

The Demand for Elections under Autocracy: Regime Approval and the Cancellation of Local Elections in Russia

Quintin Beazer*
Noah Buckley†
Ora John Reuter‡

October 27, 2025§

Abstract

Most contemporary autocracies hold elections. Does the public value these elections and, if so, do they value them enough to punish incumbents that subvert elections? We examine this question in the case of contemporary Russia by examining whether individuals withdraw support from regime leaders when local elections are cancelled. Over the past 20 years, most Russian cities have replaced their directly elected mayors with appointed chief executives. This paper uses the largest dataset on public opinion ever assembled on Russia—containing over 1.4 million polling responses drawn from two decades of polling by Russia’s top polling agencies—to analyze how the cancellation of elections in Russia’s large cities has affected public attitudes toward the authorities. Using a difference-in-differences design, we find that election cancellation reduces support for President Vladimir Putin. This effect is stronger in settings with histories of robust electoral competition. This suggests that the public is more likely to punish incumbents for canceling elections when those elections are perceived as meaningful.

KEYWORDS: Russia, Voting, Local Elections, Authoritarianism, Autocracy

*Florida State University

†Trinity College Dublin

‡University of Wisconsin–Milwaukee

§Support from the Basic Research Program of the National Research University Higher School of Economics is gratefully acknowledged.

1 Introduction

Multi-party elections are a defining feature of electoral autocracy. Does the public value such elections? If so, do they value them enough to punish leaders when elections are degraded or cancelled? Existing survey data from multiple autocracies shows that members of the public believe that leaders should be selected using elections and that electoral manipulation is inappropriate (Hale 2011; M’Cormack-Hale and Dome 2021). In addition, evidence from survey experiments using hypothetical scenarios indicates that people are willing to withdraw support from candidates who engage in fraud (Reuter and Szakonyi 2021; Svolik 2012). In this paper, we extend this line of inquiry by examining how citizens in the real world react to the wholesale cancellation of elections. Specifically, we examine whether the act of cancelling elections undermines support for authoritarian incumbents. In addition, we theorize the conditions that make individuals more likely to withdraw their support from incumbents when elections are cancelled. We argue that voters are particularly likely to punish incumbents for election cancellation in settings where elections have proven to be consequential in the past. Voters in these settings are more likely to value elections. We also argue that those who view elections as intrinsically good or instrumentally beneficial will be more likely to withdraw support from incumbents when elections are cancelled.

We address these questions using two new data sources from Russia. Specifically, we pair new data on the cancellation of mayoral elections in 212 large Russian cities between 2003 and 2018 with a collection of 394 public opinion surveys totaling over 1.4 million respondents in Russia. In 2003, almost 90% of cities directly elected their mayors, but over the past 18 years elections have gradually been eliminated in most cities, such that fewer than 10% of large cities still hold direct elections. We investigate how the cancellation of these elections affected regime support in cities where these elections were cancelled. To do so, we use pooled survey data from virtually all omnibus surveys carried out by the Levada Center and VTsIOM between 2003 and 2018, as well as survey data from the well-known

FOM GeoRating surveys.¹ This constitutes the largest assemblage of Russian survey data ever analyzed. We then match this survey database with the city-level cancellation data to examine how support for regime leaders changed in cities the cancellation of direct mayoral elections. In order to identify the effect of election cancellation, we employ a staggered difference-in-differences design that leverages the fact that elections were cancelled on an idiosyncratic schedule uncompromised by trends in public opinion. City and calendar-month fixed effects then allow us to isolate changes to regime support in cities following these unanticipated election cancellation events compared to trends in public opinion within cities where elections had not yet been cancelled.

Our findings indicate that Vladimir Putin experiences a 2.2% decline in support in cities that cancel elections. Support for regional governors—the main initiators of election cancellation—experiences a 5.4% decline in support when elections are cancelled. In line with our theoretical predictions, we show that the effect of election cancellation depends heavily on political context. In regions with a history of more vigorous electoral competition, election cancellation is associated with a larger drop in support for Putin. This suggests that the public punishes incumbents more when they have experience with meaningful elections.

We also demonstrate that our findings mask considerable heterogeneity across individuals. We find strong evidence that the effect of election cancellation on support for Putin is moderated by the age of the respondent: younger respondents are considerably more likely to withdraw support when elections are cancelled, while older individuals actually increase their support slightly. Contrary to expectations, however, we find no evidence that respondents' education levels moderates the effect of election cancellation on support for Putin.

Our findings shed light on the popular bases of regime stability under autocracy. Our results confirm descriptive survey research demonstrating that most people in electoral authoritarian regimes find value in elections. In line with findings from survey experiments, our findings suggest that the public withdraws support from authoritarian leaders who sub-

¹These three companies represent the largest and most respected polling companies in Russia over the last two decades.

vert or, in our case, eliminate elections. However, in contrast to this literature, we provide real-world, observational evidence for this claim. The recent neo-institutional scholarship on authoritarianism has outlined a number of indirect benefits that autocrats reap from the holding of elections—for example, collecting information, coopting and dividing the opposition—but the present work suggests that autocrats may also hold elections because there is popular demand for them.

These conclusions also help shed new light on the popularity of Vladimir Putin in Russia. Before the invasion of Ukraine, Vladimir Putin’s popularity was stagnating near historic lows. By all accounts, his popularity declined most precipitously among the young and the urban professional class. Some have linked this decline in support to the degradation of Russian democracy. Since the cancellation of mayoral elections has been part and parcel of this process, our findings are consistent with such conclusions.

2 Support for Elections under Autocracy

Although conceptually distinct, elections and democracy are sometimes considered as two sides of the same coin. The two are indeed intertwined in established democracies, where mass support for free elections undergirds the system’s legitimacy (Dennis 1970; Verba and Almond 1963). However, regular multiparty elections are also a part of political life in many authoritarian countries. These elections are sometimes even competitive, but electoral authoritarians tip the playing field in their favor to such an extent that these elections cannot be considered democratic (Levitsky and Way 2002). It is unclear what value citizens place on elections that are regular, but not free or fair.

Existing research on public opinion in autocracies does not have an answer to this question, but some studies have explored mass attitudes towards adjacent concepts, such as support for democracy more broadly. Some scholars have argued that many countries are autocracies precisely because popular support for democracy is low (Bell 2000; Pipes 2004). Such accounts are hard to evaluate for a variety of reasons. For one, the causal chain linking non-democratic values to regime survival is complex and contains many steps that require

additional assumptions. Endogeneity concerns also loom large; support for democracy and regime type are almost certainly co-determined by a raft of social, cultural, political and informational factors that are hard to account for in traditional research designs. What is more, regime type itself is likely to affect, in various ways, support for democracy. Finally, views of what democracy means vary dramatically across cultures, contexts and individuals. For these reasons, it is no surprise that empirical studies on the link between regime type and democratic values come to contradictory conclusions (Claassen 2020).

Given these difficulties, some scholars have focused instead on public support for specific democratic processes and rights, concepts that are less contaminated by subjective interpretation. For instance, Colton and McFaul (2002) find that while Russians have conflicting and ambivalent views about the broad concept of democracy, they are almost unanimous in their support for fundamental liberal rights such as freedom to choose one’s leaders, freedom of expression, and freedom of the press. This line of inquiry tends to find that, once asked about the constituent parts of democracy, many citizens under autocracy are not so undemocratic in their attitudes.

Existing survey work in electoral autocracies even provides some descriptive insights into public support for competitive elections, the democratic component of most concern to our study. In general, such surveys tend to show strong majority support for multi-party elections. Across 10 electoral autocracies surveyed in the 2019 AfroBarometer survey, 76% of respondents agreed with the statement that leaders should be chosen in “regular, open and honest elections” (M’Cormack-Hale and Dome 2021).² Surveys in Russia also show high levels of support for elections. Hale (2011), for instance, demonstrates that while most Russians believe that Russia needs “a strong hand” to lead it, fully 96% of respondents believed that voters should have the right to choose that candidate, and 87% said that “citizens should choose from among several candidates with different views by means of free and fair elections.” These patterns extend to the subnational level in Russia as well. After

²Respondents were asked which of two statements they agreed with: 1. We should choose our leaders in this country through regular, open, and honest elections or 2. Since elections sometimes produce bad results, we should adopt other methods for choosing the country’s leaders.

the cancellation of gubernatorial elections in Russia in 2004, surveys routinely showed that a large majority of Russians supported the return of elections. In 2008, for example, 75% of respondents in a Levada Center survey said that they supported the return of direct elections for governor. This is especially noteworthy given that surveys at that time routinely showed majority support for almost all political positions held by the government.³

Despite seemingly widespread support for the concept of elections, these surveys do not tell us how individuals living in electoral authoritarian regimes would actually react when faced with threats to elections. Respondents may voice support for the abstract ideal of elections, but such considerations may be less salient when compared to other considerations such as the economy or their standard of living. People may support elections in the abstract, but the elimination or degradation of elections may not affect how they evaluate the regime.

The existing literature on autocracies offers relatively few attempts to address such questions. While a number of survey experiments in new or developing democracies suggest that the public punishes candidates for engaging in electoral violence (Gutiérrez-Romero and LeBas 2020), vote buying (Weitz-Shapiro 2014; Mares and Young 2018) and voter intimidation (Mares and Young 2018), we are aware of only three such studies in autocracies. Using a candidate-choice experiment in Venezuela, Svolik (2019) finds that people withdrew support from candidates who supported reforms that undermined electoral integrity (such as maintaining a partisan electoral commission), but that they were more willing to overlook such norm violations if the reforms were supported by co-partisan candidates. Using a vignette experiment in Russia, Reuter and Szakonyi (2021) find that respondents withdraw their support from regime candidates who are revealed to have engaged in ballot fraud, vote buying, and voter intimidation. These results are replicated in the vignette experiments carried out in Russia by Aarslew (2023).

This nascent literature provides some hints that the public under autocracy is willing to punish incumbents for undermining the integrity of elections. However, it also suffers from

³The respective share who supported gubernatorial elections in 2009, 2010, and 2011 were 72%, 75% and 80%.

some significant shortcomings that reduce its ability to answer important questions about public support for elections under autocracy. For one, the dependent variable is narrow and conceptually distinct: these experiments focus on how individuals evaluate candidates who support or engage in fraud. It is plausible that the public transfers their disapproval of candidates who commit or support electoral manipulation to disapproval of the regime, but that is not empirically demonstrated, nor is it a given. Likewise, since these studies examine a different independent variable—electoral manipulation—their findings do not directly inform us about how people might react to the wholesale cancellation of elections. More generally, all the findings in this field come from survey experiments that rely on hypothetical scenarios. As the prevalence of such research designs grows, some critics question how well the findings from these artificial environments travel to the real world. As we discuss below, the present study attempts to gain purchase on this question by examining whether the public decreases their support for regime leaders when local elections are cancelled in Russian cities.

3 Theory

The main question that we address in this paper is whether the elimination of elections undermines popular support for authoritarian leaders. There are several theoretical reasons to believe that the public will punish authoritarian governments for eliminating elections. At one level, people may value elections as an intrinsic good and have a normative commitment to the ideal of selecting their leaders through elections. Consistent with this view, some studies have found that many people living under autocracy value democracy for its own sake (Bratton and Mattes 2001). Such individuals might then withdraw support from incumbents who deprive them of the ability to exercise the right to select their leaders.

Alternatively, the public may value elections because they provide instrumental benefits. For instance, people may value elections because they believe that elections produce more representative and accountable government. People living under autocracy may be right to be skeptical about the notion that voting is an effective method of influencing government, yet surveys nonetheless indicate that non-negligible shares of the electorate still view them

as a mechanism for removing poorly-performing leaders (M’Cormack-Hale and Dome 2021). In a similar vein, members of the public may believe that elections produce higher quality leaders (Besley and Reynal-Querol 2011). Individuals may also value elections because they think that elections produce better outcomes such as economic growth, political stability, or public goods. If voters see instrumental value in elections, then removing their access to these benefits could be a deeply unpopular decision.

One alternative argument is that individuals living under electoral autocracies do not value elections, so they will not punish regime leaders for cancelling elections. To be sure, some people hold beliefs that are the polar opposite of those described above. They may believe that elections create instability, that they lead to the selection of incompetent leaders, or that they undermine economic growth. More generally, many living under autocracy are aware that elections are not free and fair (M’Cormack-Hale and Dome 2021). If elections do not provide people with the opportunity to meaningfully influence policy, then they may not feel inclined to defend them.

While these attitudes exist, the surveys cited in the previous section indicate that they are minority viewpoints in Russia and electoral autocracies more generally. Over the period of our study, surveys indicated that majorities of respondents believed Russian elections to be mostly free and fair. However, surveys indicate that Russians support the maintenance of elections, even if they are flawed. In a 2008 survey, sixty-five percent of respondents agreed that “elections should be maintained even if they are not free and fair,” with 28 percent disagreeing (Hale 2011, p.1365). Respondents in that survey also mostly disagreed with the view that democracy is incompatible with economic growth and strong state (Hale 2011, p.1366). Thus, on balance, we expect that the average voter would be more likely to withdraw support when elections are cancelled.

Another alternative argument is that partisan biases color how people evaluate the cancellation of elections. In most contemporary autocracies, large portions of the public actively support the regime (Guriev and Treisman 2019). Supporters of the regime may forgive the cancellation of elections if it is spearheaded by regime leaders that they support. Such parti-

san biases are known to affect approval for a wide range of policy decisions (Svolik 2019)—the decision to cancel elections may be no different.

Such biases undoubtedly skew the issue evaluations of voters. However, partisanship and polarization were relatively weak in Russia before the invasion of Ukraine (see Orhan et al 2024). Indeed, survey experiments have found that regime supporters punish regime candidates who engage in undemocratic behavior just as much, if not more, than swing voters do (Reuter and Szakonyi 2021; Aarslew 2023). As noted above, large majorities of Russian voters (including 81% of Putin supporters) backed the idea of electing mayors directly when surveyed in 2021. Thus, that, on balance, incumbents should stand to lose support when cancelling elections.

Finally, and distinctly, people may decline to punish incumbents for cancelling elections because the preservation of elections is just not a salient issue for them. Individuals may well see some value in elections, but if preserving elections is not an important political issue for them, then they will not enter this factor into their assessments of incumbents. This is a particularly relevant consideration for second-order elections, which are typically viewed as less important. If this alternative hypothesis is correct, then we should expect to see no effect of election cancellation on support for the authorities.

In this paper, we test whether the public does, in fact, withdraw support from authoritarian incumbents that cancel elections. Beyond this primary goal, we also examine the conditions under which voters withdraw support. The more people value elections, the more likely they will be to punish incumbents who undermine elections. The factors that lead to support for democracy and elections are likely to vary by context, but, generally speaking, we argue that those segments of society that are more supportive of the idea of democracy will be more likely to withdraw support from incumbents who seek to cancel elections.

On this basis, we advance two propositions. First, individuals will be more likely to punish incumbents for cancelling elections if elections are seen as meaningful and consequential. Those who view elections as fraudulent and irrelevant should take less umbrage at their removal. Such beliefs may be based in partisanship, drawn from social cues, or rooted in

individual idiosyncrasies. They are also formed by experience. We expect that people who have experience with competitive elections will be more likely to punish incumbents for cancelling elections. Respondents who have experience with consequential elections—where competition was high, where turnover has occurred, where incumbents were defeated—should be more likely to view elections as worthy of preservation and thus to punish incumbents who seek to cancel them.

Second, individual level traits may condition responses to election cancellation. Those who value elections intrinsically should be more likely withdraw support from leaders when elections are cancelled. The same goes for those who believe that elections lead to good governance. Such beliefs can have many possible sources. They may be nurtured by experience with democracy or could be linked with certain value orientations. As we discuss below, certain demographic groups, such as younger voters and the better educated, are more likely to hold such beliefs (Nie, Junn and Stehlik-Barry 1996; Glaeser, Ponzetto and Shleifer 2007).

4 The Case: Cancellation of Mayoral Elections in Russia, 2003-2018

We examine these arguments in the case of Russia. Russia offers a unique opportunity to study how the cancellation of elections affects public opinion toward authoritarian leaders, as there is wide temporal and geographic variation in the cancellation of local elections in Russia. We focus specifically on the cancellation of mayoral elections in large cities—those with populations over 75,000. In the early 2000s, almost 90% of Russia’s large cities directly elected their mayors.⁴ Over the course of the 2000s, however, many Russian cities began replacing their elected chief executives with appointed chief executives, often called city managers. According to a 2003 law on local self-government, city councils were allowed to choose between having a directly elected chief executive (essentially identical to the existing model) or a model in which primary executive power would be vested in a city manager who would be selected by a commission appointed by the regional administration and city

⁴A small number of cities in Russia, mostly in ethnic republics, never held direct elections, even in the 1990s.


council.⁵ Although city councils had formal say over the model used to select the city’s mayor, the nationwide trend toward cancelling direct elections was almost universally seen as part of the Kremlin’s efforts to recentralize political authority in the 2000s and 2010s. In individual cities, the initiative for cancelling elections almost always originated either from the governor’s office or from local United Russia leaders. In some cases, governors used their immense formal and informal resources to cajole city councils into passing a cancellation decision. By the early 2010s, United Russia had come to control majorities in most large city councils. This made the cancellation process easier.

As Figure 1 shows, almost 50% of cities had changed to an appointment model by 2012. In 2014, the basic law on local self-government was amended again to allow regional legislatures to determine the method of selection for all local chief executives in a region (see also Gasparian (2021) for more on this reform). These amendments also outlined a new appointed chief executive model that regions could choose. This new model consists of a single mayor appointed by a special commission that is selected by the city council and regional legislature. Many regions availed themselves of this new opportunity to cancel elections for all cities in the region and this prompted a new wave of election cancellations in 2014–15. By January 2018, 85% of all large cities had appointed chief executives. The gradual cancellation of mayoral elections between 2003 and 2018 provides a wide range of variation that makes our analysis possible.

One intriguing aspect of the drive to cancel mayoral elections is that the Russian public appears to prefer to select their mayors via direct elections. At various points over the past

⁵The details of this process are described in Buckley et al. (2014); Reuter et al. (2016), and Minaeva et al. 2021. Three models were made possible during this period. One model was a unified executive model with a directly elected mayor. This was the model used by most Russian cities in the early 2000s. A second model was dual executive model, with an appointed city-manager (head of administration [glava administratsii]) and a head of city (glava goroda) who was selected by the members of the city council and usually served, ex officio, as the chairperson of the city-council. Under this model, the city manager was responsible for almost all policy-making decisions, while the powers of the head of city were more restricted to legislative politics and serving as the official head of the city in external relations with other cities. When cities cancelled their direct elections, this was the most commonly selected model. A third model was similar to the second model described above, except that the head of city was directly elected. This model was rare, but was used for a time in some prominent cities, such as Perm and Ekaterinburg.

Figure 1: Proportion of Russian Cities with Appointed Executive, 2003–2018



15 years, Russian survey firms have asked respondents how they think their mayors should be selected. In 2006, when the wave of cancellations was just starting, 79.9% percent said they preferred elected mayors to appointed mayors.⁶ Five years later, 68.7% of respondents supported mayors being directly elected.⁷ In a December 2013 Levada Center survey, 70% of respondents said they would disapprove of a draft bill to cancel elections in large cities.⁸ In 2017, the Levada Center asked respondents directly how they thought their local chief executives should be selected: directly elected, appointed by a special commission, or selected by the city council from among the deputies. Sixty-five percent preferred direct elections, followed by 12% preferring the special commission and 5% choosing election by city councils. Another 18% were unsure how to answer. Most recently, the 2021 Russian Election Study asked respondents whether they believe that the heads of cities should be selected directly

⁶Levada Center, cited in Reuter et al. (2016)

⁷Public Opinion Fund (FOM) public opinion survey “Georating 11mf_hse,” 2011, *questionQ342*.

⁸Levada Center public opinion survey “Courier 2013-14,” December 2013, *question Q39A*.

by voters or appointed by city councils. Eighty-seven percent said that they believed mayors should be elected, with 13% saying that they should be appointed.⁹

Looking at the election cancellation process in individual cities, we observe evidence of widespread disapproval as well. For instance, at the time Perm's city council voted in 2010 to adopt an appointment system, a Levada poll from that same period indicates that 79% of Perm residents supported direct mayoral elections (a number that, as an act of protest, opponents of reform sprayed painted on sidewalks around the city).¹⁰ Similarly, less than two months before reforms were adopted in Arkhangelsk in December 2014, the city conducted its own survey of residents' attitudes, which indicated that 75% preferred directly electing their mayor, 71% thought that moving to appointed mayors would lead to worse outcomes for the city, and 76% believed the city council should be obligated to uphold the general public's preference for direct elections.¹¹

It is also noteworthy that, while reforms unfolded with little organized public opposition in some cities, the cancellation process provoked significant challenges in others. Following the announcement in Ekaterinburg that, at the governor's request, the city council would be voting on a proposal to end direct mayoral elections, several thousand gathered on the the main square to protest the decision.¹² Protests also accompanied the cancellation of elections in other cities, such as Nizhevirtovsk, Ulan-Ude, Angarsk, Berdsk, Chita, Irkutsk, and Yuzhno-Sakhalinsk. In a few cities, such as Surgut, Biisk, and Rybtsovsk, public opposition was strong enough to quash initial attempts to cancel elections, although such efforts only delayed the inevitable for a year or two before renewed attempts were pushed through. Even after 2014, when federal law gave regional legislatures legal authority to cancel cities' direct

⁹We have do not have survey questions that ask respondents why they prefer to elect their mayors, but survey data on their views of gubernatorial elections offer some clue. In a May 2007 Levada Courier survey respondents were asked: "Which governors, those popularly elected or those appointed by the the president, would pay more attention to the needs of your region and its residents?" By a margin of 2 to 1 voters though elected governors would do a better job representing their interests (48% to 24% [24% believed neither could]).

¹⁰<https://ura.news/news/1052116650>

¹¹<https://www.arhcity.ru/?page=0/29777>

¹²<https://www.rbc.ru/politics/03/04/2018/5ac320cb9a7947ae6ea3d653>

mayoral elections, support for electing mayors was strong enough in several cities that the city council voted down the first round of proposals to change the city charter in compliance with new regional laws. In the case of Krasnoyarsk, where the ruling party's own deputies dominated the city council, it took threats from the regional prosecutor's office to dissolve the city's legislative body and almost an entire year before the city council finally changed the city charter.¹³

Moreover, there is evidence that public support for elections often persists long after the cancellation reforms have passed. Citizens and opposition figures have continued to make repeated (but so far unsuccessful) efforts to bring back mayoral elections; in 2019-2020 alone, at least thirteen cities saw popular proposals to reintroduce elections. Somewhat ironically, even some governors, such as Zabaikalskii Krai governor Aleksandr Opisov and Primorskii Krai governor Oleg Kozhemyako, have flirted publicly with the idea of returning mayoral elections to their major cities.¹⁴ While observers do not expect the governors to follow through on such talk, many see such statements as forced recognition of the public's enduring support for local elections. Thus, while support is far from universal, Russians appear to prefer mayoral elections and, in many cases, the cancellation of elections elicited significant public outcry. It is noteworthy that the regime has pursued this reform measure despite public opposition.

4.1 Hypotheses

We advance several hypotheses on the basis of the theory and case discussed above. The first, overarching hypothesis is that election cancellation will reduce support for the authorities in cities where it is carried out. As we noted above, arguments can be made that election cancellation will have no effect on public opinion, but given that direct elections are supported by sizable majorities of respondents in a wide variety of autocratic settings (including in the case of Russian mayoral elections), we expect that, on balance, election cancellation will

¹³<https://dela.ru/news/163375/>

¹⁴<https://www.chita.ru/articles/124171/>; <https://baik124.ru/text/06-08-2021/004/>

lead individuals living in electoral autocracies to withdraw support from regime leaders. This leads to our first hypothesis:

H1 The cancellation of local elections will reduce public support for regime leaders.

Surveys in Russia demonstrate that a majority of Russians support mayoral elections, but these same surveys also show that sizable portions of the electorate are not bothered by (or even support) election cancellation. This suggests that responses to the cancellation of local elections will be heterogeneous. Thus, we also advance a series of hypotheses about the conditions that make it more likely for members of the public to withdraw support when elections are cancelled.

First, consistent with the arguments presented in the previous section, we expect that when voters see elections as consequential or democratic they will be more likely to believe elections are worth defending and thus withdraw support from incumbents when those elections are then cancelled. In September 2014, the Levada Center conducted a poll in which they asked respondents who did *not* oppose cancelling local (and gubernatorial) elections why they thought that it was permissible to eliminate these elections. The most popular response (35% of respondents) was that elections were fraudulent and that they do not actually determine who gets elected.¹⁵

We expect that those who have experience with meaningful, competitive elections will be more likely to punish the authorities for cancelling elections. Specifically, we expect that people who live in regions with a history of competitive elections will be more likely to withdraw support when elections are cancelled. Such individuals have experience with observing and participating in elections that were relatively free, democratic, and, often, resulted in turnover. In order to measure experience with electoral competition, we draw on both expert-coded and more objective measures of electoral competitiveness. Specifically,

¹⁵Other responses include: regular people are not qualified to participate in elections (34%), elections don't change anything (25%), governors, mayors, etc. don't have any impact on our lives (16%) and elections undermine stability(6%). These responses reveal that election opponents were not rejecting the abstract idea of elections so much as expressing disillusionment with how they perceive to elections to operate in Russia.

we use the electoral competitiveness sub-index from the Petrov-Titkov composite index of democracy in the Russian regions (Petrov and Titkov 2013) as well as electoral data on the effective number of candidates in gubernatorial elections in the region in the 1990s and early 2000s. Since contemporaneous measures of democracy may themselves be influenced by the cancellation of elections, we examine whether the legacy of electoral politics in the 1990s and early 2000s affects how the public reacts to the cancellation of elections. We expect that individuals in regions that were more democratic in the 1990s and early 2000s should be more dissatisfied with the cancellation of elections.

H2 Individuals in regions with a history of real electoral competition will be more likely to withdraw support from the authorities when elections are cancelled.

Our framework also leads us to expect that those who intrinsically support democracy and those who are more likely to believe that elections produce better governance will be more likely to reduce their support for the regime when elections are cancelled. Unfortunately, our survey data lacks fine-grained measures on support for democracy or opinions about democracy’s effects. However, in Russia, demographic information can provide a first-cut proxy for individual attitudes since public opinion reports over the period of interest show age and education level to be correlated with attitudes towards elections and democracy. Specifically, the evidence suggests that less educated voters and older voters – those who were in the working years through the turbulent transition under Gorbachev and Yeltsin – have more skeptical views of democracy and are more supportive of the authoritarian management of the country under Putin (Volkov 2021; Levinson 2023*a,b*). Therefore, we expect that these groups will be more likely to view elections more positively.

H3 Younger individuals will be more likely than older individuals to withdraw support when elections are cancelled

H4 More educated individuals will be more likely than less educated individuals to withdraw support when elections are cancelled.

In the following section, we describe both the data we use to test these hypotheses as well as our analytical approach.

5 Data

5.1 Election Cancellation

Our main analyses draw on an original dataset that tracks the precise dates when elections were cancelled in the 212 largest cities in Russia. This data builds on other datasets used to study the causes and consequences of local election cancellation in Russia (Reuter et al. 2016; Beazer and Reuter 2019, 2022). We build on this data by collecting and verifying fine-grained information on the exact date of election cancellation in these cities. Existing datasets contain year-level information on the type of selection model used in a given city. However, since our main empirical goal in this paper is to examine how changes in mayoral selection model affect public opinion—and since public opinion can react immediately to the cancellation of elections—it is important for us to establish a clear timeframe for when elections were cancelled.

In order to do this, we read media reports and the texts of laws passed by city councils to determine when elections could be considered cancelled in a given city. Substantively, we seek to establish the date on which it became clear—to observers, key political actors, and, most importantly, the public—that elections have been cancelled in a city. For several reasons, this is more difficult than it may initially appear. The actual process of cancelling elections consists of multiple steps. In some cases, the idea is initially floated in the press by a United Russia city council deputy or by the mayor himself. The first formal step in the process is a bill introduced to the city council in a first reading. If the bill passes the first reading, a public hearing and period of commentary sometimes then set. After this, the bill advances to a second and final reading, at which point the changes are formally entered into the city charter(*ustav*). In many instances, however, this entire legislative process is condensed into a single day and the bill is passed quickly in three readings.

In coding the data, our task was to determine at which point in this process voters would

reasonably conclude that mayoral elections in their city had been eliminated. We rely on close readings of press accounts to determine which development is the most resonant and the most likely to be perceived by the public as the date of cancellation. In the moment, observers sometimes portray a city's elections as cancelled on the first reading or after the public hearing period, treating charter changes as a foregone conclusion once the initial legislative and public relations work is completed. In some instances, this seems accurate, but that coding rule would obscure cases in which the process drags out over a long period or when there is uncertainty about whether the initiative will succeed.¹⁶ In any case, it is much more common for press reports to characterize the moment of final passage as the date when elections are cancelled. Thus, in most cases, we code the final third reading passage (i.e., when the changes are entered into the charter) as the official cancellation date. For completeness, we also collected secondary dates that might reasonably be considered as the cancellation date. Below, we carry out a series of robustness checks that use these secondary dates to demonstrate our results are not sensitive to this choice of dates. In addition, we also show that our results are robust when we use the date (where available) when the proposal to cancel first came on the agenda.

A 2014 federal law further complicates the task of determining the precise date of subset of cancellations. This reform gave regional authorities the right to cancel elections on behalf of municipal governments, and many regional governments used this authority to pass legislation cancelling mayoral elections in their large cities. Although the regional legislation's date is a reasonable contender for cancellation in these cities, reforms only took full effect once cities brought their charters into compliance with regional legislation, and many cities dragged their feet for months before finally complying.¹⁷ In most instances, we use the date

¹⁶For example, several cities have abandoned the cancellation process part-way through (sometimes to be restarted months or years later). This alone makes it hard to conclude that the outcomes are foreordained once the initiative is first introduced.

¹⁷In fact, some cities, such as Arkhangelsk and Krasnoyarsk, put up such active opposition that it was not immediately clear when (or even if) the changes would be formally adopted by the city. In Krasnoyarsk, for example, the regional bill cancelling elections passed in December 2014, but the local legislature only brought the changes to a vote in April 2015. The city council rejected the bill, and then rejected it a second time in October 2015. The changes were finally adopted and took effect only in November 2015,

that the city charter was brought into conformity with the regional law as the official date of cancellation. When the city date is unavailable, we use the regional date. We also use the date of regional law passage as a robustness check for all cities where this applies.¹⁸

5.2 Public Opinion Data

In our analyses, we examine the shifts in public attitudes generated by abrupt changes to the selection model of local executives in Russia. To capture these shifts, we use a massive assemblage of Russian public opinion data. Collected, standardized, and merged into a single database over a period of years, this dataset represents over 1.4 million survey responses in 394 survey waves from 2003 through 2018. To our knowledge, this is the largest and most comprehensive collection of Russian public opinion data ever produced.

The dataset consists of over 1.4 million individual-level entries from Levada Courier, FOM GeoRating, and VTsIOM Express survey waves starting in March 2003. Run by the three leading Russian polling firms, these surveys were conducted every month or every week for most of the period under study. The Levada and VTsIOM surveys are ‘omnibus,’ nationally representative surveys while the FOM GeoRating surveys are regionally representative surveys.¹⁹ Aside from demographics, each wave only asked about a small set of questions.²⁰

To create the primary data frame for testing our hypotheses, we merge our pooled public opinion data with our city-level data on cancellation dates. This process begins with multiple, painstaking rounds of cleaning and collating the many survey waves. Within the pooled public opinion data, the authors recoded each survey wave to ensure that common ques-

nearly one year after the regional law was passed.

¹⁸A further, though rare, complication is that some cities have a dual-executive model where the main policy making figure—the city manager—is appointed but voters still directly elect a symbolic head of the city. It is unclear whether transitions toward or away from this model should be considered as cancellations. Therefore, we drop the 10 cities that used this model from our analyses. We test for robustness to their inclusion in the robustness section.

¹⁹Each GeoRating survey includes approximately 500-600 respondents per region in approximately 65 regions. Thus, each GeoRating survey is much larger than the Levada and VTsIOM surveys, but we only have 2-4 GeoRating surveys per year for the period 2003-2011.

²⁰The survey data underlying this database was provided by the Levada Center, the Higher School of Economics data repository, VTsIOM, and other sources. We are grateful for these organizations’ cooperation.

tions’ responses are identically constructed before being merged into one pooled database. When necessary, we aggregated categorical demographic variables to the smallest common categories across survey waves to ensure compatibility. We then extracted the respondents in these surveys who lived in the 212 largest cities in Russia and matched these observations to the appropriate cities in our city-level, cancellation dataset. This creates a final, merged dataset of approximately 450,000 survey responses in large Russian cities and a total of 127 cancellation events.²¹

The core of our analyses rests on measuring general support for the regime. As described in Hypotheses 1–4, our theory asserts that cancellation of mayoral elections in an electoral authoritarian regime like Russia will decrease the public’s support for the regime and prominent regime leaders. We focus on approval of Vladimir Putin, as he has consistently served as the key figurehead of the regime since the early 2000s. Surveys frequently ask respondents to state their approval or disapproval of Putin—this dichotomous question serves as the main dependent variable in our analyses.²² Surveys over this time period only rarely ask about approval or support of United Russia or the authorities more broadly, so our focus remains on Putin himself. In some analyses, we also look for changes in approval of regional governors. Using approval of one’s governor as a dependent variable, we can examine the effects of election cancellation on those who are implicated most directly in that process.

5.3 Modeling Strategy

By combining the survey data and election cancellation information, we gain an unprecedentedly fine-grained view into how support for an electoral authoritarian regime changes when popular elections are eliminated. We are able to leverage individual-level, geographically-diverse, over-time survey data with monthly resolution. This represents a much more power-

²¹Approximately 5% of cities never had elections and approximately 5% had not cancelled elections by 2018, the end of dataset. We are also missing the precise cancellation date for some cities.

²²FOM GeoRating surveys occasionally ask an alternative formulation of this question. Here, they attempt to capture both levels of support for Putin and changes to that support. We aggregate responses along the ‘change’ dimension in order to keep the focus squarely on levels of support and maintain standardization across surveys. Survey firm fixed effects help eliminate any remaining heterogeneity. See Table 5 Model 3 for a test of the robustness of our results to controlling for these question wording differences.

ful tool for studying election cancellation than the typical case of a rare, one-shot national-level cancellation. Instead of relying on a pair of measurements around one single cancellation event, we bring to bear a full 127 election cancellation events along with hundreds of thousands of survey responses in the weeks and months before and after elections are eliminated.

In addition to providing highly precise temporal resolution, we offer a focused look at the *city-level* effects of this event. In our data structuring process, we link the cities that each survey response originates from with the 212 cities in our cancellation dataset. This allows us to eliminate rural respondents and those who live in cities for which we do not have data on the selection mechanism.

Our main estimating strategy is a two-way fixed effect regression (TWFE), which includes fixed effects for city and month-year, along with a binary treatment indicator at the city-month level. This variable takes a value of one in the month after a city’s elections are cancelled and zero in all preceding months. Our base model estimating equation is the following:

$$Putin\ Approval_{itr} = \alpha + \beta_1 C_{itr} + \gamma X_{itr} + \theta_i + \eta_t + \phi_{itr} + \varepsilon_{itr}$$

where i indexes each city, t indexes the month-year and r indexes the respondent; $PUTIN\ APPROVAL_{itr}$ is the measure of each respondent’s approval of Putin; C is the measure of whether the city has cancelled elections (i.e., uses an appointment system); $\mathbf{X}$ is a vector of relevant, individual level control variables; α , β , and γ are parameters to be estimated; θ , η , and ϕ are fixed effects parameters for city, month-year, and survey company, respectively; and ε is the error term. Here, our main quantity of interest is β_1 , the estimated effect of cancelling elections on support for Vladimir Putin.

This TWFE approach is equivalent to generalized difference-in-differences research design, where, for a given month, cities with elected mayors serve as a control group and

counterfactual for comparisons with cities that switch to an appointment system.²³ We use the linear probability model via OLS estimation in all analyses. We include demographic controls for age, gender, and a four-point educational attainment variable in all models. Standard errors are clustered at the city level.

6 Results

Table 1 presents results from our baseline models. Model 1 shows the effect of election cancellation on approval for Vladimir Putin. The results indicate that election cancellation decreases support for Putin by a small but significant degree: 2.2%. Given the very large number of considerations that inform evaluations of Putin, this effect is sizable. For example, in our models we see that the effect of gender on support for Putin is 6.5%, which is generally considered a large and significant effect. An average effect of mayoral election cancellation of more than one-third that of gender is substantively large. Model 2 looks at the effect of election cancellation on support for regional governors. Governors are political figures who are usually directly implicated in the process of cancelling elections, so it makes sense that the public would adjust their opinion of the governor in response to election cancellation. Model 2 indicates that this is the case and that election cancellation reduces support for the regional governor by 5.2%. However, this survey question is missing in almost three quarters of our surveys, so findings must be interpreted with caution.

Model 3 indicates that individuals also withdraw support from the government when elections are cancelled, but this effect is smaller and less precisely estimated. Support for the government in Russia is almost always lower than support for Putin and trends in support for the government do not move in lock step with support for Putin. The government is perceived as less political and more of a technocratic policy making entity. More to the point, rational voters would be correct in seeing little linkage between the actions of the

²³For discussion and exploration of potential sources of bias in the TWFE specifications we employ, see Sections 7 and 8. The additional tests and analyses in these sections demonstrate that: 1.) our data display very few markers of bias-inducing heterogeneity, and 2.) our results hold when using modeling strategies to counteract any such potential bias.

Table 1: Mayoral Election Cancellation Reduces Regime Approval

	DV: Putin approval	DV: Governor approval	DV: Gov't approval
	(1)	(2)	(3)
Elections cancelled	−0.022** (0.007)	−0.054* (0.023)	−0.019+ (0.010)
Age	−0.001** (0.000)	−0.001** (0.000)	−0.002** (0.000)
Educational attainment	−0.005** (0.001)	−0.006** (0.002)	−0.009** (0.001)
Male	−0.065** (0.002)	−0.058** (0.004)	−0.062** (0.003)
Number of Observations	401 343	109 113	331 579
R^2 Adj.	0.073	0.062	0.109

Coefficient estimates from OLS regression with standard errors in parentheses clustered by city. All models include fixed effects for city, month-year, and survey firm. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

government and the cancellation of elections at the local level. In contrast, Putin is Russia’s primary political actor and governors are key political representatives of the regime in the regions.

Table 2 explores how the effect of election cancellation varies across political contexts depending upon the level of political competition. In Model 1, we interact our treatment (election cancellation) with a measure of the effective number of candidates in the most recent gubernatorial elections held in the region (before all such elections were cancelled in 2004). Model 2 interacts treatment with a measure of the effective number of candidates averaged across all elections in the 1990s. Lastly, Model 3 interacts treatment with Petrov and Titkov’s (2013) expert-coded assessment of electoral competitiveness in the regions in the 1990s. Across all three measures, we observe the same result: the negative effect of election cancellation on Putin support is considerably larger in regions with a history of robust electoral competition. Figure 2 demonstrates how the marginal effect of election cancellation on Putin support changes with levels of electoral competitiveness. In panel 1, for instance, we see that mayoral election cancellation reduces support by about 7 percentage

Table 2: Effects of Mayoral Election Cancellation on Regime Approval, Moderated by Political Competitiveness


	DV: Putin approval		
	(1)	(2)	(3)
Elections cancelled	0.016 (0.016)	0.038 (0.023)	0.018 (0.016)
ENC of most recent govr election pre-2004	0.011** (0.004)		
ENC last govr elec pre-2004 $\times$ Elections cancelled	-0.011** (0.004)		
Avg ENC of govr elections pre-2004		0.038** (0.005)	
Avg govr ENC pre-2004 $\times$ Elections cancelled		-0.019** (0.006)	
Regional electoral competitiveness 1990s			0.002 (0.002)
Reg electoral compet 1990s $\times$ Elections cancelled			-0.013* (0.005)
Age	-0.002** (0.000)	-0.002** (0.000)	-0.001** (0.000)
Educational attainment	-0.008** (0.001)	-0.008** (0.001)	-0.005** (0.001)
Male	-0.070** (0.002)	-0.070** (0.002)	-0.065** (0.002)
Number of Observations	295 129	295 129	401 343
R^2 Adj.	0.073	0.073	0.073

Coefficient estimates from OLS regression with standard errors in parentheses clustered by city. All models include fixed effects for city, month-year, and survey firm. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

points in the most democratic regions, while the effect is null in the least democratic regions. These analyses make a compelling case that, in places with a history of more vibrant electoral contests, people place greater value on elections. When confronted with regime efforts to remove elections, these individuals experience greater frustration and dissatisfaction than individuals living in places where elections have a less-impressive track record.

Table 3 investigates individual-level heterogeneous effects of election cancellation on Putin approval. In Model 1, we interact our treatment variable with respondent age. The results indicate that age significantly moderates the effect of election cancellation on support for Putin. Figure 3 plots the marginal effect of treatment by respondent age. For a young

Figure 2: Marginal Effects by Regional Political Competitiveness



Marginal effect estimates based on Model 2 (Avg ENC) and Model 3 (90s Electoral Competitiveness) in Table 2. Bands represent 95% confidence intervals.


Table 3: Effect of Mayoral Election Cancellation on Regime Approval, Moderated by Age and Education

	DV: Putin approval	
	(1)	(2)
Elections cancelled	-0.077** (0.011)	-0.015* (0.007)
Age	-0.002** (0.000)	-0.001** (0.000)
Age $\times$ Elections cancelled	0.001** (0.000)	
Educational attainment	-0.005** (0.001)	-0.004** (0.001)
Education $\times$ Elections cancelled		-0.003 (0.002)
Male	-0.065** (0.002)	-0.065** (0.002)
Number of Observations	401 343	401 343
R^2 Adj.	0.074	0.073

Coefficient estimates from OLS regression with standard errors in parentheses clustered by city. All models include fixed effects for city, month-year, and survey firm. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

respondent (22 years old), the effect of election cancellation is strongly negative and statistically significant. The estimated reduction in support for Putin is roughly 5 percentage points. By contrast, for the oldest respondents, the effect is actually positive and also statistically significant (but smaller in magnitude). This supports our hypothesis that younger respondents—who are more likely to hold democratic attitudes in Russia—are the most likely to punish the authorities for cancelling elections. It is also noteworthy that heterogeneity in treatment effect is very large and is even positive for sizable subsets of the sample. This is important to keep in mind when evaluating the baseline treatment effects presented in Table 3; while cancelling elections appears to have been an unpopular move that reduced regime support on average, there are individuals who responded to these reforms with approbation and enthusiasm.

Figure 3: Marginal Effects by Age and Education



Marginal effect estimates based on Model 1 (age) and Model 2 (education) in Table 3. Bands represent 95% confidence intervals.

Model 2 interacts election cancellation with the education level of the respondent. Figure 3 displays the conditional coefficients from this interaction. Contrary to our expectations, there is only a weak indication that more educated respondents are more likely to punish

Putin for election cancellation. In contrast to the strong interactive results for age, this difference (between the educated and less educated) is not statistically significant. Ultimately, our data provide no clear evidence that the negative effects of election cancellation on Putin support are stronger among Russians with higher levels of education.

6.1 Alternative Mechanisms and Robustness Checks

In Table 4, we explore the robustness of our main results to alternative data formulations. In Model 1, we drop all cities located in Moscow oblast. As much of Moscow oblast is a densely-populated commuter area for Moscow city proper, attitudes in these cities may be affected by the Moscow metropole or by events in neighboring settlements. Our results remain robust to dropping these cities. Model 2 omits all survey responses from FOM GeoRating. Since the GeoRating survey waves contain very large sample sizes from the period 2003-2011 only, one could be concerned about our results being inadvertently and unduly weighted towards the cities and time periods that happened to be included in this survey. Once we drop these observations, the estimated treatment effect is somewhat larger.

Table 4: Negative Effects of Election Cancellation on Regime Approval are Robust to Alternative Specifications

	DV: Putin approval				
	No Msk Obl	No FOM	No demogr	Day FEs	Alt cancel dates
	(1)	(2)	(3)	(4)	(5)
Elections cancelled	−0.021** (0.007)	−0.034** (0.012)	−0.021** (0.007)	−0.022** (0.006)	−0.022** (0.007)
Age	−0.001** (0.000)	−0.001** (0.000)		−0.001** (0.000)	−0.001** (0.000)
Educational attainment	−0.005** (0.001)	−0.002 (0.001)		−0.005** (0.001)	−0.005** (0.001)
Male	−0.065** (0.002)	−0.067** (0.004)		−0.065** (0.002)	−0.065** (0.002)
Number of Observations	389 801	174 817	404 993	401 343	402 070
R^2 Adj.	0.072	0.075	0.064	0.075	0.073

Coefficient estimates from OLS regression with standard errors in parentheses clustered by city. All models include fixed effects for city, month-year, and survey firm. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

Model 3 in Table 4 omits the three demographic control variables that we normally

include in all models; results are unchanged. Model 4 substitutes time fixed effects at the day level for the month-level FEs included in other models. Finally, Model 5 replaces our primary cancellation date with an alternate, ‘secondary’ date in those cases when our coding process uncovered uncertainty about which date better represents the moment of election cancellation.²⁴ Our results are robust to these alternatives as well.

7 Addressing Potential Bias from Heterogeneous Treatment Effects

In recent years, a fast-growing literature in econometrics has reconsidered the unbiasedness of coefficients produced by two-way fixed effects (TWFE) estimators. These studies have concluded that TWFE estimates may be biased when treatment effects are heterogeneous across time or units (Goodman-Bacon 2021; Callaway and Sant’Anna 2021; De Chaisemartin and d’Haultfoeuille 2020). These problems arise because TWFE coefficients are calculated as a naively-weighted composite of “clean” comparisons between treated observations and untreated (or yet-to-be-treated) observations on the one hand and, on the other hand, “forbidden” (“dirty”) comparisons that use already-treated observations as the counterfactual for the current treatment group (see Roth et al. (2022) for a good overview).²⁵ The bias arises when treatment effects are heterogeneous because the estimates from those “forbidden” comparisons will be determined by the comparison between the late-treated units’ post-treatment values and the early-treated units’ values *after* the late treated group is treated. If treatment effects are homogeneous (and the parallel trends assumption holds) then this difference will be zero, but if treatment effects are heterogeneous, then this quantity will be non-zero and will reflect variation in the treatment effect over time. Traditional TWFE estimators naively

²⁴See Table A2 for additional robustness checks using alternative cancellation dates. These analyses produce similarly robust results. In addition, Table A7 demonstrates that results are consistent if, instead of our TWFE specification, we adopt a multilevel modeling approach that replaces fixed effects with varying intercepts for city, month-year, and survey firm.

²⁵The coefficients in TWFE models are a weighted average of all the potential 2×2 difference-in-difference comparisons that can exist between different treatment timing groups (Goodman-Bacon 2021). The weights are proportional to the group sample sizes and the variance of the treatment within each 2x2 DID.

incorporate the quantities from these “forbidden” comparisons into the coefficient estimates and thus introduce the potential for biased coefficient estimates.

This section discusses our efforts to address the potential for such biases in our research design. One primary concern is that our treatment effects might not be homogeneous across time. For instance, the treatment effect’s magnitude might decline over time as those living within cities where elections have been canceled become resigned to living under an appointment system. Alternatively, the treatment effect might increase over time, perhaps because quality of governance declines without elections.

Heterogeneous treatment effects can create bias as a result of the negative weights that are often applied to treated units. As noted by Jakiela (2021), “negative weights are not a pathology, but a desired and natural consequence of the (implicit) modeling assumption we make when we difference out country and year means and estimate a single (implicitly homogeneous) treatment effect.” However, if the assumption of homogeneous treatment effects is violated and some treated units receive negative weights, then TWFE does not necessarily provide an unbiased estimate of the treatment effect (see Roth et al. (2022) and Jakiela (2021)).

Therefore, we begin by following Jakiela (2021) and apply diagnostics for calculating the proportion of treated observation weights that are negative. These diagnostics indicate that approximately 25.4% of treated observations in our main model have negative regression weights. As predicted by econometric theory, these negatively-weighted observations mostly occur in the final years of our data. Following Jakiela (2021), we show that our results are robust to dropping any of several sets of the final years in our dataset. These results are available in Appendix Figure A4.

To investigate whether our treatment effects are heterogeneous, we also conduct the diagnostic proposed by Jakiela (2021). This test entails examining the relationship between loess and OLS regression lines for residualized treatment and outcome. Potentially-troubling treatment effect heterogeneity would be evident in loess or OLS estimates that vary significantly between different values of the residualized treatment in the untreated and treated

groups, resulting in diverging slopes. As Figure A5 shows, these slopes do not vary appreciably in our data. We also conduct a formal statistical test of treatment effect heterogeneity by regressing the residualized outcome (Putin approval) on an interaction between the residualized treatment and our treatment indicator. The coefficient on this interaction term is statistically indistinguishable from zero, indicating no evidence of treatment effect heterogeneity. Based on these tests, heterogeneous treatments effects are unlikely to be a substantial concern in our analyses.

Nonetheless, we also explore how our results perform when using an estimator that is robust to heterogeneous treatment effects. Specifically, we employ the heterogeneity-robust estimator developed by Borusyak, Jaravel and Spies (2021).²⁶ One important advantage of the Borusyak, Jaravel and Spies (2021) estimator is that, unlike most heterogeneity-robust estimators, it allows for estimation with multiple observations per unit-time (in our case, multiple respondents per city-month) and also allows for interrupted time series, which is important given gaps in our survey coverage. However, the estimator also has significant limitations when applied to our data. Since it uses an imputation method to predict counterfactual outcomes of untreated observations, the procedure sometimes fails to impute values for certain classes of observations (e.g., low-variance times and groups) and drops large numbers of observations from the sample. In addition, convergence is sometimes problematic. Thus, we interpret these results with caution.

Appendix Table A3 reports the coefficients on our treatment variable (*Elections cancelled*) across the baseline, non-interactive specifications using the Borusyak, Jaravel and Spies (2021) estimator. Since this method does not allow for interaction terms, we split the sample at the median values of key modifying variables (e.g., the region’s democratic history). Coefficient estimates are highly consistent with the estimates in our main regressions, although, given the imputation and sometimes reduced sample size in these regressions, the standard errors are larger. On the whole, however, the coefficient estimates from the

²⁶This estimator is very similar to the estimators produced by Callaway and Sant’Anna (2021) and Sun and Abraham (2021). Zhang, Li and Liu (2021) and Gardner 2021 have independently proposed estimators identical to Borusyak, Jaravel and Spies (2021).

heterogeneity-robust estimator are consistent with the coefficient estimates from our TWFE specifications. Since the main purpose of these estimators is to correct for biased coefficient estimates, this gives us confidence in the results from our primary models, which are more efficient and make use of the full sample.

8 Assessing Threats to Inference


Our estimation strategy relies on a parallel trends assumption for identification. We assume that, absent election cancellation, Putin’s approval rating within treatment cities would have trended similarly to Putin’s approval rating within control group cities (i.e., cities that had not yet eliminated elections). This assumption establishes control observations as valid counterfactuals for treatment observations. If leaders eliminated elections precisely because Putin’s popularity was declining in a city, this would violate the parallel trends assumption and potentially bias our inferences. While this assumption is never fully testable, we employ several additional analyses to probe the plausibility of parallel trends for our study.

First, we look for parallel trend violations by comparing trends in Putin’s approval between treatment and control groups in the pre-treatment period. Focusing on observations of treated cities before treatment occurs and never-treated cities, we regress Putin approval on our control variables and fixed effects, plus a suite of “pre-treatment” dummies representing monthly intervals prior to treatment. Figure 4 plots the coefficient estimates for these “pre-treatment” dummies.²⁷ If elections were eliminated in places where Putin’s popularity was declining faster than elsewhere, we would expect to see a downward trend and negative, statistically significant coefficients on the pre-treatment dummies in the months leading to cancellation. Figure 4 shows no evidence of such trends.

We also evaluate the parallel trends assumption using an additional placebo test of pre-treatment cross-sectional differences. Using never-treated cities and only the pre-treatment years for treated cities, we regress Putin’s approval on an indicator equal to one for cities

²⁷Results are similar when using the heterogeneity-robust estimator of Borusyak, Jaravel and Spies (2021) discussed above.

Figure 4: No Evidence of Differing Trends in Regime Approval Between Treated and Control Cities in the Lead-up to Cancellation



Investigating the parallel trends assumption by comparing Putin approval in never-treated cities and treated cities before treatment occurs. Points represent coefficient estimates for “pre-treatment” dummies at various monthly intervals prior to treatment; significant coefficients would indicate non-parallel trends between treatment and control groups. Bands represent 95% confidence intervals. Model includes all control variables, and fixed effects for city, month-year, and survey firm; standard errors clustered by city.

“treated in the future” and our usual vector of controls and fixed effects.²⁸ Here, a significant coefficient on future cancellation status would suggest treatment cities are fundamentally different from untreated cities in their approval of Putin even before election cancellation, and our results could plausibly stem from unmodeled differences between treated and never-treated cities. Appendix Table A5 presents this test’s results; we find no evidence of such pre-treatment differences.

Another possible concern is that our results are not caused by election cancellation, but rather by unobserved region-level events that correlate with city-level election cancellation. We can test this by examining whether election cancellation affects Putin approval among respondents who live in the region but not in the city under question. In Appendix Table A6, we look for an effect of election cancellation on respondents in *suburban areas* around the cities in our dataset. For these analyses, we hand-coded all respondent localities within 50km of our main cities as being either: 1.) nearby, obvious suburbs of those cities, 2.) localities more distant or more tenuously geographically connected to the cities, or 3.) distant and very questionably functioning as suburbs. If our results capture city- rather than region-level political events, we expect weaker or null results in these areas, since these residents are not affected by election cancellation in the urban core. We find no significant effect of election cancellation on these localities, whether we consider nearby suburban areas only (Model 1), suburban areas at a medium distance from our cancellation cities (Model 2), or medium and distant (or likely independent) localities (Model 3). This demonstrates that the effect of election cancellation is limited to the places where cancellation actually occurs.

9 Conclusions

Multiparty elections are a fundamental feature of many modern autocracies. How does the public react when these elections are subverted or cancelled? We address this question by examining how city-dwellers in Russia change their assessments of Vladimir Putin in response to the cancellation of direct mayoral elections in their cities. Using a large assemblage of

²⁸Since over 95% of our cities are treated before our sample frame ends, we note that there is limited variation on this treated-in-the-future variable and encourage readers to consider results accordingly.

survey data across Russia’s 212 largest cities, we demonstrate that the cancellation of direct mayoral election is associated with decreases in support for Vladimir Putin. This effect is especially large in settings where elections might be perceived as more meaningful, such as in regions with a history of democracy and cities with more competitive elections. Our findings are limited by the fact that they focus only on municipal elections, which are much less important than national elections, but our results suggest that the regime would stand to lose even more support if it cancelled presidential or national legislative elections.

Our findings point to a trade-off that dictators face between control and popularity. By degrading democracy and undermining the electoral process, dictators can reduce uncertainty and obtain greater political control. But our findings suggest that these same actions can harm their popularity. Indeed, our findings point to under-appreciated reason why many autocrats choose to retain elections: because there is demand for elections among citizens. Of course, dictators often reason that they can absorb the reputational damage from cancelling elections and thus prefer to expand their political control nonetheless. This appears to be the calculation that regime leaders in Russia ultimately made in the case of mayoral elections. Our findings should not be interpreted to suggest that dictators will always be dissuaded from cancelling elections by their concerns about popular backlash; we do however demonstrate that such backlash exists.

As regards Russian politics, our findings are consistent with a number of well-known trends. Over the past fifteen years, the regime has lost much of its urban support, especially among the middle class and young, educated voters. We do not presume to suggest that the regime’s loss of support in cities is primarily due to the cancellation of elections in large cities, but our findings suggest that this—and the more general decline in Russian democracy of which it is a part—have been a contributing factor.

10 Bibliography

- Aarslew, Laurits F. 2023. “Why Don’t Partisans Sanction Electoral Malpractice?” *British Journal of Political Science* 53(2):407–423.
- Beazer, Quintin H and Ora John Reuter. 2019. “Who is to blame? Political centralization and electoral punishment under authoritarianism.” *The journal of politics* 81(2):648–662.
- Beazer, Quintin H and Ora John Reuter. 2022. “Do Authoritarian Elections Help the Poor? Evidence from Russian Cities.” *The Journal of Politics* 84(1):000–000.
- Bell, Daniel A. 2000. *East meets West: Human rights and democracy in East Asia*. Princeton University Press.
- Besley, Timothy and Marta Reynal-Querol. 2011. “Do democracies select more educated leaders?” *American political science review* 105(3):552–566.
- Borusyak, Kirill, Xavier Jaravel and J Spies. 2021. “Revisiting Event Study Designs: Robust and Efficient Estimation.” *University College London, mimeograph* .
- Bratton, Michael and Robert Mattes. 2001. “Support for Democracy in Africa: intrinsic or instrumental?” *British Journal of Political Science* 31(3):447–474.
- Buckley, Noah, Guzel Garifullina, Ora John Reuter and Alexandra Shubenkova. 2014. “Elections, appointments, and human capital: The case of Russian mayors.” *Demokratizatsiya* 22(1):87.
- Callaway, Brantly and Pedro HC Sant’Anna. 2021. “Difference-in-differences with multiple time periods.” *Journal of Econometrics* 225(2):200–230.
- Claassen, Christopher. 2020. “In the mood for democracy? Democratic support as thermostat opinion.” *American Political Science Review* 114(1):36–53.
- Colton, Timothy J and Michael McFaul. 2002. “Are Russians Undemocratic?” *Post-Soviet Affairs* 18(2):91–121.
- De Chaisemartin, Clément and Xavier d’Haultfoeuille. 2020. “Two-way fixed effects estimators with heterogeneous treatment effects.” *American Economic Review* 110(9):2964–96.
- Dennis, Jack. 1970. “Support for the Institution of Elections by the Mass Public.” *American Political Science Review* 64(3):819–835.
- Gasparyan, Olga. 2021. “Devolution, Local Efficiency, and Resource Allocation: Evidence from Russian Cities.”.
- Glaeser, Edward L, Giacomo AM Ponzetto and Andrei Shleifer. 2007. “Why does democracy need education?” *Journal of economic growth* 12:77–99.
- Goodman-Bacon, Andrew. 2021. “Difference-in-differences with variation in treatment timing.” *Journal of Econometrics* 225(2):254–277.

- Guriey, Sergei and Daniel Treisman. 2019. “Informational autocrats.” *Journal of economic perspectives* 33(4):100–127.
- Gutiérrez-Romero, Roxana and Adrienne LeBas. 2020. “Does electoral violence affect vote choice and willingness to vote? Conjoint analysis of a vignette experiment.” *Journal of peace research* 57(1):77–92.
- Hale, Henry E. 2011. “The myth of mass Russian support for autocracy: The public opinion foundations of a hybrid regime.” *Europe-Asia Studies* 63(8):1357–1375.
- Jakiela, Pamela. 2021. “Simple diagnostics for two-way fixed effects.” *arXiv preprint arXiv:2103.13229* .
- Levinson, Aleksei. 2023a. “Putinskie Signal.” *Moscow: Levada Center* .
- Levinson, Aleksei. 2023b. “What is Really Behind Putin’s Approval Ratings.” [Online; posted 3 Nov 2023].
URL: <https://russiapost.info/politics/ratings-pu>
- Levitsky, Steven and Lucan A Way. 2002. “Elections without democracy: The rise of competitive authoritarianism.” *Journal of democracy* 13(2):51–65.
- Mares, Isabela and Lauren E Young. 2018. “The core voter’s curse: Clientelistic threats and promises in Hungarian elections.” *Comparative Political Studies* 51(11):1441–1471.
- M’Cormack-Hale, Fredline and Mavis Zupork Dome. 2021. “Support for elections weakens among Africans; many see them as ineffective in holding leaders accountable.”.
- Nie, Norman H, Jane Junn and Kenneth Stehlik-Barry. 1996. *Education and democratic citizenship in America*. University of Chicago Press.
- Petrov, Nikolai and Aleksei Titkov. 2013. “Democracy rating of Russian regions by Carnegie Moscow Center: 10 years in Service.” *Moscow: Carnegie Moscow Center* .
- Pipes, Richard. 2004. “Flight from freedom: What Russians think and want.” *Foreign Affairs* pp. 9–15.
- Reuter, Ora John and David Szakonyi. 2021. “Electoral manipulation and regime support: survey evidence from Russia.” *World Politics* 73(2):275–314.
- Reuter, Ora John, Noah Buckley, Alexandra Shubenkova and Guzel Garifullina. 2016. “Local elections in authoritarian regimes: An elite-based theory with evidence from Russian mayoral elections.” *Comparative political studies* 49(5):662–697.
- Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski and John Poe. 2022. “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature.” *arXiv preprint arXiv:2201.01194* .
- Sun, Liyang and Sarah Abraham. 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” *Journal of Econometrics* 225(2):175–199.

- Svolik, Milan W. 2012. *The politics of authoritarian rule*. Cambridge University Press.
- Svolik, Milan W. 2019. “Polarization versus democracy.” *J. Democracy* 30:20.
- Verba, Sidney and Gabriel Almond. 1963. “The civic culture.” *Political attitudes and democracy in five nations* .
- Volkov, Denis. 2021. “Demobilizatsiya i Polyarizatsiya Parlamentskie Vybery v Zerkale Oprosov Obshhestvennogo Mneniya.” *Moscow: Levada Center* .
- Weitz-Shapiro, Rebecca. 2014. *Curbing clientelism in Argentina: Politics, poverty, and social policy*. Cambridge University Press.
- Zhang, Weijia, Jiuyong Li and Lin Liu. 2021. “A unified survey of treatment effect heterogeneity modelling and uplift modelling.” *ACM Computing Surveys (CSUR)* 54(8):1–36.

Appendix Material for “The Demand for Elections under Autocracy: Regime Approval and the Cancellation of Local Elections in Russia”

Table of Contents

A	Alternative survey sample	APP-1
B	Additional alternative cancellation dates	APP-2
C	Negative weights and heterogeneity-robust estimation	APP-3
D	Additional placebo tests	APP-5
E	Multilevel modeling	APP-7

A Alternative survey sample

Table A1: Robustness to excluding FOM GeoRating surveys

	DV: Putin approval				
	(1)	(2)	(3)	(4)	(5)
Elections cancelled	−0.034** (0.012)	−0.075** (0.016)	−0.026+ (0.014)	0.017 (0.034)	−0.014 (0.034)
Avg ENC of govr elections pre-2004				−0.826 (33 749.390)	
Avg govr ENC pre-2004 x Elections cancelled				−0.022+ (0.011)	
Regional electoral competitiveness 1990s					0.020* (0.008)
Reg electoral compet 1990s x Elections cancelled					−0.006 (0.012)
Age x Elections cancelled		0.001** (0.000)			
Educ x Elections cancelled			−0.003 (0.003)		
Age	−0.001** (0.000)	−0.001** (0.000)	−0.001** (0.000)	−0.002** (0.000)	−0.001** (0.000)
Educational attainment	−0.002 (0.001)	−0.002 (0.001)	−0.001 (0.002)	−0.008** (0.002)	−0.002 (0.001)
Male	−0.067** (0.004)	−0.067** (0.004)	−0.067** (0.004)	−0.093** (0.005)	−0.067** (0.004)
NObs	174 817	174 817	174 817	68 603	174 817
R2 Adj.	0.075	0.075	0.075	0.054	0.075

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

OLS regression coefficients; SEs clustered by city.

All models include fixed effects for city, month, and survey firm.

B Additional alternative cancellation dates

Table A2 tests additional alternative municipal election cancellation dates resulting from our data-collection process. In Model 1, we include those few cities that used a dual-executive model with an appointed city manager and directly elected head of city. As described in the Election Cancellation section above, we drop these cities from most analyses. When we include them here in Model 1, our results are slightly weakened, which is to be expected given the blurred lines of accountability resulting from this executive structure, but remain in line with our main results. Model 2 in Table 4 adds a fixed effect for the structure of the Putin approval question that serves as the dependent variable. As a few waves of the FOM Georating survey asked a version of this question that incorporates change in opinion, we include a fixed effect for this difference. Our results remain unchanged.

Table A2: Additional robustness checks

	DV: Putin approval			
	Dual-exec included	Question type	No regional dates	Regional date only
	(1)	(2)	(3)	(4)
Elections cancelled	−0.020** (0.006)	−0.018** (0.006)	−0.022** (0.007)	−0.022** (0.006)
Age	−0.001** (0.000)	−0.002** (0.000)	−0.001** (0.000)	−0.001** (0.000)
Educational attainment	−0.005** (0.001)	−0.006** (0.001)	−0.004** (0.001)	−0.005** (0.001)
Male	−0.064** (0.002)	−0.063** (0.002)	−0.065** (0.002)	−0.065** (0.002)
NObs	424 321	299 447	372 375	402 070
R2 Adj.	0.073	0.071	0.074	0.073

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

OLS regression coefficients; SEs clustered by city.

All models include fixed effects for city, month, and survey firm.

Model 3 addresses another difference in the coding of regional vs. local cancellation dates after the 2014 reform. For those cities where we have both a region-level election cancellation date and a local-level implementation date, we use the local date instead of the regional date in Model 3. Model 4 implements the reverse—using the regional date rather than local date,

when we have both dates available. Our results lose some statistical significance in Model 4 but remain substantively unchanged.

C Negative weights and heterogeneity-robust estimation

In this section, we explore the heterogeneity-robust estimator from Borusyak, Jaravel and Spies (2021), as described in Section 7. As noted there, since this method does not allow for interaction terms, we split the sample at the median values of key modifying variables (e.g., the region’s democratic history). Coefficient estimates are highly consistent with the estimates in our main regressions, although, given the imputation and sometimes reduced sample size in these regressions, the standard errors are larger. One possible exception to this pattern is the results on the individual-level interactions. Here the results are significantly different from our main results: we observe a moderating effect for education, such that more educated voters are more likely to reduce their support for Putin after cancellation, but the moderating effect of age is actually the reverse of that observed in the main regressions. This suggests that these subgroup-level findings should be interpreted with special caution.

Table A3: Heterogeneity-Robust Estimator: Baseline

VARIABLES	(1) Baseline	(2) Low ENC	(3) High ENC	(4) Low ENC Avg.	(5) High ENC Avg.	(6) Uncompet. Elec	(7) Compet. Elec
Elections Cancelled	-0.017* (0.009)	-0.004 (0.010)	-0.027** (0.008)	-0.015* (0.009)	-0.019* (0.011)	-0.011 (0.012)	-0.032** (0.008)
Observations	397,745	141,967	232,115	186,855	185,049	316,952	80,793

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

SEs clustered by city.

All models include demographic controls

Dropping negative weights

As discussed elsewhere in Section 7, we also show that our results are robust when dropping the final years in our dataset (≥ 2016), which is where the negative weights are most likely to occur. This approach, from Jakiela (2021), helps to eliminate the vast majority of observations with negative weights that could cause bias due to heterogeneous treatment

Table A4: Dropping observations likely to have negative weights (Jakiela 2021)

	DV: Putin approval				
	Pre-2015	Pre-2016	Pre-2017	Max 3yr treatment	Max 5yr treatment
	(1)	(2)	(3)	(4)	(5)
Elections cancelled	−0.028** (0.008)	−0.027** (0.007)	−0.024** (0.007)	−0.021** (0.007)	−0.022** (0.007)
Age	−0.002** (0.000)	−0.002** (0.000)	−0.002** (0.000)	−0.002** (0.000)	−0.002** (0.000)
Educational attainment	−0.006** (0.001)	−0.005** (0.001)	−0.005** (0.001)	−0.005** (0.001)	−0.005** (0.001)
Male	−0.071** (0.002)	−0.069** (0.002)	−0.065** (0.002)	−0.066** (0.002)	−0.066** (0.002)
NObs	338 724	362 234	386 877	342 309	364 557
R2 Adj.	0.075	0.077	0.075	0.075	0.075

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

OLS regression coefficients; SEs clustered by city.


All models include fixed effects for city, month, and survey firm.

effects in our staggered TWFE design.

Models 1–3 in Table A4 show our main results (Table 1, Model 1) but dropping all observations from 2015 on, 2016 on, and 2017 on, respectively. Models 4 and 5 drop observations that have been treated for more than 3 years and 5 years, respectively. Our results are robust to all of these procedures for eliminating observations likely to have negative weights.

In Figure A5, we show loess and OLS regression lines for residualized treatment and outcome. As shown by Jakiela (2021), the presence of potentially-troubling treatment effect heterogeneity would appear here as loess or OLS estimates that vary dramatically between different values of the residualized treatment in the untreated and treated groups, resulting in diverging slopes. We do not observe evidence of this heterogeneity, which provides additional evidence that such treatment effect heterogeneity is unlikely to be a substantial concern in our analyses.

Figure A5: Looking for treatment effect heterogeneity (Jakiela 2021)



D Additional placebo tests

Table A5: Placebo Test

VARIABLES	(1)
	Putin Approval
Male	-0.067*** (0.002)
Age	-0.002*** (0.000)
Education	-0.005*** (0.001)
Treated in Future	-0.005 (0.003)
Constant	0.930*** (0.011)
Observations	289,282
R-squared	0.076

Cell entries are OLS Regression Coefficients. Standard errors clustered by city.
All models include fixed effects for city, calendar month, and survey firm.
* p<.1; **p<.05; ***p<.01

In Table A6, we provide results from another type of placebo test. Here, we look for an effect of our main treatment variable, Elections cancelled, on respondents in *suburban areas* around the cities in our dataset. For these analyses, we hand-coded all respondent localities within 50km of our main cities as being either nearby, obvious suburbs of those cities, localities more distant or more tenuously geographically connected to the cities, or distant and very questionably functioning as suburbs. We expect weaker or null results in all of these areas, since these residents are not being directly treated by election cancellation.

Table A6: Placebo tests with suburbs

	DV: Putin approval		
	Near suburbs	Medium suburbs	Med and distant suburbs
	(1)	(2)	(3)
Elections cancelled	−0.005 (0.014)	0.010 (0.023)	0.003 (0.018)
Age	−0.001** (0.000)	−0.001** (0.000)	−0.001** (0.000)
Educational attainment	−0.006** (0.002)	−0.004* (0.002)	−0.005* (0.002)
Male	−0.064** (0.004)	−0.076** (0.004)	−0.073** (0.003)
NObs	121 076	117 576	132 881
R2 Adj.	0.073	0.077	0.076

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

OLS regression coefficients; SEs clustered by city.

All models include fixed effects for city, month, and survey firm.

We find no significant effect of election cancellation on these localities, whether we consider nearby suburban areas only (Model 1), suburban areas at a medium distance from our cancellation cities (Model 2), or medium and distant (or likely independent) localities (Model 3). This indicates that the effect of cancellation is truly geographically limited to the places where cancellation is taking place—Russians’ attitudes are affected by the loss of their own ability to vote for their mayors, not simply observing this occurring in nearby areas.

E Multilevel modeling

In this section, we reproduce our main models from Tables 1–3, principally Model 1 from Table 1, shown in Appendix Table A7 as Model 1, using non-nested two-way linear multilevel models instead of two-way fixed effects models. Here, city, month, and survey firm are represented as random effects (varying intercepts). These results strongly support our primary findings, albeit with the expected much stronger statistical significance.

Table A7: Effect of Mayoral Election Cancellation on Regime Approval, linear multilevel models

	DV: Putin approval					
	(1)	(2)	(3)	(4)	(5)	(6)
Elections cancelled	−0.020** (0.002)	−0.075** (0.004)	−0.013** (0.004)	0.013* (0.007)	0.028** (0.010)	0.017+ (0.009)
ENC of most recent govr election pre-2004				0.003 (0.002)		
ENC last govr elec pre-2004 x Elections cancelled				−0.010** (0.002)		
Avg ENC of govr elections pre-2004					0.003 (0.003)	
Avg govr ENC pre-2004 x Elections cancelled					−0.015** (0.003)	
Regional electoral competitiveness 1990s						0.006 (0.004)
Reg electoral compet 1990s x Elections cancelled						−0.012** (0.003)
Age x Elections cancelled		0.001** (0.000)				
Educ x Elections cancelled			−0.003+ (0.001)			
Age	−0.001** (0.000)	−0.002** (0.000)	−0.001** (0.000)	−0.002** (0.000)	−0.002** (0.000)	−0.001** (0.000)
Educational attainment	−0.005** (0.001)	−0.005** (0.001)	−0.004** (0.001)	−0.008** (0.001)	−0.008** (0.001)	−0.005** (0.001)
Male	−0.065** (0.001)	−0.065** (0.001)	−0.065** (0.001)	−0.070** (0.001)	−0.070** (0.001)	−0.065** (0.001)
NObs	401 343	401 343	401 343	295 129	295 129	401 343

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Linear multilevel regression coefficients.

All models include random effects for city, month, and survey firm.