sell products to specific audiences, survey researchers can also use this technology to recruit specific samples. To date, researchers in political science have already taken advantage of this targeting to conduct cluster-randomized experiments on national U.S. samples (Ryan 2012; Broockman and Green 2014), and to collect survey data in developing countries such as Brazil (Samuels and Zucco 2014). Yet no study has fully leveraged these targeting methods to recruit participants into local election surveys— perhaps because the extent of Facebook's reach is not generally realized. According to data cited by Samuels and Zucco (2013), there are currently 1.7 billion active monthly users on Facebook worldwide, with 84% outside the United States and Canada.[2] In January 2015, *Adweek* reports a total of 152 million US-based users, or a little less than half of the total population of 320 million (Adweek 2016).[3]

Facebook allows any user to register as an advertiser and display advertisements, consisting of a short amount of text and a small image, to target audiences. When Facebook users click on the ad, they can be re-directed to an online survey hosted on an external site, such as the Qualtrics platform.[4] My surveys targeted residents in Memphis, TN; Nashville, TN; and smaller Illinois cities. Following past online survey recruitments (Berinsky et al. 2012; Samuels and Zucco 2014), my ad text in Memphis read "[Institutional name redacted] researchers want to know your views on the mayor election; take a 10 minute survey!" In Nashville, it read "Researchers want your opinion on the mayor's race. You could win a $100 Amazon gift card!"[5], and in Illinois, it read, "Researchers want YOUR opinion on Illinois issues. You could win a $100 Amazon gift card." Subjects from Nashville and Illinois were incentivized to participate with a chance of winning a single Amazon.com gift card worth $100; no added incentive was provided in Memphis.[6]

---

[2] These figures are cited in Samuels and Zucco (2013, p. 5), and I have updated them by visiting http://www.socialbakers.com/statistics/facebook/ (figures current as of May 4, 2016).

[3] In Brazil, in contrast, Samuels and Zucco (2013) report 40 million users out of a population of 200 million, or 20%.

[4] See http://www.qualtrics.com.

[5] For the Nashville ad, subjects could intuit the institutional sponsor of the survey by reading the accompanying link.

[6] Advertisers have the choice of paying Facebook per ad view (called *impressions*), or per click. In either case, Facebook presents advertisers with a range of possible prices (called *bids*) that might actually be paid, depending on the demand for ads. Following Samuels and Zucco (2013), I chose to pay per click, as the goal of my ad campaign was to have users click on the survey link. Actual costs per click for my surveys varied from $0.30 in Memphis, to $0.38 in Illinois, to $0.98 in Nashville. Actual costs per survey completion were $5.48 in Memphis, $4.45 in Nashville, and $1.31 in Illinois. The variation in cost per completion may be due to the vastly greater number of Facebook respondents in the suburban Illinois pool (5.9 million) versus in the Tennessee cities (about 300,000 each). I show screenshots of the ads used in Sect. 2 of the Online Appendix.

A useful feature of the advertising interface is the specification of demographic targets. In Memphis and Nashville, I restricted each audience to adults living within a zero mile radius of the city (thus excluding suburban areas) and who expressed interest in the following areas: local government; their own city (Memphis, TN or Nashville, TN); and politics and social issues. Prior to purchasing any ads, the Ads Manager returns an estimated "potential reach" based on these targets; these figures were 310,000 in Memphis and 330,000 in Nashville.[7] In Illinois, I restricted the audience to adults 18 and older living in small cities, and did not limit the audience based on interests; for this audience, Facebook returned a "potential reach" of 5.9 million users.[8] While selecting voters with an expressed interest in local politics may make a finding of ideological voting more likely, it seemed necessary in order to recruit both likely voters and a sufficiently large sample to make statistical inferences, given that those without such interests would be less likely to click on the recruitment ad. Of course, such a restriction is not so different from an exit poll design, which by construction only includes local voters. Additionally, I will show below that the effects are similar in the Illinois sample, which was recruited without any interest restrictions.

The Memphis election took place October 8, 2015. The Facebook ad campaign and Qualtrics survey were in the field from September 28th to October 6th. Over the course of the ad campaign, roughly 120,000 users viewed the ad, 4625 clicked on the survey link, 482 passed the initial consent screen, and 253 completed the survey. The Nashville election took place August 6th, 2015, and the pre-election survey was in the field from May 7th to June 2nd. About 196,000 users viewed the ad, and 3043 clicked to start the survey. Of these, 1029 consented, and 662 completed the survey. The Illinois survey was fielded from October 21 to October 29, 2016; 140,936 users viewed the ad, 3631 clicked, 1499 consented, and 1035 completed.[9]

---

[7] In Memphis, I restricted age to 18 and older. In Nashville, I restricted the population to 23 and older to avoid sampling college students. On the consent screen of each survey, respondents affirmed that they were eligible to vote in their respective mayoral election. For comparison, the Census estimates of the 18+ population for Memphis and the 23+ population in Nashville are 482,000 and 472,000, respectively.

[8] I limited the audience to small cities by targeting based on zip codes. Details on the generation of small-city zip codes are provided in Sect. 5 of the Online Appendix.

[9] I provide a tabular summary of the recruitment process in Sect. 1 of the Online Appendix. Survey break-off rates [partial responses/partial & complete responses; Callegaro and DiSogra 2008) were 50% in Memphis, 35% in Nashville, 32% in Illinois. These rates were not that out of line with median break-off rates found in meta-analyses of online surveys (Peytchev 2009, p. 75)], and were only slightly higher than the median break-off rate (27%) in my own analysis of eight studies I previously fielded via Survey Sampling International (detailed results available on request). However, break-off rates are significantly higher than surveys fielded by GFK (formerly Knowledge Networks); see, for example, Hainmueller and Hiscox (2010, p. 67) who report a break-off rate of 4.5%. Additionally, Peytchev (2009) reports break-off rates of 16 and 9% in two SSI samples fielded in 2003 and 2004. Section 3 of the Online Appendix tests whether break-off and attrition are correlated with any observable characteristics, or any particular moment in the survey. The major determinant of break-off appears to be fatigue, with most break-off occurring before around the tenth question. Those who start but fail to complete a single wave are slightly less likely to be Barry voters in Nashville, but there are no consistent predictors in the other two surveys. Those who fail to continue to the second wave in Nashville are slightly less likely to be Barry voters and slightly less likely to be white; while significant, these differences are substantively small.

At the end of the Nashville survey, I asked respondents for their e-mail addresses so as to recontact them for a second wave, for which they would be entered in a second gift card drawing. I fielded a post-election survey between August 7th and August 12th. Of the 662 respondents who completed the survey, 590 provided their e-mail addresses, 332 responded to the post-election invitation, and 291 completed the second wave survey.

How Representative Are Facebook-Recruited Samples?

A concern with any non-probability sample is how well it approximates the target population (Berinsky et al. 2012). With Facebook, in particular, we might be concerned that potential respondents are disproportionately young. In Brazil, Samuels and Zucco (2013) find that Facebook samples over-represent the young, the wealthy, and (in one study) men. However, they stress that their sample is more representative than other convenience samples.[10]

How well do my local samples match population characteristics? Figure 1 compares the age distribution of respondents (thick black lines) to known population distributions from the 2010 Census (gray areas).[11] According to the top panel, the Memphis sample recruited via Facebook is in fact somewhat *older* than the population. In general, however, there is considerable overlap between the two densities. The middle panel shows that the Nashville sample is also older than the population, though here the two densities are more similar to one another; the bottom panel shows that the Illinois small-city sample is similarly older than the population.

Table 1 presents additional comparisons between the three samples and their respective population figures from the 2010 Census. Again, the Facebook samples are older: the median age in the Memphis sample is 54 versus 44 in the Census; the corresponding figures for Nashville are 48 and 44, and in Illinois they are 49 and 58. On gender, the large-city samples approximate the population figures: 56% of the sample and 53% of the population in Memphis are female, and 52% of both the sample and the population in Nashville are female; yet the Illinois sample is only 38% female. All of the samples diverge from their populations on race, education, and economic status: Memphis, for instance, is 31% white, while the sample is 72%. The Nashville sample similarly over-represents whites, though to a lesser extent, with 61% of the population and 83% of the sample, while for Illinois, the sample is 92% white versus 76% in the population. Finally, both samples over-represent homeowners and those with a four-year college degree or more, though the Illinois sample appears less biased than the large city samples.

In short, while the samples recruited via Facebook do not perfectly match the underlying population, they can be characterized as diverse, and more representative

---

[10] Samuels and Zucco (2013) also note that Facebook surveys are amenable to quota sampling, and attempt to over-sample certain subpopulations in their second study. See Zhang et al. (2017) for an implementation of quota sampling in the United States.

[11] The age densities from the Census are computed as the number of persons in each age group, divided by the total number of persons 18 and older (22 and older in Nashville).

Memphis

Nashville

Illinois
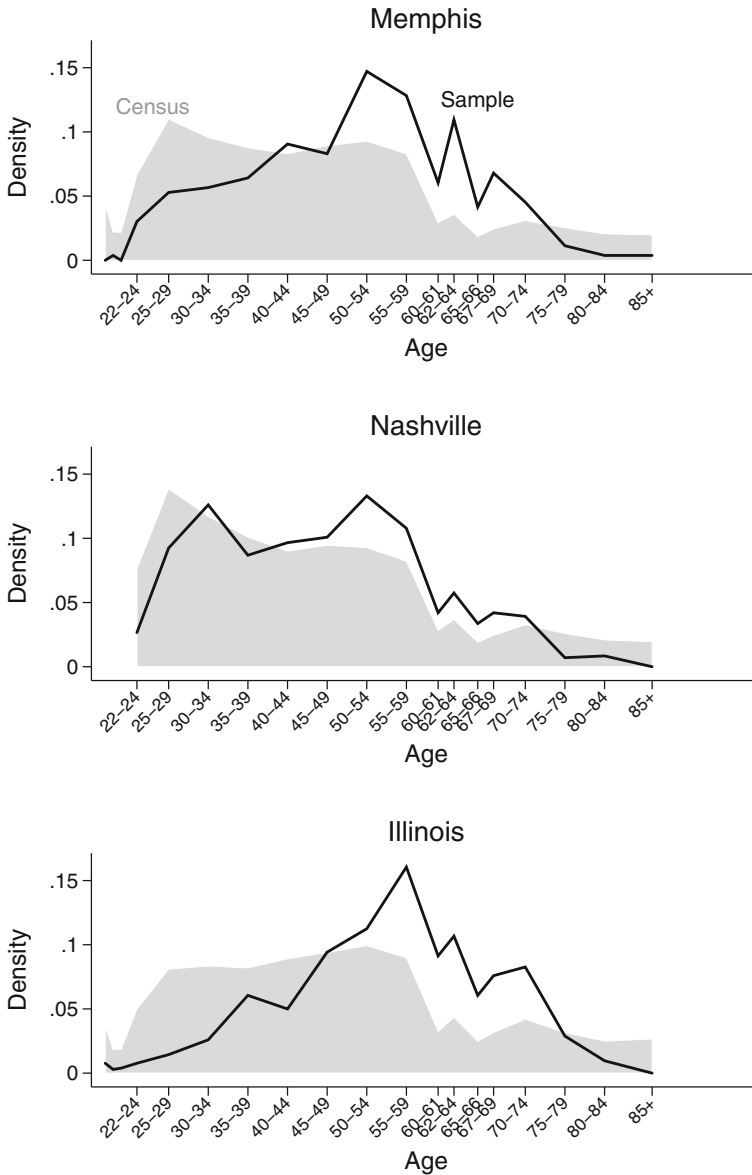
Fig. 1 Age distribution of Facebook survey samples compared to Census. *Notes Black lines* represent the distribution of age in each survey; *gray* areas represent the 2010 Census age distribution. The Census comparison is limited to persons aged 18 and older for Memphis and Illinois, and proportions are shown but not labeled for the categories 18–19, 20, and 21. The Census comparison is limited to persons aged 22 and older in Nashville (while the Facebook ad was targeted at those 23 and older, the closest Census age category is 22–24)

Table 1 Demographics of Facebook survey samples compared to population figures

	Memphis		Nashville		Illinois	
	Census	Sample	Census	Sample	Census	Sample
Median age	44	54	44	48	49	58
Percent female	53	56	52	52	51	38
Percent white	31	72	61	83	76	92
Percent college	22	66	34	81	29	48
Percent homeowner	50	76	54	73	72	85

Population statistics are from the 2010 Census and the 2010–2014 American Community Survey. The Census *Median age* is computed among those age 18 and older (in Nashville, 22 and older)

than any existing online method for recruiting voters from local contexts. Moreover, it is not the case that the samples over-represent the young. In fact, the samples are on average anywhere from four to ten years older than their populations. Finally, while the population of local *voters* is unknown, it likely diverges from the Census in ways that are similar to the Facebook samples. For instance, Oliver and Ha (2007, p. 393) conclude that the suburban elections are "dominated by a nonrepresentative group of 'stakeholders'" who are more educated, more likely to own a home, older, and more likely to be white (406). Additionally, Boudreau et al. (2015) report that their exit poll of San Francisco voters is similarly more educated (79% B.A. in sample vs. 51% in population) and more likely to own a home (42 vs. 37%).

The 2015 Memphis and Nashville Mayoral Races

Memphis and Nashville are both located in Tennessee, and both have populations around 650,000. Aside from these facts, however, the two cities are very different: Memphis is a majority-minority city, while Nashville is predominantly white; Memphis exists independently of its parent county and competes with the surrounding communities for resources, while Nashville and Davidson County are a single governmental entity; Memphis, located in the southwestern corner of Tennessee, borders the Mississippi River and Arkansas directly to the west and Mississippi directly to the south, while Nashville, the state capitol, is located near the center of the state.

Both cities also had mayoral elections in 2015, Memphis in October and Nashville in August, though the particular circumstances of each election varied. In Memphis, incumbent A.C. Wharton, an African American male, lost to challenger and City Councilor Jim Strickland, who is white. Wharton, first elected in 2009, was the first sitting incumbent to lose in 24 years. Strickland received 42 % of the vote to Wharton's 22%, while two other African American candidates—Memphis Police Association President Mike Williams and City Councilor Harold Collins—received a combined total of 34 percent. As Memphis does not hold runoff elections, Strickland was declared the winner and became mayor in January 2016.

The Nashville race, in contrast, was open, as incumbent mayor Karl Dean was term limited. City Councilor Megan Barry received a plurality of the vote, 24%, in the general election on August 6th, followed by businessman David Fox with 23%; the remaining 54% of the vote was split between five other candidates who received between 5 and 22%. Because no candidate received a majority in the general election, a runoff election, in which Barry won with 55% of the vote to Fox's 45%, was held on September 10th.

Neither city holds partisan elections, and both are majority-Democratic. Thus, similar to state primary elections (Hirano et al. 2015), these mayoral elections are well-suited to studying the impact of ideology on vote choice. In races with partisan labels, there is a risk that voters first choose their preferred candidate based on ideology, and then change their ideology to conform to this choice. This "following" behavior will bias estimates of the relationship between ideology and vote choice. In contrast, non-partisan mayoral elections allow for a cleaner test of the impact of ideology (Boudreau et al. 2015). These particular races are also notable in that they paired a white versus a black candidate in Memphis, and a male versus a female candidate in Nashville. Given the theoretical power of group identities in city elections, these contests present an interesting test of the power of ideology in shaping vote choice.

During the campaigns in each city, the leading candidates did work to differentiate themselves from one another in terms of policy. In Memphis, a key issue was how to respond to a perceived rise in violent crime. According to the Memphis *Commercial Appeal*, crime was "perhaps the most pointed issue differential of the campaign" (Veazey 2015a). Challenger Strickland presented himself as tougher on crime, supporting harsher penalties for violent criminals, including the detention of juvenile offenders. Despite Strickland's constant raising of the crime issue, incumbent Wharton proposed no new crime initiatives, arguing for a more "strategic" and level-headed approach (Veazey 2015b). Further differentiating the two candidates, Strickland—who like Wharton is, in fact, a Democrat—received the endorsement of the Shelby County Republican Party (Veazey 2015a).

No single policy issue defined the Nashville campaign, but local observers frequently cast the race in ideological terms, featuring a liberal, in Megan Barry, pitted against a conservative, in David Fox. As the *Tennessean* reported the day after the general election, the campaign "clearly shows the final race is between a liberal candidate and a conservative candidate" (Boucher 2015). During the campaign, "Barry consistently touted her views on social issues," pointing to her support of abortion rights and her officiation of the first ever same-sex marriage in the city's history. Fox, in turn, "proclaimed himself the only true fiscal conservative in the race, promising to cut spending and control the size of government" (Boucher 2015). So ideologically charged was the ensuing runoff campaign between Fox and Barry—featuring, among other barbs, accusations that Barry was an atheist and that Fox was a segregationist—that it received coverage in the *New York Times* (Fausset 2015).

Press reports therefore suggest that both the Memphis and Nashville races featured a more liberal candidate facing off against a more conservative candidate. My own survey data are consistent with this interpretation. On each of my city surveys, I asked respondents to place themselves, and the candidates, on a seven-
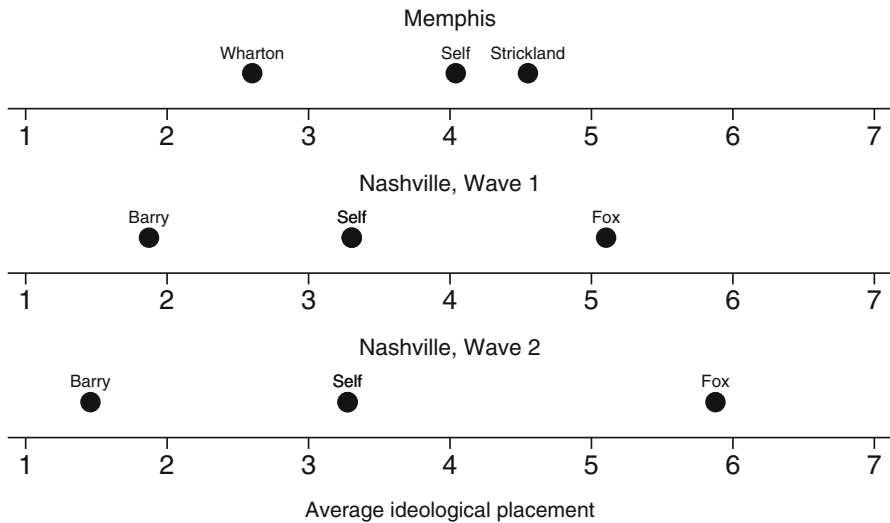
Fig. 2 Ideological placements of voters and major candidates in Memphis and Nashville elections. *Notes* This figure plots the average ideological placement of the two top vote-getters in each election, as well as voters' own self-placement. The figure only includes data from "expert" respondents, defined as meeting the following criteria: participated in both waves of the study in each city, and offered a ranking for all three items on each wave. For the Memphis data, the sample size is 98. For the Nashville data, the sample size is 203 for each wave

point "scale on which the political views that people might hold are arranged from strongly liberal to strongly conservative."[12] In Fig. 2, I show the average placements from highly informed respondents—those who participated in both waves of each study, and who placed both themselves and the two leading candidates on each wave.[13] The top panel shows rankings from Memphis, where respondents rated the incumbent Wharton at about 2.6, and the challenger and winner Strickland at about 4.6, and themselves at about 4.1. The second panel shows that even in the first wave, informed respondents perceived Barry as more liberal, at about 1.9, than Fox, at about 5.1. Finally, the third panel shows that two months of campaigning later, these expert respondents perceived an even wider ideological gap—about 1.5 for Barry versus 5.9 for Fox.

While no two cities can claim to be representative, Memphis and Nashville do not appear to be outliers on observable characteristics. Section 4 of the Online Appendix provides comparisons to other large (populations over 20,000) cities using data from Tausanovitch and Warshaw (2014). While certainly larger than most cities, Memphis and Nashville are typically near the central tendency (income, partisan elections, the initiative) or at different points along the distribution (for instance, Memphis is 60% black while Nashville is around 25% black; Nashville has term limits while Memphis does not). That said, certain features of these cities—such as their scheduling

---

[12] I show full question wordings in Sect. 10 of the Online Appendix.

[13] There are 98 respondents in the Memphis data in this figure, and 203 in the Nashville data. I use the responses from the full sample, including non-responses, later in the analysis.

elections in odd years, without any upper-level races to distract voters, and their active local media (Hopkins and Pettingill 2017)—may overstate the extent of ideological voting in all cities. However, later in the paper I show that similar voting behavior occurs in a sample of small-city voters in another state.

Ideology and Voting in Big-City Elections

If the candidates in each campaign may have staked out distinct ideological visions for the future of their cities, did voters then use this information in their voting decision? Figure 3 provides a preliminary answer. In this figure, I show the relationship between a vote for the winner in each city sample, and voters' own ideological placement on a seven-point scale. For these analyses, I use only respondents from the first wave of each survey. First, I calculate the average support for the winning candidate for respondents in each of the seven ideological bins (gray circles), as well as a 95% confidence interval (gray vertical lines). I then fit a linear regression line to these averages (dashed black line). Because the Memphis race featured a conservative winner and the Nashville race featured a liberal winner, the
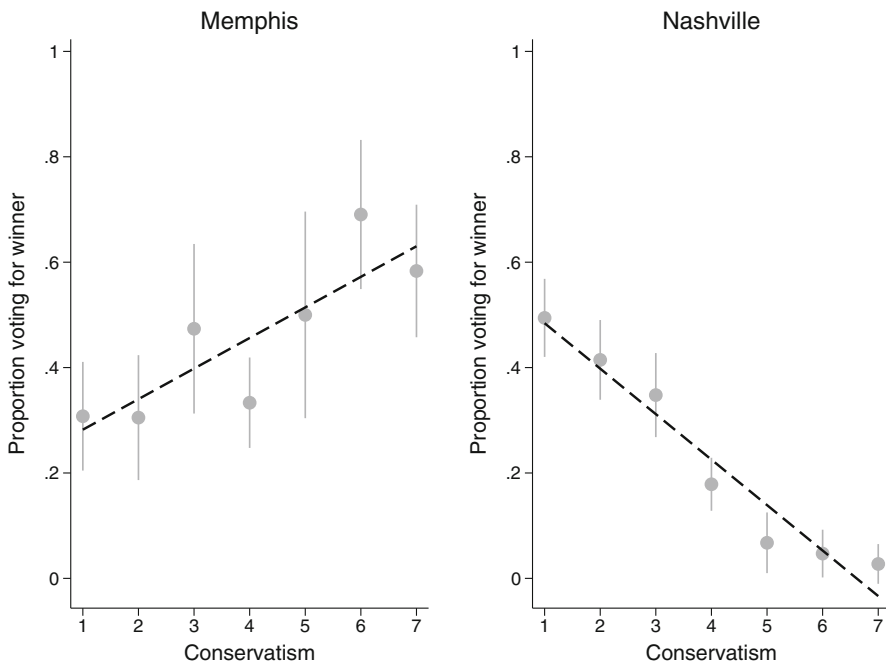
Fig. 3 Ideological voting in the 2015 Memphis and Nashville mayoral elections. *Notes* Points represent the proportion of respondents in each ideological category voting for the winner (Jim Strickland in Memphis and Megan Barry in Nashville), with *vertical lines* spanning 95% confidence intervals. *Dashed lines* represent linear regressions of proportions voting for the winner on conservatism. The sample size used to generate the *left panel* is 420, and the sample size in the *right panel* is 934

expectation is that we would see a positive relationship in the left panel, and a negative relationship in the right panel.

Figure 3 shows this is indeed what we see: in Memphis, the proportion supporting the more conservative Jim Strickland was about 30% for the most liberal respondents, while it was about 60 percent for the most conservative respondents. In Nashville, the proportion supporting liberal Megan Barry was about 50% for the most liberal respondents, and about 5% for the most conservative. Both of these patterns are evident whether we examine the data non-parametrically, using the binned averages, or when fitting a linear regression.

Of course, theories of mayoral voting suggest that group identity is of equal, if not greater importance, in city elections. And as Hajnal and Trounstine (2014) note, group identities are often themselves powerful predictors of ideologies, making it difficult to isolate the pure effect of ideology. As a first attempt to separate the effect of ideology from other factors, I estimate the following multivariate regression,

$$\text{Vote}_i = \alpha + \beta * \text{Conservatism}_i + \sum_{k=1}^{K} \pi * x_{k,i} + \epsilon_i$$

where Vote_i is an indicator for whether respondent i voted for the winner (a conservative in Memphis and a liberal in Nashville), Conservatism is seven-point ideology where higher values indicate greater conservatism, and the K x_k variables include indicators for race (white or non-white), gender (female or not), home-ownership, education (a four-year degree or higher, or not), age, and whether the respondent self-identifies as a Democrat. All predictor variables are re-scaled such that a zero represents the sample minimum and one represents the sample maximum, and all specifications use robust standard errors.

I show estimates and standard errors in Table 2, beginning with the results from Memphis. In column (1), I show results from a bivariate regression of vote choice on ideology. Similar to Fig. 3, the regression suggests that the most conservative voters are thirty-three percentage points more likely to vote for the conservative candidate, relative to the most liberal voters. The estimate is highly statistically significant, with a standard error of seven percentage points.

Because many respondents began but did not complete the survey, missingness is higher on the demographic control variables than on the vote choice and ideology variables. Before including controls, I show that the result in column (1), which includes 420 respondents, replicates among the 252 respondents who also provided answers to the demographics. As shown in column (2), the point estimate and level of significance are unchanged. In column (3), I show the relationship between the various control variables and vote choice, without adjusting for ideology. This regression shows that, as found in previous studies, race and class are powerful predictors of mayoral vote choice: white respondents are 17 points (standard error = 6 points) more likely to vote for Strickland, while college-educated respondents are 26 points (standard error = 6 points) more likely.

In column (4), I add conservatism back to the specification. Interestingly, none of the estimates from columns 1–3 change significantly (the exception is the Democrat dummy). The estimate on conservatism is now 0.30, and is only slightly more

Table 2 Ideological voting in the 2015 Memphis and Nashville mayoral elections

	Memphis				Nashville			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Conservatism	0.33***	0.33***		0.30**	−0.54***	−0.57***		−0.55***
	(0.07)	(0.09)		(0.12)	(0.04)	(0.04)		(0.07)
White			0.17**	0.18**			0.15***	0.15***
			(0.06)	(0.06)			(0.04)	(0.04)
Female			−0.15*	−0.14*			0.00	−0.01
			(0.06)	(0.06)			(0.03)	(0.03)
Homeowner			0.01	0.03			−0.05	−0.02
			(0.08)	(0.07)			(0.04)	(0.04)
BA or higher			0.26***	0.25***			0.11**	0.07
			(0.06)	(0.06)			(0.04)	(0.04)
Age			0.21	0.16			−0.09	−0.10
			(0.15)	(0.15)			(0.08)	(0.08)
Democrat			−0.13*	0.01			0.25***	0.01
			(0.06)	(0.08)			(0.03)	(0.05)
Constant	0.27***	0.24***	0.13	−0.07	0.49***	0.51***	0.00	0.39***
	(0.04)	(0.05)	(0.11)	(0.12)	(0.03)	(0.03)	(0.07)	(0.08)
Sample size	420	252	252	252	934	695	695	695

The dependent variable is voting for the winning candidate. Columns 1 and 5 include all respondents who reported their vote choice and their own ideology; columns 2–4 and 6–8 include only respondents who provided answers to all the included demographic controls. Cell entries are estimated coefficients from linear regressions, with robust standard errors in parentheses. All predictors are rescaled such that zero is the sample minimum and one is the sample maximum. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

uncertain, with a standard error of 0.12. Meanwhile, whites and college-educated respondents are still 18 and 25 percentage points more likely to vote for Strickland, and both these estimates retain their levels of statistical significance. While race and class still appear to be important factors in the election, it is striking that the estimate on conservatism is still the largest in substantive terms.

The next four columns present the results from Nashville, where ideology has an even more powerful effect. As shown in column (5), the most conservative voters were 54 points less likely (standard error = 0.04), relative to the most liberal voters, to vote for liberal Megan Barry. As in Memphis, in column (6) I show the result for the 934 voters who answered the ideology and vote choice questions also replicates among the 695 who answered the demographic questions: the point estimate is −0.57 with a standard error of 0.04. Column (7) shows that race also predicts voting, with white respondents 15 points (standard error = 4 points) more likely to vote for Barry as opposed to the alternatives.[14] However, despite the presence of a male-female contest between the two leading candidates, gender appeared to play

---

[14] Howard Gentry, an African American and Clerk of the City Criminal Court, received 11.6% of the vote in the general election. Dropping the 40 Gentry voters from the sample has no substantive impact on the estimates.

little role in voters' decisions: the estimates suggest no advantage for Barry among female voters without adjusting for ideology, and an insignificant one point disadvantage adjusting for ideology. Education has a positive and significant relationship with liberal vote choice when ideology is not taken into account (estimate = 0.11, standard error = 0.04), and a positive and insignificant relationship when adjusting for ideology (estimate = 0.07, standard error 0.04).

In two disparate cities and two very different election campaigns, we see a robust association between ideology and vote choice. This relationship is substantively larger than any other predictor, including group identities such as race and gender, which by all accounts should have been activated in these elections. As noted previously, however, cross-sectional analyses of voting behavior may mislead due to voters' tendency to adopt candidates' positions. In the next section, I account for this possibility by leveraging the panel nature of my Nashville survey.

Voter Learning and the Causal Effect of Ideology

The regressions in the previous section seek to isolate the impact of ideology from other factors, yet they are unable to account for potential simultaneity between vote choice and ideology. If vote choice is in fact a cause of ideology—which would be the case if voters first choose a candidate, and then adopt their issue positions—then the coefficients shown above will be biased upward, even if we adjust for observable factors. Given the tendency of voters to first choose candidates and then adopt their issues positions (Lenz 2012), this bias is, *a priori*, likely in this case. Although these elections lacked partisan labels, theoretical accounts of mayoral elections emphasize the importance of other group identities, such as race. If voters first choose candidates based on race, they may then adjust their self-reported ideology to conform to their preferred candidate.

Panel surveys are one way for researchers to account for the reverse causality explanation and further isolate the independent variable of interest (Lenz 2012; Hirano et al. 2015). By examining respondents both earlier and later in an election campaign, we can observe what comes first: ideology or vote choice. Further, the campaign itself can be used as an exogenous shock to voter ideology. If ideology truly has a causal impact on voter choice, then liberal respondents who learn about the ideology of the candidates over the course of the campaign should become more likely to vote for the liberal candidate, and conservative respondents should become less likely. In contrast, if ideology is merely a proxy for racial identity, there should be no change in vote choice induced by voter learning, as voter and candidate identities are both fixed over the course of the campaign.

While panels are impossible to conduct using exit poll surveys, and are prohibitively costly when using traditional telephone surveys, I conducted a second wave of my Nashville study by simply requesting respondents' e-mail addresses at the conclusion of the first wave.[15] I specifically chose the dates of the surveys to be

---

[15] Respondents were offered the chance to win an additional $100 Amazon.com gift card for their participation in the second wave.

as close together as possible, while maximizing the amount of campaign exposure between waves. Based on an analysis of the volume of election-related press coverage in 2007, the most recent open-seat race, I fielded the pre-election survey in June, and the post-election survey immediately after the election in August.[16]

To analyze learning, I construct the following variable:

$$\text{Knowledge}_i^w = \begin{cases} 1 & \text{if placed Barry to the left of Fox} \\ 0 & \text{if placed Fox equal or to the left of Barry,} \\ 0 & \text{if could not place either Fox or Barry} \end{cases}$$

where i indexes respondents and w survey waves. By this measure, 56% of the 331 respondents who participated in both waves were knowledgeable in June, compared to 89% in August, for a net learning difference of 33%.[17] To measure learning at the respondent level, I follow Hirano et al. (2015) and construct a learning indicator as follows:

$$\text{Learned}_i = \begin{cases} 1 & \text{if Knowledge}_i^1 = 0 \text{ and } \text{Knowledge}_i^2 = 1 \\ 0 & \text{if Knowledge}_i^1 = 0 \text{ and } \text{Knowledge}_i^2 = 0 \\ 0 & \text{if Knowledge}_i^1 = 1 \text{ and } \text{Knowledge}_i^2 = 0 \\ . & \text{if Knowledge}_i^1 = 1 \text{ and } \text{Knowledge}_i^2 = 1 \end{cases}$$

Thus respondents "learned" if they could not place Barry to the left of Fox in wave 1 but could in wave 2, and "did not learn" if they either failed the knowledge test in both waves, or answered correctly in June but failed in August. As in Hirano et al. (2015, p. 104), I exclude respondents who correctly placed candidates in both waves, as this group may still have increased their knowledge over the course of the campaign. By this measure, 114 respondents—76%—learned, while 36 did not.

Figure 4 presents the average support for the more liberal candidate by wave, learning status, and ideology. For this analysis, I code respondents as liberal if they place themselves below a four on a seven-point ideological scale, and conservative if they place themselves above four. The figure shows that among liberals who learned, support for Barry increased from about 47 points in wave 1 to 62 points in wave 2, while support among liberal non-learners actually decreased slightly, from 53 to 47. Among conservatives, learners decreased their support from 9 to 6 points, while non-learners actually increased their support, from 0 to 15 points.

To formally test whether learning occurs, I re-estimate the regression presented earlier in a sample that pools respondents from both waves, interacting ideology with time. That is, I estimate

---

[16] I also fielded a third wave in September, after the runoff election. Because there was little learning between the second and third wave, I omit this sample from the analysis. Results are substantively similar if I compare the first and third wave instead of the first and second wave. I provide details on the third wave, including sample size, in Sect. 1 of the Online Appendix. Sect. 6 of the Online Appendix presents an analysis of election-related press coverage in the 2015 election.

[17] Among the full sample of wave 1 participants who answered the relevant items, wave 1 knowledge was 50%.
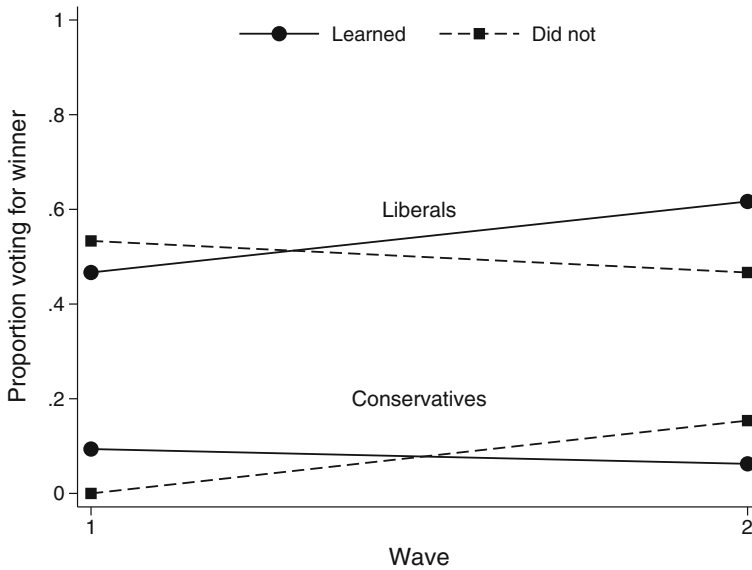
Fig. 4 Learning and opinion change in the Nashville election. *Notes* This figure compares the proportion voting for the winning candidate (Megan Barry) by wave, ideology, and learning status. *Liberals* are those respondents who placed themselves at a 1, 2, or 3 on a seven-point ideological scale on Wave 1; *conservatives* are those who placed themselves at 5, 6, or 7. Respondents are coded as knowledgeable if they placed Barry to the *left* of David Fox, and as not knowledgeable if they either placed Barry to the *right* of Fox, or answered "don't know" when asked to place either Fox or Barry. *Learners* are respondents who failed the knowledge test in Wave 1, but passed in Wave 2. Those who *did not* learn failed the knowledge test in both waves, or passed in Wave 1 but failed in Wave 2. There are 120 respondents in total in each wave

$$\text{Vote}_i^w = \alpha + \beta_1 * \text{Conservatism}_i^w + \beta_2 * \text{Post}_i^w + \beta_3 * \text{Conservatism}_i^w \times \text{Post}_i^w + \epsilon_i^w$$

where w indexes survey waves, Post is an indicator for the August wave, Conservatism is on a seven-point scale (rescaled, such that zero is most liberal and one is most conservative) and I cluster standard errors at the respondent level. If respondents adjust their voting decisions in accord with the candidates' ideology, it should be the case that either the Post indicator is positive, or the interaction between Conservatism and Post is negative, either of which would reflect liberals increasing support of Barry over the course of the campaign.

Column (1) of Table 3 shows this is the case. The coefficient on Conservatism is −0.54, with a standard error of 0.04, as it is exactly the same comparison as performed in Table 2, column (5): in the first wave, the most conservative voters were 54 points less likely to support Barry relative to the most liberal voters. The coefficient on Post indicates that the most liberal voters increased their support of Barry by 19 points (standard error of 4 points) over the course of the campaign, while incorporating the interactive term suggests conservatives decreased their already low support by one point (calculating the change in support for conservatives requires adding the coefficients for Post and the interaction). In

Table 3 Learning and opinion change in the Nashville election

	(1)	(2)	(3)	(4)	(5)
Conservatism	−0.54***	−0.57***			
	(0.04)	(0.06)			
Post	0.19***	0.13***			
	(0.04)	(0.04)			
Conservatism X post	−0.20**	−0.16*			
	(0.07)	(0.07)			
Prior conservatism			−0.38***	−0.37**	0.18
			(0.07)	(0.12)	(0.24)
Learned					0.38*
					(0.16)
Prior conservatism X learned					−0.70**
					(0.27)
Prior vote			0.59***	0.48***	0.65**
			(0.05)	(0.09)	(0.20)
Prior vote X learned					−0.19
					(0.22)
Constant	0.49***	0.54***	0.36***	0.38***	0.09
	(0.03)	(0.04)	(0.05)	(0.08)	(0.14)
Sample size	1262	648	324	144	144

The dependent variable is voting for the winning candidate. Cell entries are estimated coefficients from linear regressions, with robust standard errors in parentheses. All predictors are rescaled such that zero is the sample minimum and one is the sample maximum. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

sum, liberals increased their support for Barry by 19 points, while conservatives barely shifted.

To rule out the possibility that vote intention determines ideology, and not the other way around, I conduct an additional regression that adjusts for both wave 1 conservatism and wave 1 vote intention. That is, I estimate

$$\text{Vote}_i^w = \alpha + \beta_1 * \text{Conservatism}_i^{w-1} + \beta_2 * \text{Vote}_i^{w-1} + \epsilon_i^w$$

The adjustment for wave 1 conservatism ensures that my measure of ideology cannot be caused by changes in vote intent, and the inclusion of prior vote intent adjusts for regression to the mean (Finkel 1995). Additionally, the lagged outcome variable means that the effects can be interpreted in terms of changes in vote intent between June and August. Because this regression can only include respondents who answered the vote intent and ideology questions on both waves, I first present the pooled regression estimate for this subsample in column (2). The estimates are similar to column (1), and imply that liberals increased their support for Barry by 13 points over the campaign, and that conservatives decreased their support by 3 points.

Column (3) shows the results controlling for lagged vote and ideology. The coefficient on prior conservatism is −0.38, with a standard error of 0.07. This

estimate implies that, relative to the most liberal voters, the most conservative voters in June decreased their support for Barry by 38 points by August. This estimate is over twice as large as those reported in the first two columns, but is also more credible, as it is unlikely that the June measure of ideology is causally impacted by the change in vote intent between June and August.

Finally, I examine changes in voting among learners specifically, estimating a regression of the form,

$$
\begin{aligned}
\text{Vote}_i^w = \alpha + \beta_1 * \text{Conservatism}_i^{w-1} + \beta_2 * \text{Learned}_i + \beta_3 * \text{Conservatism}_i^{w-1} \\
\times \text{Learned}_i + \beta_4 * \text{Vote}_i^{w-1} + \beta_5 * \text{Vote}_i^{w-1} \times \text{Learned}_i + \epsilon_i^w
\end{aligned}
$$

In this regression, the coefficient on Conservatism represents the change in Barry support among the most conservative voters who did not learn (relative to the most liberal who did not learn), the Learned coefficient represents the change in Barry support among the most liberal voters who learned (relative to the most liberal who did not learn), and the interaction between Conservatism and Learned represents the change in support among the most conservative learners, relative to the most conservative non-learners. Because this regression further restricts the sample (see the definitions of Learning and Knowledge above), I first estimate the preceding regression among this group in column (4), which shows the estimates are basically unchanged.

Column (5) presents the results from the learning regression. It shows that the most conservative voters who did not learn did not alter their vote for Barry over the course of the campaign; the estimate is actually positive, at 18 points, but statistically insignificant with a standard error of 24 points. Among the most liberal learners, Barry's support increased by 38 points relative to the most liberal non-learners, similar in magnitude to the change in vote intent seen in columns (3) and (4). Finally, the most conservative learners decreased their support for Barry by 33 points (standard error of 16 points) (calculated by summing the coefficient on learned and the interaction), relative to the most conservative non-learners.

In sum, the effect of ideology is robust to panel methods that control for reverse causality. Voters in the Nashville election not only learned about the candidates' relative ideological positions between June and August, but they also acted on this knowledge by shifting their vote according to their own ideological positions in June. In contrast, voters who did not learn actually moved away from the ideologically "correct" candidate, a normatively troubling result.

In addition to ruling out reverse causality due to voter following, these results have implications for ideology's impact relative to group identity. While I cannot rule out the possibility that respondents' wave 1 ideology is itself influenced by race, that respondents who learned shifted their vote suggests ideology does have a causal effect independent of other group identities, at least for the subset of voters who did not know the candidates' ideological positions in both waves. For these voters, the election campaign constituted an exogenous shock to their knowledge of the candidates' ideologies, while the identities of the candidates and voters were held constant.

If ideology does have a causal impact on vote choice, does it do so because voters prefer the candidate who is more liberal or conservative, or does the exact ideological placement of the candidates matter (Jessee 2009; Joesten and Stone 2014; Shor and Rogowski 2016)? In Memphis and Nashville, I lack a precise measure of how liberal or conservative each candidate is, which means I can not calculate an objective measure of ideological distance. However, evidence presented in Sect. 7 of the Online Appendix suggests a role for proximity. There I show that voters' perceived ideological distance between themselves and the candidate is strongly correlated with vote choice. Although this subjective measure of distance may be biased by voters' "projecting" their own ideology onto their preferred candidates (Krosnick 1990), the results also hold when using the change in ideological distance that occurs across waves. Finally, the results in the next section also suggest a role for proximity.

Ideology and Voting in Small-City Elections

Two original surveys of voters in large Southern cities show evidence of ideological voting in mayoral elections. However, questions remain regarding both internal and external validity. Regarding the former, the complexity of real election campaigns means we can never be completely sure that the impact of ideology has been disentangled from other factors. Regarding the latter, it is not clear whether the patterns observed in these two cities would extend to cities that are smaller or are located in a different region, or to samples that are not limited to those who express an interest in local politics.

To address both these issues, I conducted a third survey of voters residing in small cities and villages in Illinois. Rather than intervene in a real election campaign, I presented these respondents with a conjoint experiment (Hainmueller et al. 2014; Carnes and Lupu 2016; Kirkland and Coppock 2017) designed to elicit the relative importance of different attributes in their selection of local leaders.[18] As explained by Hainmueller et al. (2014), a conjoint experiment allows researchers to estimate causal effects in multi-dimensional choice settings, such as the selection of political candidates. In my case, I am interested in estimating the impact of candidate ideology, independent of other candidate attributes. In addition to ideology, past research suggests local voters may select candidates on the basis of race, as well as valence issues such as geographic proximity (i.e. "friends and neighbors" voting; Key 1949; Meredith 2013) and experience (Trounstine 2013; Schaffner et al. 2001).

In the experiment, subjects were asked to suppose they were about to cast a vote for the next mayor or president of their local government.[19] They were then shown

---

[18] While a direct replication of the prior studies would be valuable, it would be logistically more difficult as small IL cities vary in election timing, and sampling from a particular small city would likely yield too few respondents.

[19] In Illinois, small municipalities are known as cities or villages. Generally, the chief executive of a city is referred to as a mayor, while in a village she is referred to as the president. I show a screenshot of the conjoint experiment in Sect. 8 of the Online Appendix.

two pairs of candidates, and asked which of these two candidates they would prefer to see as their next local mayor or president.[20] Each candidate in each pair was randomly assigned a set of attributes, including race (Asian American, Black, Hispanic, or White), policy position (supports cutting property taxes, supports maintaining property taxes, or supports increasing taxes), whether the candidate lived in the respondent's neighborhood (no or yes), and experience in local government (2, 5, 10, or 20 years). Following Hainmueller et al. (2014), I analyze the results of the roughly (1000 respondents * 2 hypothetical contests * 2 candidates per contest =) 4000 candidate evaluations, clustering standard errors at the respondent level. For the analysis, I transform the candidate issue position to a measure of issue distance, taking the difference between the respondent's own position on raising property taxes and those of the candidate.[21]

Note the experiment varies candidates' policy positions, rather than seven-point ideology, as I believed actual positions would be less abstract to voters than a seven-point scale. I also wished to avoid lumping in "symbolic" conservatives—those who say they are conservative but fail to endorse conservative policy positions—with "operational conservatives"—those who both self-identify as conservative and endorse conservative positions (Ellis and Stimson 2012). Of course, one drawback to this choice is that the estimates from the conjoint study are not perfectly comparable to those from the field studies, as the key independent variables are on different scales. Nonetheless, evidence of a positive and substantively large impact of issue agreement in this more controlled context should boost the credibility of the results shown earlier.[22]

Figure 5 shows the marginal effects of each candidate attribute on vote choice.[23] The first four estimates show there is no detectable impact of candidates' race in this

---

[20] To be precise, each respondent was actually shown five pairs of candidates. Due to a programming error, only the attributes of the first two pairs were recorded in the survey data, and so only the responses from the first two pairs can be analyzed.

[21] Formally, the measure of policy distance is

$$\text{Issue distance}_{ic} = -|x_i - x_c|,$$

where i indexes respondents, c indexes candidates, and x_i and x_c are the respondent's and candidate's views on raising property taxes, where $x_i \in \{0 = \text{supportscuts}, .5 = \text{supportsmaintainingcurrentlevels}, 1 = \text{supportsincreases}\}$.

[22] Section 9 of the Online Appendix replicates the analysis presented here using a measure of distance based on voter's seven point ideology (making the assumption that those less than four on the scale are the same as voters who wish to raise taxes, those that are equal to four wish to maintain taxes, and those greater than four wish to cut taxes). The results are substantively very similar.

[23] As in Hainmueller et al. (2014), this figure is constructed by estimating four regressions (one for each set of attributes) where vote choice is regressed on indicator variables for each attribute value. For instance, the first four estimates are from a regression of the form:

$$\text{Vote}_{ic} = \alpha + \beta_1 \text{Black}_{ic} + \beta_2 \text{Hispanic}_{ic} + \beta_3 \text{White}_{ic} + \epsilon_{ic}$$

where i indexes respondents and c indexes candidate pairs. The three β coefficient estimates and their 95% confidence intervals are then plotted in Fig. 5. As the attributes are independent of one another, there is no need to include attributes in a single regression; doing so, however, yields substantively similar results.
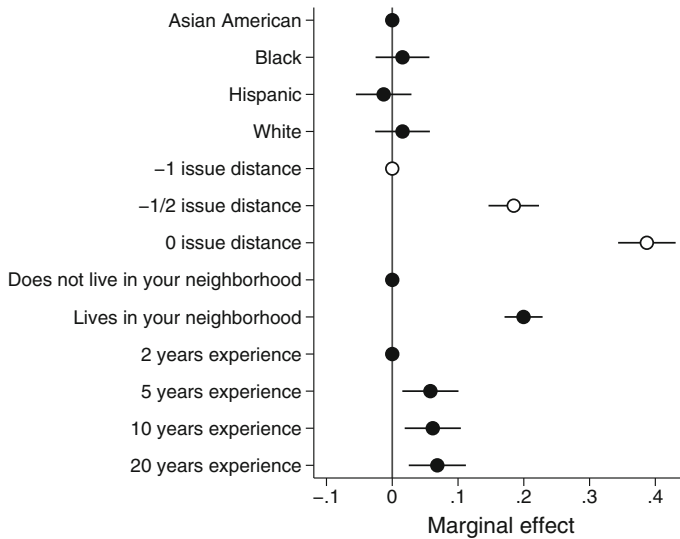
Fig. 5 Effect of mayor candidate attributes on vote choice among small-city Illinois voters. *Notes* This plot shows estimates of the effects of randomly assigned local candidate attributes on the probability of choosing that candidate. Estimates are based on regressions with standard errors clustered by respondent; *horizontal lines* represent 95% confidence intervals. Points without *horizontal lines* denote the reference value for the attribute category. There are 4350 candidate observations and 1109 unique respondents

sample. The next three estimates, however, show a large and significant impact of ideology. Respondents prefer candidates who share their policy position (and thus have zero issue distance) about 39 percentage points (standard error of 2 percentage points) more than they prefer candidates with the opposite position; they also prefer candidates who are 1/2 points away from their own position by about 18 percentage points (standard error of 2 percentage points) more than they prefer candidates who are 1 point away. Only the estimate for geographic proximity comes close in terms of magnitude: compared to candidates who do not live in their neighborhood, voters prefer candidates who do by 20 percentage points (standard error of 1 percentage point). Voters also prefer candidates with more than two years experience relative to those with just two years, though more experience is not necessarily better. The point estimates for the experience coefficients are all between 6 and 7 points, and while statistically different from the reference category, they are not different from one another.

Consistent with the big-city studies, the survey of small-city voters reveals ideology to be a powerful, independent predictor of candidate choice. In a more controlled setting, and in a sample of voters from a different region and from much smaller communities, we see a similar impact of policy positions on voting. Indeed, as in the studies reported previously, we see ideology—here measured as the distance between subjects' and candidates' policy positions—has the largest impact on vote choice compared to any other factor.

Conclusion

A major justification for elections is to provide voters a voice in the policy process. While national elections have long been seen as more or less fulfilling this purpose, city elections have been seen as divorced from policy. This separation is both normatively concerning, given the impact of city policies on voter welfare, as well as theoretically puzzling, given recent evidence that cities do, in fact, provide voters with the policies they say they want.

In this paper, I show that the presumed substantive gap between local and national politics is smaller than once thought. In two original studies of mayoral elections, I show that candidates stake out independent ideological positions, that voters learn these positions, and that voters use what they learn about ideology to choose between candidates. The effect of ideology rivals that of other observable factors, and is robust to adjusting for race, class, and potential reverse causality. The effect also replicates in a sample of small-city voters who are asked to evaluate hypothetical mayoral candidates. Substantively, these findings show that account-ability for policy is possible in U.S. cities, a result that is both normatively positive and theoretically helpful in interpreting recent evidence of city responsiveness.

I uncover the impact of ideology using a novel methodological tool that can now be used by other scholars of local political behavior. Using geographically targeted advertisements on the Facebook social media platform, I recruited three diverse samples of local voters from two separate cities. Both the technical and financial barriers to recruitment via this method are far lower than the alternatives, such as in-person exit polls or telephone surveys, and samples recruited over Facebook can easily be recontacted using respondent-provided e-mail addresses.

Of course, no method is perfect, and questions remain about the external validity of these samples. In particular, two issues may lead these results to overstate the extent of ideological voting in city elections. First, those who volunteer for surveys about local politics may be more ideological than the broader population. If the broader population is taken to be all residents in a city, these volunteers are clearly not representative. However, if the broader population is the local electorate, the extent to which this is an issue is less clear. Several existing studies have noted that local electorates over-represent whites, homeowners, and the elderly, and so these volunteers may simply be an accurate reflection of the biased local electorate. Second, most of the results presented here come from off-cycle elections. In local races held concurrently with national elections, voters may be distracted from learning local candidates' positions, hindering ideological voting.

With these caveats in mind, the results presented here suggest that city elections are, at least in one sense, similar to races that occur at higher levels. Specifically, voters in state primary elections also appear to vote ideologically (Hirano et al. 2015), as do independents in national elections (Jessee 2009; Shor and Rogowski 2016). One commonality between these cases and (most) city elections is the lack of a party label, which might otherwise overpower ideological considerations if voters follow party cues instead of learning candidates' positions. Rather than depressing accountability, then, the widespread adoption of the nonpartisan ballot in cities may

have facilitated greater policy voting (Schaffner et al. 2001)—though at the cost of a smaller and less representative city electorate.

References

Abrajano, M., Nagler, J., & Michael Alvarez, R. (2005). A natural experiment of race-based and issue voting: The 2001 city of Los Angeles elections. *American Politics Quarterly*, *58*(2), 203–218.

Adweek. (2016). Number of Facebook users in the United States as of January 2015, by age group (in millions). In *Statista—The Statistics Portal*. Retrieved May 12, 2016 from http://www.statista.com/statistics/398136/us-facebook-user-age-groups/.

Alvarez, R. M., & Nagler, J. (1995). Economics, issues and the Perot candidacy: Voter choice in the 1992 presidential election. *American Journal of Political Science*, *39*(3), 714–744.

Arceneaux, K. (2006). The federal face of voting: Are elected officials held accountable for the functions relevant to their office? *Political Psychology*, *27*(5), 731–754.

Arnold, R. D., & Carnes, N. (2012). Holding mayors accountable: New York's executives from Koch to Bloomberg. *American Journal of Political Science*, *56*(4), 949–963.

Barreto, M. (2007). Si Se Puede! Latino candidates and the mobilization of Latino voters. *American Political Science Review*, *101*(3), 425–441.

Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, *20*(3), 351–368.

Berry, C. R., & Howell, W. G. (2007). Accountability and local elections: Rethinking retrospective voting. *Journal of Politics*, *69*(3), 844–858.

Boucher, D. (2015). Nashville mayoral election: Five takeaways. In *The Tennessean* August 7. Retrieved May 12, 2016 from http://www.tennessean.com/story/news/politics/2015/08/07/5-takeaways–nashvilles-mayoral-election/31102247/.

Boudreau, C., Elmendorf, C. S., & MacKenzie, S. A. (2015). Lost in space? Information shortcuts, spatial voting, and local government representation. *Political Research Quarterly*, *68*(4), 843–855.

Broockman, D. E., & Green, D. P. (2014). Do online advertisements increase political candidates' name recognition or favorability? Evidence from randomized field experiments. *Political Behavior*, *36*(2), 263–289.

Burnett, C. M., & Kogan, V. (2017). The politics of potholes: Service quality and retrospective voting in local elections. *Journal of Politics*, *79*(1), 302–314.

Callegaro, M., & DiSogra, C. (2008). Computing response metrics for online panels. *Public Opinion Quarterly*, *72*(5), 1008–1032.

Campbell, A., Converse, P. E., Miller, W., & Stokes, D. E. (1960). *The American Voter*. Ann Arbor: University of Michigan Press.

Carnes, N., & Lupu, N. (2016). Do voters dislike working-class candidates? Voter biases and the descriptive underrepresentation of the working class. *American Political Science Review*, *110*(4), 832–844.

de Benedictis-Kessner, J., & Warshaw, C. (2016). Mayoral partisanship and the size of municipal government. *Journal of Politics* (in press).

Downs, A. (1957). *An economic theory of democracy*. New York: Harper Collins.

Einstein, K. L., & Kogan, V. (2016). Pushing the city limits: Policy responsiveness in municipal government. *Urban Affairs Review*, *52*(1), 3–32.

Ellis, C., & Stimson, J. A. (2012). *Ideology in America*. New York: Cambridge University Press.

Fausset, R. (2015). In Mayoral race, Nashville politics forgets its manners. The New York Times. Retrieved May 12, 2016 from http://www.nytimes.com/2015/09/10/us/mudslinging-in-race-for-nashville-mayor-shakes-citys-political-scene.html?_r=0.

Finkel, S. E. (1995). *Causal analysis with panel data*. Thousand Oaks: Sage.

Gerber, E. R., & Hopkins, D. J. (2011). When mayors matter: Estimating the impact of mayoral partisanship on city policy. *American Journal of Political Science*, *55*(2), 326–339.

Hainmueller, J., & Hiscox, M. J. (2010). Attitudes toward highly skilled and low-skilled immigration: Evidence from a survey experiment. *American Political Science Review*, *104*(1), 61–84.

Hainmueller, J., Hopkins, D. J., & Yamamoto, T. (2014). Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis*, *22*(1), 1–30.

Hajnal, Z. (2006). *Changing white attitudes toward black political leadership*. Cambridge: Cambridge University Press.

Hajnal, Z., & Trounstine, J. (2014). What underlies urban politics? Race, class, ideology, partisanship, and the urban vote. *Urban Affairs Review*, *50*(1), 63–99.

Hetherington, M. J. (2001). Resurgent mass partisanship: The role of elite polarization. *American Political Science Review*, *95*(3), 619–631.

Hirano, S., Lenz, G. S., Pinkovskiy, M., & Snyder, J. M. (2015). Voter learning in state primary elections. *American Journal of Political Science*, *59*(1), 91–108.

Hopkins, D. J., & Pettingill, L. M. (2017). Retrospective voting in big-city US mayoral elections. *Political Science Research and Methods* (forthcoming).

Jacoby, W. G. (2009). Ideology and vote choice in the 2004 election. *Electoral Studies*, *28*(4), 584–594.

Jessee, S. A. (2009). Spatial voting in the 2004 presidential election. *American Political Science Review*, *103*(1), 59–81.

Joesten, D. A., & Stone, W. J. (2014). Reassessing proximity voting: Expertise, party, and choice in congressional elections. *Journal of Politics*, *76*(3), 740–753.

Kaufmann, K. M. (1998). Racial conflict and political choice a study of mayoral voting behavior in Los Angeles and New York. *Urban Affairs Review*, *33*(5), 655–685.

Kaufmann, K. (2004). *The urban voter: Group conflict and mayoral voting in American cities*. Ann Arbor: University of Michigan Press.

Key, V. O. (1949). *Southern politics in state and nation*. New York: A. A. Knopf.

Kirkland, P. A., & Coppock, A. (2017). Candidate choice without party labels: New insights from conjoint survey experiments. *Political Behavior* (forthcoming).

Knight, K. (1985). Ideology in the 1980 election: Ideological sophistication does matter. *Journal of Politics*, *47*(3), 828–853.

Krosnick, J. A. (1990). Americans' perceptions of presidential candidates: A test of the projection hypothesis. *Journal of Social Issues*, *46*(2), 159–182.

Lenz, G. S. (2012). *Follow the leader?*. Chicago: University of Chicago Press.

Marschall, M., Shah, P., & Ruhil, A. (2011). The study of local elections. *Political Science & Politics*, *44*(1), 97–100.

Meredith, M. (2013). Exploiting friends-and-neighbors to estimate coattail effects. *American Political Science Review*, *107*(4), 742–765.

Oliver, J. E. (2012). *Local elections and the politics of small-scale democracy*. Princeton: Princeton University Press.

Oliver, J. E., & Shang, E. H. (2007). Vote choice in suburban elections. *American Political Science Review*, *101*(3), 393–408.

Peterson, P. E. (1981). *City limits*. Chicago: University of Chicago Press.

Peytchev, A. (2009). Survey breakoff. *Public Opinion Quarterly*, *73*(1), 74–97.

Ryan, T. J. (2012). What makes us click? Demonstrating incentives for angry discourse with digital-age field experiments. *Journal of Politics*, *74*(4), 1138–1152.

Samuels, D. J., & Zucco, C. (2013). Using Facebook as a subject recruitment tool for survey-experimental research. Working paper, Department of Political Science, University of Minnesota. Retrieved May 12, 2016 from http://ssrn.com/abstract=2101458.

Samuels, D., & Zucco, C. (2014). The power of partisanship in Brazil: Evidence from survey experiments. *American Journal of Political Science*, *58*(1), 212–225.

Schaffner, B. F., Streb, M., & Wright, G. (2001). Teams without uniforms: The nonpartisan ballot in state and local elections. *Political Research Quarterly*, *54*(1), 7–30.

Shor, B., & Rogowski, J. C. (2016). Ideology and the US congressional vote. *Political Science Research and Methods* (forthcoming).

Stein, R. M., Ulbig, S. G., & Post, S. S. (2005). Voting for minority candidates in multiracial/multiethnic communities. *Urban Affairs Review*, *41*(2), 157–181.

Tausanovitch, C., & Warshaw, C. (2014). Representation in municipal government. *American Political Science Review*, *108*(3), 605–641.

Trounstine, J. (2013). Turnout and incumbency in local elections. *Urban Affairs Review*, *49*(2), 167–189.

Vavreck, L., & Rivers, D. (2008). The 2006 cooperative congressional election study. *Journal of Elections, Public Opinion and Parties*, *18*(4), 355–366.

Veazey, K. (2015a). A closer look: Strickland's crime plan heavy on juveniles, which concerns some. The Commercial Appeal, September 20. Retrieved May 12, 2016 from http://www.commercialappeal.com/news/government/city/a-closer-look-stricklands-crime-plan-heavy-on-juveniles-which-concern-some-1f3d8262-f920-3927-e053-0-328432381.html.

Veazey, K. (2015b). Wharton, Strickland release detailed crime plans. The Commercial Appeal, September 9. Retrieved May 12, 2016 from http://www.commercialappeal.com/news/government/city/wharton-strickland-release-detailed-crime-plans-ep-1266377385-327921431.html.

Zhang, B., Mildenberger, M., Howe, P. D., Marlon, J., Rosenthal, S., & Leiserowitz, A. (2017). Quota sampling using Facebook advertisements can generate nationally representative opinion estimates. Working paper, Department of Political Science, Yale University. Retrieved July 5, 2017 from http://qssi.psu.edu/new-faces-papers-2017/zhang-quota-sampling-using-facebook-advertisements.