# Ballot Stuffing: a Structural Approach

Anastasia Burkovskaya

Kadir Atalay

October 2024

#### Abstract

We examine ballot stuffing, a form of electoral fraud involving the illegal addition of extra ballots. While common in autocratic regimes, this practice is increasingly cited in democracies as grounds for challenging election results. We propose a non-parametric structural model that links unobservable polling station characteristics to instances of ballot stuffing. By leveraging less fraudulent data—such as data from independent observers or exit polls—our model can identify ballot stuffing at the individual polling station. Applying this model to the 2011 Russian Parliamentary election, we find substantial evidence of ballot stuffing, with fraud rates varying across and within regions.

Keywords: electoral fraud, ballot stuffing, nonparametric identification, electoral

turnout, nonclassical measurement error

**JEL codes:** C14, D72

Anastasia Burkovskaya: Corresponding Author, School of Economics, University of Sydney, Australia email: anastasia.burkovskaya@sydney.edu.au); Kadir Atalay: School of Economics, University of Sydney, Australia email: kadir.atalay@sydney.edu.au

## 1 Introduction

Many modern nondemocratic leaders prefer maintaining a facade of democracy, with elections being a key part of that image (Guriev and Treisman, 2022). To guarantee the electoral outcome that projects their strength, they often engage in electoral manipulations (Gehlbach et al., 2016). Recent examples of electoral fraud have been documented in Russia (Enikolopov et al., 2013; Kobak et al., 2018), Turkey (Aksoy, 2016), and Mozambique (Leeffers and Vicente, 2019). Autocrats have devised countless ways to tamper with election outcomes, from ballot stuffing—the illegal addition of extra ballots—to direct overwriting of the final count. While statistical traces of interference often appear at the aggregate level, accurately assessing electoral fraud at the local level—such as determining where it occurs and to what extent—remains challenging. This requires understanding unobserved factors like ideological preferences and voting costs at individual polling stations. However, local-level analysis is essential to ensure electoral integrity and to guide the strategic allocation of resources, such as independent observers or vote recounts, in future elections.

In this paper, we examine ballot stuffing, a specific form of electoral fraud often seen in autocratic regimes and increasingly cited in recent democratic elections as grounds for contesting results (BBC News, 2023; Bloomberg News, 2024). We propose a non-parametric structural model of ballot stuffing and explain how it can be empirically estimated from the data. Our findings show that with the availability of a less fraudulent sample (*cleaner data*)—such as polling stations with independent observers or exit polls—we can accurately detect ballot stuffing or establish its bounds at each polling station. This enables us to measure the level of fraud in specific areas and the election overall. We estimate our model using data from 2011 Russian parliamentary election, a case often scrutinized for its lack of transparency and integrity (Golos, 2011; OSCE, 2011).

The existing research on electoral fraud explores the following main directions: the analysis of different statistical irregularities in the data, such as Benford's law (Breunig and Goerres, 2011; Mebane, 2008; Skovoroda and Lankina, 2017), unusual kurtosis of the distribution of electoral data (Klimek et al., 2012), and spikes in the distribution of votes (Kobak et al., 2016a, 2018; Rozenas, 2017); the evaluation of fraud through both natural experiments (Cantu, 2014; Casas et al., 2017) and randomized assignment of independent observers (Asunka et al., 2019; Enikolopov et al., 2013); fitting a parametric model (Levin et al., 2009; Mebane et al., 2022) and training a machine-learning algorithm (Cantu and

<sup>&</sup>lt;sup>1</sup>A quote commonly attributed to Joseph Stalin, though its authenticity is debated.

Saiegh, 2011). The existing papers concentrate on either identifying the fact of electoral manipulations through statistical irregularities or evaluating the average effect on the overall result; however, most approaches cannot predict the effect on the result in a specific polling station in the absence of a suitable natural experiment.

This paper focuses on elections with ballot stuffing being the main source of fraud. We assume only the incumbent government has access to ballot stuffing, but not the opposition. In this form of electoral fraud, polling station administrators insert additional ballots into the ballot box and manipulate the official voter lists and turnout records, thereby artificially lowering abstention rates to favor the incumbent party. As a result, both the turnout and the number of votes for the incumbent increase. However, even though the voter shares are affected as well, the number of votes for the opposition candidate stays unchanged. We exploit these legitimate votes together with a simple structural model to non-parametrically identify unobservable polling station random term by following Matzkin (2003)'s approach for inference of a nonadditive unobservable variable. This random term is associated with unobservable citizens' characteristics that we cannot control for such as unobservable component of ideological preferences, personal costs of voting, information, civic duty, etc. Next, we evaluate how this unobservable polling station random term relates to ballot stuffing by using a sample of less fraudulent data such as polling stations with independent observers that reduce fraud (Enikolopov et al., 2013) or reliable exit polls. Because we are able to identify the random term in each polling station, knowing its relationship with ballot stuffing allows us to estimate it everywhere else. If the additional data contains polling stations where there is certainty that no violations occurred or there is a reliable exit poll, then we are able to evaluate ballot stuffing in every polling station. Otherwise, if we rely on independent observers only reducing some amount of fraud, we are able to evaluate the lower bound or what would have been the final election result if independent observers were present at every polling station. It is worth noting that the need for an additional sample is met as long as independent observers or reliable exit polls are present in at least one region, which are increasingly common in today's elections.<sup>2</sup>

Our model and the estimation procedure offered in this paper have a number of advantages. First, to our knowledge, this paper is the first to propose a structural electoral model that avoids parametric assumptions and assumptions on preference formation. Bypassing restrictive assumptions makes the identification procedure more robust to different underlying specifications. Moreover it provides a flexible framework in which the model can be extended to include a specific preference formation and develop further analysis of the detailed preference components if needed. Second, the approach proposed in this paper does

<sup>&</sup>lt;sup>2</sup>For instance, a comprehensive list of upcoming elections attended by OSCE independent observers can be found at https://www.osce.org/odihr/elections.

not require prior electoral data, and hence is applicable not only to regular elections, but also to unique events such as referenda or for studying electoral boycotts. Third, we are able to make inference about any polling station in the country rather than averages across regions. This feature of our method is particularly useful for officials in charge of guaranteeing free and fair elections, because our analysis can detect polling stations with questionable results requiring additional attention with potential cancellation of the results and/or re-voting.

Our approach also contributes to the econometric literature on evaluation of the nonclassical measurement error in non-linear models (see Schennach (2016) and Schennach (2022) for the most recent literature review). Our work falls in the stream of papers that use a validation sample for the identification of the variables of interest. The validation sample in our case is the cleaner, less fraudulent, sample such as exit poll or data from independent observers. Most literature in this stream either focuses on the measurement error in a regressor or an independent variable (Bound et al., 1989; Carroll et al., 2010; Carroll and Wand, 1991; Hausman et al., 1991; Hu and Ridder, 2012; Pepe and Fleming, 1991); or uses a parametric model (Hsiao, 1989; Katz and Katz, 2010); or studies the problem with additive noise (Chen et al., 2005; Lee and Sepanski, 1995). Ballot stuffing affects the legitimate votes, which is a dependent variable, and it cannot be negative, implying a non-zero mean and systematic biases in the observations. Additionally, we allow ballot stuffing to correlate with the legitimate votes, hence, we face a mismeasured dependent variable that is correlated with the measurement error. What sets us apart from the existing papers is that we work with a non-parametric model with non-additive random term (unobserved citizens' characteristics in our context). Even though our assumptions are weak yet we are able to identify the distributions of interest because we exploit the legitimate number of votes for the opposition.

To demonstrate our estimation strategy, we applied it to the 2011 Russian parliamentary election, notorious for widespread ballot stuffing. We use the data from a randomized control trial by Enikolopov et al. (2013), where independent observers were deployed to specific polling stations in Moscow, resulting in reduction of fraud at those locations. By analyzing these cleaner polling stations, we identified the relationship between unobservable polling station random term and ballot stuffing. Then, studying the other polling stations, our estimates suggest ballot stuffing affected approximately 13% of the overall votes. Our results indicate considerable variation across regions, while also highlighting substantial differences in fraud rates within regions.

This paper proceeds as follows: Section 2 introduces a structural model of ballot stuffing and discusses its identification. In Section 3, we discuss the estimation procedure. Section 4 offers the empirical application.

## 2 Model and identification

Ballot stuffing is directly connected to turnout. That is why our approach to modeling it is closely related to the models of electoral competition and voter turnout, which often hinge on individual characteristics such as the cost of voting and preferences (see Downs (1957); Riker and Ordeshook (1968) and the literature deriving from them). However, due to the anonymity of voting, these models struggle to empirically connect the individual characteristics directly to voting behavior (King et al., 2004). As a result, researchers must impose additional assumptions, such as a specific structure of utility and cost of voting together with a specific parametric distribution that a voter would vote (Degan and Merlo, 2011; Kawai et al., 2021) or rely on unexpected and transitory shocks to voting costs, such as rainfall (Fujiwara et al., 2016).<sup>3</sup>

Our approach takes a different direction. Instead of modeling individual voting behavior, we focus on the aggregated electoral variables without making any assumptions on how these variables have been formed. Although the aggregated variables we use can be connected to the standard models of voting behavior (see an example of how a simplified version of calculus of voting relates to our model in the Online Appendix C), our estimation results are robust to the choice of a specific underlying model.

### 2.1 Electoral model

Two candidates, incumbent and opposition, are running for office. In every election, each polling station reports data on the number of votes for each candidate and the voter turnout. In each polling station, our observable electoral variables will be Y and Z: Y denotes the proportion of votes for the incumbent calculated from the entire voter population, whereas 1-Z denotes the proportion of votes for the opposition calculated from the entire voter population as well. Hence, voter turnout is Y + 1 - Z or the abstention rate is t = Z - Y.

In the absence of electoral fraud, the building block of our model is as follows:

incumbent votes: 
$$Y = m(X, \varepsilon) - t(X, \varepsilon)$$
  
opposition votes:  $1 - Z = 1 - m(X, \varepsilon)$ ,

where t and m:  $A \times E \to \mathbb{R}$  are continuous in  $(X, \varepsilon)$  and m is strictly increasing in  $\varepsilon$ ,  $A \subset \mathbb{R}^L$  is the support of observable regional characteristics X, L is the dimension of X,  $E \subset \mathbb{R}$  is the support of  $\varepsilon$ , X has a continuous density  $f_X$ , and  $\varepsilon$  is an unobservable random term related to a polling station, it has continuous distribution  $F_{\varepsilon}$  and it is independent on

<sup>&</sup>lt;sup>3</sup>Although these studies theoretically model the cost of voting for individual voters, their empirical estimations, often based on aggregated data (e.g., county-level voting), still rely on variations in costs at the aggregate level rather than the individual level.

X. This random term explains the difference between two polling stations where all observable characteristics (income, proportion of young voters, etc.) are the same yet the results are different. Term  $\varepsilon$  captures unobservable characteristics of the polling station after accounting for observable factors. These may include polling station-specific elements, such as unobservable components of ideological preferences, voting costs, access to information, and voters' sense of civic duty. The fact that m is strictly increasing in  $\varepsilon$  means that the greater  $\varepsilon$  the smaller the number of votes for the opposition. Hence, stronger ideological leaning towards the incumbent, greater voting costs, lower sense of civic duty, lack of information about the candidates would all result in a greater  $\varepsilon$ .

In any clean election, this model is identified subject to a normalization and can be non-parametrically estimated. Additionally, because the used functions are fully non-parametric, this model can cover most underlying preference formation processes. See Online Appendix C for an example.

Next, we turn our attention to ballot stuffing. Formally, ballot stuffing is defined as the illegal practice of one person submitting multiple ballots during a vote in which only one ballot per person is permitted. In non-democratic elections, ballot stuffing happens when the official staff that runs the voting in a polling station puts a number of prefilled ballots into the ballot box either before the polls open or after they close. In addition, after the polling station closes, the staff illegally corrects the official turnout to be consistent with the number of votes in the ballot box. Hence, ballot stuffing operates through lowering the abstention rate.

We assume only incumbent candidate has an opportunity to rig the election. In this case, the observable variables are affected as follows: Variable Y denotes the number of votes for the incumbent in a polling station. If  $q(X,\varepsilon)$  is a number of additional ballots due to fraud, then instead of Y we observe variable  $Y^f = m(X,\varepsilon) - t(X,\varepsilon) + q(X,\varepsilon)$ . We assume that  $q: A \times E \to \mathbb{R}$  is continuous in  $(X,\varepsilon)$ . Variable Z will not be affected, because 1-Z represents the number of votes for the opposition. To summarize, the observable variables in our case are  $Y^f$ , Z and controls X; the unobservable variables are Y and  $\varepsilon$ ; and the model is

$$Y^{f} = m(X, \varepsilon) - t(X, \varepsilon) + q(X, \varepsilon)$$

$$Z = m(X, \varepsilon),$$
(1)

where all functions are unknown.

### 2.2 Identification in nonadditive models

The only truthful observable electoral variable in this context is Z, which is generated by the model  $Z = m(X, \varepsilon)$  in Eq. (1). Generally, given unobservable function m, the random term  $\varepsilon$  is not uniquely identified. This relationship is a well-known nonadditive model that can be identified subject to a normalization and estimated as discussed in Matzkin (2003). In the empirical application, we use the midpoint of support of X for normalization of the random term, which follows Specification I in the above reference.

**Assumption 1** (Normalization). There exists known  $X = x^n$  such that  $m(x^n, \varepsilon) = \varepsilon$ .

With this normalization, Matzkin (2003) obtains expressions (3.1') and (3.3') in her paper, which are

$$F_{\varepsilon}(e) = F_{Z|X=x^n}(m(x^n, e))$$
  

$$m(x, e) = F_{Z|X=x}^{-1} \left( F_{Z|X=x^n}(e) \right),$$

where  $F_{Z|X}$  is a conditional distribution of Z given X.

This assumption requires choosing characteristics  $X=x^n$  that will be "the standard" for polling station uncertainty. We might choose one particular region or a vector of characteristics without connection to any particular region as long as  $x^n \neq x$  in the regions with X=x for which we conduct the estimation. The main property of the chosen value  $x^n$  should be the quality of the estimator of  $F_{Z|X=x^n}(\cdot)$ , hence, the vector should be well inside of the support of X.

# 2.3 Identification of unobservable components

Without any additional structure and/or data, the identification of  $t(\cdot)$  and  $q(\cdot)$  is not possible. Indeed, functions  $\tilde{q}(X,\varepsilon) = 0.5q(X,\varepsilon)$  and  $\tilde{t}(X,\varepsilon) = t(X,\varepsilon) - 0.5q(X,\varepsilon)$  produce the exactly same observable variables  $Y^f$  and Z. To proceed further we need additional assumptions either on functions such as turnout or additional observations. This paper explores evaluation of ballot stuffing through additional cleaner, less fraudulent, data.

Before proceeding further, note that the realization of the random term  $\varepsilon$  is identifiable from observation Z=z as

$$\varepsilon = m^{-1}(x, z) = F_{Z|X=x^n}^{-1} \left( F_{Z|X=x}(z) \right),$$
 (2)

for each polling station.

Suppose also we observe a cleaner sample in the region with X = x'. Such data could be exit polls, surveys or reports from randomly assigned independent observers in one region.

In the empirical application, we explore the latter. The independent observers might report violations, however, their presence reduces the amount of ballot staffing (as in Enikolopov et al. (2013)). This data allows us to obtain the truthful (or cleaner) distribution of the abstention rate in this region: t = Z - Y. Given that we also can identify  $\varepsilon$  and its distribution, hence we can identify the abstention rate  $t(x', \varepsilon)$  but only conditional on X = x'. So we need more assumptions to carry the inference onto other regions.

We assume that ballot stuffing is reduced by  $d(\varepsilon) + \eta$  in a polling station with random term  $\varepsilon$  in the cleaner sample, so it does not depend on any observable variables and is only a function of the random terms such that  $\varepsilon \perp \!\!\! \perp \eta$  and  $\eta$  is a mean-zero random noise.

Denote by  $Y^{obs}$  the value of Y in the cleaner sample. Then we have the following relationships in the region with X = x':

$$\mathbb{E}(Z - Y^f | X = x', \varepsilon = e) = t(x', e) - q(x', e)$$

$$\mathbb{E}(Z - Y^{obs} | X = x', \varepsilon = e) = t(x', e) - q(x', e) + d(e).$$

Given that  $\varepsilon$  is identified in each polling station, hence, d(e) is identified as well:

$$d(e) = \mathbb{E}(Z - Y^{obs}|X = x', \varepsilon = e) - \mathbb{E}(Z - Y^f|X = x', \varepsilon = e), \tag{3}$$

which also represents the lower bound for ballot staffing in a polling station with  $\varepsilon = e$ . Thus, we can also identify the lower bound of the abstention rate in each polling station

$$\tilde{t}(Y^f, Z, e) = \mathbb{E}(Z - Y^{obs}|X = x, \varepsilon = e) = Z - Y^f + d(e).$$

The function  $d(\cdot)$  will have a different interpretation depending on the type of the cleaner sample. If the additional data is perfectly clean, such as a reliable exit poll or observers reporting no violations, and we believe that ballot stuffing  $q(X,\varepsilon) = q(\varepsilon)$  does not depend on the observable variables and it depends only on the random term, then we are able to evaluate the amount of fraud  $q(\varepsilon) = d(\varepsilon)$ . On the other hand, if the additional sample is the sample with randomized independent observers (with different reports of polling station quality) and we assume that the effect from an observer present at a polling station  $obs(\varepsilon)$  depends only on the random term, then we have  $obs(\varepsilon) = d(\varepsilon)$ , which gives us the lower bound. The counterfactual in this case can be interpreted as a scenario where independent observers are present at all polling stations. Finally, note that this approach can also be applied to biased exit polls.<sup>4</sup>

<sup>&</sup>lt;sup>4</sup>For example, some opposition voters might be reluctant to say publicly that they vote for the opposition, so the exit poll will be biased towards the incumbent. If the exit poll still provides a more accurate measurement of the incumbent's votes than the official (fraudulent) result, and the difference is driven by a

## 3 Estimation

The estimation procedure follows the identification strategy and involves two steps. We assume the overall sample size N, the cleaner sample size  $N_c$  in region X = x' and the rest non-cleaner sample size  $N_f$  in region X = x'. The econometrician observes variables  $X, Y^f$  and Z for each polling station.

First, we estimate Eq. (2) and obtain the unobserved citizens' characteristics  $\hat{\varepsilon}$  in every polling station with observable variables X=x and Z=z. To do so, by following the identification strategy, we have  $\hat{\varepsilon}(x,z)=\hat{F}_{Z|X=x}^{-1}\left(\hat{F}_{Z|X=x^n}(z)\right)$ , where  $\hat{F}_{Z|X=x}(z)$  is an estimator for the cumulative distribution function of Z given X=x. We follow Matzkin (2003) and use a kernel estimator for  $\hat{F}_{Z|X=x}(\cdot)$  that under the standard assumptions has the normal asymptotics and converges with the rate  $\sqrt{Nh^L}$ , where h is the normalized bandwidth and L is the dimension of X. Alternatively, one could use a series estimator as in Imbens and Newey (2009).

Second, we estimate Eq. (3) and obtain the difference between the original sample and the cleaner one,  $d(\cdot)$ , that we assume depends only on  $\varepsilon$  estimated in the first step. Denote  $\mathbb{E}(Z-Y^f|X=x',\varepsilon=e)=d^f(e)$  and  $\mathbb{E}(Z-Y^{obs}|X=x',\varepsilon=e)=d^{obs}(e)$  in the fixed region with X=x'. Then  $d(e)=d^{obs}(e)-d^f(e)$ . Depending on the size of the cleaner sample, we might use either parametric (OLS) or non-parametric (Nadaraya-Watson kernel regression or series regression) approach for estimation of  $d^f(e)$  and  $d^{obs}(e)$ , delivering estimators with rates of convergence  $\sqrt{N_c}$  and  $\sqrt{N_f}$ . However, we do not have the true value of  $\varepsilon$  but its estimate. Thus, to be able to meaningfully use it in a parametric or non-parametric regression, generally,  $\hat{\varepsilon}$  has to converge faster than the corresponding regression estimators. Hence, we must have the cleaner sample's size small in comparison with the actual sample size and that of the overall region's sample. If these and other standard assumptions hold, then often  $\hat{d}(\hat{\varepsilon})$  will have the normal asymptotics.<sup>5</sup>

Note that  $\widehat{d}(\widehat{\varepsilon})$  represents the estimator for the lower bound of ballot stuffing if the cleaner, less fraudulent, sample is not of perfect quality. If the additional data is non-fraudulent, then  $\widehat{d}(\widehat{\varepsilon})$  estimates the actual ballot stuffing based on the random term. The estimator of the turnout (or its bound) follows directly from the estimator  $\widehat{d}(\widehat{\varepsilon})$ . Additionally, Online Appendix B presents a simulation study that evaluates the proposed estimator.

random term, we can use these exit polls to estimate the lower bound of ballot stuffing.

<sup>&</sup>lt;sup>5</sup>See Online Appendix A for a discussion on the asymptotic properties of this estimator when using kernel estimators in both steps.

# 4 Empirical Application

We provide an empirical application of our approach using the 2011 Russian parliamentary election, which transparency and integrity are often questioned (Golos, 2011; OSCE, 2011). The election took place on December 4, 2011, with seven parties participating. The official results reported a significant victory 49% for the incumbent party, United Russia (UR), led by Vladimir Putin.

### 4.1 Dataset

Our data comes from three sources. The primary dataset is from the Central Electoral Commission of the Russian Federation, including the election results, votes each party received, reported turnout, and the number of registered voters for all polling stations.<sup>6</sup>

The second dataset helps us identify polling stations with minimal ballot stuffing. This dataset is sourced from Enikolopov et al. (2013), who conducted a randomized control trial during the same election. They assigned independent observers to 156 random polling stations in Moscow, showing that the presence of observers significantly reduced potential fraud and the incumbent party's voter share. Out of 156 randomly selected polling stations, observers reported no violations in 76 of them. These polling stations form our cleaner sample.

Lastly, since both datasets lack socio-economic characteristics of voters in the registered districts, we use data from the 2010 Russian Census. This dataset includes demographic and economic information (average salaries, occupations), as well as housing and ethnicity details. We observe this information at the *Selsoviet* level, the smallest administrative division in the country, with 21,659 Selsoviets recorded in the 2010 Census. We matched each polling station to its corresponding Selsoviet. After excluding polling stations that could not be precisely matched due to a lack of accurate address information, we retained a total of 71,064 polling stations. Among these, 138 had observers, and in 71 of those, no violations were reported.

Our model assumes a two-party race. To employ our technique, we group all opposition parties and invalid ballots<sup>8</sup> into one variable. Specifically,  $Y^f$  represents the officially reported proportion for UR among registered voters at a given polling station, 1-Z denotes the official proportion for the opposition parties (comprising the remaining six parties and invalid ballots) among registered voters. For regional characteristics (X), we use principal factor analysis to create three factors from 14 variables that capture the socio-economic

<sup>&</sup>lt;sup>6</sup>The data was taken from the website of the Central Electoral Commission when it was publicly available.

<sup>&</sup>lt;sup>7</sup>See the Online Appendix D for details on matching.

<sup>&</sup>lt;sup>8</sup>Ballot spoiling is a way of protest voting that is common in Russia (Panyushkin, 2011; Smyth, 2022).

characteristics of registered voters in each polling station.<sup>9</sup>

Panel A of Table 1 presents the official election results from our sample, which closely align with the full sample results. The incumbent UR received 48.8% of the votes in our sample, compared to 49.3% in the full sample. Among registered voters, the UR's proportion is 28.4%, while the proportion for other parties, including invalid ballots, totals 30.8%, indicating an overall turnout rate of 59.2%.

The second column shows the results for Moscow, the largest region in Russia, with 7.1 million voters. The percentages are similar to those in column (1), except for a larger share of votes for the non-UR group. Columns (3) and (4) present results for Tatarstan and Dagestan, two of the ten most populous regions. In these regions, UR secured a significant majority, with approximately 80% of the votes in Tatarstan and 92% in Dagestan, both of which had turnout rates exceeding 80%.

Panel B shows the mean values of regional characteristics derived from the principal factor analysis. These characteristics are normalized to the sample, resulting in a mean of 0 in column (1). Although they may not have direct interpretations, they can be viewed as socio-economic indexes that capture regional differences. The next three columns display the variations across regions.

### 4.2 Estimation and Results

We first estimate Eq. (2) and recover the individual polling station random terms,  $\varepsilon$ , from the data on the opposition votes using the non-parametric kernel estimators for the distributions involved. We use the midpoint of the support of X as the standard for the normalization.<sup>10</sup> We chose the midpoint based on the expected quality of the estimator of  $F_{Z|X=x^n}(\cdot)$  at such  $X=x^n$ . Figure 1 Panel A plots the density of estimated polling station random terms,  $\varepsilon$ , for the entire sample. The estimated values of  $\varepsilon$  range from 0.25 to 1, with 99% of the sample being above 0.5. The range of  $\varepsilon$  is determined by the normalization and does not affect fraud estimation.<sup>11</sup> These polling station random terms capture unobserved citizens' characteristics (e.g., unobservable component of ideological preferences, voting costs,

<sup>&</sup>lt;sup>9</sup>This technique is similar to the methods used to create Socio-Economic Advantage/Disadvantage indexes for localities by statistical agencies, with the data guiding the number of factors (see the Online Appendix E for details on the principal factor analysis and variables).

<sup>&</sup>lt;sup>10</sup>In the Online Appendix F.1, we use one of the regions for robustness check and show the results remain robust regardless of the choice of normalization.

<sup>&</sup>lt;sup>11</sup>For an example of relationship between the normalization and range of  $\varepsilon$ , consider the following scenario. Suppose X is a one-dimensional variable and the true function  $m(X,\varepsilon) = X(1+\varepsilon)$ , where true  $\varepsilon \in [0,1]$ . We pick a value  $X = x^n$  for normalization, implying that normalized  $\varepsilon^n = x^n(1+\varepsilon)$ , with its range being  $[x^n, 2x^n]$ . Similarly, if instead of value  $x^n$ , we use a different value  $X = \bar{x}$ , then the range of the corresponding  $\bar{\varepsilon}$  would be  $[\bar{x}, 2\bar{x}]$ . However, the normalization does not affect the second step of our estimation, because conditional expectation with the condition  $\varepsilon = c$  is equivalent to the condition  $\varepsilon^n = x^n(1+c)$  or  $\bar{\varepsilon} = \bar{x}(1+c)$  depending on the choice of the normalization and the corresponding range.

civic duty), with higher values of  $\varepsilon$  indicating stronger pro-incumbent preferences or higher voting costs.

The second step is to estimate Eq. (3), for what we use Moscow as a region with our cleaner sample. To obtain the estimator of  $d(\cdot)$ , we compare the conditional means  $E(Z-Y^f|X=x',\epsilon=e)$  without observers and  $E(Z-Y^{obs}|X=x',\epsilon=e)$  with observers, where Z-Y represents the polling station abstention rate. This can be interpreted as the first difference in abstention rates between cleaner polling stations and the rest of Moscow, accounting for observables and  $\varepsilon$ . This is our estimate for "ballot stuffing". We use the Nadaraya-Watson kernel regression for estimation of these conditional means. For the cleaner sample, we first restrict the data from independent observers only to those that reported no violations. We are left with 71 data points that we use to evaluate the conditional mean,  $E(Z-Y^{obs}|X=x',\epsilon=e)$ . The rest of the Moscow sample has 2,882 data points, which are used to evaluate  $E(Z-Y^f|X=x',\epsilon=e)$ . This information provides us with an estimate of ballot stuffing based on the random term  $\varepsilon$ . We plot the estimator of the function  $d(\cdot)$  in Figure 1 Panel B1 across  $\varepsilon$ . We use bootstrap to obtain the 95% confidence interval constructed from 150 iterations.

The estimated function  $d(\cdot)$  indicates that ballot stuffing is more prevalent in polling stations with the highest  $\varepsilon$ . Because  $\varepsilon$  is a residual that represents the aggregation of all unobserved effects, multiple underlying reasons might be contributing to this result. Higher  $\varepsilon$  can be interpreted as reflecting areas with stronger support for the incumbent UR or weaker support for the opposition. This implies that ballot stuffing is more frequent in proincumbent regions, likely because it is easier to manipulate votes there. Higher voting costs, lack of information and lower sense of civic duty all contribute to lower turnout. Hence, when  $\varepsilon$  proxies some of these variables, it would likely lead to increased ballot stuffing, in line with the intuition that areas with lower turnout rates offer more opportunities for fraud, as the "cost of fraud" is lower (i.e., there is a greater margin to add extra ballots).

Given that ballot stuffing depends solely on the random term for each polling station, and not on regional characteristics, we apply the obtained estimator to the remaining polling stations in Russia in the final step. Table 1 Panel C shows that the turnout would have been 46.4% rather than the reported 59.2% if no fraud had occurred. This suggests that electoral fraud across Russia accounts for approximately 12.8% of all registered votes, reducing the proportion of the registered voters for the incumbent UR to 26.8%. In Moscow, the second column shows a slightly higher fraud rate, with the UR's estimated vote proportion at 24.8%. Columns (3) and (4) report higher rates of ballot stuffing in Tatarstan and Dagestan, around

<sup>&</sup>lt;sup>12</sup>Appendix F.3 presents summary statistics comparing the control group (rest of Moscow) with the treatment group (polling stations with no reported violations), showing their similarity. We also construct a pseudo-control group using propensity score matching, which yields results consistent with our baseline estimates.

18%. Figure 2 shows the mean fraud estimates across all 83 regions. Estimates suggest the largest ballot stuffing in Karachay-Cherkess, Mordovia, Dagestan, Chechnya, and Tatarstan, ranging from 18% to 21.6%. All of these regions are on the list of top 15 most fraudulent Russian regions identified in Kobak et al. (2016a) using federal elections data from 2000 to 2012. The least ballot stuffing is observed in Kostroma, Yaroslavl, Oryol, Mari El, and Novgorod Oblasts, ranging from 9.2% to 9.9%. While these estimates represent regional averages, we can also derive the polling station-specific estimates. Figure 3 presents the box plots of fraud estimates for the top 30 most populous regions. The figure also highlights significant variations within the regions, showing that the estimated fraud is not uniformly distributed across polling stations in each region.

A practical question arises: what if we do not know the cleaner sample precisely, but can reasonably guess which polling stations are less likely to be affected by ballot stuffing? To check this, we modify our cleaner sample to include all polling stations with independent observers. While presence of independent observers reduces potential fraud in polling stations, it does not eliminate it entirely. For instance, in our data, serious violations were reported in a quarter of the polling stations with independent observers (Enikolopov et al., 2013). In this case, the estimator provides a lower bound for ballot stuffing. Hence, we are able to evaluate how the presence of observers influences ballot stuffing based on the polling station's random term. This relationship is described by the function  $d(\cdot)$  which is presented in Figure 1 Panel B2. As expected, the estimator flattens, indicating that as cleaner sample (now including reported fraud) becomes more similar to other polling stations, fraud detection becomes less accurate.

Table 1 Panel D presents estimates showing that electoral fraud decreases by approximately 5% across regions, with a notably larger reduction in areas where support for UR is higher. The second column indicates that the UR's estimated voter share in Moscow is around 35% under this specification. These findings align with previous studies, such as (Enikolopov et al., 2013) and (Kobak et al., 2016b), which suggest that the estimated UR's vote share in Moscow should be between 30% and 36%.

Our results suggests that ballot stuffing contributed somewhere from 8.5 million to 14.6 million of votes for UR, which is consistent with 11 million votes detected by Kobak et al. (2012). In terms of parliamentary representation, UR should have secured between 171 (38%) and 203 (45%) parliamentary seats, rather than the 238 (53%) they officially obtained. This indicates that, if relying only on legitimate votes, then UR would have lost its parliamentary majority in the 2011 election.

<sup>&</sup>lt;sup>13</sup>Refer to Online Appendix Figure F4 for the complete set of regions.

## 5 Conclusion

We present a non-parametric structural model that leverages a cleaner sample of electoral data—either from independent observers or exit polls—to identify ballot stuffing or its bounds at each polling station. This methodology addresses the challenge of evaluating fraud on a granular level, which previous studies could not achieve due to limitations in methodology.

While focusing on elections where the prevalent manipulation is ballot stuffing is an obvious limitation of our approach, it allows us to exploit the legitimate votes for the opposition to non-parametrically estimate unobservable polling station random terms associated with unobserved citizens' characteristics such as unobservable component of ideological preferences, cost of voting, information, etc. The cleaner sample extends the procedure to evaluating the relationship between the unobservable terms and ballot stuffing, allowing us to obtain fraud or its bound in each polling station. Hence, while avoiding parametric or statistical assumptions, we are able to conduct a detailed analysis at the polling station level rather than relying solely on regional or aggregate data.

The empirical application of our model to the 2011 Russian parliamentary elections illustrates its practical utility and shows that fraud is most prevalent in polling stations with stronger ideological leaning and/or voting costs.

Overall, our approach provides an additional tool for election monitors, researchers, and policymakers to better assess and address ballot stuffing. It offers a novel way to leverage cleaner data to uncover hidden fraud and make informed decisions about the strategic allocation of resources, such as independent observers or analysts reviewing election results, to mitigate electoral fraud in future elections.

# References

Aksoy, O. (2016). Regression discontinuity diagnostics reveal statistical anomalies in Turkish elections. *Electoral Studies* 44, 284–292.

Asunka, J., S. Brierley, M. Golden, E. Kramon, and G. Ofosu (2019). Electoral fraud or violence: The effect of observers on party manipulation strategies. *British Journal of Political Science* 49(1), 129–151.

BBC News (2023, October 5). Former US congressman guilty of ballot-stuffing in Philadelphia. BBC News. https://www.bbc.com/news/world-us-canada-68025894.

Bloomberg News (2024, July 19). France's lawmakers shocked by irregular ballot at

- session start. *Bloomberg*. https://www.bloomberg.com/news/articles/2024-07-19/france-s-lawmakers-shocked-by-irregular-ballot-at-session-start.
- Bound, J., C. Brown, G. Duncan, and W. Rodgers (1989). Measurement error in cross-sectional and longitudinal labor marker surveys: results from two validation studies. *NBER Working Paper 2884*.
- Breunig, C. and A. Goerres (2011). Searching for electoral irregularities in an established democracy: Applying Benford's law tests to Bundestag elections in unified Germany. *Electoral Studies* 30, 534–545.
- Cantu, F. (2014). Identifying irregularities in Mexican local elections. *American Journal of Political Science* 58(4), 936–951.
- Cantu, F. and S. Saiegh (2011). Fraudulent democracy? An analysis of Argentina's infamous decade using supervised machine learning. *Political Analysis* 19(4), 409–433.
- Carroll, R., X. Chen, and Y. Hu (2010). Identification and estimation of nonlinear models using two samples with nonclassical measurement errors. *Journal of Nonparametric Statistics* 22(4), 379–399.
- Carroll, R. and M. Wand (1991). Semiparametric estimation in logistic measurement error models. *Journal of the Royal Statistical Society B* 53, 573–585.
- Casas, A., G. Diaz, and A. Trindade (2017). Who monitors the monitor? Effect of party observers on electoral outcomes. *Journal of Public Economics* 145, 136–149.
- Chen, X., H. Hong, and E. Tamer (2005). Measurement error models with auxiliary data. *Review of Economic Studies* 72, 343–366.
- Degan, A. and A. Merlo (2011). A structural model of turnout and voting in multiple elections. *Journal of the European Economic Association* 9(2), 209–245.
- Devroye, L. (1978). The uniform convergence of the Nadaraya-Watson regression function estimate. The Canadian Journal of Statistics 6(2), 179–191.
- Downs, A. (1957). An economic theory of political action in a democracy. *Journal of Political Economy* 65(2), 135–150.
- 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.

- Fujiwara, T., K. Meng, and T. Vogl (2016). Habit formation in voting: Evidence from rainy elections. *American Economic Journal: Applied Economics* 8(4), 160–188.
- Gehlbach, S., K. Sonin, and M. Svolik (2016). Formal models of nondemocratic politics.

  Annual Review of Political Science 19, 565–584.
- Golos (2011). Preliminary report on the State Duma election.
- Guriev, S. and D. Treisman (2022). Spin Dictators: The Changing Face of Tyranny in the 21st Century. Princeton University Press.
- Hausman, J., H. Ichimura, W. Newey, and J. Powell (1991). Identification and estimation of polynomial errors-in-variables models. *Journal of Econometrics* 50, 273–295.
- Hsiao, C. (1989). Consistent estimation for some nonlinear errors-in-variables models. *Journal of Econometrics* 41, 159–185.
- Hu, Y. and G. Ridder (2012). Estimation of nonlinear models with mismeasured regressors using marginal information. *Journal of Applied Econometrics* 27, 347–385.
- Imbens, G. and W. Newey (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica* 77(5), 1481–1512.
- Imbens, G. and J. Wooldridge (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature* 47(1), 5–86.
- Katz, J. and G. Katz (2010). Correcting for survey misreports using auxiliary information with an application to estimating turnout. *American Journal of Political Science* 54(3), 815–835.
- Kawai, K., Y. Toyama, and Y. Watanabe (2021). Voter turnout and preference aggregation. American Economic Journal: Microeconomics 13(4), 548–586.
- King, G., M. A. Tanner, and O. Rosen (2004). *Ecological inference: New methodological strategies*. Cambridge University Press.
- Klimek, P., Y. Yegorov, R. Hanel, and S. Thurner (2012). Statistical detection of systematic election irregularities. *Proceedings of the National Academy of Sciences* 109(40), 16469–16473.
- Kobak, D., S. Shpilkin, and M. Pshenichnikov (2012). Statistical anomalies in 2011-2012 russian elections revealed by 2d correlation analysis. pp. arXiv:1205.0741v2 [physics.soc-ph].

- Kobak, D., S. Shpilkin, and M. Pshenichnikov (2016a). Integer percentages as electoral falsification fingerprints. *Annals of Applied Statistics* 10, 54–73.
- Kobak, D., S. Shpilkin, and M. Pshenichnikov (2016b). Statistical fingerprints of electoral fraud? Significance 13(4), 20–23.
- Kobak, D., S. Shpilkin, and M. Pshenichnikov (2018). Putin's peaks: Russian election data revisited. Significance 15(3), 8–9.
- Lee, L. and J. Sepanski (1995). Estimation of linear and nonlinear errors-in-variables models using validation data. *Journal of the American Statistical Association* 90, 130–140.
- Leeffers, S. and P. Vicente (2019). Does electoral observation influence electoral results? Experimental evidence for domestic and international observers in Mozambique. *World Development* 114, 42–58.
- Leon, G. (2017). Turnout, political preferences and information: Experimental evidence from Peru. *Journal of Development Economics* 127, 56–71.
- Levin, I., G. Cohn, P. Ordeshook, and R. Alvarez (2009). Detecting voter fraud in an electronic voting context: An analysis of the unlimited reelection vote in Venezuela.
- Matzkin, R. (2003). Nonparametric estimation of nonadditive random functions. *Econometrica* 71(5), 1339–1375.
- Mebane, W., D. Ferrari, K. McAlister, and P. Wu (2022). Measuring election frauds. *Working paper*.
- Mebane, W. J. (2008). Election forensics: The second-digit Benford's law test and recent American presidential elections. pp. 162–181.
- OSCE (2011). Statement of preliminary findings and conclusions: Russian Federation State Duma elections.
- Panyushkin, V. (2011, December 12). A united russia? far from it. *The New York Times*, A38. Retrieved from https://www.nytimes.com.
- Pepe, M. and T. Fleming (1991). A general nonparametric method for dealing with errors in missing or surrogate covariate data. *Journal of the American Statistical Association* 86, 108–113.
- Persson, T. and G. Tabellini (2000). *Political Economics : Explaining Economic Policy*. MIT Press Ltd.

- Riker, W. and P. Ordeshook (1968). A theory of the calculus of voting. *American Political Science Review* 62(1), 25–42.
- Rozenas, A. (2017). Detecting election fraud from irregularities in vote-share distributions. *Political Analysis* 25(1), 41–56.
- Schennach, S. (2016). Recent advances in the measurement error literature. *Annual Reviews of Economics* 8, 341–377.
- Schennach, S. (2022). Measurement systems. *Journal of Economic Literature* 60(4), 1223–1263.
- Skovoroda, R. and T. Lankina (2017). Fabricating votes for Putin: new tests of fraud and electoral manipulations from Russia. *Post-Soviet Affairs* 33(2), 100–123.
- Smyth, R. (2022). *Elections, protest, and authoritarian regime stability*. Cambridge University Press.

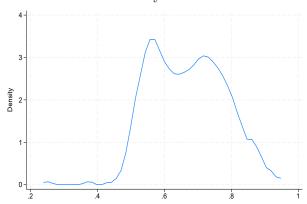
Table 1: Descriptive Results and Empirical Estimation of the 2011 Russian Parliamentary Election

	All Russia [1]	Moscow [2]	Tatarstan [3]	Dagestan [4]			
Panel A: Official Results							
Voter Share for UR	48.8%	47.3%	79.6%	91.6%			
Proportion for UR among Reg Voters	28.3%	28.3%	62.9%	83.6%			
Proportion for Non-UR among Reg Voters	30.8%	32.6%	16.8%	7.8%			
Turnout Rate	58.2%	59.9%	79.1%	91.4%			
Number of Polling Stations	71,064	2,952	2,631	1,661			
Panel B: Region Characteristics							
Number of Unique Selsoviet	13,964	127	867	644			
Mean of First Factor	0	-0.73	0.27	-0.11			
Mean of Second Factor	0	2.32	-0.39	-0.59			
Mean of Third Factor	0	-0.22	0.11	0.74			
Panel C: Estimation Results I: Cleaner Sa	mple = Polling	g Stations R	teporting No V	violations violations			
Estimated Turnout Rate	46.4%	47.5%	61.7%	72.6%			
Estimated Fraud/Registered Voters	12.8%	13.4%	18.0%	18.7%			
Estimated UR Vote /Registered voters	15.6%	14.9%	44.9%	64.8%			
Estimated Voter Share for UR	33.6%	31.4%	72.7%	89.2%			
Panel D: Estimation Results II: Cleaner Sample = Polling Stations with Independent Observers							
Estimated Turnout Rate	51.6%	53.3%	70.6%	81.8%			
Estimated Fraud/Registered Voters	7.5%	7.6%	9.12%	9.55%			
Estimated UR Vote /Registered voters	20.8%	20.7%	53.8%	74.1%			
Estimated Voter Share for UR	40.3%	38.8%	76.2%	90.4%			

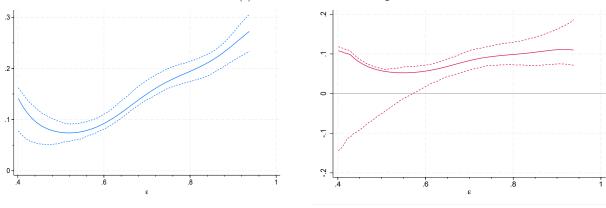
Notes: Table 1 Panel A presents the official results reported by the Central Electoral Commission of the Russian Federation. Panel B summarizes regional characteristics, with principal factor numbers derived from the CENSUS data (see Appendix E for details). Panel C shows estimation results using a cleaner sample of polling stations with independent observers who reported no violations. Panel D provides estimation results based on cleaner polling stations with all observers. See Figure F4 for the fraud estimates across the remaining regions.

Figure 1: Estimates of  $\varepsilon$  and  $d(\cdot)$ 

Panel A: Density of estimated  $\varepsilon$ 



Panel B: Estimator of  $d(\cdot)$  with 95% bootstrap confidence intervals



**B1:** No Violations **B2:** All Observers

Notes: Panel A presents the estimation results of Eq. (2) and shows the kernel density of estimated  $\varepsilon$ . The median of the distribution is 0.66, and less than 1% of the density lies below 0.4. Panel B consists of two sub-panels. In both sub-panels, the solid line represents the estimated function  $d(\cdot)$  across  $\varepsilon$ , while the dashed lines show the 95% bootstrapped confidence intervals. Panel B1 estimates  $d(\cdot)$  using polling stations with independent observers reporting no violations. Panel B2 estimates  $d(\cdot)$  using all polling stations with independent observers.

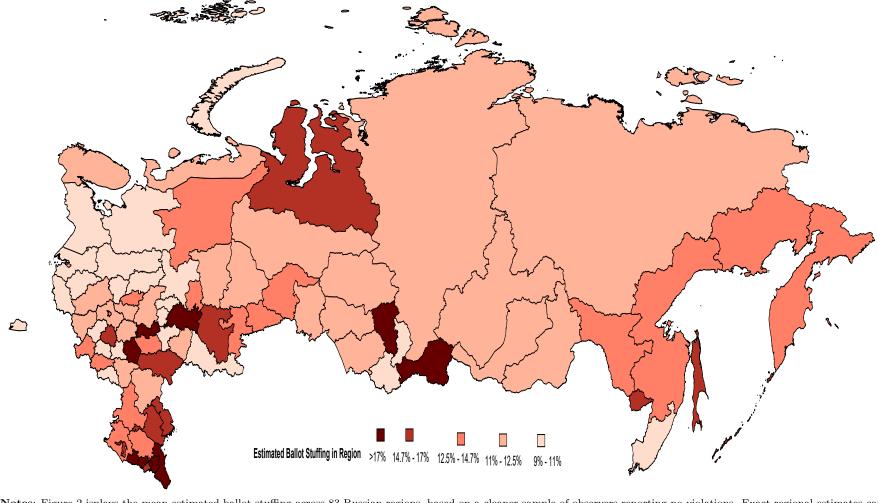
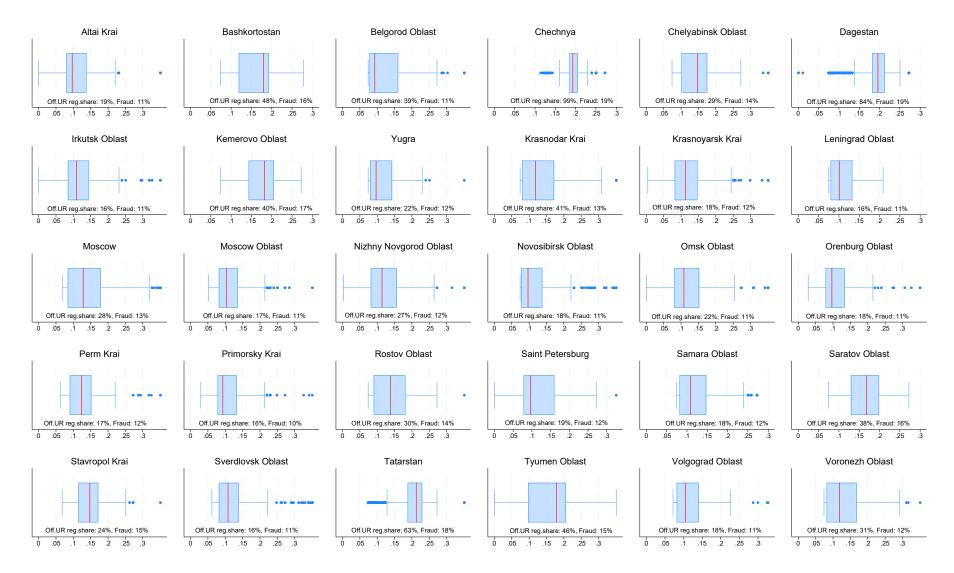


Figure 2: Ballot Stuffing Estimates Across Regions

Notes: Figure 2 isplays the mean estimated ballot stuffing across 83 Russian regions, based on a cleaner sample of observers reporting no violations. Exact regional estimates can be found in Appendix Figure F4.

Figure 3: Estimated Ballot Stuffing Across the Top 30 Most Populous Regions



Notes: Figure 3 shows the estimated ballot stuffing for the top 30 most populous regions in Russia, using the cleaner sample consisting of observers who reported no violations. For estimates of the remaining regions, refer to Appendix Figure F4. The box represents the interquartile range (IQR), covering the middle 50% of the data. The red line inside the box marks the median. The whiskers extend to the minimum and maximum values, excluding outliers, which are shown as individual points beyond the whiskers. The first number at the bottom of the figure, "Off. UR reg share," represents the official proportion of votes for UR among registered voters. The second number, "Fraud," indicates the total estimated fraudulent votes as a proportion of the registered votes.

# Online Appendices - Not For Publication

Appendix A: Asymptotics of  $\widehat{d}(\widehat{\varepsilon})$ Appendix B: Simulations of  $\widehat{d}(\widehat{\varepsilon})$ 

Table B1: Simulation Results

Figure B1: Simulation Results

Appendix C: An Underlying Structural Model

Appendix D: Data Appendix

Table D1: Summary of Sample Sizes at Different Thresholds

Figure D1: Census Data Russia Map (ADR2)

Figure D2: Moscow Polling Stations

Figure D3: Russia Polling Stations

Appendix E: Principal Component Analysis

Table E1: Eigenvalues and Randomly Generated Eigenvalues

Table E2: Factor Analysis/Correlation: Iterated Principal Factor

Table E3: Full Factor Loadings Using Principal Factor Analysis and Varimax Rotation

Figure E1: Parallel Analysis

Appendix F: Additional Results and Robustness Checks

Table F1: Sample Robustness Checks

Table F2: Summary Statistics for Treatment and Control Groups

Table F3: Results using New Control Group

Figure F1: Estimates of  $\varepsilon$  using the midpoint vs Altai Krai normalization

Figure F2: Estimator of  $d(\cdot)$  using the Altai Krai normalization

Figure F3: Estimated function  $d(\cdot)$  across  $\varepsilon$  for the robustness samples

Figure F4: Estimated Ballot Stuffing Across Russia

# A Asymptotics of $\widehat{d}(\hat{\varepsilon})$ for kernel estimators

Let  $\hat{\varepsilon}(x,z) = \widehat{F_{Z|X=x}^{-1}}\left(\widehat{F_{Z|X=x}}(z)\right)$  be the estimator of the unobserved term, using the standard kernel estimators of the cumulative distribution function and its inverse. We will abuse the notation by dropping the dependence of  $\varepsilon$  on x and z in the below. Under the set of standard assumptions, this estimator converges uniformly with rate  $\sqrt{Nh^L}$  (see Theorem 2 in Matzkin (2003)). Denote the asymptotic variance by  $V_{\varepsilon}$  and the normalized bandwidth by h.

Let  $\hat{d}^{obs}(e) = \mathbb{E}(Z - Y^{obs}|X = x, \varepsilon = e)$  and  $\hat{d}^f(e) = \mathbb{E}(Z - Y^f|X = x, \varepsilon = e)$  be the Nadaraya-Watson kernel regression estimators. Under a set of standard assumptions, they converge uniformly with rates  $\sqrt{N_c}$  and  $\sqrt{N_f}$ , correspondingly (see, for example, Devroye (1978)). Denote the corresponding normalized bandwidths by  $h_c$  and  $h_f$ .

**Theorem 1.** Suppose the following set of assumptions holds

- 1.  $X_i$  are iid,  $\{Y_i^f, Z_i\}$  are iid conditional on X, and  $Y_i^{obs}$  are iid in region with  $X_i = x$ ;
- 2. Density  $f_{X,Y,Z}$  has compact support in  $\mathbb{R}^{L+2}$  and is continuously differentiable up to the order s' > 0;
- 3. The kernel function K is differentiable of order s'', the derivatives of K of order s'' are Lipschitz, K vanishes outside a compact set, integrates to 1, and is of order s''', where  $s'' + s''' \le s'$  and  $s'', s''' \ge 2$ ;
- 4. As  $N \to \infty$ ,  $N_c \to \infty$ ,  $N_f \to \infty$ ,  $h \to 0$ ,  $h_c \to 0$ ,  $h_f \to 0$ ,  $\frac{N_c}{Nh^L} \to 0$ ,  $\frac{N_c}{N_f} \to 0$ ,  $\frac{\ln N_c}{Nh^{L+1}} \to 0$ ,  $\frac{\ln N_c}{N_c h_c^3} \to 0$ ,  $\frac{\ln N_f}{N_f h_f^3} \to 0$ ,  $\sqrt{Nh^L} \to \infty$ ,  $\sqrt{N_c h_c^2} \to \infty$ ;  $\sqrt{N_f h_f^2} \to \infty$ ,  $\sqrt{N_f h_f^2} \to \infty$ ,  $\sqrt{N_f h_f^2} \to 0$ , and  $\sqrt{Nh^L} \left(\sqrt{\frac{\ln N}{Nh^{L+1}}} + h^{s'''}\right)^2 \to 0$ ;
- 5. The values,  $x^n$  and x, are different at least at one coordinate;

then  $\widehat{d}(\widehat{\varepsilon}) \stackrel{p}{\to} d(\varepsilon)$  and  $\sqrt{N_c}(\widehat{d}(\widehat{\varepsilon}) - d(\varepsilon)) \stackrel{d}{\to} N(0, V_d(\varepsilon))$ , where  $V_d(\varepsilon) > 0$  is the asymptotic variance.

Proof. Denote  $\hat{d}(e) - d(e) = h_d(e)$  and  $\hat{\varepsilon} - \varepsilon = h_{\varepsilon}$  such that the differences are sufficiently small. Then we have that there exist constants  $0 < a_1, a_2 < \infty$  such that  $|h_d(e)| \le a_1||H_d||$  and  $|h_{\varepsilon}| \le a_2||H_{\varepsilon}||$ , where  $||\cdot||$  denotes the supnorm. Because of uniform convergence of the estimators, there exist constants  $0 < a_3, a_4 < \infty$  such that  $|\hat{d}^{obs}(\varepsilon) - d^{obs}(\varepsilon)| \le a_3||H_{d^{obs}}||$  and  $|\hat{d}^f(\varepsilon) - d^f(\varepsilon)| \le a_4||H_{d^f}||$ . Also denote by  $||H|| = \max(||H_d||, ||H_{\varepsilon}||, ||H_{d^{obs}}||, ||H_{d^f}||)$ . Hence,  $|h_d(e)| \le a_1||H||$ ,  $|h_{\varepsilon}| \le a_2||H||$ ,  $|\hat{d}^{obs}(\varepsilon) - d^{obs}(\varepsilon)| \le a_3||H||$  and  $|\hat{d}^f(\varepsilon) - d^f(\varepsilon)| \le a_4||H||$ .

Next, note the following:

$$\begin{split} \widehat{d}(\widehat{\varepsilon}) - d(\varepsilon) &= \widehat{d}(\widehat{\varepsilon}) - \widehat{d}(\varepsilon) + \widehat{d}(\varepsilon) - d(\varepsilon) \\ &= \frac{\partial d(\varepsilon)}{\partial \varepsilon} (\widehat{\varepsilon} - \varepsilon) + \frac{\partial h_d(\varepsilon)}{\partial \varepsilon} (\widehat{\varepsilon} - \varepsilon) + \widehat{d}(\varepsilon) - d(\varepsilon) + O(||H||^2) \\ &= \frac{\partial d(\varepsilon)}{\partial \varepsilon} (\widehat{\varepsilon} - \varepsilon) + \widehat{d}^{obs}(\varepsilon) - d^{obs}(\varepsilon) - \widehat{d}^f(\varepsilon) + d^f(\varepsilon) + Res, \end{split}$$

where  $Res = O(||H||^2)$ .

Thus, there exist constants  $0 < a_5, a_6 < \infty$  such that  $|\widehat{d}(\hat{\varepsilon}) - d(\varepsilon)| \le a_5||H|| + a_6||H||^2$ . Given that  $||H|| \stackrel{p}{\to} 0$ , we have  $|\widehat{d}(\hat{\varepsilon}) - d(\varepsilon)| \stackrel{p}{\to} 0$ .

Next, we have

$$\sqrt{Nh^L} \frac{\partial d(\varepsilon)}{\partial \varepsilon} (\hat{\varepsilon} - \varepsilon) \stackrel{d}{\to} N \left( 0, \left( \frac{\partial d(\varepsilon)}{\partial \varepsilon} \right)^2 \times V_{\varepsilon} \right).$$

Also,  $Res = o(h_{\varepsilon})$ , so  $\sqrt{N_c}(\frac{\partial d(\varepsilon)}{\partial \varepsilon}h_{\varepsilon} + o(h_{\varepsilon})) \xrightarrow{p} 0$ , implying that the resulting asymptotics is determined by the slowest term. By the assumption on convergence rates, the slowest term is  $\hat{d}^{obs}(\varepsilon) - d^{obs}(\varepsilon)$ .

# **B** Simulations

In this section, we provide a small simulation experiment of the two-step estimation procedure. We generate regional characteristic  $X \sim U[0,1]$  for  $n_r + 1$  regions, where  $n_r$  of them have only one polling station and one large region has  $n_m$  polling stations. We also generate  $n_r + n_m$  unobservable random terms  $\varepsilon \sim U[0,1]$  that are independent of X. We obtain the electoral variables Y and Z as follows:

$$Y = \frac{1}{2}X^2 + \frac{3}{8}\varepsilon^2$$
$$Z = \frac{1}{2}X^2 + \frac{1}{2}\varepsilon.$$

We assume that ballot stuffing depends only on  $\varepsilon$  as  $d(\varepsilon) = 0.3 + 2(\varepsilon - 0.5)^3$ , implying that the observable  $Y^f = Y + d(\varepsilon)$ . For half of the large region polling stations  $(n_m/2)$  we observe the truthful Y, so they represent the cleaner polling stations.

We run 1,000 simulations for different combinations of values for  $n_r$  and  $n_m$ , where  $n_r$  takes values of 500 and 1,000 and  $n_m$  takes values of 250 and 500. For each simulation, we use the value of X = 0.5 for normalization and estimate  $\hat{\varepsilon}$  for each polling station. Next, we use the sample from the large region and estimate  $d(\cdot)$  at 100 fixed points, which we drew from a uniform distribution with the support [0, 1]. We use kernel estimation for both steps. For each simulation, we calculated the absolute value of bias, variance, and mean squared error. We report averages of these values in Table B1.

Figure B1 also demonstrates the average over 1,000 simulations of the estimated function  $d(\cdot)$  (the solid line) together with its true value (the dashed line) and the 2.5% and 97.5% percentiles (the dot-dashed lines) for  $n_r = 1,000$  and  $n_m = 500$ .

# C An Underlying Structural Model

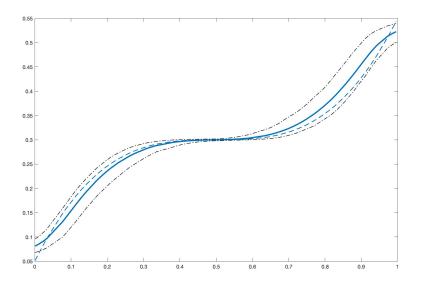
This section presents a simple example of an underlying structural model related to the calculus of voting (Downs, 1957; Riker and Ordeshook, 1968) that could have produced the

Table B1: Simulation Results

Specification	Bias	Variance	MSE
$n_r = 500, n_m = 250$	0.0120	0.0117	0.0120
$n_r = 500, n_m = 500$	0.0114	0.0115	0.0118
$n_r = 1,000, n_m = 250$	0.0115	0.0101	0.0104
$n_r = 1,000, n_m = 500$	0.0110	0.0094	0.0097

**Notes**: Table B1 presents the simulation results for different values of  $n_r$  ( $n_r + 1$  being the number of regions) and  $n_m$  (the number of polling stations in the large region). The metrics include the absolute bias, variance, and mean squared error (MSE) for each scenario.

Figure B1: Simulation Results



**Notes**: Average behavior of the estimator for function  $d(\varepsilon)$  based on 1,000 simulations with  $n_r = 1,000$  and  $n_m = 500$ . True value is shown by the blue dashed line,  $d(\cdot)$  is the blue solid line. Black dot-dash lines present the 2.5% and 97.5% percentiles confidence interval.

data we study.

Two candidates, A (incumbent) and B (opposition), are running for office. Each voter has preferences over candidates: Similarly to probabilistic voting in Persson and Tabellini (2000), voter i in a polling station j in region K chooses candidate A over B if

$$\sigma_A^{ijK} + \mu_A^K > \sigma_B^{ijK} + \mu_B^K,$$

where  $\sigma_t^{ijK}$  is a parameter of individual "pure" preferences toward candidate t of voter i, and  $\mu_t^K$  is a regional effect in the popularity of candidate t and is a function of some observable characteristics of the region X, such as average income, level of education, share of old population, etc. Thus,  $\mu_t^K = \tilde{h}_t(X)$ .

To obtain the reduced form of the model, we define parameters of difference in preferences

between candidates  $\sigma = \sigma_B^{ijK} - \sigma_A^{ijK}$  and  $\mu^K = \mu_B^K - \mu_A^K = \tilde{h}_B(X) - \tilde{h}_A(X) \equiv h(X)$ . Therefore, voter i in a polling station j in region K with observable characteristics X chooses candidate A over B if

$$\sigma + h(X) < 0,$$

where  $\sigma$  is personal "pure" preference for the candidate B and h(X) is the regional effect on preferences.

To include turnout in the model, we follow the empirical evidence that suggests voting costs affect electoral participation (Fujiwara et al., 2016; Leon, 2017). Similar to Kawai et al. (2021), a voter chooses to participate in elections if the difference in her preferences from other candidates is higher than the costs of participation:

$$|\sigma + h(X)| \ge c$$
.

Participation costs c are random and the same for all voters in the same polling station, but might be different across different polling stations. Costs might represent the length of the line to vote, the weather, the difficulty of obtaining a voter card, etc. Such representation of participation implies that if a voter's preferences are close to indifference between the candidates, she does not attend elections. And, by contrast, if a person has very strong preferences toward one candidate or another, she comes to the polling station even when costs are high.<sup>14</sup>

In addition, we assume individual "pure" preferences  $\sigma$  are independent identically distributed variables with density  $g(\cdot)$  and cumulative distribution  $G(\cdot)$ . Costs of voting c are independent identically distributed variables with density  $f_c$ .

Next, we introduce "swing voters",  $\sigma_A$  and  $\sigma_B$ , in every polling station, who are indifferent between participating and not participating in elections:

$$\sigma_A = -h(X) - c$$
 and  $\sigma_B = \sigma_A + 2c$ .

Notice people with "pure" preferences  $\sigma < \sigma_A$  will vote for candidate A, people with  $\sigma > \sigma_B$ will vote for candidate B, and everybody in between the swing voters will abstain from elections. As a result, the number of people who vote for A in a polling station is  $n_A =$  $\int_{-\infty}^{\sigma_A} dG(x) = G(\sigma_A)$ . The same number for candidate B is  $n_B = \int_{\sigma_B}^{+\infty} dG(x) = 1 - G(\sigma_B)$ . Thus, turnout in the polling station is  $\tau = 1 - G(\sigma_B) + G(\sigma_A)$ , and A's share of votes is  $\pi_A = \frac{n_A}{n_A + n_B} = \frac{G(\sigma_A)}{1 - G(\sigma_B) + G(\sigma_A)}$ . Hence, this model produces our variables of interest Y and Z as follows:

$$Y = G(\sigma_A)$$
 and  $Z = G(\sigma_B)$ 

or in the following form:

$$G(-h(X) - c) = m(X, \varepsilon) - t(X, \varepsilon)$$
  
$$G(-h(X) + c) = m(X, \varepsilon).$$

<sup>&</sup>lt;sup>14</sup>We recognize this model does not account for "marginal voter" thinking, i.e., when voters believe their vote does not matter and abstain from elections. The data show people vote, and we do not attempt to contribute to the divisive question of why.

Note that when we use normalization  $m(X^n, \varepsilon) = \varepsilon$ , it has the following connection with the cost of voting in this model:

$$G(-h(X^n) + c) = \varepsilon.$$

Thus,  $\varepsilon$  is a positive transformation of the cost of voting c that could be further identified if necessary.

Note that this model is only an example of an underlying data generating process.

# D Data Appendix

## D.1 Creating the Dataset

To construct our final dataset used in Section 4, we draw on three sources: the Russian Census (2010), the Russian Presidential Election (March 2012), and the Russian Parliamentary Election (December 2011). The primary task with these three datasets is to obtain demographic controls for the 2011 parliamentary elections. The main challenge is that the 2011 parliamentary election data does not include precise polling station locations. Instead, it provides information on the Region, Territorial Electoral Commission (TIK), and polling station number (UIK). In larger regions, an additional layer of geographic information is available. Because the difference between the parliamentary and presidential elections is only 3 months, the majority of the polling stations are expected to stay the same.

We begin by matching polling stations between elections. This enables us to align the address information of polling stations available from the 2012 presidential election data. The process of matching 2011 parliamentary polling stations to 2012 presidential polling stations is based on the Region and Territorial Electoral Commission (TIK). Out of 86 unique Region values, 83 were successfully matched across the two elections. The remaining three correspond to foreign locations outside of Russia. Additionally, by grouping the two datasets by Region, all unique TIK values could be matched across the elections. From there, we matched UIK numbers within TIKs, resulting in a dataset of 100,321 observations. However, 10,396 observations lack a corresponding polling station in either the presidential or parliamentary election datasets. This is likely due to new polling stations being added during the presidential election or polling stations from the parliamentary election being discontinued in the presidential election.

We also assessed fluctuations in voter registration across the two elections for the matched polling stations. The thresholds range from  $\pm$  10% to  $\pm$  100%, reflecting the extent to which registration could fluctuate between the two elections. Table D1 provides the corresponding sample sizes.

For our estimation sample, we restrict the data to a  $\pm$  20% threshold. After matching polling stations, we use the location of each UIK to align with the 2010 census data at the Selsoviet level (ADR3). We exclude 762 polling stations where the census matching is deemed unreliable, such as cases where the population count is smaller than the number of registered voters. The final sample includes 71,064 polling stations, representing 79% of the total sample.

Threshold	Sample Size
$\pm 10\%$	59,366
$\pm 20\%$	71,826
$\pm 50\%$	81,062
$\pm 80\%$	84,841
$\pm 100\%$	86,377
No Threshold	89,925

Table D1: Summary of Sample Sizes at Different Thresholds

## D.2 Visual Representation of Matching

The census data provides two levels of administrative district classifications: ADR2 and ADR3. ADR3 represents a more localized administrative region nested within ADR2. While Figure D1 shows the ADR2 regions on a map of Russia, we use the more granular ADR3 data, which is more challenging to visualize. To illustrate the matching process, Figure D2 presents an example of the ADR2 districts around central Moscow, with polling stations from the presidential dataset overlaid. Figure D3 shows the matched polling station locations in the 2011 parliamentary election.



Figure D1: Census Data Russia Map (ADR2)

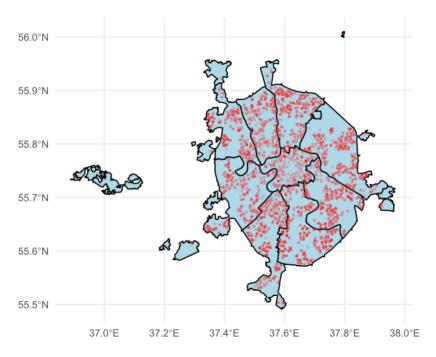


Figure D2: Moscow Polling Stations

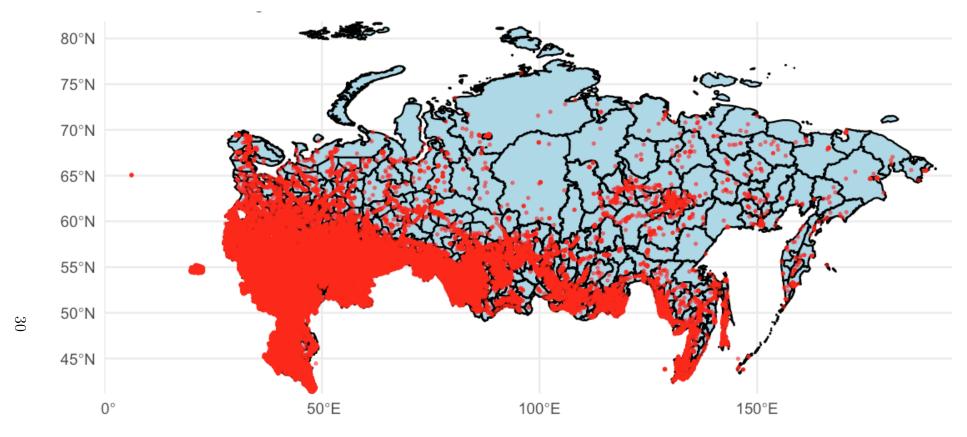


Figure D3: Russia Polling Stations

#### D.3 Census Information

The Russian census data contains information pertaining to demographic, economic, housing, migration, transportation and social factors. Note that not all information is available for every region (ADR2 and ADR3). Therefore, when selecting variables, we restrict our analysis to data that is widely available across the regions. As a result, we focus on the following variables:

- Total Population
- Percentage of Rural Population
- Percentage of Female Population
- Percentage of Population Change Compared to the Last Census
- Average Salary
- Average Government Salary
- Percentage of Families Receiving Government Subsidies
- Number of Government Employees
- Percentage of Disabled People
- Percentage of School-Aged Children
- Number of Schools
- Number of Kindergartens
- Average Housing in Square Metres

Note that in section F.2, we present results using all available information from the CENSUS, while accounting for missing data in certain regions.

# E Principal Component Analysis

Principal Component Analysis (PCA) is a statistical technique used to summarise a dataset containing numerous correlated variables into a smaller set of uncorrelated components, each being a linear combination of the original variables. There are as many principal components as there are variables.

In situations where the original variables are highly correlated, PCA enables a significant reduction in dimensionality by capturing most of the variation in a few principal components, thus simplifying further analysis. The first principal component is responsible for the largest proportion of variance in the dataset, and each subsequent component explains a progressively smaller portion of the variance.

The PCA process assigns an eigenvalue to each component, which indicates the amount of variance that component captures from the original data. The proportion of variance

explained by any principal component is calculated by dividing its eigenvalue by the sum of all eigenvalues.

Additionally, the "loading" for a variable on a principal component is determined by multiplying the eigenvector associated with that component by the square root of its eigenvalue. This loading measures the strength of the relationship between the variable and the component.

Section D.3 lists the variables included in PCA analysis. Below, we provide a step-bystep explanation of how we conducted the PCA analysis and why we chose to focus on three factors.

(1) Decide whether it is appropriate to perform exploratory factor analysis. This is done by examining the 'factorability' of the correlation matrix.

Barlett test of sphericity: The null hypothesis for this test is an identity matrix, so rejecting the null means you infer non-zero correlations among the variables.

P-value of the test: 0.000. This result suggests the matrix is factorable.

Kaiser-Meyer-Olkin Measure of Sampling Adequacy: The Kaiser-Meyer-Olkin Measure (KMO) is another index for assessing the appropriateness of factor analysis. KMO ranges from 0 to 1, with higher values providing greater justification for factor analysis of the correlation matrix. As described in the Stata manual.

KMO = 0.903, which provides additional evidence for the appropriateness of conducting factor analysis on our variables.

(2) How many factors to retain?

We use parallel analysis to determine how many factors to retain. Parallel analysis compares eigenvalues from the PCA with randomly generated eigenvalues. As shown in Table E1, the eigenvalues associated with first 3 principal components are larger than randomly generated ones. This suggests retention of three factors.

Figure E1 shows an overlay of the randomly generated eigenvalues and shows a 'break-point' at three factors. This helps to visualize where the randomly generated eigenvalues exceed the eigenvalues computed from the data.

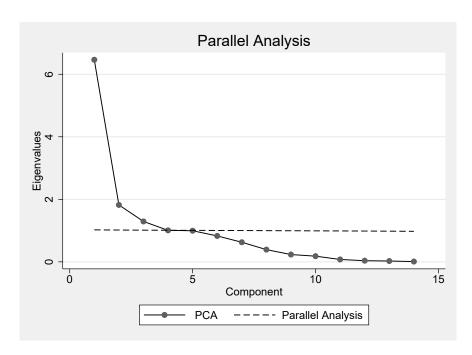


Figure E1: Parallel Analysis.

Table E1: PCA Eigenvalues and Randomly Generated Eigenvalues.

	Principal Comp. Analysis	Parallel Analysis	Difference
	Eigenvalues	Random Eigenvalues	Difference
1	6.46429	1.023207	5.441083
2	1.821781	1.018152	0.8036298
3	1.292837	1.014183	0.2786534
4	1.00793	1.010738	-0.0028085
5	0.995452	1.007542	-0.0120899
6	0.8300372	1.004558	-0.1745212
7	0.6277536	1.001498	-0.3737449
8	0.3916343	0.9985071	-0.6068728
9	0.2340445	0.9954785	-0.7614341
10	0.1825673	0.9923968	-0.8098295
11	0.0769674	0.9892223	-0.912255
12	0.0380089	0.9856929	-0.947684
13	0.0272562	0.9818299	-0.9545738
14	0.0094408	0.976993	-0.9675523

### (3) Extract the factors and varimax rotation (a type of orthogonal rotation)

Interpreting a factor analytical solution can be challenging, largely due to the inherent indeterminacy in factor analysis. This indeterminacy is often addressed through orthogonal transformations of the common factors and their associated factor loadings. This approach preserves the property that common factors remain uncorrelated and does not affect the reconstructed (fitted) correlation matrix. Table E2 shows that Factor 1 explains 65.14% of the total observed variance in our factor model. Table E3 presents the final loadings used to predict the factor loadings for each factor.

Table E2: Factor Analysis/Correlation: Iterated Principal Factor.

Factor	Variance	Difference	Proportion	Cumulative
Factor1	5.55572	3.26436	0.6514	0.6514
Factor2	2.29136	1.60986	0.2687	0.9201
Factor3	0.68150		0.0799	1.0000
N = 71,	064; LR tes	st: independe	ent vs. saturat	ted: P-value=0.0000

Table E3: Full Factor Loadings Using Principal Factor Analysis and Varimax Rotation.

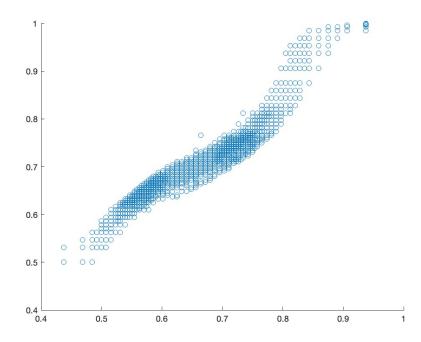
Variable (Selsoviet level)	Factor1	Factor2	Factor3	Uniqueness
Total Population	0.9711	0.2040	-0.0982	0.0056
% of Rural	-0.3726	-0.6112	0.1589	0.4623
% of Female	0.2448	0.2071	-0.3482	0.7760
% Pop Change	0.2066	0.2779	0.6144	0.5026
Average Salary	0.1199	0.8948	-0.0224	0.1844
Average Government Salary	0.1499	0.8439	0.0067	0.2653
% of Families Receiving Gov. Subsidies	-0.0038	0.0010	0.0094	0.9999
Averag Tax Income	0.8323	0.1816	-0.0431	0.2725
No. of Government Employees	0.9311	0.2047	-0.0833	0.0841
% of Disabled People	-0.1550	-0.1976	-0.1479	0.9151
% of School Kids	-0.0355	-0.1022	0.3010	0.8977
No. of Schools	0.9611	0.2001	-0.0803	0.0299
No. of Kindergartens	0.9508	0.2021	-0.1020	0.0448
Average Housing in Square Metres	0.9601	0.1948	-0.0956	0.0312

## F Additional Results and Robustness Checks

### F.1 Robustness Check: Normalization

To conduct a robustness check, we change the normalization used to estimate the polling station random term in Eq. (2). Instead of the midpoint of the support of X we use the characteristics X in uik 1136 in Altai Krai.

Figure F1: Estimates of  $\varepsilon$  using the midpoint vs Altai Krai normalization



**Notes**: Figure F1 shows the scatterplot of the estimators of  $\varepsilon$  in Eq.(2) based on the midpoint and Altai Krai normalizations.

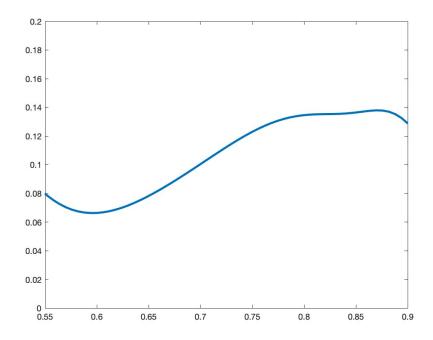
Figure F1 shows the scatterplot of the estimators of  $\varepsilon$  in Eq.(2) based on the midpoint and Altai Krai normalizations. When the estimates from the two normalizations are identical, they should align along the 45-degree diagonal line. The scatterplot shows that although the Altai normalization produces slightly different estimates, they remain very close to one another.

Figure F2 shows the estimator of  $d(\cdot)$  under the robustness normalization. While there are some differences with the original estimator, the overall shape is the same. Hence, we conclude that the choice of the normalization does not affect the results.

# F.2 Robustness Check: Sample and PCA analysis

In the main results, we estimated the model using restricted data. In this subsection, we perform two robustness checks: (1) using the full polling station dataset without excluding stations with low matching rates (see subsection D), and (2) incorporating additional

Figure F2: Estimator of  $d(\cdot)$  using the Altai Krai normalization



**Notes**: The graph presents the estimated function  $d(\cdot)$  across  $\varepsilon$ , estimated using the Altai Krai normalization and all polling stations with independent observers as the cleaner sample.

variables from the CENSUS to construct factors via PCA (see subsection E). Panel A of Table F1 presents the estimated fraud rates for the full sample, similar to Table 1. In our final sample, the fraud rate is slightly higher, at 13.5%, compared to 12.8% when using the cleaner sample of polling stations with no violations. It is also 8% compared to 7.5% when using the cleaner sample with all independent observers. The results for Moscow and the other regions are similarly consistent, confirming the robustness of our findings.

In Panel B, we expand the set of information used in PCA to generate factors. While the base estimation was restricted to data available for all Selsovets from the CENSUS, here we include all available information, marking any missing data with an indicator variable. The additional data includes variables on age distribution, occupation, and ethnic composition. Comparing these results with Table 1, we again observe similar patterns, reinforcing the robustness of our conclusions.

Table F1: Sample Robustness Checks

Panel A: Full Sample									
	Clean	Clean Sample: No Violations				Clean Sample: All Observers			
	All Regions	Moscow	Tatarstan	Dagestan	All Regions	Moscow	Tatarstan	Dagestan	
Estimated Turnout (%)	46.3	47.3	60.1	72.1	51.6	53.3	70.1	81.4	
Estimated Fraud / Reg. Voters (%)	13.3	14.0	19.5	19.0	8.0	8.0	9.7	9.7	
Estimated UR Vote / Reg. Voters (%)	15.6	14.5	43.1	64.3	20.9	20.5	53.1	73.6	
Estimated UR Voter Share (%)	26.2	23.7	54.1	70.6	35.1	33.5	66.6	80.8	
Official UR Vote / Reg. Voters (%)	28.9	-28.6	62.7	83.3	28.6	28.6	-62.7	83.3	
Official UR Voter Share (%)	48.5	46.6	78.7	91.5	48.0	46.6	78.7	91.5	
Official Turnout (%)	59.6	61.3	79.7	91.1	59.6	61.3	79.8	91.1	
Number of Polling Stations	88,326	3,361	2,749	1,877	88,326	3,361	2,749	1,877	

Panel B: New Factor Variables

	Clean Sample: No Violations			Clean Sample: All Observers				
	All Regions	Moscow	Tatarstan	Dagestan	All Regions	Moscow	Tatarstan	Dagestan
Estimated Turnout (%)	46.5	47.7	61.0	72.0	51.2	53.0	70.4	81.4
Estimated Fraud / Reg. Voters (%)	12.7	13.3	18.7	19.3	7.9	8.0	9.3	10.0
Estimated UR Vote / Reg. Voters (%)	15.7	15.1	44.2	64.3	20.5	20.4	53.6	73.6
Estimated UR Vote Share (%)	26.6	24.7	55.5	70.4	34.6	33.4	67.2	80.6
Official UR Vote / Reg. Voters (%)	-28.4	$-\bar{28}.\bar{3}$	62.9	83.6	28.4	28.3	62.9	83.6
Official UR Vote Share (%)	48.0	46.5	78.9	91.5	48.0	46.5	78.9	91.5
Official Turnout (%)	59.2	61.0	79.7	91.4	59.2	61.0	79.7	91.4
Number of Polling Stations	71,064	2,952	2,631	1,661	71,064	2,952	2,631	1,661

**Table Note**: The table provides two robustness checks for the main results. Panel A presents fraud rate estimates using the full polling station dataset, without excluding stations with low matching rates. Panel B incorporates additional variables from the CENSUS to generate factors via PCA. For each panel, results are reported for the cleaner sample of polling stations: those with "no violations" (similar to Table 1, Panel C) and those with independent observers (similar to Table 1, Panel D). Figure F3, Panels A and B presents the estimated functions  $d(\cdot)$  for the sample used in Panel A and Panel B.

## F.3 Alternative Control Group

In this section, we assess the robustness of using Moscow polling stations without independent observers as the control group in Eq. (3). Table F2 presents summary statistics for polling stations with independent observers and no reported violations, showing that Moscow polling stations exhibit similar characteristics. To further ensure the robustness of this control group, we construct a pseudo-control group using propensity score matching with nearest-neighbor matching based on Mahalanobis distance. This matched sample is then used as the control group to re-estimate Eq. (3), minimizing selection bias from observable covariates (Imbens and Wooldridge, 2009). The summary statistics for the matched group are shown in the third panel of Table F2, and Table F3 provides the estimation results, which remain consistent with those in Table 1.

Table F2: Summary Statistics for Treatment and Control groups

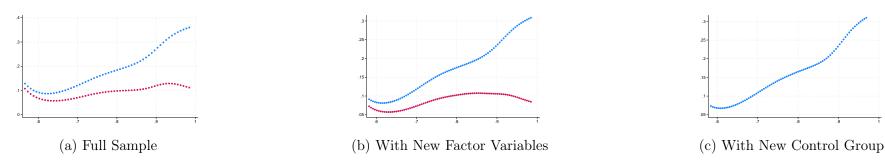
	Treat Gr: No Violations	Control: Rest of Moscow	Pseudo Control Group
Factor 1	-0.7298	-0.7249	-0.7370
Factor 2	2.2684	2.3224	2.2387
Factor 3	-0.2250	-0.2178	-0.2019

Table F3: Results using New Control Group

Results using New Control Group								
All Regions Moscow Tatarstan Day								
Estimated Turnout (%)	47.8	49.1	63.6	75.8				
Estimated Fraud Rate (%)	11.4	11.9	16.1	15.6				
Estimated UR Vote / Reg. Voters (%)	17.0	16.5	46.8	68.0				
Estimated UR Voter Share (%)	28.8	27.0	58.7	74.4				
Official UR Vote / Reg. Voter (%)	28.4	28.3	62.9	83.6				
Official UR Voter Share (%)	48.0	46.5	78.9	91.6				
Official Turnout (%)	59.2	61.0	79.7	91.4				
Number of Polling Stations	71,064	2,952	2,631	1,661				

Table Note: This table presents estimates using a new control group to assess robustness. We use Moscow polling stations without independent observers as the control group for Eq. (3), and construct a pseudocontrol group through propensity score matching with nearest-neighbor matching based on Mahalanobis distance. The results are comparable to Table 1 Panel C.Figure F3, Panel C presents the estimated function  $d(\cdot)$  for this specification.

Figure F3: Estimated function  $d(\cdot)$  across  $\varepsilon$  for the robustness samples.



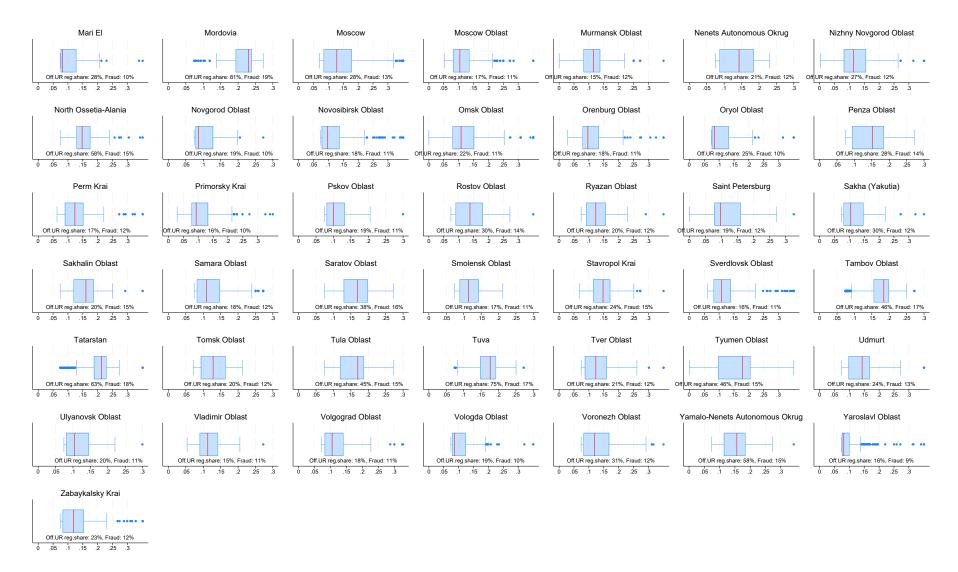
**Note:** This figure presents the estimated function  $d(\cdot)$  across  $\varepsilon$  for the samples used in Tables F1 and F3. The estimates are based on the mid-point normalization. Red points represent estimates from all polling stations with independent observers (cleaner sample), while blue points show our preferred estimates, where the cleaner sample consists of polling stations that reported no violations. Panel A corresponds to Table F1, and Panel B corresponds to Table F3.

Figure F4: Estimated Ballot Stuffing Across Regions



Notes: Figure F4 shows the estimated ballot stuffing using the cleaner sample consisting of observers that reported no violations. The box represents the interquartile range (IQR), covering the middle 50% of the data. The red line inside the box marks the median. The whiskers extend to the minimum and maximum values, excluding outliers, which are shown as individual points beyond the whiskers. The first number at the bottom of the figure, "Off. UR reg share," represents the official proportion for UR among registered voters. The second number, "Fraud," indicates the total estimated fraudulent votes out of the registered votes.

Figure F4: Estimated Ballot Stuffing Across Regions (continued)



Notes: Figure F4 shows the estimated ballot stuffing using the cleaner sample consisting of observers that reported no violations. The box represents the interquartile range (IQR), covering the middle 50% of the data. The red line inside the box marks the median. The whiskers extend to the minimum and maximum values, excluding outliers, which are shown as individual points beyond the whiskers. The first number at the bottom of the figure, "Off. UR reg share," represents the official proportion for UR among registered voters. The second number, "Fraud," indicates the total estimated fraudulent votes out of the registered votes.