Jul 2019 No.433

Patronage and Election Fraud: Insights from Russia's Governors 2000–2012

Christoph Koenig

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

Patronage and Election Fraud: Insights from Russia's Governors 2000–2012

Christoph Koenig

University of Bristol and CAGE Warwick *

July 19, 2019

Abstract

Theory and empirics suggest that patronage fosters election fraud. But why does fraud vary within autocracies where patronage's incentives to manipulate should be uniformly high? In this paper, I explore whether information asymmetries can explain this phenomenon. I study the introduction of a patronage system which allowed Russia's president to discretionarily appoint all 89 regional governors. After December 2004, all national elections were organized by governors facing removal but, crucially, only some were actually patronage-appointed with lower need to signal their qualities. I estimate the effect of the reform's introduction and its staggered implementation on a new and verified regional fraud indicator for 7 national elections

*Department of Economics, University of Bristol. Email: Christoph.Koenig@bristol.ac.uk

This paper supersedes a previous version entitled "Competence vs. Loyalty: Political survival and electoral fraud in Russia's regions 2000–2012". I would like to thank Sascha O. Becker, Peter Buisseret, Mirko Draca, Jon Eguia, Ruben Enikolopov, Francesco Giovannoni, David Gomtsyan, Sergei Guriev, Mark Harrison, Giovanni Mastrobuoni, Sharun Mukand, Maria Petrova, Kirill Pogorelskiy, Fabian Waldinger and Yanos Zylberberg for important discussions and useful insights and Maxim Krukov of GOLOS and Graeme Robertson for sharing their data. The paper further benefited from seminars and participants at Warwick, RES Bristol and RES Warwick. Financial support by the ESRC DTC Warwick is gratefully acknowledged. All remaining errors are mine.

from 2000–2012. Results show that patronage increased overall levels of rigging but less so with patronage-appointed, connected governors. Appointments had no effect on actual election results and regional economic performance, which makes reduced uncertainty about governors' loyalty the most plausible explanation.

In Bashkortostan, President Rakhimov delivered 92 per cent of the vote for Putin; Dagestan, 94 per cent; Kabardino-Balkaria, 96; Ingushetia, 98. Were they running a competition?

Anna Politkovskaya, *"A Russian Diary"*, 2007

I think you need to pay attention to those areas where our people have denied United Russia serious trust. Not because it is a tragedy, but because it is a signal for the authorities.

Dimitri Medvedev, 2011

1 Introduction

Election fraud is an almost proverbial and extensively studied feature of imperfect democracies and dictatorships. Scholars have been investigating the techniques and motivations behind the manipulation of elections since more than a century by now.[1] Until recently, the literature predominantly understood electoral fraud as the centralized action of a non-democratic regime or a single dictator. This view has been challenged by a growing empirical and theoretical body of work on the local dimensions of rigging and, crucial to this paper, on the role of public officials in carrying out manipulation.[2] As a result, there is now extensive evidence from across the globe that bureaucrats' involvement matters for electoral accountability (Ziblatt, 2009; Folke et al., 2011; Martinez-Bravo, 2014; Callen and Long, 2015;

[1] According to a review by Lehoucq (2003), the earliest work covering electoral fraud is Seymour's study on England and Wales (1915). One of the topic's first quantitative accounts is provided by Cox and Kousser (1981).

[2] Mares and Young (2016) provide an extensive overview of the quantitative work on extralegal voter mobilisation including those by local electoral officials. More closely related empirical studies are discussed at the end of this section. A short summary of the theoretical literature on election fraud is provided in Gehlbach et al. (2016).

Martinez-Bravo et al., 2017; Cantú, 2019). Yet, we still know comparatively little about what actually motivates officials to engage in the manipulation of elections and which role institutional settings play in their decision-making process. Giving an answer to these questions is important since it can inform international actors which institutional changes to lobby for in order to increase electoral accountability in a given context.

One such institutional feature of particular interest is *patronage*. If bureaucrats in charge of organizing an election are themselves discretionarily selected by the state or a dictator, there are doubtlessly clear incentives for them to ensure that their appointer (or her candidate) also wins the ballot. According to theoretical work by Martinez-Bravo (2014) and Rundlett and Svolik (2016), the main incentives are avoiding reprisal by a disgruntled dictator (or a victorious opposition) and rewards such as keeping office or other material gains. Both models also predict that the motivation of appointed bureaucrats to manipulate will increase in the probability of incumbent survival and should thus lead to uniformly high fraud levels in ballots with near-certain outcomes. The considerable sub-national variation in manipulation observed in some autocracies with non-competitive elections can thus not be fully explained (Myagkov et al., 2009).[3] One explanation for this phenomenon could be the following: first, non-competitive elections primarily matter for *"more*

[3] Rundlett and Svolik (2016) explain this phenomenon through agents' varying beliefs in the incumbent's survival probability which they infer from electoral support in their vicinity. This is supported by the abnormal occurrence of United Russia vote shares ending in 0 or 5 where the party fared best in the 2011 and 2012 elections. Despite the persuasiveness of their general argument, there are two important caveats to applying this to the case of Russia: first, as argued in this article and others (Reuter, 2011; Little, 2015), the victory was expected across the country in those elections which should have resulted in uniformly high levels of fraud. Secondly, it is unclear whether high results which have been trimmed to match a specific share can be interpreted as a signal of genuine popularity or are just a by-product of the manipulation itself.

than winning" (Simpser, 2013) such as generating information about local officials. Second, patronage and career concerns induce bureaucrats to use fraud in order to signal their *loyalty* and *competence* to the dictator (Gandhi and Lust-Okar, 2009; Egorov and Sonin, 2011; Zakharov, 2016). Finally, even under patronage, local officials may be quite heterogenous in their characteristics and the need to use fraud as a signal.[4] Unfortunately, there is so far no evidence how patronage and bureaucrat heterogeneity affect election fraud under autocratic rule. From an empirical viewpoint, this is also particularly challenging since changes in autocracies' organisational structures are very rare and often coincide with large-scale purges which presumably result in near-uniform levels of loyalty and competence among public officials.

In this paper, I empirically study the effect of patronage on election fraud in a non-competitive setting by exploiting the abolition of Russia's governor elections in late 2004. Until then, the leaders of Russia's 89 regions who are in charge of organizing national elections at the sub-national level were elected by the population of their respective regions.[5] Following a major terrorist attack, however, President Vladimir Putin signed a law which ended governors' appointment via elections in an attempt to strengthen federal authority. Thereafter, regional leaders had to be re-appointed by the president in order to stay in office after their term but could also be replaced without any legal barriers (Hill, 2012). This newly introduced patron-

[4] Several other scholars have also discussed fraud as a signalling tool in the context of non-competitive elections (Magaloni, 2006; Simpser, 2013; Little, 2015; Rozenas, 2016). While my paper shares important insights from these models, I focus on the role of the bureaucracy in the manipulation process rather than that of the dictator. My notion of fraud as a signal also assumes that it can be observed or inferred by the central administration which is shared by parts of this literature (Simpser, 2013; Little, 2015).

[5] Technically, those elections are called *federal elections*. To avoid confusion, I refer to them as *national elections* instead. The word *federal* is still used to describe several other Russian institutions.

age system severely altered officials' motivation to please the central government and deliver high results in national elections. Crucially, the new law was not accompanied by a uniform purge but rather implemented in a staggered fashion leading to a transition period during which *connected* (i.e. screened and hand-picked) and *non-connected* (i.e. inherited and re-appointed) governors coexisted and organized 4 national ballots until the year 2012. This unique setup offers an opportunity to separate the impact of patronage's *introduction* which induces agents to signal loyalty and competence (*incentive effect*) and its actual *implementation* which reduces again the need to signal (*information effect*). To the best of my knowledge, these concurring effects of patronage have not been addressed so far in formal or empirical work on political economy.

My paper also tries to make a methodological advancement in identifying a valid and reliable measure of fraud in the studied context. The chosen indicator is a version of the turnout-vote share correlation (henceforth TVSC) as outlined in Myagkov et al. (2009) which is first constructed at the sub-regional level and then aggregated up. The idea behind the TVSC is that, absent fraud, the correlation between the number of total votes and those for a particular candidate across voting stations in sufficiently small and homogeneous areas should be equal to the candidate's true level of support. Turnout-inflating fraud such as ballot stuffing increases this correlation to unreasonable levels or, in the extreme, above 1. To construct this measure, I use unique voting station level for all 7 national elections – parliamentary and presidential – held in the Russian Federation from 2000 to 2012. Parts of this data have been used in previous studies on Russia discussed further below using a variety of indicators for rigged elections. I improve on this pre-existing work in two ways: first, I use additional data on fraud reports from a crowd-sourcing project to benchmark existing methods against each other and identify an abnormally high TVSC as the most credible measure to detect actual manipulation. Second, I use this measure to quantify the extent of rigging across both regions *and* elections.

I then use this indicator in a Differences-in-Differences (DID) framework to study how the introduction and implementation of patronage differentially affected fraud levels across space and time.

My empirical results show that, in line with the information effect, replacing a non-connected governor with a connected one reduces the share of suspicious votes by 6 percentage points on average. Given a baseline mean of 20 percentage points, this is a substantial effect. This finding is also noteworthy since aggregate data show that fraud levels were actually *increasing* in the aftermath of the December 2004 reform which can be interpreted as evidence for the incentive effect of patronage. These baseline findings are robust to a number of additional tests and checks for the identifying assumptions including flexibly controlling for pre-reform institutional quality and economic performance. In spite of this, one may still doubt that the identified effect can be solely attributed to changes in the need to signal. For instance, the government could have mainly dismissed notorious cheaters or governors with low track records on popularity and regional development and a heavy reliance on fraud to meet *target* election outcomes (White and Feklyunina, 2011). I tackle this concern by using an event-study design which disaggregates the effect of receiving a connected governor into the periods immediately prior and after the actual replacement. Using this methodology, my results indicate that fraud levels, incumbent vote shares as well as GDP growth and unemployment rates were not statistically different prior to replacement.

While these findings are consistent with lower information asymmetries, I also explore alternative explanations by looking into the effect of connected governors on other electoral, political and economic outcomes as well as governor characteristics. In line with a purely symbolic signalling purpose of fraud, I do not find any significant effects on turnout and vote shares of the incumbent and other parties. There is also no evidence on changes in levels of oppression, governor popularity and additional economic outcomes. Connected governors are thus unlikely to have been systemat-

ically selected based on their competence as policy-makers. Additional results on leader characteristics show that the reform was largely used to remove members of the unpopular, old Communist elite and replace them with young outsiders who recently joined the United Russia party and did not have a lot of political experience. A simultaneous drop in rigging capacity is, however, unlikely to drive my results since there is no evidence for learning or catching up by connected governors. The most plausible explanation for my findings is thus that patronage's implementation must have reduced fraud because it removed uncertainty about governors' loyalty.

This paper contributes to three areas in the political science literature. The first is the literature on election forensics (Myagkov et al., 2009; Shpilkin, 2009; Beber and Scacco, 2012; Rozenas, 2017). The TVSC method used in this study has been adopted by a large amount of papers studying the case of Russia (Filippov and Ordeshook, 1997; Myagkov et al., 2009; Lukinova et al., 2011; Enikolopov et al., 2013). Other studies on the same country have focussed on abnormal digit distributions in absolute (Kalinin and Mebane Jr., 2011; Mebane Jr., 2013; Skovoroda and Lankina, 2016) or relative counts (Klimek et al., 2012; Kobak et al., 2016; Rozenas, 2017).[6] My paper adds to this work by leveraging actual fraud reports to provide further evidence on the TVSC's reliability and by proposing a metric which can be used for comparisons across areas and elections.[7] To the best of my knowledge, my study also provides the first panel data analysis of election fraud.[8]

Second, this article ties in with empirical and theoretical work on the nexus between state bureaucrats, patronage and political accountability (Ziblatt, 2009; Folke

[6] Few studies also regard high turnout per se as an indication of fraud (Bader and van Ham, 2014; Moser and White, 2016).

[7] This is not to generally dismiss the other methods but rather to say that they perform less well in measuring visible acts of fraud at the sub-national level.

[8] The closest to such an analysis is Ziblatt (2009) who uses pooled cross-sectional data for Imperial Germany without accounting for area-specific fixed-effects.

et al., 2011; Martinez-Bravo et al., 2017; Cantú, 2019). Studies on the relationship between patronage and fraud in nascent democracies with competitive elections are particularly closely related (Martinez-Bravo, 2014; Callen and Long, 2015). Yet, I depart from this literature by shifting the focus to authoritarian, non-competitive elections and the role of information asymmetries. This paper also complements formal models on patronage and fraud by Martinez-Bravo (2014) and Rundlett and Svolik (2016) in exploring whether bureaucrat heterogeneity can explain varying degrees of fraud within the same country. My empirical results reinforce the view that the positive effect of patronage on fraud is strongly driven by officials' need to overcome uncertainty about their qualities. The actual implementation of patronage by appointing pre-screened officials, can thus lead to a comparative reduction in rigging.

Finally, my paper also directly links to recent work on the role of clientelism in Russian elections (Frye et al., 2014; Reuter and Szakonyi, 2019). A part of this literature is also specifically concerned with the December 2004 law change and the way in which Russian governors were selected in the reform's aftermath. Work by Reuter and Robertson (2012), Moraski and Reisinger (2013) and Rochlitz (2016) suggests that loyalty rather than competence was the main appointment criterion of the government. My findings arrive at similar conclusions in the sense that connected leaders do not perform better economically but still engaged in less fraud *as if* they had lower need to signal their loyalty to the regime. However, I do not find that election results or economic performance were driving governor replacement or were affected by it. Closest to this article is a working paper by Kalinin and Mebane Jr. (2011) which anticipates the idea that Russia's governors could be using fraud as a signalling tool to assure political survival. My study differs from theirs by focussing on the period around the abolishment of governor elections 2000 to 2012 and assuming a simple top-down relationship between the central administration and governors. Using a verified fraud indicator and a DID

approach, my results also cannot replicate their finding that connected governors engage in more rigging but rather imply the exact opposite.

The paper starts with a brief description of Russia's political and electoral system in Section 2. After introducing the data in Section 3, I present the main fraud indicator with a special emphasis on its reliability and comparison with other potential alternatives in Section 4. Section 5 starts by discussing the DID approach and then presents the baseline results, robustness checks and the effect of connected governors on other outcomes. Section 6 concludes.

2 Institutional background

2.1 The organisation of national elections

The focus of this paper is on *national* elections for the Russian president and the *Duma* parliament in which regional leaders are not participating themselves but rather organize ballots where the incumbent party or its candidate are running. Presidential elections are carried out by simple majority which proceeds to a second round if no candidate attains 50% in the first. In order to rule out any strategic interactions, this paper focusses exclusively on the results in the first round. For the parliamentary elections, half of the 450 seats are allocated proportionally through regional party lists and the other half via 225 majoritarian single-member districts. This system was suspended from 2005 to 2014 in favour of a pure proportional system (Moraski, 2012). In order to establish comparability with the 2007 and 2011 results, I use only the proportional votes for the 2003 ballot. The actual organisation of national elections in Russia strongly corresponds to its administrative divisions, both horizontally and vertically. This aspect of the institutional context is crucial in order to understand at what stage manipulation through governors could technically occur.

The highest authority is the Central Election Commission (CEC), a permanent body whose members are nominated by the president and the two chambers of the legislative – the *Duma* parliament and the *Council of Regions*. The CEC's main tasks are the coordination of the 89 Regional Election Commissions (RECs) and the organisation of the national elections. An REC fulfils the same role as the CEC at the regional level and coordinates the territorial election commissions (TECs). Unlike the CEC, however, the members of an REC are appointed by recommendation of the federal government which gives a region's leader very limited power in influencing results at this level of aggregation. The next administrative level below the region is the district (*rayon*). Like regions, the districts can vary considerably in size and population but unlike the former there may be several TECs within the same rayon (e.g. in larger cities or former *closed towns*). The members of the roughly 3,000 TECs are permanent delegates by the regional executive, legislative and parties. At the lowest level, about 95,000 precinct election commissions (PECs) are responsible for the local organisation of all elections. They are formed only one month before the ballot and through nomination by the electorate and it remains unclear to what extent regional authorities can exert control over their composition (OSCE, 2000, 2004a,b, 2012b,a). Their ad-hoc nature and sheer size in numbers, however, requires more in-depth knowledge and makes them an unlikely target of direct intervention by a region's leader.

In sum, governors' main power over outcomes in national elections lies in controlling the TECs. The fraud measure presented below therefore implicitly assumes that governors have the power to induce TECs to engage in rigging to varying degrees. Visualisations of fraud across regions and districts shown below are supportive of this assumption.

2.2 The abolition of governor elections in December 2004

As in most federal states, each Russian region has its own legislative and executive, the latter being led by a governor.[9] The members of the regional assemblies are chosen in local elections. Similar to the national level, these assemblies are dominated by the executive making governors the most important political institution in the regions. From 1996 until December 2004 these were chosen in local elections. However, in 2004 the constitution was changed in favour of a direct appointment of governors by the president (Slider, 2012). This drastic constitutional change was decided in the aftermath of the *Beslan Massacre*. On the 1st of September 2004, a multinational Islamist terror squad took over 1000 hostages in a school in Beslan, a town in the Republic North Ossetia-Alania close to Chechnya. When security forces attempted to free the hostages, more than 300 people were killed. This national tragedy demonstrated the increased power of Chechen insurgents and their allies but also showed the lack of coordination between federal and regional authorities.[10] Very soon after the attacks, president Vladimir Putin initiated a law which re-introduced the appointment of governors. The draft passed both chambers of the Federal Assembly and came into effect in December 2004.

What may seem puzzling is that the new law was accepted by both the population and the governors without any major opposition. Goode (2007), who analysed the parliamentary debates in late 2004, concludes that a combination of rally-around-the-flag effects and an appeal to Soviet legacies made it impossible to reject the new law. Additionally, being independent of the local electorate and depending only on the central executive was in the interest of many governors. Only 7 years after its introduction, however, the central government concluded that the abolition had

[9] Many regions use different titles such as *President*, *Head of the Republic*, or *Head of the administration*. For the sake of simplicity, I refer to all these in this paper simply as *governors*.

[10] This was apparent even though a lot of information about the Beslan hostage crisis was actually withheld from the Russian public (Haraszti, 2004).

been a mistake which led to public discontent and eventually re-introduced governor elections in 2012. The new law was accepted at first reading in late February 2012 just before the last presidential election in my sample and became effective in May 2012. Importantly for this study, the exact timing of when the law would become effective remained unclear to most of the elite even after it passed the parliament with one representative even calling for a six-year postponement (Teague, 2014).

3 Data

3.1 Elections and fraud reports

All voting data used in this project was provided by the organisation GOLOS, an independent Russian NGO concerned with election monitoring. The data is compiled from the official CEC website and covers PEC-level results for the entire country and all presidential and parliamentary election since 2000.[11] For each election, I assigned each PEC to a TEC over which I calculate the measures of electoral fraud explained in Section 4.1. The main variables of interest are absolute number of votes for specific candidates and parties, size of the electorate as well as valid and invalid votes in order to calculate the turnout for each election-precinct cell. Vote shares are calculated for the incumbent party United Russia (including presidential candidates Putin and Medvedev) and three main opposition parties. These are the communist *KPRF*, the ultra-national *LDPR* and the democratic *Yabloko* party.[12]

[11] The two exceptions are the Republic of Sakha and the Nenets Autonomous District for which no PEC-level data is provided in 2000.

[12] The presidential candidates for *KPRF* are Gennady Zyuganov (2000, 2008, 2012) and Nikolay Kharitonov (2004) and for the *LDPR* Vladimir Zhirinovsky (2000, 2008, 2012) and Oleg Malyshkin (2004). The *Yabloko* party did not contest in all presidential elections either due to boycott (2004) or because of bureaucratic barriers (2008 and 2012). While party leader Grigory Yavlinsky ran himself for office in 2000, I use the votes for Irina Khakamada (2004), Andrey

From GOLOS I also obtained direct indicators of election rigging. During the 2011 and 2012 elections in my sample, the association ran the *Karta Narusheniy* (map of violations) project which provided a platform for citizens to anonymously report incidents of fraud in national, regional, and local elections and send detailed reports of observed electoral law violations via phone, internet, and text message. The users could also give information whether the action happened during campaigning or on election day and which type of fraud had taken place.[13] Of particular interest for this research project are the categories *distortion of results* and *exclusion of observers, committee members, or media*. Lastly, the observers could also provide information on the location where the action was witnessed. I used this information to match reports to their respective district after checking and correcting for informational consistency (about 5,600 out of 7,100 for 2011 (= 78.9%) and 3,700 out of 4,800 (= 77.2%) for 2012).[14]

3.2 Governor characteristics, other variables and sample

In order to accurately define this, I first compiled a panel dataset of all governors serving between 2000 and 2012 with the start and end date of their terms. This information was taken from the website *rulers.org* and cross-checked with *wikipedia.ru*.[15] The switch from a *non-connected* governor to a *connected* governor is coded in the main treatment variable $ConnectedGovernor_{it}$. It has value 1 if

Bogdanov (2008) and Mikhail Prokhorov (2012) in the remaining elections. This choice of candidates is arguably not perfect but reflects the main alternative to the three other parties in the respective presidential elections.

[13] A detailed description of the *Karta Narusheniy* data is provided in Bader (2013a).

[14] I am only aware of two other studies who used this type of data – each, however, with a slightly different topical focus (Bader, 2013a; Skovoroda and Lankina, 2016)

[15] Acting governors, for which often only very limited information exists, were ignored for simplicity and it was assumed that their successors took over office immediately from the last governors.

the governor ruling region i at election or year t started his term after the 12^{th} of December 2004 – and therefore had to be selected by the Russian president – and 0 otherwise. From *wikipedia.ru*, I also extracted the exact geo-location for each governors' place of birth. I then matched all governors in my sample to the *ICSID Russian Governors bios* database which provides me with important information on job backgrounds and their pre-term place of living as well as their party affiliations with the Communist Party and United Russia. I refine and correct the entry dates into United Russia for all governors serving during the interim period 2003 to 2007 using data reported in Reuter (2010).

An obvious way to evaluate governors' performance would be indicators of economic prosperity. I therefore use the *ICSID Social and economic database* which contains panel data from the Russian Federal State Statistics Service (GKS) on regional GDP and income per capita, average wages, unemployment and poverty rates as well as regional consumer price indices. In addition I also use data on the amount of yearly federal financial transfers to each region extracted from the official budget plans 1999 to 2012. Another set of indicators is related to political outcomes. First, I use democracy ratings by region reported in Petrov and Titkov (2013). These scores are based on expert assessments and always given over a specified time horizon. The ranking of 1991-2001 is used as a control variable and baseline measure of institutional quality at the start of the sample period. The ones from 1999-2003 to 2006-2010 are converted into a panel dataset using the last year as a time variable. In addition, I use yearly regional estimates of governor popularity which comes originally from the *Public Opinion Foundation* and was also used in Reuter and Robertson (2012). This yearly measure is the % of survey respondents in a Russian regions answering positively to the question whether their governor is doing a good job or not. The survey had limited coverage across Russia, omitting almost the entire North Caucasus, and could only be used for 60 regions in my sam-

ple.[16] Finally, I extract information from maps provided by the *Glasnost Defense Foundation* to construct a dummy whether a region had relatively free media or not for the years 2000, 2001, 2002, 2006, 2008 and 2010.[17]

The final panel dataset covers 77 out Russia's 89 regions over all 7 national elections during the period 2000 to 2012 apart from Republic of Sakha and the Nenets Autonomous District where no detailed election data is available for 2000. During the sample period there were 5 mergers between 2 or 3 regions which reduced the initial amount of 89 subdivisions to 83. Such mergers are likely to fundamentally change the power structure of a governor and make it difficult to compare the new and old units especially since all mergers took place after the abolishment of governor elections. For this reason I excluded all 11 regions affected by a merger. In addition to that, Chechnya lacked information on economic outcomes until the mid 2000s and was hence dropped from the dataset. The standard sample size is thus 539 for electoral outcomes and 924 for yearly variables with full availability from 2000 to 2012. As noted above, particular variables may also vary additionally in their regional and yearly coverage. Summary statistics of the final sample are reported in the Appendix in Tables A1 and A2.

[16] I would like to thank Graeme Robertson for kindly sharing this data with me. Occasional missing values were linearly interpolated.

[17] The main reason for using a dummy is that the maps change their classification scale over time, i.e. from 2000 to 2002 it ranges between *High*, *Medium* and *Low* and from 2006 to 2010 between *Free*, *Relatively Free*, *Relatively Unfree* and *Unfree*. My *FreeMedia* dummy variable takes a value of 1 if a region falls into the *High*, *Free* and *Relatively Free* categories and 0 otherwise if it was included in the ranking. Occasional missing values were linearly interpolated. The series for the Chukotka and Yamalo-Nenetsky Autonomous Districts, however, could not be included due to insufficient data available.

4 Measuring election fraud

4.1 The turnout/vote share correlation

The turnout/vote share indicator was first applied by Sobyanin for the 1993 constitutional referendum and is probably the most widely used tool for detecting election rigging in Russia. It is most suited for turnout-inflating cases of fraud, most importantly ballot-stuffing, and relies on the assumption that within a given entity and absent manipulation there should be no correlation between how many people vote and their choice across lower-tier areas. Figure 1a illustrates this with a brief example similar to Myagkov et al. (2009): assume there are 24 voting stations in an area with a given homogenous support of 75% for candidate i. Half of the stations are in high-turnout areas where 60% of the electorate casts their ballot, whereas the remaining ones only have a turnout of 40%. Absent fraud, a 1% higher turnout T is thus associated with an increase of 0.75% in votes for i out of the total electorate, V/E. A simple OLS regression of V/E on T thus yields a turnout/vote-share coefficient (henceforth $TVSC$) equal to the average support of the candidate – 0.75. This relation, however, would not hold in the case of ballot stuffing or other turnout inflating methods of manipulating the outcome as can be seen from Figure 1b. In this scenario, eight out of the twelve low-turnout voting stations see their turnout artificially increased to 80% with all additional votes going to candidate i. The TVSC thus changes from 0.75 to 1.07 which cannot be equal to i's natural support in that area anymore.[18]

Following Myagkov et al., one can distinguish between two degrees of suspicious results. In the first, the TVSC exceeds the candidates vote share in the respective

[18] One could imagine a scenario in which fraud is conducted in such a way that turnout and vote share are identical in each voting station. In this case there would be no variation and a TVSC could not be calculated. While this is theoretically possible, it is very difficult to implement in reality. I did not encounter such a case during the construction of my fraud data.

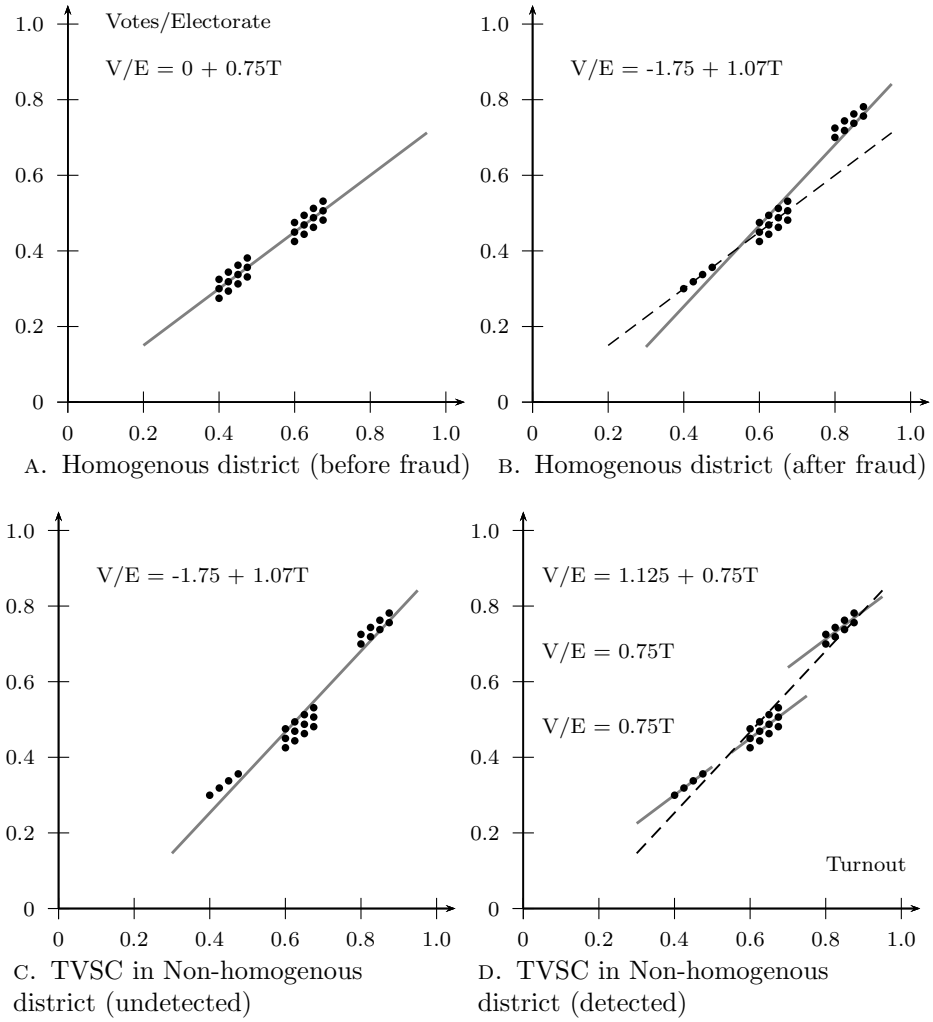

A. Homogenous district (before fraud)

B. Homogenous district (after fraud)

C. TVSC in Non-homogenous district (undetected)

D. TVSC in Non-homogenous district (detected)

FIGURE 1: EXAMPLE OF TVSC ABSENT BEFORE (TOP LEFT) AND AFTER FRAUD (TOP RIGHT). EXAMPLE OF BIASED TVSC IN THE CASE OF NON-HOMOGENEOUS AREAS (BOTTOM LEFT AND RIGHT).

area but is smaller than one. In the second, the TVSC is equal or even bigger than one. In scenario 1, the conclusion is ambiguous and will only be a safe detector if one can rule out that the favoured candidate – absent fraud – would have fared particularly well in lower-tier areas of high turnout – a premise which is quite difficult to check. Scenario 2, i.e. TVSC ≥ 1, appears to be a stronger indicator of manipulative turnout inflation, but it is also not immune to fallacies as depicted in Figures 1c and 1d. One may think of a region with uniform support of 75% for candidate i across its three districts with 4, 12, and 8 voting stations respectively. The voting stations, however, are not homogeneous since districts differ substantially in

their average turnout level. A regression of V/E on T across the whole region will thus suffer from aggregation bias and yield a TVSC larger than one even though fraud did not take place. While this error cannot be ruled out entirely, it can be mitigated by using highly disaggregated data and calculating the TVSC for reasonably homogeneous areas. Enikolopov et al. (2013), for instance, have shown that the random allocation of election observers across voting stations within Moscow's 125 TECs during the 2011 Duma election significantly decreased the TVSC calculated for the United Russia party.

Further estimates of election fraud in Russia's regions using the TVSC have been scarce so far and mostly relied on district aggregates (e.g. Myagkov et al., 2009). In these cases the assumption of homogeneity is more difficult to defend and the amount of districts/observations to calculate the TVSC may be very low. The availability of election results at the voting stations level since 2000 allows me to calculate TVSCs in each district of a given region and to construct new and more reliable estimates of rigging intensity across Russia's regions. As a new measure of regional fraud intensity, I propose the *share of votes from districts with a TVSC* ≥ 1. This indicator has the main advantage of using data from comparatively small areas but simultaneously providing a regional aggregate. It is supposed to capture the intensity of rigging rather than its mere existence which has been observed in virtually every region across the country and therefore does not provide much information. Another interpretation of the indicator is the percentage of votes (for either party/candidate, valid or invalid) likely to be affected by manipulation or, simply, the *share of suspicious votes*.

Figure 2 plots the district and regional measure of fraud described above for the four presidential elections in my sample.[19] The left-hand side figures show whether a district was classified as suspicious based on the TVSC≥ 1 criterion and visualizes

[19] The corresponding graphs for the parliamentary elections are not shown here for space constraints. They are displayed in Figure 5 in the Appendix.

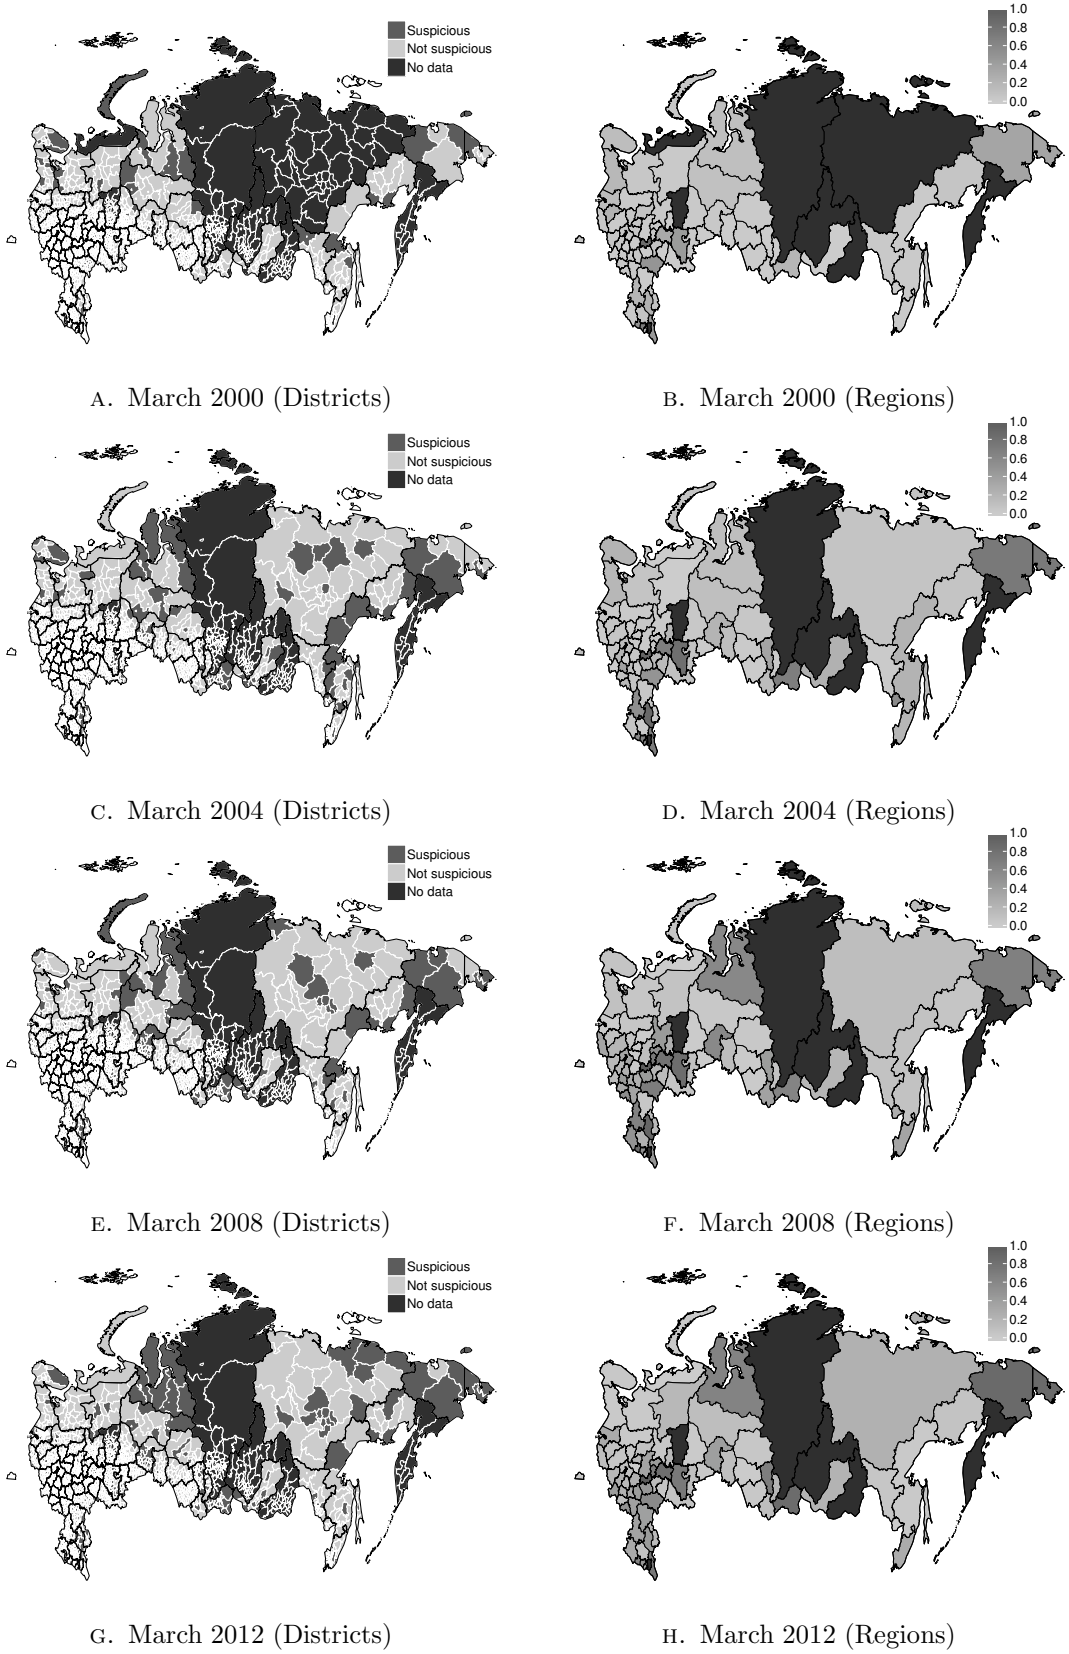

A. March 2000 (Districts)

B. March 2000 (Regions)

C. March 2004 (Districts)

D. March 2004 (Regions)

E. March 2008 (Districts)

F. March 2008 (Regions)

G. March 2012 (Districts)

H. March 2012 (Regions)

FIGURE 2: CLASSIFICATION AND AGGREGATION OF SUSPICIOUS VOTES BEFORE/AFTER THE ABOLITION OF GOVERNOR ELECTIONS IN DECEMBER 2004

Notes: Map of the division of the Russian Federation into regions and districts. Panels on the left are at the district levels and show their classification into *Suspicious* (dark grey) and *Not suspicious* (light grey) based on their voting station-level results according to the $TVSC \geq 1$ criterion. Panels on the right show the share of valid ballots in a specific region/election coming from *Suspicious* districts. Dark grey states are not included in the sample for this particular election.

how fraud evolved at the extensive margin. In 2000, suspicious districts are strongly concentrated in particular regions, most notably in the Republics of Tatarstan and Mordovia in Western Russia. In the 2004 election, suspicious votes start showing up in a number of formerly *clean* districts and entire regions exhibit almost uniformly suspicious ballot counts – Republics of Tuva in South-Central and Bashkortostan in the South-West. Other areas, particularly in the West and conflict-ridden South-Western Causcasus territory, are joining in during the election in 2008. The right-hand side figures display my preferred measure of fraud, the *share of suspicious votes*, based on the corresponding left-hand side classification. The graphs help understanding the variation in the main outcome variable and at the same time also show how rigging changed at the intensive margin within regions over the time period studied. While always present to some degree in few subdivisions, suspicious results seem to take off during the 2004 election from an average of 10.5 to 21.2%. During the 2008 election it kept on rising to about 28.1% and fell slightly to 26.4% in 2012.[20] The strong concentration suggests that election- as well as region-specific characteristics are important drivers of the prevalent fraudulent election outcomes since 2000.

4.2 Validity and reliability checks

The choice of the particular indicator used in the previous section to estimate regional fraud intensity may appear arbitrary at first. Before proceeding with this methodology, one should therefore first assure that it is an actually valid and reliable measure of rigging. I assess the performance of my fraud indicator $TVSC \geq 1$ by looking at its relation with reported incidents of fraud in the 2011 and 2012 elections provided by *GOLOS*. This is possible because roughly 80% of reports also

[20] The corresponding values for the Duma elections in 2003, 2007 and 2011 are 7.8%, 25.1% and 25.4%.

Table 1: Partial correlation between reported fraud and TVSC-based indicators

	(TVSC $\geq$ 1) = 1	(TVSC − Incumbent vote $\geq$ 0) = 1	TVSC	TVSC − Incumbent vote
	(1)	(2)	(3)	(4)
Any report = 1	0.052***	0.009	0.037	0.026
	(0.019)	(0.017)	(0.025)	(0.024)
Improper counting = 1	0.064***	0.018	0.063**	0.054*
	(0.023)	(0.020)	(0.030)	(0.029)
Exclusion of voters = 1	0.018	0.008	0.010	0.005
	(0.023)	(0.020)	(0.030)	(0.029)
Illegal campaigning = 1	−0.019	−0.021	−0.016	−0.027
	(0.021)	(0.019)	(0.028)	(0.027)
Observers excluded = 1	0.088***	0.036*	0.083***	0.078***
	(0.024)	(0.021)	(0.031)	(0.030)
Faulty ballot box = 1	0.040	−0.012	0.012	0.018
	(0.024)	(0.021)	(0.032)	(0.031)
Secrecy violated = 1	0.078***	0.002	0.010	−0.009
	(0.024)	(0.021)	(0.031)	(0.030)
Illegal voting = 1	0.050**	0.018	0.045*	0.029
	(0.021)	(0.018)	(0.027)	(0.026)
Other violations = 1	0.085***	−0.028	0.041	0.032
	(0.023)	(0.020)	(0.029)	(0.029)

Notes: Each cell represents the standardized coefficient from a single regression where the row denotes the *dependent* variable and the column denotes the *independent variable*. Observations are at the district-level. The sample period is 2011 to 2012 and includes 2 national elections (1 Presidential, 1 Parliamentary). Robust standard errors in parentheses: *p<0.1; **p<0.05; ***p<0.01. Included control variables are District FEs, Election FEs and Region×Elections FEs.

allow identifying the district where election irregularities were witnessed. In Table 1, I therefore estimate the correlation between dummy variables for having any particular fraud report in a specific district and another dummy indicating whether a district reported a TVSC in excess of 1. Apart from this validity check , I also benchmark the performance of the $TVSC \geq 1$ dummy against that of some related alternative measures: the first one is another dummy variable for the TVSC exceeding the candidate's vote share which was also mentioned in Myagkov et al. (2009). The remaining two candidate measures are actual TVSC and TVSC in excess of the incumbent vote share.

Table 1 presents 36 (standardized) coefficients from regressing a dummy for each 9 types of reported election irregularity on my preferred indicator, the $TVSC \geq 1$ dummy, and the 3 variants mentioned above. Observations are at the district-election level for 2011 and 2012 and results are conditional on district, election, and region by election dummies. The results indicate that the dummy for TVSC exceeding 1 is by far the most reliable correlate of having actual fraud reports. Most

notably, it is strongly and significantly correlated with reports on those irregularities associated with fraud such as *Observers excluded*, *Improper counting* and *Illegal voting*. Also other violations like *Secrecy violated* seem to be correlated with the first indicator. Using the first alternative measure in column 2 produces no significant correlations apart from *Observers excluded* which is about half the size and only significant at the 10% level. The continuous indicators in columns 3 and 4 are also significantly correlated with *Observers excluded* and *Improper counting* but do not perform as well for the other types of fraud. Taken together, having a TVSC larger than 1 emerges as the most reliable tool at detecting relevant types of election fraud for the Russian context in this benchmark check.[21]

Next, I assess the reliability of the *regional* fraud measure through its correlation with changing opportunities for ballot stuffing. A significant decrease in such opportunities was marked by the start of electronic vote counting via optical scanners during the mid-2000s across Russia in about 5% of all voting stations (Bader, 2013b). The introduction was carried out in a staggered fashion, starting with the 2007 Duma elections, which makes it unlikely to be correlated with other incentives for fraud. From official government reports documented in Central Election Commission of the Russian Federation (2014), I collected information on the numbers of voting stations equipped with such a device in each region across national elections and calculated the share by dividing through all stations in a given region-election cell. This variable captures the reduction in scope for ballot stuffing and would be ex-ante expected to correlate negatively with turnout-inflating fraud. Table 2 shows the results from regressing the share of suspicious votes based on $TVSC \geq 1$ and TVSC exceeding incumbent votes on the share of regions' voting stations equipped with an electronic ballot box over the time period 2000 to 2012. Once election and region fixed effects are controlled for, one can see that having 1% more stations with

[21] Appendix Section C.1 discusses further results showing that also number-based fraud indicators do not constitute a promising alternative to my measure.

Table 2: The effect of electronic ballot boxes on regional fraud measures

	Share of suspicious votes from districts with							
	TVSC ≥ 1				TVSC > Incumbent Vote			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
% Electronic ballot boxes	0.288	0.653***	−0.386***	−0.387***	1.112***	1.371***	−0.259	−0.265
	(0.246)	(0.203)	(0.106)	(0.106)	(0.308)	(0.318)	(0.204)	(0.228)
Region FE	N	Y	Y	Y	N	Y	Y	Y
Election FE	N	N	Y	Y	N	N	Y	Y
Controls	N	N	N	N	N	N	N	N
Regions	77	77	77	74	77	77	77	74
Observations	537	537	537	516	537	537	537	516
Mean DV	0.2	0.2	0.2	0.19	0.74	0.74	0.74	0.74
R^2	0.002	0.665	0.760	0.727	0.038	0.322	0.544	0.574

Notes: Observations are at the region-level. The sample period is 2000 up to 2012 and includes 7 national elections (4 Presidential, 3 Parliamentary). Reported standard errors in parentheses are clustered at the region-level: *p<0.1; **p<0.05; ***p<0.01. Included control variable is Log(GDP p.c.).

electronic ballot boxes in a region significantly reduces the suspicious vote share based on $TVSC \geq 1$ by 0.3%. The findings for the alternative indicator are qualitatively similar but not significant which is presumably due to a large amount of false positives. All in all, this section has provided strong evidence that my indicator is reliable and valid. The share of suspicious votes based on TVSC$\geq$ 1 is arguably not a perfect measure of rigging intensity but the findings above suggest that it is strongly correlated with what it is supposed to measure and can rule out many important concerns.

5 Empirical analysis

5.1 Identification

The main purpose of the empirical analysis is to test whether there is evidence for the *information effect* described in Section 1. Based on this, one would expect that connected governors appointed by patronage have less incentives to rig elections since there is less need to cast signals about their uncertain qualities. Regions receiving a connected governor should therefore see their overall fraud levels decrease. I model

the relationship between $ShareSuspicious_{it}$, i.e. fraud in region i at time t, with the forced change to a governor chosen by the central administration, captured in $ConnectedGovernor_{it}$, in a standard DID specification:

$$ShareSuspicious_{it} = \alpha + \gamma_i + \lambda_t + \beta ConnectedGovernor_{it} + \boldsymbol{\mu}\, \boldsymbol{X}_{it} + \epsilon_{it} \qquad (1)$$

The addition of time- and region-specific fixed effects λ_t and γ_i restricts the focus only to variation in suspicious votes within regions off any election- or year-specific trend. Time FEs can, for instance, account for the strong upward shift in suspicious votes over time documented in Figure 3a or macro-economic shocks which would be mechanically correlated with the arrival of connected governors after December 2004. Region-specific fixed effects further control for permanently strong political machines originating from times of the Soviet Union (Hale, 2003). Further controls $\boldsymbol{X}_{it}$ include regions' aggregate democracy rating between 1991 and 2001 as well as the 2004 values of log of population and log GDP per capita. All controls are interacted with election fixed effects. Standard errors are clustered at the region level to account for autocorrelation of region-specific unobservables which may downward bias conventional robust estimates of the residuals' variance.

In order to consistently estimate the effect of having a connected governor, one needs to assure that treated regions did not systematically differ from non-treated ones and that replacement did not result in other simultaneous changes in other governor characteristics or performance which could explain the results. The first assumption is casually checked by Figure 3a which displays the mean share of potentially fraudulent votes over time for regions with and without a replacement of governor until the very first post-reform election in December 2007.[22] As can be

[22] I use this initial treatment allocation as a split to account for the changes in group composition over time. The event-study checks in Figure 4 test the common trends assumption more rigorously and use the correct treatment assignment in each period.

seen, the two groups follow roughly similar trends before the new law, even during the first major increase of suspicious votes in the presidential elections of March 2004.[23] In the following elections, the patterns start diverging with regions headed by a *non-connected* governor displaying notably higher levels of suspicious votes. After an initial peak in 2008 with an average difference of 6.4%, regions with a *connected* leader remain about 3% below the level of their counterparts in 2011 and 2012.

Unlike the common trends assumption, the absence of confounding events is not straightforward to check. This is because replacing a region's leader may also affect many other factors potentially correlated with rigging incentives and capabilities. Two such confounders may be that the central government was removing particularly unsuccessful or unpopular governors. In terms of *electoral* success, Figure 3b reveals that incumbent votes in national elections were at a similar level and that replacement does not appear to have been a punishment for pre-2005 election results. It is important to note that this pattern also holds when connected governors are in charge for the subsequent elections. Given that different levels of fraud were used to produce those results, one cannot infer much about the true popularity of governor's or the central government from those results.

More reliable measures are provided in Figures 3c and 3d which use regions' yearly GDP growth and unemployment rates as proxies for economic success across the two subgroups. Again, the raw data does not provide any evidence for substantial differences across regions before and after the appointment of governors.

[23] This initial jump is most likely due to other laws passed in 2000 and 2001 which removed immunity of governors as members of the Council of Regions and allowed the president to replace a governor in case of criminal activities (Sharafutdinova, 2010; Hill, 2012). With this opportunity to prosecute governors, it is not surprising the they wanted to please the president in the 2004 elections while not making much use of manipulation in order to help the recently founded United Russia party in 2003. Importantly, these events seem to affect governors in the same way before December 2004 since none of them had been hand-picked by then.

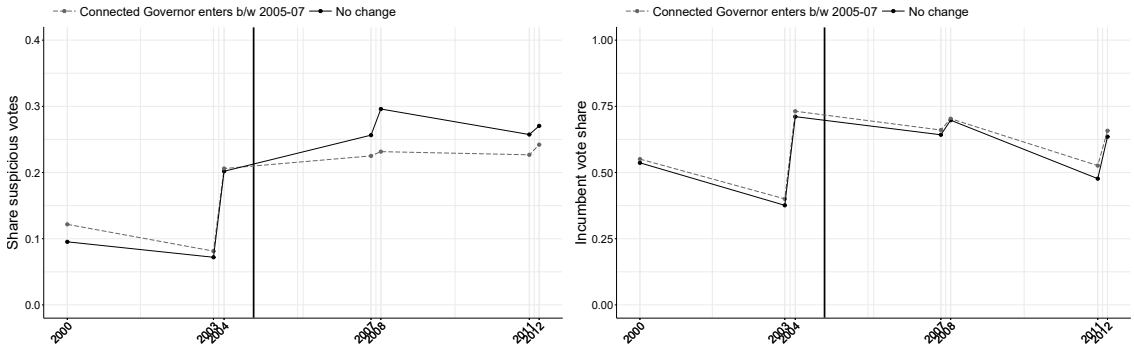

A. Average suspicious vote share

B. Average incumbent vote share

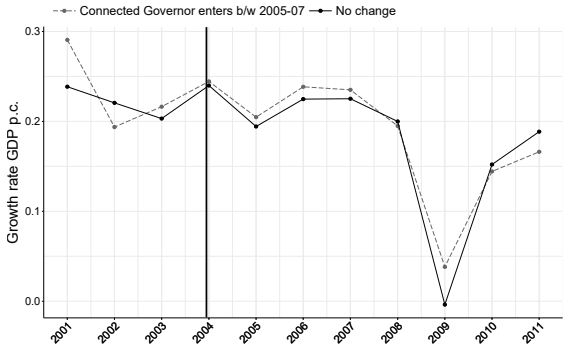

C. Avg. 1-year growth rate GDP per capita

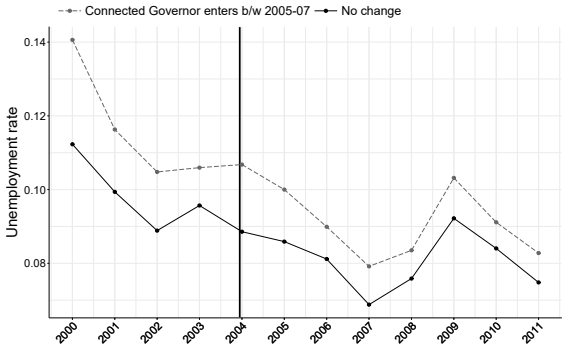

D. Avg. unemployment rate

FIGURE 3: THE EFFECT OF RECEIVING A CONNECTED GOVERNOR BEFORE THE FIRST POST-REFORM ELECTION

Notes: Plot of variable means over time in regions receiving a connected governor before the first post-reform elections in December 2007 or not.

This finding is important for several reasons: First, governor replacement is most likely not driven by low economic performance. Second, Figures 3b–3d indicate that connected governors are not chosen for their political or their economic *competence* which, in turn, makes *loyalty* the more likely selection criterion. Third, suppose that connected governors needed to manipulate less since their arrival somehow increases government support. Also this unlikely since the results on GDP growth indicate that economic performance – by far the most important issue for Russian voters (White, 2012) – does not play any role here. A final worry could be that voters were rewarding the government for replacing deeply disliked governors and in this way reduced the need for manipulation. While this is generally plausible, it is not in line with the fact that the largest difference in Figure 3a is detected in 2008

when Vladimir Putin, who initiated the law change, was not running for election. Section 5.3 will look at variables capturing this aspect in more detail.

5.2 The effect of connected governors on election fraud

Table 3 reports the main results of the DID estimation. As can be seen from the first 2 columns, *ConnectedGovernor* is positively correlated with the share of suspicious votes even when controlling for time-invariant regional characteristics. This, however, is due to the general rise in replaced governors and fraud over time which leads to a notable upward bias. Once election fixed effects are accounted for in column 3 the coefficient remains of similar magnitude and significance but flips sign. According to this specification, having a connected governor reduces suspicious votes by 6.5%. This is of a similar magnitude as the difference observed in the raw data in Figure 3a and equivalent to a quarter of a standard deviation or moving from the median to first tercile of the distribution. Specification 4 accounts for the possibility that connected governors could have been placed systematically in more prosperous or democratically minded regions by including a set of appropriate control variables. This makes the coefficient decrease marginally in magnitude and has no strong effect on the significance of the estimates. The results thus lend tentative support to the hypothesis that patronage may reduce information asymmetries and the incentives to engage in ballot rigging.

As a first test for the stability of the baseline estimates, I investigate the possibility that simple pre-trends could be driving my findings. The two setups I am consider are election-wise fixed-effects for each of Russia's 8 Federal Districts and region-specific linear time trends. Federal Districts were created in 2000 by President Putin as an intermediary subdivision and cover between 6 to 18 of these. The corresponding *plenipotentiaries* are directly appointed by the president and were intended to tighten control over territories' leaders (Hill, 2012). A particularly progressive district leader could therefore affect the replacement of corrupt governors

Table 3: Baseline results and different FE specifications

	Share of suspicious votes						
	Baseline				Robustness Checks		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ConnectedGovernor	0.047*	0.073***	−0.065**	−0.062**	−0.050*	−0.049	−0.052
	(0.027)	(0.017)	(0.027)	(0.029)	(0.030)	(0.040)	(0.040)
Region FE	N	Y	Y	Y	Y	Y	Y
Election FE	N	N	Y	Y	Y	Y	Y
Controls	N	N	N	Y	Y	Y	Y
Fed.Distr.×Election FE	N	N	N	N	Y	N	Y
Region FE×t	N	N	N	N	N	Y	Y
Regions	77	77	77	77	77	77	77
Observations	537	537	537	537	537	537	537
Mean DV	0.2	0.2	0.2	0.2	0.2	0.2	0.2
R^2	0.008	0.672	0.763	0.785	0.814	0.852	0.870

Notes: Observations are at the region-level. The sample period is 2000 up to 2012 and includes 7 national elections (4 Presidential, 3 Parliamentary). Reported standard errors in parentheses are clustered at the region-level: *p<0.1; **p<0.05; ***p<0.01. Included control variables are Democracy rating 1991-2001, Log(pop) 2004 and Log(GDP p.c.) 2004. All controls are interacted with Election FE.

and simultaneously dis-incentivize ballot rigging or introduce manipulation techniques that the TVSC cannot capture which would give similar results to the ones in the preferred specification of column 4.[24] Region-specific linear time trends, on the other hand, provide a test whether the effect could be driven by diverging linear trends in fraud between treated and non-treated region regardless of governor replacements. Specifications 5 to 7 in Table 3 show how the baseline estimates react to applying these two additional fixed effect specifications by themselves and jointly. The estimates for the main effect are slightly smaller between -0.049 and -0.052 while the standard errors increase substantially with the inclusion region-specific linear trends in column 6 and 7. As a result, the coefficients on *ConnectedGovernor* in those specifications are insignificant. Reassuringly, however, the drop in coefficients' magnitude is small and suggests that the baseline estimates were only

[24] This specification also serves as an additional check whether results are driven by the North Caucasian federal district which hosts a large amount of the ethnic republics notorious for high levels of fraud. In Table A3 I test the impact of Republics on my baseline estimates more formally but do not find my findings to be driven by these regions.

slightly biased downwards. Taken together, the evidence is not entirely conclusive but also does not support the view that pre-trends are driving the main results.

An important shortcoming of the checks above is that they do not allow for non-linear pre-trends. For instance, new laws in 2000 and 2001 removed governors' immunity and made it possible to prosecute them for criminal activities (Sharafutdinova, 2010). If under-performing, to-be-removed governors were decreasing election fraud in response to this law or any other policy change or if fraud reduction was in fact an anticipatory behaviour, this may still yield results similar to the baseline but would not be captured by linear pre-trends. In order to provide a more rigorous test of the common trends assumption, I use a Granger causality test which checks whether the decrease in fraud actually coincided with receiving a connected leader or if these regions already diverged *before* governor replacement. In order to do this I first construct a dummy for *ever having a connected governor*, $EverConnectedGovernor_i$. I then center elections in each region around the first election organized by a connected governor. This gives values between -6 and 3 to regions which had their governor exchanged and 0 to regions which did not. I collapse these into six bins for values -2 through to 2 and one for 3 and above and interact them with $EverConnectedGovernor_i$. Values below -2 are summarized in a residual bin which is also the omitted reference category. I then substitute $ConnectedGovernor_{it}$ with the aforementioned six interaction terms to trace the timing of the baseline effect.

The results of the Granger causality tests described above for different outcomes are summarized in Figure 4 for different outcomes. Figure 4a shows that levels of election fraud in the 2 ballots prior to the first election under a connected governor were marginally smaller but not statistically different from other regions. With the inaugural election at point 0 and beyond, the share of suspicious votes drops by about 8 percentage points and marginally decreases further for subsequent elections. These results reinforce the assumption that the reduction in fraud indeed coincided

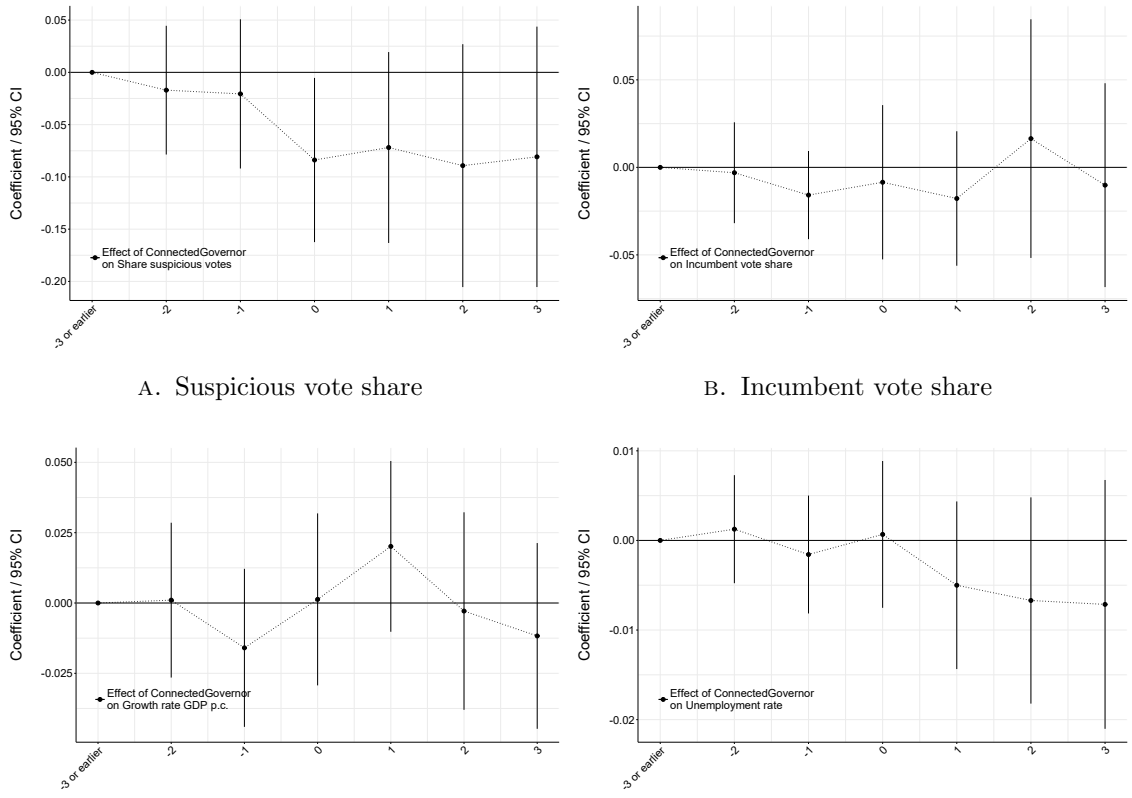

A. Suspicious vote share

B. Incumbent vote share

C. 1-year growth rate GDP per capita

D. Unemployment rate

FIGURE 4: EVENT STUDY GRAPHS FOR RECEIVING A CONNECTED GOVERNOR

Notes: Coefficients and 95% confidence intervals for the effect of receiving a connected governor on respective variable before/during/after the governor's first national election.

with the appointment of connected governors and was not anticipated. While the interactions with 0 is significant at the 5% level, respectively, the coefficients for 1, 2 and 3 are less precisely estimated and only significant at the 12-20% level. The main reason behind this is that about half of the ever treated regions have a connected governor only for the last two elections of the sample. The relative stability of the coefficients after the inaugural election is both reassuring but also informative about the type of treatment. For instance, it indicates that there are no learning effects for newly assigned leaders on how to conduct fraud or reduce it even further which I will discuss later in more detail.

Figures 4b–4d apply the same methodology to the other outcomes displayed in Figure 3. Figure 4b shows that replacement seems to be neither driven by past

national election outcomes nor does it have any impact on them. This confirms the analysis of the raw data and reinforces the view that under-performers in national elections were not systematically replaced. The non-results for the post-period, however, may seem surprising at first since one would expect lower fraud to translate into lower incumbent vote shares. There are two potential explanations to this: First, reports suggest that the central government may impose specific *targets* for each election which regional leaders have to hit (White and Feklyunina, 2011). When natural government support is low, these targets may only be achievable by manipulation and fraud would then serve as a substitute for government popularity. While this could, in principle, explain why official vote shares and turnout are nearly identical across regions, it implies that government support must have *increased* with a receiving a connected leader since fraud is falling. The final two graphs give a first idea about the plausibility of such a scenario. Figure 4c indicates a slight upward trend in GDP per capita growth around the time of replacement but none of these effects is statistically significant. Finally, Figure 4d reveals that unemployment is very similar before replacement across the two groups and then decreases starting in the first year after a connected governor assumes office. Even though these coefficients are never significant, they still raise the question whether this mild reduction in unemployment could just be a by-product of other, more significant, changes in policy outcomes caused by appointed governors.

5.3 The effect of connected governors on election results and policy outcomes

Section 5.1 already briefly discussed the possibility of correlated shocks, i.e. that receiving a connected governor could be accompanied by other events which may reduce the need to commit fraud. In this part of the analysis, I therefore look at the effect of connected governors on other outcomes related to elections, the polit-

ical environment and the economy. While the event-study graphs presented above already cast some doubt on this, I now look at a number of additional outcomes to obtain a clearer picture. Table 4 starts by investigating the impact of connected governors on turnout and vote shares in national elections. Specification 1 shows that turnout is decreasing by 1.1% on average in regions with a connected governor. Even though this finding is generally in line with a turnout-inflating type of fraud such as ballot-stuffing, it is not significant at conventional levels and very small compared to the mean of 0.65. The results in column 2 are confirming the pattern in Figure 4b that votes for the incumbent party or candidate are neither significantly nor meaningfully different for regions with a connected leader. The estimates for the main competing parties are higher in relation to the variables' mean but never significant. Closest to significance is the effect on the Ultranationals which, coincidentally, are also perceived as the major satellite party of the central government (Gel'man, 2008).

Next, I check a couple of indicators related to the political sphere. First, I test whether connected governors fostered government support by oppression which should be captured in the democracy rating by Petrov and Titkov (2013) and the *FreeMedia* dummy compiled from the Glasnost reports. Columns 6 and 7 show that democracy ratings marginally increased while the probability of free media fell. Both estimates are insignificant and small compared to their mean values. Specification 8 then looks at the regional percentage of survey respondents approving of their governor's work.[25] The results show that governor popularity increased by a marginally insignificant 7.1 percentage points in response to receiving a connected leader. Another possibility could be that stronger government ties enabled new leaders to attract pork-barrel projects and in this way increased government support

[25] This data was originally used in Reuter and Robertson (2012) and kindly shared by the authors. Apart from covering only 6 sample years, a major drawback is that the poll was only conducted in 61 regions out of which 60 are used in this study.

Table 4: Results on election results and policy outcomes

	Electoral (%)					Political				Economic					
	Turn-out	In-cum-bent	Com-munist	Ultra-national	Demo-cratic	Demo-cracy rating	Media Free-dom = 1	Gover-nor Popu-larity	Fin. trans fers p.c.	1-year GDP p.c. Growth	% Un-em-ploy-ment	Log Income p.c.	Avg. wage	% Po-verty	CPI
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
ConnectedGovernor	−0.011	−0.004	0.003	−0.005	0.001	0.513	−0.063	0.071	0.019	0.011	−0.003	−0.003	0.246	−0.003	0.297
	(0.010)	(0.012)	(0.007)	(0.003)	(0.002)	(0.332)	(0.079)	(0.044)	(0.943)	(0.010)	(0.003)	(0.021)	(0.312)	(0.010)	(0.275)
Region FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Election/Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Regions	77	77	77	77	77	77	75	60	77	77	77	77	77	77	77
Observations	539	539	539	539	539	616	450	360	924	847	924	924	924	924	924
Mean DV	0.65	0.59	0.18	0.08	0.03	29.41	0.49	0.42	3.33	0.19	0.09	8.8	10.56	0.24	112.61
R^2	0.811	0.890	0.860	0.907	0.867	0.956	0.639	0.807	0.813	0.517	0.903	0.990	0.984	0.901	0.830

Notes: Observations are at the region-level. The sample period is 2000 up to 2012 and includes 7 national elections (4 Presidential, 3 Parliamentary) for all electoral outcomes. For non-economic outcomes, the sample period is 2000 up to 2011 apart from Democracy Rating (2003-2010), Media Freedom (2000-2002, 2006, 2008 and 2010), Gov. Popularity (2003-2008) and 1-year Growth GDP p.c. (2001-2011). Reported standard errors in parentheses are clustered at the region-level: *p<0.1; **p<0.05; ***p<0.01. Included control variables are Democracy rating 1991-2001, Log(pop) 2004 and Log(GDP p.c.) 2004. All controls are interacted with Election FE for electoral outcomes and Year FE for all others.

and reduced the need for fraud. In column 9, I therefore look at the effect on government transfers per capita as a proxy for such pork-barrel spending. The results suggest that there is no empirical support for this story.

Finally, I revisit the effect on yearly growth in GDP per capita and unemployment, albeit this time in the standard DID setting, and perform the same regressions for a number of other economic indicators. The first two specifications 10 and 11 essentially confirm the results in Figure 4. Columns 12 to 16 investigate log income per capita, average wages, the share of people living below the subsistence level as well as the consumer price index. The results on these outcomes is mixed and imply both positive and negative effects on economic well-being. However, as with the other economic, political and electoral outcomes, none of them seems to respond significantly to receiving a connected governor.

5.4 Changes in governor characteristics

The previous results have highlighted two important results. First, lower fraud is unlikely to be the result of lower rigging incentives due to an increase in government support. Second, given that there is no impact of connected governors on economic and political performance, competence seems a far less likely appointment criterion than loyalty. In Table 5, I therefore check the effect of connected governors on various metrics of leader characteristics related to party affiliations, ties with their region and professional background. According to specification 1, connected governors had about 10 years less experience in office than their counterparts in other regions. Given the type of treatment, this results may appear somewhat mechanic but also shows that governors who were *not* replaced had predominantly already assumed their position in the 1990s. Compared to their new colleagues, leaders in regions without replacement were thus far more likely to have well-functioning electoral machines at their disposal which could in principle be an alternative explanation beyond loyalty for the baseline findings. This goes back to Reuter (2013) who

argues that post-2004 governor replacements often wiped out well-functioning political machines and thereby undermined (official) political support for United Russia in regional elections. Two important points speak against this reasoning: First, if capacity was critical to conducting fraud, one would expect this to also show up in lower incumbent vote shares in regions with connected governors which was not the case according to Table 4. Second, one should then also expect new governors to slowly build their own electoral machines over time. Again, the stable coefficients in Figure 4a speak against such learning effects.

Columns 2 and 3 show that governor appointments led to a change in leader generations. Newly appointed governors were on average 10 years younger and less likely to have ever been affiliated with the Communist party. Absent any tangible improvements, this ousting of the "old guard" is the most plausible explanation for the marginally improved popularity ratings in Table 4. Next, I investigate whether Communists were systematically replaced by individuals who had *ever* been members of the incumbent party *United Russia*. The results in column 4 show that this was actually not the case which is most likely due to the fact that many governors only joined United Russia after their removal. When looking at *current* membership, however, specification 5 shows that the odds of having a current United Russia member as governor increased by 23.7% with a connected governor which is close to the sample mean. The actual length of their membership was, however, shorter as shown in column 6. Given that governors did not differ in their membership length before replacement, this last result can only be explained if connected governors were predominantly joining the party only very close to their appointment.[26] Taken together, this means that incumbent party membership was an important prerequisite for governors who were explicitly chosen by the central government but that

[26] See also Figure 7 in relation to this argument.

Table 5: Effect of connected governors on leader characteristics

	Tenure	Age	Ever Communist = 1	Ever United Russia = 1	Currently United Russia = 1	UR Membership Length	Born in Region = 1	Lived in Region = 1	Security or Army background = 1	Distance St. Petersburg–Place of Birth	Business background = 1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
ConnectedGovernor	−9.827***	−10.209***	−0.368***	−0.020	0.237***	−1.814***	−0.091	−0.164**	−0.023	−0.380	0.096*
	(0.614)	(1.299)	(0.094)	(0.100)	(0.068)	(0.383)	(0.089)	(0.074)	(0.058)	(1.577)	(0.051)
Region FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Regions	77	77	77	77	77	77	77	77	77	77	77
Observations	924	924	924	924	924	924	924	924	924	924	924
Mean DV	6.41	55.14	0.6	0.74	0.26	1.57	0.49	0.83	0.09	32.29	0.04
R^2	0.825	0.761	0.724	0.546	0.578	0.704	0.759	0.635	0.617	0.784	0.619

Notes: Observations are at the region-level. The sample period is 2000 up to 2011 for all variables. Reported standard errors in parentheses are clustered at the region-level: *p<0.1; **p<0.05; ***p<0.01. Included control variables are Democracy rating 1991-2001, Log(pop) 2004 and Log(GDP p.c.) 2004. All controls are interacted with Year FE.

experience or a high position therein was not an important criterion. Entry could thus be interpreted as one way to make loyalty to the central government official.

Next, I check whether new leaders were systematically selected on their geographical origins. First, column 7 reveals that connected governors are 9.1% less likely to have been born in the region they govern. The estimate is, however, not significant and in fact only every second governor fulfils this condition. Having lived in the governed region before assuming office is far more common but 16.4% less likely to be the case for connected governors. This finding, which is also statistically significant, implies that the central government was using its powers to systematically replace natives with outsiders, a development which was also noted by Reuter (2013). Finally, I statistically check whether there was a systematic selection of governors with ties to the secret service or military, to the city of St. Petersburg or experience in business management (Blakkisrud, 2011). While the first two cannot be confirmed according to the results in columns 9 and 10, specification 11 provides indicates that connected governors had a 9.6% higher chance of having a business background than non-connected ones.

In sum, I find that governor replacements were predominantly used to replace long-serving, older leaders with a background in the Communist party. Their successors, however, were predominantly recent United Russia members who had not lived in the region before their term and who were likely to have served in a management position before without much political experience. None of these findings constitutes an alternative reason for why connected governors should have engaged in less election fraud, if it was not for their loyalty. This loyalty became official through their party membership but was presumably assured in a screening process before. This claim can, however, not be further explored with the data available.

6 Conclusion

In this paper, I study the importance of incentive structures of local officials on the dynamics of election fraud in Russia. I exploit a radical law change in December 2004 which allowed the central government to remove governors without any constraints and thus created strong motivation for the latter to use rigging as a signal in order to stay in office. I hypothesize that governors connected to the central government had less need to engage in fraud than non-connected ones since their loyalty was assured and there was lower need to signal. The paper develops and extensively tests a new indicator of electoral fraud for Russian regions between 2000 and 2012 which is created from a unique micro-level dataset of election results at the voting station level. The effect of having a connected governor on the share of suspicious votes in a region is estimated using a DID model. The baseline results support the hypotheses above and show that regions with a connected governor had on average 6% less suspicious votes than those without. The effect is highly significant and passes several robustness checks concerning the validity of the common trends assumption and placebo treatments.

While the share of suspicious votes is affected by the law change and connected governors, I also show that other electoral, political and economic outcomes did not respond. This serves as evidence against the claim that there was lower need for fraud because connected governors were improving government popularity or otherwise reducing the need for manipulation in order to meet target election outcomes. Looking at governor characteristics, I also find that the central government systematically used the law change to replace long-serving governors with ties to the Communist party by young, outsider candidates with a management background and little experience in politics. Connected governors are more likely to be a member of the United Russia party but join very close to the start of their office term.

Loyalty thus must have been assured at a previous stage of the screening process for prospective governors.

Despite the focus on Russia, my findings provide interesting insights into the functioning of authoritarian systems. In particular, elections can actually still function as an arena for political competition, albeit only among lower-tier officials. Contrary to common sense, I argued conceptually and empirically that incentives to manipulate under patronage may be lower among a dictator's cronies due to lower uncertainty about their qualities. From a policy perspective, this means that advocating for the co-optation of non-cadres to government positions in an authoritarian regime may actually have counter-productive results and lead to – comparatively – higher levels of election fraud.

References

Bader, M. (2013a). Crowdsourcing election monitoring in the 2011–2012 Russian elections. East European Politics 29(4), 521–535.

Bader, M. (2013b). Do new voting technologies prevent fraud? Evidence from Russia. USENIX Journal of Election Technology and Systems 2(1), 1–8.

Bader, M. and C. van Ham (2014). What explains regional variation in election fraud? Evidence from Russia: a research note. Post-Soviet Affairs 31(6), 514–528.

Beber, B. and A. Scacco (2012). What the Numbers Say: A Digit-Based Test for Election Fraud. Political Analysis 20(2), 211–234.

Blakkisrud, H. (2011). Medvedev's New Governors. Europe-Asia Studies 63(3), 367–395.

Callen, M. and J. D. Long (2015). Institutional Corruption and Election Fraud: Evidence from a Field Experiment in Afghanistan. American Economic Review 105(1), 354–381.

Cantú, F. (2019). The Fingerprints of Fraud: Evidence from Mexico's 1988 Presidential Election. American Political Science Review 113(03), 710–726.

Central Election Commission of the Russian Federation (2014). VBPC: The History of Creation and Use: Collection of Materials.

Cox, G. W. and J. M. Kousser (1981). Turnout and Rural Corruption: New York as a Test Case. American Journal of Political Science 25(4), 646.

Egorov, G. and K. Sonin (2011). Dictators And Their Viziers: Endogenizing The Loyalty–Competence Trade-Off. Journal of the European Economic Association 9(5), 903–930.

Enikolopov, R., V. Korovkin, M. Petrova, K. Sonin, and A. Zakharov (2013). Field experiment estimate of electoral fraud in Russian parliamentary elections. Proceedings of the National Academy of Sciences 110(2), 448–452.

Filippov, M. and P. C. Ordeshook (1997). Who stole what in Russia's December 1993 elections. Demokratizasiya 5(1), 36–52.

Folke, O., S. Hirano, and J. M. Snyder (2011). Patronage and Elections in U.S. States. American Political Science Review 105(3), 567–585.

Frye, T., O. J. Reuter, and D. Szakonyi (2014). Political Machines at Work Voter Mobilization and Electoral Subversion in the Workplace. World Politics 66(02), 195–228.

Gandhi, J. and E. Lust-Okar (2009). Elections Under Authoritarianism. Annual Review of Political Science 12(1), 403–422.

Gehlbach, S., K. Sonin, and M. W. Svolik (2016). Formal Models of Nondemocratic Politics. Annual Review of Political Science 19(1), 565–584.

Gel'man, V. (2008). Party Politics in Russia: From Competition to Hierarchy. Europe-Asia Studies 60(6), 913–930.

Goode, J. P. (2007). The puzzle of Putin's gubernatorial appointments. Europe-Asia Studies 59(3), 365–399.

Hale, H. (2003). Explaining Machine Politics in Russia's Regions: Economy, Ethnicity, and Legacy. Post-Soviet Affairs 19(3), 228–263.

Haraszti, M. (Ed.) (2004). Report on Russian media coverage of the Beslan tragedy: Access to information and journalists' working conditions. The Representative on Freedom of the Media.

Hill, R. J. (2012). The Putin era. In G. Gill and J. Young (Eds.), Routledge handbook of Russian politics & society, pp. 13–26. London: Routledge.

Kalinin, K. and W. R. Mebane Jr. (2011). Understanding Electoral Frauds through Evolution of Russian Federalism: from "Bargaining Loyalty" to "Signaling Loyalty". Presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL, March 31–April 2, 2011.

Klimek, P., Y. Yegorov, R. Hanel, and S. Thurner (2012). Statistical detection of systematic election irregularities. Proceedings of the National Academy of Sciences 109(41), 16469–16473.

Kobak, D., S. Shpilkin, and M. S. Pshenichnikov (2016). Statistical fingerprints of electoral fraud? Significance 13(4), 20–23.

Lehoucq, F. (2003). Electoral Fraud: Causes, Types, and Consequences. Annual Review of Political Science 6(1), 233–256.

Little, A. T. (2015). Fraud and Monitoring in Non-Competitive Elections. Political Science Research and Methods 3(1), 21–41.

Lukinova, E., M. Myagkov, and P. C. Ordeshook (2011). Metastasised Fraud in Russia's 2008 Presidential Election. Europe-Asia Studies 63(4), 603–621.

Magaloni, B. (2006). Voting for Autocracy. Cambridge: Cambridge University Press.

Mares, I. and L. Young (2016). Buying, Expropriating, and Stealing Votes. Annual Review of Political Science 19(1), 267–288.

Martinez-Bravo, M. (2014). The Role of Local Officials in New Democracies: Evidence from Indonesia. American Economic Review 104(4), 1244–1287.

Martinez-Bravo, M., P. Mukherjee, and A. Stegmann (2017). The Non-Democratic Roots of Elite Capture: Evidence From Soeharto Mayors in Indonesia. Econometrica 85(6), 1991–2010.

Mebane Jr., W. R. (2013). Using Vote Counts' Digits to Diagnose Strategies and Frauds: Russia. Presented at the 2013 Annual Meeting of the American Political Science, Chicago, IL, August 29–September 1, 2013.

Moraski, B. J. (2012). The Duma's Electoral System: Lessons in Endogeneity. In G. Gill and J. Young (Eds.), Routledge handbook of Russian politics & society, pp. 102–114. London: Routledge.

Moraski, B. J. and W. M. Reisinger (2013). Deference or Governance? A Survival Analysis of Russia's Governors Under Presidential Control. In W. M. Reisinger (Ed.), Russia's Regions and Comparative Subnational Politics, pp. 40–62. Routledge.

Moser, R. G. and A. C. White (2016). Does electoral fraud spread? The expansion of electoral manipulation in Russia. Post-Soviet Affairs 33(2), 85–99.

Myagkov, M. G., P. C. Ordeshook, and D. Shakin (2009). The forensics of election fraud: Russia and Ukraine. Cambridge: Cambridge University Press.

OSCE (2000). Russian Federation, Elections to the State Duma, 19 December 1999. OSCE/ODIHR Election Observation Mission Final Report. Warsaw.

OSCE (2004a). Russian Federation, Elections to the State Duma, 7 December 2003. OSCE/ODIHR Election Observation Mission Final Report. Warsaw.

OSCE (2004b). Russian Federation, Presidential Election, 14 March 2004. OSCE/ODIHR Election Observation Mission Final Report. Warsaw.

OSCE (2012a). Russian Federation, Elections to the State Duma, 4 December 2011. OSCE/ODIHR Election Observation Mission Final Report. Warsaw.

OSCE (2012b). Russian Federation, Presidential Election, 4 March 2012. OSCE/ODIHR Election Observation Mission Final Report. Warsaw.

Petrov, N. and A. Titkov (2013). Democracy Rating of Russian Regions by the Carnegie Moscow Center: 10 Years in Service. Moscow: Carnegie Moscow Center.

Reuter, O. J. (2010). The Politics of Dominant Party Formation: United Russia and Russia's Governors. Europe-Asia Studies 62(2), 293–327.

Reuter, O. J. (2011). United Russia and the 2011 Elections. Russian Analytical Digest 102, 2–6.

Reuter, O. J. (2013). Regional patrons and hegemonic party electoral performance in Russia. Post-Soviet Affairs 29(2), 101–135.

Reuter, O. J. and G. B. Robertson (2012). Subnational Appointments in Authoritarian Regimes: Evidence from Russian Gubernatorial Appointments. Journal of Politics 74(4), 1023–1037.

Reuter, O. J. and D. Szakonyi (2019). Elite Defection under Autocracy: Evidence from Russia. American Political Science Review 113(2), 552–568.

Rochlitz, M. (2016). Political Loyalty vs Economic Performance: Evidence from Machine Politics in Russia's Regions. Higher School of Economics Research Paper WP BRP 34/PS/2016.

Rozenas, A. (2016). Office Insecurity and Electoral Manipulation. The Journal of Politics 78(1), 232–248.

Rozenas, A. (2017). Detecting Election Fraud from Irregularities in Vote-Share Distributions. Political Analysis 25(01), 41–56.

Rundlett, A. and M. W. Svolik (2016). Deliver the Vote! Micromotives and Macrobehavior in Electoral Fraud. American Political Science Review 110(01), 180–197.

Seymour, C. (1915). Electoral Reform in England and Wales: The Development and Operation of the Parliamentary Franchise, 1832-1885. New Haven and CT: Yale University Press.

Sharafutdinova, G. (2010). Subnational Governance in Russia: How Putin Changed the Contract with His Agents and the Problems It Created for Medvedev. Publius: The Journal of Federalism 40(4), 672–696.

Shpilkin, S. (2009). Statistical Investigation of the Results of Russian Elections in 2007-2009. Troitskij Variant (21), 2–4.

Simpser, A. (2013). Why Governments and Parties Manipulate Elections: Theory, Practice, and Implications. Cambridge: Cambridge University Press.

Skovoroda, R. and T. Lankina (2016). Fabricating Votes for Putin: New Tests of Fraud and Electoral Manipulations from Russia. Post-Soviet Affairs 33(2), 100–123.

Slider, D. (2012). Regional governance. In G. Gill and J. Young (Eds.), Routledge handbook of Russian politics & society, pp. 153–163. London: Routledge.

Teague, E. (2014). Russia's Return to the Direct Election of Governors: Re-Shaping the Power Vertical? Region: Regional Studies of Russia, Eastern Europe, and Central Asia 3(1), 37–57.

White, S. (2012). Public Opinion and Voting Behaviour. In G. Gill and J. Young (Eds.), Routledge handbook of Russian politics & society, pp. 353–363. London: Routledge.

White, S. and V. Feklyunina (2011). Russia's Authoritarian Elections: The View from Below. Europe-Asia Studies 63(4), 579–602.

Zakharov, A. V. (2016). The Loyalty-Competence Trade-Off in Dictatorships and Outside Options for Subordinates. The Journal of Politics 78(2), 457–466.

Ziblatt, D. (2009). Shaping Democratic Practice and the Causes of Electoral Fraud: The Case of Nineteenth-Century Germany. American Political Science Review 103(01), 1.

Appendix

A Tables

Table A1: Descriptive statistics (Regional data)

Variable	Obs	Mean	Std.Dev.	Min	Max
Fraud and Elections					
Share of ballots w/ TVSC≥1	537	0.20	0.24	0.00	1.00
Share of ballots w/ TVSC≥Incumbent %	537	0.74	0.21	0.00	1.00
% Turnout	539	0.65	0.11	0.44	0.98
% Incumbent vote	539	0.59	0.16	0.25	0.99
% Communist vote	539	0.18	0.08	0.00	0.47
% Ultranational vote	539	0.08	0.05	0.00	0.23
% Democratic vote	539	0.03	0.03	0.00	0.20
% voting stations w/ electronic ballot box	539	0.02	0.04	0.00	0.29
Political and economic variables					
Democracy rating 1991-2001	77	27.61	6.04	14	45
Democracy rating	616	29.41	5.86	16.00	46.00
Free Media = 1	450	0.49	0.50	0	1
Governor Popularity	360	0.42	0.18	0.05	0.89
Financial transfers per cap.	924	3.33	6.40	−0.01	53.90
Population in 100,000	924	1.76	1.68	0.04	11.82
1-year Growth GDP p.c.	921	0.20	0.11	−0.32	1.02
Unemployment rate	924	0.09	0.06	0.01	0.63
Log(Income p.c.)	924	8.80	0.86	6.38	10.91
Average wages	924	10.56	8.75	0.88	59.10
% below subsistence	924	0.24	0.13	0.06	0.94
Consumer price index	924	112.61	4.59	101.40	138.70
Treatment and governor variables					
Connected Governor	924	0.26	0.44	0	1
Connected Governor b/w 2004/2007 = 1	924	0.30	0.46	0	1
Tenure	924	6.41	4.57	0.00	20.21
Age (in years)	924	55.14	8.09	34.05	75.95
Ever Communist = 1	924	0.60	0.49	0	1
Ever United Russia = 1	924	0.74	0.44	0	1
Currently United Russia = 1	924	0.26	0.44	0	1
Length UR Membership	924	1.57	2.27	0.00	11.00
Born in region = 1	924	0.49	0.50	0	1
Lived in region = 1	924	0.83	0.38	0	1
Security/Army background = 1	924	0.09	0.28	0	1
Distance St. Petersburg-Place of Birth	924	32.29	10.79	12.26	78.72
Business background = 1	924	0.04	0.20	0	1

Notes: The unit of observation is one of the 77 regions in the sample at time t. t represents an election for *Fraud and Elections* and a year for all other data. Variables provided at the cross-sectional level only (i.e. with only 77 observations). Coverage for specific variables can vary and is explained in detail in Section 3

Table A2: Descriptive statistics (District data)

Variable	Obs	Mean	Std.Dev.	Min	Max
Fraud indicators					
TVSC≥1 = 1	4,984	0.32	0.47	0	1
TVSC≥Incumbent % = 1	4,984	0.86	0.34	0	1
TVSC	4,984	0.88	0.58	−3.60	26.92
TVSC - Incumbent %	4,984	0.30	0.54	−4.29	26.25
Fraud reports					
Any fraud report = 1	4,984	0.03	0.17	0	1
Improper counting = 1	4,984	0.01	0.12	0	1
Exclusion of voters = 1	4,984	0.01	0.09	0	1
Illegal campaigning = 1	4,984	0.00	0.03	0	1
Observers excluded = 1	4,984	0.01	0.11	0	1
Faulty ballot box = 1	4,984	0.01	0.07	0	1
Secrecy violated = 1	4,984	0.00	0.05	0	1
Illegal voting = 1	4,984	0.01	0.11	0	1
Other violations = 1	4,984	0.02	0.13	0	1

Notes: The unit of observation is one of the 2,492 districts in the sample at election t (2011 or 2012).

Table A3: Baseline effect and heterogeneity for ethnic republics

	Share of suspicious votes				
	(1)	(2)	(3)	(4)	(5)
ConnectedGovernor	0.047*	0.073***	−0.051**	−0.062**	−0.061**
	(0.027)	(0.017)	(0.022)	(0.029)	(0.030)
ConnectedGovernor×Republic					−0.004
					(0.043)
Region FE	N	Y	Y	Y	Y
Election FE	N	N	N	N	N
Controls	N	N	N	Y	Y
Regions	77	77	77	77	77
Observations	537	537	537	537	537
Mean DV	0.2	0.2	0.2	0.2	0.2
R^2	0.008	0.672	0.737	0.781	0.781

Notes: Observations are at the region-level. The sample period is 2000 up to 2012 and includes 7 national elections (4 Presidential, 3 Parliamentary). Reported standard errors in parentheses are clustered at the region-level: *p<0.1; **p<0.05; ***p<0.01. Included control variables are Democracy rating 1991-2001, Log(pop) 2004 and Log(GDP p.c.) 2004. All controls are interacted with Election FE.

B Figures

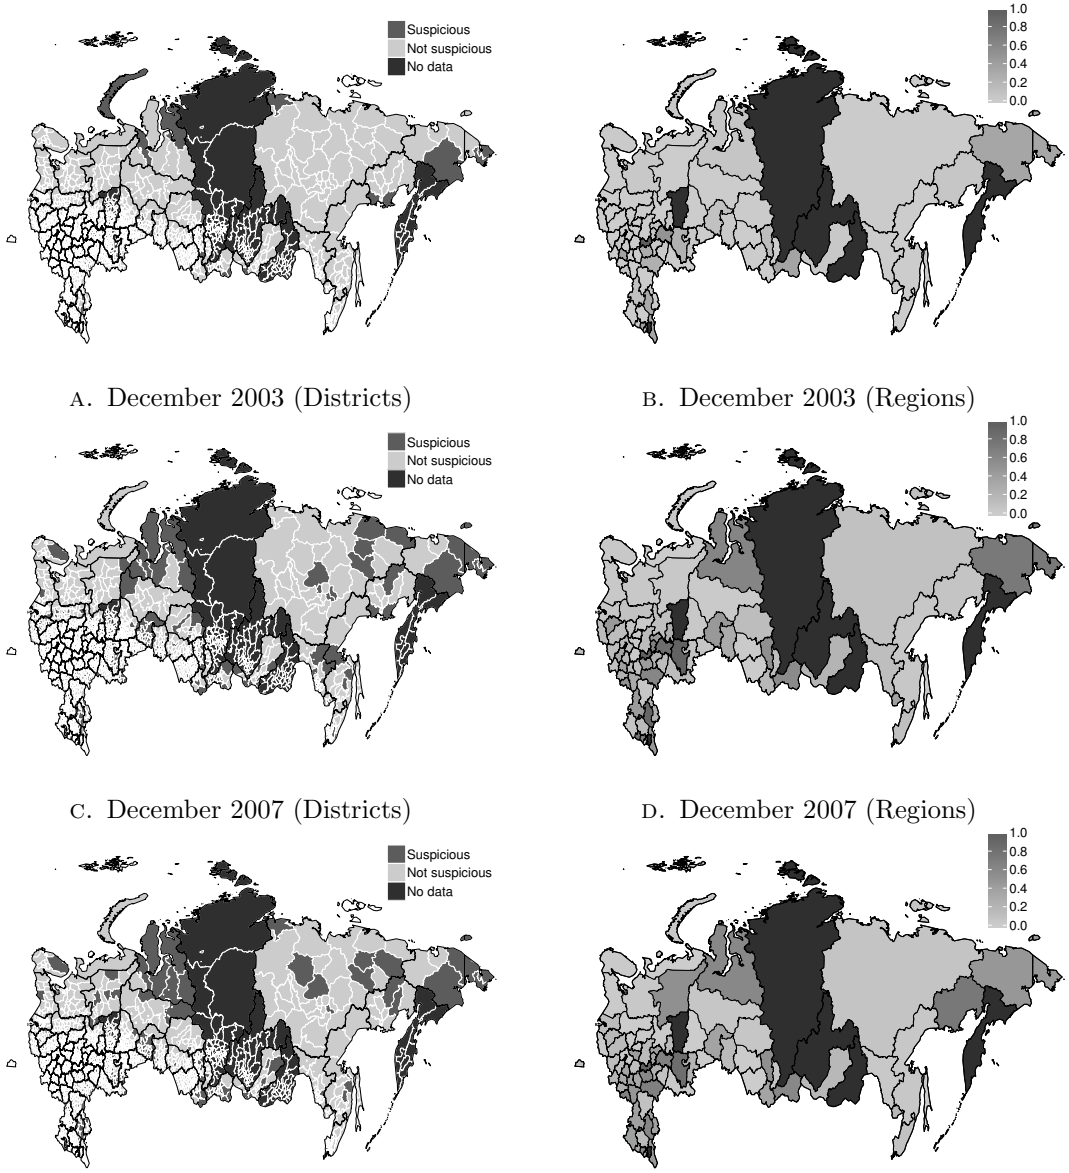

A. December 2003 (Districts)

B. December 2003 (Regions)

C. December 2007 (Districts)

D. December 2007 (Regions)

E. December 2011 (Districts)

F. December 2011 (Regions)

FIGURE 5: CLASSIFICATION AND AGGREGATION OF SUSPICIOUS VOTES BEFORE/AFTER THE ABOLITION OF GOVERNOR ELECTIONS IN DECEMBER 2004 (DUMA ELECTIONS)

Notes: Map of the division of the Russian Federation into regions and districts. Panels on the left are at the district levels and show their classification into *Suspicious* (dark grey) and *Not suspicious* (light grey) based on their voting station-level results according to the $TVSC \geq 1$ criterion. Panels on the right show the share of valid ballots in a specific region/election coming from *Suspicious* districts. Dark grey states are not included in the sample for this particular election.

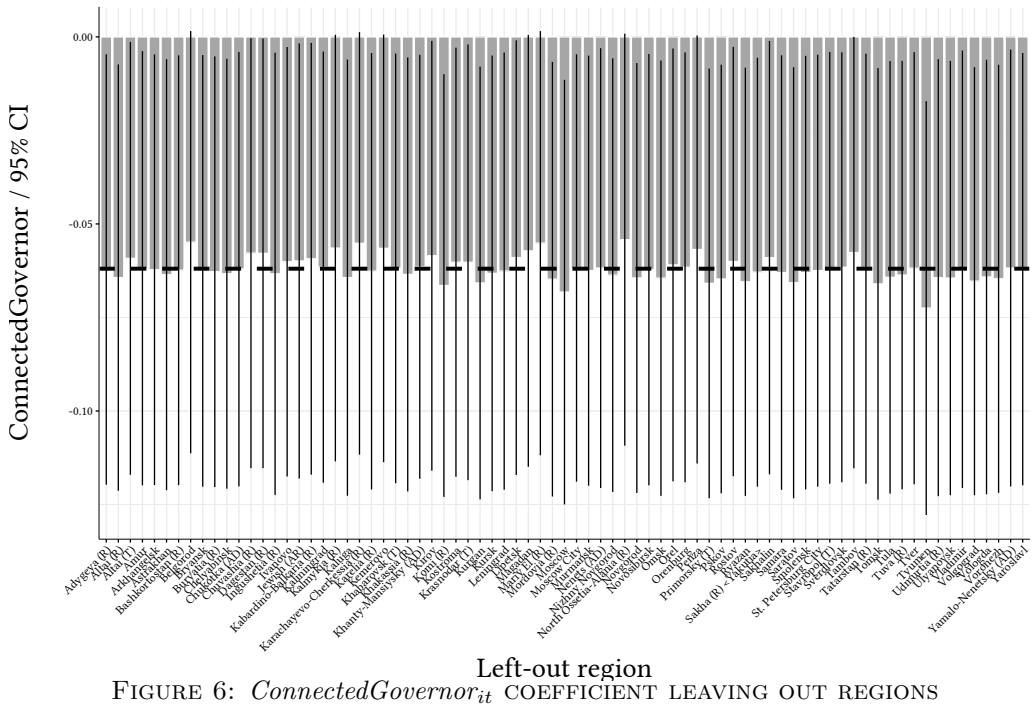

FIGURE 6: *ConnectedGovernor*$_{it}$ COEFFICIENT LEAVING OUT REGIONS

Notes: Coefficients for *ConnectedGovernor*$_{it}$ after removing a single region (denoted on the x-axis) from the sample and corresponding 95% confidence intervals. The dashed line indicates the baseline, i.e. the magnitude when excluding no region.

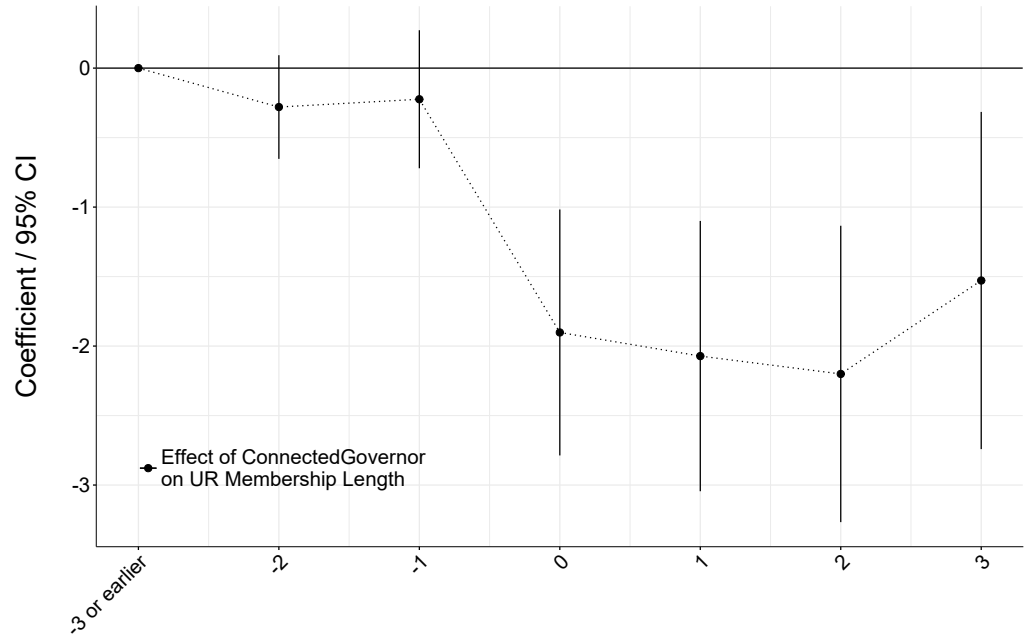

FIGURE 7: EVENT STUDY GRAPH FOR UR MEMBERSHIP LENGTH

Notes: Coefficients and 95% confidence intervals for the effect of receiving a connected governor on UR Membership Length before/during/after the governor's first national election.

C Additional analyses

C.1 Alternative indicators

The measures discussed in Section 4.2 solely investigated tools for detecting *turnout-inflating* types of rigging which was motivated by the vast amount of evidence suggesting these to be the most widely used manipulation techniques in the Russian context (Filippov and Ordeshook, 1997; White and Feklyunina, 2011; Enikolopov et al., 2013). Yet, anecdotal evidence from the Russian republics – Tatarstan, Ingushetia and Dagestan in particular – suggests that election results in some areas may not only be manipulated but entirely fabricated (Myagkov et al., 2009; Lukinova et al., 2011). In detecting this kind of fraud I follow the methodology of Beber and Scacco (2012) who rely on human preferences for specific numbers and biases in number generation. The main argument is that, under fairly generous assumptions, the final digit as well as the distance between the last and second-last digit of the vote count should follow a uniform distribution.

In order to create alternative indicators of the share of suspicious votes, I adapt the methodology of Beber and Scacco (2012) to identify fraud at the district level and then aggregate this to the regional level using the share of affected votes as for the TVSC. In detail, I proceeded as follows: first, I calculated for each district and 2011/2012 election the p-values of a Pearson's chi-squared test of uniform distribution of the last digit and the distance between last and second-last digits and repeated this procedure for the reported absolute counts of valid votes as well as incumbent votes. In a second step, I calculated dummy variables whether the hypothesis of uniform distribution could be rejected at a significance level of either 5 or 1%. This procedure has the advantage of being completely agnostic about the kind of bias, i.e. whether there is a bias towards fives in one district vs. eights in another, and only assumes whether votes for the incumbent or the number of valid votes/turnout were affected.

TABLE A4: PARTIAL CORRELATION BETWEEN REPORTED FRAUD AND NUMERIC ANOMALIES

Election outcome	Valid votes				Incumbent votes			
Test uniform distribution of	last digit		Δ last 2 digits		last digit		Δ last 2 digits	
Criterion	$p \leq 0.05 = 1$	$p \leq 0.01 = 1$	$p \leq 0.05 = 1$	$p \leq 0.01 = 1$	$p \leq 0.05 = 1$	$p \leq 0.01 = 1$	$p \leq 0.05 = 1$	$p \leq 0.01 = 1$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any report = 1	−0.012	0.011	−0.025	−0.002	0.009	0.018	−0.024	0.001
	(0.014)	(0.014)	(0.016)	(0.015)	(0.015)	(0.017)	(0.015)	(0.015)
Improper counting = 1	0.008	0.029*	−0.018	−0.003	−0.010	0.026	−0.043**	−0.040**
	(0.017)	(0.016)	(0.019)	(0.018)	(0.018)	(0.020)	(0.017)	(0.017)
Exclusion of voters = 1	0.013	0.019	−0.038**	−0.000	−0.021	−0.001	0.002	0.053***
	(0.017)	(0.016)	(0.019)	(0.018)	(0.018)	(0.020)	(0.017)	(0.017)
Illegal campaigning = 1	0.002	−0.001	−0.032*	−0.002	0.003	0.002	0.019	−0.004
	(0.016)	(0.015)	(0.018)	(0.016)	(0.017)	(0.018)	(0.016)	(0.016)
Observers excluded = 1	−0.005	0.001	−0.021	−0.005	−0.008	0.000	−0.021	0.002
	(0.018)	(0.017)	(0.020)	(0.018)	(0.019)	(0.020)	(0.018)	(0.018)
Faulty ballot box = 1	−0.045**	−0.001	−0.020	−0.005	−0.025	−0.005	−0.044**	−0.067***
	(0.018)	(0.017)	(0.020)	(0.019)	(0.019)	(0.021)	(0.018)	(0.018)
Secrecy violated = 1	0.002	0.000	−0.023	−0.005	−0.000	−0.000	0.013	0.002
	(0.018)	(0.017)	(0.020)	(0.018)	(0.019)	(0.021)	(0.018)	(0.018)
Illegal voting = 1	−0.005	0.000	−0.019	−0.006	0.001	0.002	0.010	−0.001
	(0.016)	(0.015)	(0.017)	(0.016)	(0.016)	(0.018)	(0.016)	(0.016)
Other violations = 1	−0.051***	0.000	−0.024	−0.003	−0.011	0.026	−0.039**	−0.036**
	(0.017)	(0.016)	(0.019)	(0.018)	(0.018)	(0.019)	(0.017)	(0.017)

Notes: Each cell represents the standardized coefficient from a single regression where the row denotes the *dependent* variable and the column denotes the *independent variable*. Observations are at the district-level. The sample period is 2011 to 2012 and includes 2 national elections (1 Presidential, 1 Parliamentary). Robust standard errors in parentheses: *p<0.1; **p<0.05; ***p<0.01. Included control variables are District FEs, Election FEs and Region×Elections FEs.

Table A4 repeats the analyses of Table 1 for each of the 9 measures of electoral violations. Only two out of 72 cells report a significant, positive coefficient. The first is the $p \leq 0.01$ indicator for last digits of valid votes and improper counting, the second is the $p \leq 0.01$ indicator for the distance in last two digits of incumbent votes and the exclusion of voters. While the results on these two number anomalies are indicative, their correlations with a single fraud report dummy do not hold for the remaining measures. Most types of fraud report dummy variables are in fact negatively affected by the existence of numeric anomalies – sometimes even significantly. One possible interpretation of these results is that fabricating votes may be a complement to turnout inflating methods which does apparently not require actions widely observable to the public and hence results in lower amounts of reported fraud. More generally, voters' inability to detect and report entirely invented results makes it hard to verify the reliability of any fraud indicator built such on numeric anomalies. The results in Table A4 should thus also be regarded as a reminder that some fraud may still go undetected by the fraud indicator employed in the remainder of this paper. Bearing this caveat in mind, I now turn towards the extent of potentially fraudulent election outcomes and its changes over time.