WP 1716 – April 2017

Electoral fraud and voter turnout: An experimental study

Vardan Baghdasaryan, Giovanna Iannantuoni, Valeria Maggian

Abstract:

In this paper we experimentally investigate the consequences of electoral fraud on voter turnout. The experiment is based on a strategic binary voting model where voters decide whether to cast a costly vote in favour of their preferred candidate or to abstain. The minority candidate can illicitly influence the electoral process by applying ballot-box stuffing. In the experiment we implement two different framings: we compare voter turnout in a neutral environment and with framed instructions to explicitly replicate elections. This approach enables to both test the model's predictions and to estimate the framing effects of voting and fraud. Comparison of experimental results with theoretical predictions reveals over-voting, which is exacerbated when fraud occurs. Turnout increases as predicted with moderate level of fraud while, with higher electoral fraud, voters fail to recognize that the existence of a relatively larger number of "agents" voting with certainty considerably decreases the benefits of voting. Importantly, framing matters, as revealed by the higher turnout of those in the majority group, against which the fraud is applied.

Keywords:

Laboratory experiment, Framing, Voting, Electoral fraud, Ballot box stuffing and Voter turnout

JEL codes:

D72, C52, C91, C92

Electoral fraud and voter turnout: An experimental study*

Vardan Baghdasaryan[†], Giovanna Iannantuoni[‡], Valeria Maggian[§]

Abstract

In this paper we experimentally investigate the consequences of electoral fraud on voter turnout. The experiment is based on a strategic binary voting model where voters decide whether to cast a costly vote in favour of their preferred candidate or to abstain. The minority candidate can illicitly influence the electoral process by applying ballot-box stuffing. In the experiment we implement two different framings: we compare voter turnout in a neutral environment and with framed instructions to explicitly replicate elections. This approach enables to both test the model's predictions and to estimate the framing effects of voting and fraud. Comparison of experimental results with theoretical predictions reveals over-voting, which is exacerbated when fraud occurs. Turnout increases as predicted with moderate level of fraud while, with higher electoral fraud, voters fail to recognize that the existence of a relatively larger number of “agents” voting with certainty considerably decreases the benefits of voting. Importantly, framing matters, as revealed by the higher turnout of those in the majority group, against which the fraud is applied.

JEL classification: D72, C52, C91, C92

Keywords: Laboratory experiment, Framing, Voting, Electoral fraud, Ballot box stuffing and Voter turnout.

***Acknowledgements:** The authors gratefully acknowledge financial support from PRIN (2012-2015) “New approaches to political economics : positive political theories, empirical evidence and experiments in laboratory”. Valeria Maggian’s contribution to this research was also performed within the framework of the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d’Avenir (ANR-11-IDEX-007) operated by the French National Research Agency (ANR). We thank Natalia Montinari, Marie Claire Villeval, the participants at the 2016 Behavioral Models of Politics conference at the University of Pittsburgh and at the 2015 ESA European Meeting at the University of Heidelberg for useful comments. We are grateful that David K. Levine and Thomas Palfrey provided us with their data from the laboratory experiment “The paradox of voter participation? A laboratory study” (2007).

[†]American University of Armenia and Affiliate Fellow at CERGE-EI, Prague. E-mail: vbaghdasaryan@aau.am.

[‡]University of Milano-Bicocca. E-mail: giovanna.iannantuoni@unimib.it.

[§]Corresponding author. Univ Lyon, CNRS, GATE L-SE UMR 5824, F-69130 Ecully, France. E-mail: maggian@gate.cnrs.fr

1 Introduction

Since the 1990s, with the last large wave of democratization of post-Soviet and eastern European countries, the number of elections in the world has witnessed a sharp increase. At the same time, the quality of elections has deteriorated over time: while the share of elections reported to be illicitly influenced was around 15% out of around 180 national level elections during the period 1980-1990, it increases to almost 40% out of around 480 elections in the period 2001-2010.¹

In this paper we investigate, both theoretically and experimentally, the effects of electoral fraud on voter turnout. Our objective is twofold. First, high turnout rate is often quoted as a by-side indicator of properly organized and conducted elections. We thus provide the first evidence on whether a lower participation rate is effectively associated with a higher degree of fraudulent intervention in the elections. Second, our experimental design enables us to test whether the behavioural reaction to an electoral procedure that is considered to be unfair also affects voter decision to participate in the elections. Citizens may shy away from voting when the election is characterized by a fraudulent mechanism or, on the opposite, they may be rather more likely to mobilize and go to the polls in order to counteract it.

Evidence of electoral fraud has been provided worldwide, both in developed and in developing countries. Weak electoral identity-checking schemes are responsible for at least 42 convictions for electoral fraud in the UK in the period 2000-2007 (Wilks-Heeg, 2008) and for promoting electoral fraud in Japanese municipal elections, as shown by the natural experiment by Fukumoto and Horiuchi (2011). News reports, highlighting the risk of voter fraud in several EU countries,² have become a fairly routine part of electoral campaign. Douglas (2013) examines the multiplicity of election contest provisions and

¹ Authors' calculations based on data from Institute for Democracy and Electoral Assistance (IDEA), database on political institutions (DPI) and the Cingranelli-Richards human rights dataset (CIRI).

² EUobserver (2014): Threat of voter fraud haunts EU vote in Romania, Bulgaria and Croatia. See <http://euobserver.com/eu-elections/123485>

the procedural mechanism to manage them among the US states: the notorious hanging chads in 2000 in Florida is an example of US not being immune to (potential) electoral malpractice. In developing countries, electoral fraud is usually more evident and takes a wider variety of forms: from vote buying to intimidation, from ballot-box stuffing and carousel voting to explicit violence.

Whenever electoral fraud is pervasive and eliminates any notion of political competition, any effort made to analyse its impact on voters' behaviour will hardly be of any use. Differently, when illicit electoral intervention leaves some room for candidates or parties to compete, voters' strategic and behavioural responses to fraud will be a key determinant of the election outcome.

In order to investigate the effects of fraud on voter turnout, we modify the conventional binary elections model with costly voting to allow for fraud, where illicit intervention is promoted only by one out of two parties. Assuming that the implementation of fraud is costly (and thus limited in scope), we show that in equilibrium fraud might increase turnout rate. In particular, a low amount of fraud exerted in favor of the competitor with the ex-ante smaller support should increase turnout in both groups, whereas higher fraudulent intervention should suppress voting motives and decrease the turnout rates to a level lower than in equilibrium without fraud.

We test our predictions in a laboratory experiment, where we compare voting behavior in a baseline condition with respect to a situation where, respectively, a low and a high amount of fraud is introduced in the elections. The baseline condition is designed in a way that allows direct comparisons with the paper by Levine and Palfrey (2007). In their laboratory experiment, they test for the effect of changing the (relative and absolute) size of the electorate of two competing groups of supporters on voter turnout: while a more numerous electorate lowers turnout rate, having closer elections increases it. Finally, they show that voters supporting the less popular alternative have higher turnout rates. We thus provide an additional validity test for their results. In order to

analyse the effect of fraud on voting behaviour, we introduce two levels of fraudulent intervention in the elections. When fraud is low our model predicts a higher turnout than in the baseline, whereas a higher intervention should have the opposite effect. In the experiment, we implement two different treatments: in the *Neutral* treatment, the wording in the instructions is designed to induce a neutral environment. There is no mention of voting, electoral fraud or voting costs. In the *Framing* treatment the instructions are framed as to replicate elections. We explicitly refer to elections and ballot-box stuffing in order to investigate the importance of social and psychological factors in affecting voting behaviour. In particular, if voters fear that polls are corrupted, they have less incentive to bother casting a vote; however, the electorate may also decide to heavily participate in the elections when fraud is acknowledged, in order to punish the malfeasant politician.

The experiment has provided a number of noteworthy results. First, the comparison of experimental results with theoretical predictions reveals over-voting (analogous results were obtained by Levine and Palfrey (2007)). Second, as expected, limited intervention increases the turnout and this result is observed both in the Neutral and Framing treatment. Third, extensive fraud does not have a significant negative impact on turnout rates as suggested by the theoretical model. Moreover, its effect is qualitatively different when comparing the Neutral and the Framing treatments. While in the Neutral treatment there is no significant variation in turnout between a limited and a more extensive fraudulent intervention, in the Framing treatment we do observe such differences. In particular, the majority group, against which the fraud is undertaken, responds with significantly higher participation in the extensive intervention case. The observed differences between the Framing and Neutral treatments suggest that voting behavior cannot be solely explained by the (in)ability of agents to correctly infer their expected pay-offs when deciding to vote or to abstain. In this paper we provide evidence of behavioural aspects that cannot be disregarded when considering voting in a fraudulent environment.

2 Related literature

In spite of its frequent occurrence, fraud has rarely been studied within the theoretical literature on voting. Fraud may affect elections via numerous different practises. Among them, vote buying has been mainly analysed as a legal phenomenon adopted by both competitors in binary elections. Dekel et al. (2008) present a model in which vote buying is free of stigma and show that politicians obtain higher rents when electoral competition rests on pre-electoral vote-buying rather than when it relies on campaign promises. Ellman and Wantchekon (2000) study the case of post-electoral violence and demonstrate that, in a situation without policy commitment and asymmetry of information about the probability of unrest, voters select the strong party to avoid disruptions. In a more recent paper, Collier and Vicente (2012) analyse the difference of electoral fraud and electoral violence in Sub-Saharan Africa, showing that when political competitors are weak they are more likely to resort to violence, whereas when they are strong they prefer to use bribery and ballot box stuffing in order to win elections. The closest study to our own is provided by Vorobyev (2016), who investigates the mechanism of making participation decisions in the presence of electoral fraud. In his theoretical model electoral fraud is designed as ballot rigging, so that when a voter abstains, her unused ballot may be transformed into a vote for the Incumbent with known probability. According to his analysis, two equilibria exist: an abstention one, where none of the voters vote and the Incumbent always wins, and a coordination one, where some challenger's supporters vote and the candidate preferred by the majority is likely to win. With respect to his paper, there are a number of differences in the way we set up both the model and the fraud technology, as shown in section 3. In particular, we do not model electoral fraud as stolen votes, but rather consider the number of fake ballots introduced independently from the number of abstaining voters. An important outcome difference is that in our model the probability of the Incumbent winning the election is monotonously increasing with the number of fake ballots. Most importantly, we test our predictions by means

of an experiment and provide evidence that voter turnout may increase, rather than decrease, as a consequence of electoral fraud.

In the empirical literature electoral fraud has been studied from various perspectives. Enikolopov et al. (2013) show that, during the 2011 Russian parliamentary elections, electoral fraud was sufficient to impact substantially the outcome of the elections. In line with this study, Collier and Hoeffler (2009) investigate the causal relationship between electoral fraud and the probability of the Incumbent to win the election (i.e. efficiency of fraud) in a cross-country context. According to the results obtained by analysing 786 elections in 155 countries, fraud during elections increases the probability of the Incumbent to win from 62% up to 84%. Vergne (2009) is one of the few papers to examine the role of electoral fraud in influencing voter turnout finding it to be a significant negative predictor of turnout. Her analysis covers 60 developing countries during the period 1980-2005 and the microeconomic foundation of fraud is that it increases the cost of participation by making the outcome less predictable. Birch (2010) shows, in an analysis conducted on aggregate-level data from 31 countries, that perceptions of electoral integrity are positively associated with the propensity to vote.

Carreras and İrepoğlu Yasemin (2013) found that the irregularities (i.e. vote buying) that still characterize the electoral process in Latin America negatively affect citizens' trust in elections, finally reducing their willingness to participate in national elections (for Mexico, see also McCann and Dominguez (1998)). Differently, by using data from a voters' representative survey, Bratton (2008) shows that vote buying has no statistically significant impact on voter turnout, while the threat of violence has a negative effect on it.

Overall, previous empirical studies suggest that electoral fraud has a negative impact on turnout, even if there is no clear-cut result. However, the intrinsic characteristics of electoral fraud - difficult to observe and measure - make it difficult to study the relationship between fraud and turnout using empirical data. The impact of electoral integrity on

voter turnout may not be as straightforward, since different types of electoral malpractices (threat of violence, coercion, vote buying, etc.) and magnitudes of electoral fraud may lead to different and contrary effects, as well as both complex micro and macro conditions affecting voter turnout having to be controlled for in empirical studies (Norris, 2014). These studies are often based on the perception of fairness of elections, an opinion that also depends on people's candidate preference and campaign competitiveness (Wolak, 2014). Moreover, the consequences of fraud on voters' participation is not clear also from a *behaviorial* point of view. In the electoral context, the procedure by which the candidate is elected is of extreme importance. In particular, in the social psychology literature, Deutsch (1975) states that "*people are more apt to accept decisions and their consequences if they have participated in making them*" (p.139): when the electoral process is unfair, in the sense that it allows fraudulent interventions to limit people's participation, voters may be more likely to cast a vote for their preferred, but unfairly sabotaged, candidate with respect to a situation when there is no fraud at all. In order to define an election as fair, procedural characteristics are fundamental, as shown by the experiment by Wilking (2011): participants were asked to judge a series of hypothetical scenarios and their opinion about fairness in the elections was largely independent of whether they resided in a democracy or non-democracy, or whether they were politically engaged. Previous research on the effect of procedural injustice in Economics has shown that people are prone to impose a cost, both on self and others, to resist procedures that they value as biased (Bolton et al., 2005). Individuals intrinsic sense of (procedural) fairness, activated by framing the experiment as an election characterized by a fraudulent voting mechanism, will thus act as a motivating force in individual choice behavior. Whenever procedural injustice affects emotions such as anger, we expect that it can affect turnout in turn (Tarrow, 1998). People may be willing to bear a higher cost of voting in order to take revenge when being treated unfairly in an election. However, it is also possible that voters, when facing electoral fraud, may be rather less likely to

cast a vote for their preferred party, shying away from the unfairness of the electoral procedure.

In this paper we examine how these two forces affect voting behavior of the electorate by comparing a Neutral Treatment with a Framing one, where the election environment is explicitly stated and the fraudulent intervention in the electoral process is made common knowledge. In this regard, the laboratory represents the ideal environment for testing the predictions of the strategic costly voting model with fraud, a controlled environment where *ceteris paribus* analysis is possible. Given the complexity of the real world, especially in the political context, we aim at analyzing whether framing the experiment as an election where a fraud in the voting system is introduced, makes people to react differently with respect to an identical situation, where actions and payoffs are expressed in neutral terms. In order to test for the external validity of our results we thus examine the role of framing the experiment as an election, which has rarely been systematically tested. Previous researchers have studied how media frames affect voter turnout and vote choice (Schuck et al., 2013, Valentino et al., 2001): by using empirical data, Schuck et al. (2014) show that, when campaign news coverage was framed in terms of conflict rather than strategy, voters were more likely to turn out and vote in 2009 European Parliament elections. The natural voting experiment by Butler and Marchal (2007) provides evidence of how different emphasis in the initiatives' titles of two virtually identical popular initiatives, both demanding a decrease in the legal age of retirement in Switzerland, led to different approval rates. Druckman (2001), moreover, put into evidence how source credibility is a prerequisite for successful framing, when elites want to manipulate public opinion. Our paper extends this literature by testing for the existence and consequences of the framing effect of ballot rigging on voting behavior within a laboratory experiment. As suggested by Druckman (2004), investigating in which contexts framing effects occur allows for a deeper understanding of the applications of the rationality assumptions.

Previous researchers have investigated voter turnout in the laboratory, under many different institutions: Schram and Sonnemans (1996), Grosser et al. (2005), Grosser and Schram (2010) focus on participation games and coordination problems such as the volunteer’s dilemma.³ One of the most influential experimental papers on voter turnout is provided by Levine and Palfrey (2007), who investigate voter turnout by addressing comparative statics questions about the effect of the electorate size, relative party size, and voting cost on subjects’ decision to vote.⁴ Their results support theoretical predictions with the exception that turnout probabilities are higher than those predicted when considering larger electorates (i.e. with 27 and 51 voters). The design of our experiment allows us not only to compare our results with those obtained by Levine and Palfrey (2007) but, most importantly, also to investigate the effect of introducing fraud in the electoral system on voting behaviour of the electorate.

3 The model

We define electoral fraud as “*the corruption of the process by which votes are cast and counted*”, as suggested by Minnite and Callahan (2003) and, similarly, by Lehoucq (2002). In the present study we apply ballot-box stuffing type of fraud, such as repeated voting by the same voter, elimination of valid ballots or voting instead of abstaining. It is worth emphasizing that qualitatively similar theoretical results are obtained when asymmetries in costs are introduced between the supporters of the two groups.

In our model we include electoral fraud in the framework developed by Levine and Palfrey (2007). There are N voters who can choose to vote or to abstain. They support one of the two candidates: the Incumbent (I) or the Challenger (C) with $N_I < N_C$, so that supporters of the Incumbent are in the minority. The sizes of the two groups are common knowledge. Voting is assumed to be costly and every voter i knows her

³Palfrey (2009)) provides a survey of experimental studies of voting behavior.

⁴Levine and Palfrey (2007) run their experiment in a completely neutral framing: we replicate their design in (the first part of) our Neutral treatment, as explained in Section 4.

own cost c_i . She also knows the single-peaked distribution function $f(c)$ from which the voting costs of the other voters are independently drawn. The candidate who receives more votes wins the elections and ties are broken by a fair coin toss. If candidate I wins, all voters supporting him receive a reward H and all voters supporting candidate C receive a reward L , with $L < H$, while the opposite happens when candidate C wins. The size of the rewards is common knowledge.

The Incumbent can interfere with the elections by introducing fake ballots. The strategy of the Incumbent is defined as the (integer) number $\delta \in [0, \Delta]$ of fake ballots which are placed in the polling box. The electoral fraud applied by the Incumbent is supposed to be costly. Interventions must increase his probability of winning the election, given the electoral rule in force. Let $\Phi(\delta)$ be the convex cost function of the fraud. Also, let n_I^* be the expected number of the Incumbent's supporters who vote in equilibrium and n_C^* be the expected number of the Challenger's supporters who vote in equilibrium.

We define the expected plurality EP as the difference between the expected number of votes in favour of the Incumbent and in favour of the Challenger - $EP(\delta) = \delta + n_I(\delta) - n_C(\delta)$. The falsifier's problem would be⁵:

$$\max_{\delta} -(EP(\delta) - \mathcal{M})^2 \quad (1)$$

$$\text{s.t. } \Phi(\delta) \leq \mathcal{B} \quad (2)$$

where $\mathcal{M}$ is the difference in votes that the Incumbent wants to achieve in equilibrium. Thus, $\mathcal{M} = 1$ describes a situation where the Incumbent is only interested in winning the elections, whereas $\mathcal{M} > 1$ corresponds to a situation where a larger expected margin of victory is required. Though theory suggests that in a binary election a

⁵An alternative formulation where the win probability is directly incorporated into the Incumbent's preferences would be the following: $\max_{\delta} R \cdot \text{prob}[n_I^*(\delta) = n_C^*(\delta) + \mathcal{M}] - \Phi(\delta)$, where R is the benefit associated with winning and $\Phi(\delta)$ is the cost function of the fraud.

candidate must be just interested in getting more votes than the opponent, there exists wide evidence that Incumbents apply wider manipulations to win elections with a larger margin. Benefits of excessive fraud during the elections can be of various types, such as discouraging opponents from joining or supporting rival parties, from voting, or from participating in the competition in other ways or, alternatively, it may motivate supporters to participate more actively (Simpser, 2013). $\mathcal{B}$ is the exogenously given budget of the falsifying Incumbent.

Players move simultaneously in the above described game. While in turnout models simultaneity is considered a natural modelling approach for voters, the actions of the Incumbent require an additional clarification. An alternative view would be to assume that the Incumbent can actually monitor the process of the election (e.g. by applying exit polls) or even falsify the electoral results after the vote count. Differently, we assume simultaneity of choices believing that the fraudulent interventions on behalf of the Incumbent need to be planned and designed in advance and thus cannot be considered a “last minute” response to negative developments on the election day.

In the following, we look for group symmetric equilibria of the voting game, which are known to be in cut-point strategies. A cut-point strategy for voter i specifies a cost level c_i^* such that voter i abstains if and only if $c_i > c_i^*$. This implies an aggregating voting probability for each group, (p_C^*, p_I^*) given by:

$$p_C^* = \int_{-\infty}^{c_C^*} f(c)dc = F(c_C^*) \quad (3)$$

$$p_I^* = \int_{-\infty}^{c_I^*} f(c)dc = F(c_I^*) \quad (4)$$

Given that in equilibrium every voter in the group follows the same rule, this means that a voter with a voting cost exactly equal to c^* must be indifferent between voting and abstaining. It is easy to demonstrate that these indifference conditions are given by

the following equations:

$$c_I^* = \frac{H-L}{2}(\text{prob}[n_I^i + \delta + 1 = n_C] + \text{prob}[n_I^i + \delta = n_C]) \quad (5)$$

$$c_C^* = \frac{H-L}{2}(\text{prob}[n_C^j + 1 = n_I + \delta] + \text{prob}[n_C^j = n_I + \delta]) \quad (6)$$

Superscript ‘i’ (‘j’) indicates that the number of those who voted for the Incumbent (Challenger) is n_I (n_C) excluding voter ‘i’ (‘j’).

Given the preferences of the falsifier, the equilibrium of the game is characterised by the following conditions:

$$c_I^* = \frac{H-L}{2}\pi_I^*(p_I^*, p_C^*, \delta^*, N_I, N_C) \quad (7)$$

$$c_C^* = \frac{H-L}{2}\pi_C^*(p_I^*, p_C^*, \delta^*, N_I, N_C) \quad (8)$$

$$\arg \max_{\delta} -(\delta^* + p_I^*(\delta)N_I - p_C^*(\delta^*)N_C - \mathcal{M})^2 \quad (9)$$

$$\Phi(\delta^*) \leq \mathcal{B} \quad (10)$$

where π_I and π_C are the pivot probabilities and have the following representation:

$$\begin{aligned} \pi_I^* = & \sum_{k=0}^{N_I-1} \binom{N_I-1}{k} \binom{N_C}{k+\delta} (F(c_I^*))^k (1-F(c_I^*))^{N_I-1-k} \\ & \times (F(c_C^*))^{k+\delta} (1-F(c_C^*))^{N_C-k-\delta} + \sum_{k=0}^{N_I-1} \binom{N_I-1}{k} \binom{N_C}{k+\delta+1} (F(c_I^*))^k \times \\ & (1-F(c_I^*))^{N_I-1-k} (F(c_C^*))^{k+\delta+1} (1-F(c_C^*))^{N_C-k-1-\delta} \end{aligned} \quad (11)$$

$$\begin{aligned}
\pi_C^* = & \sum_{k=0}^{N_I} \binom{N_I}{k} \binom{N_C-1}{k+\delta} (F(c_I^*))^k (1-F(c_I^*))^{N_I-k} \times \\
& (F(c_C^*))^{k+\delta} (1-F(c_C^*))^{N_C-k-\delta-1} + \sum_{k=0}^{N_I} \binom{N_I}{k} \binom{N_C-1}{k+\delta-1} \times \\
& (F(c_I^*))^k (1-F(c_I^*))^{N_I-k} (F(c_C^*))^{k+\delta-1} (1-F(c_C^*))^{N_C-k-\delta} \quad (12)
\end{aligned}$$

We do not make any explicit assumptions about the cost function $\Phi(\delta)$ and we are not going to solve the Incumbent's problem explicitly. As it will be made clear below, in the current framework fraud monotonically increases the expected plurality (or winning probability) in favour of the Incumbent, hence limited fraud can be rationalized only by assuming that the falsifier is cost constrained. Since we are interested in the reaction of the electorate to different levels of electoral fraud, we assume that the parameters of the Incumbent's objective function and budget constraint are such that he can introduce a sufficient level of fraud in the election to win it (in expectation). When deriving the equilibrium of the voting game, we will thus focus only on the strategies of the voters. In the general case, the exogenously given absolute margin of victory $\mathcal{M}$ derives the equilibrium of the the game which determines the optimal fraud level endogenously.

3.1 Predictions

Numerical estimation techniques are applied to obtain the exact predictions for the turnout rate in the game. Given that we do not have a falsifier as a subject in the lab, we fix the fraud level exogenously. As mentioned above, the Incumbent has ex ante less support. For the purposes of our experiment, we fix $N_I = 9$ and $N_C = 18$. This corresponds to the landslide treatment with 27 voters in Levine and Palfrey (2007). Other parameters of the experiments are chosen in a way to guarantee direct comparability of their experiment with our Baseline condition without fraud.

In particular, benefits for winning and losing the elections are $H=105$ and $L=5$, respectively. The voting cost is assumed to be distributed uniformly from 0 to 55 and this distribution is the same for both groups. On the one hand, the choice of such parametrization is equivalent to assuming that the voting costs are sufficiently high for some voters, so that abstention is their dominant strategy (Palfrey and Rosenthal, 1985). On the other hand, a zero voting cost implies that, with positive probability, for some voters voting will be a weakly dominant strategy, irrespective of the number of fake ballots introduced. This, along with the assumption of a single-peaked cost distribution, ensures a unique group-symmetric equilibrium in our setup. Moreover, from a practical perspective, this framework is easy to be explained to subjects.

Table 1 presents the results of our numerical estimations of the equilibria. Though we have used only two levels of fraud for the experiment, here we present the estimations for a wider range of possible values in order to explicitly underline the inverse U-shaped relationship between fraud and turnout rates. In Figure 1, in order to verify the robustness of our predictions, we provide a graphical view of our numerical estimations for different electorate sizes, while keeping the Incumbent's relative support constant (Left Panel), as well as for a range of support for the Incumbent, keeping the total number of voters fixed (Right Panel).

Table 1: Theoretical predictions.

Incumbent group size	9	(fraud in favour of Incumbent)				
Challenger group size	18					
	Number of fake ballots					
	0	1	2	3	4	5
Group turnout probability (Incumbent)	0.27	0.29	0.30	0.30	0.25	0.11
Group turnout probability (challenger)	0.23	0.26	0.29	0.30	0.28	0.19
Expected Turnout (Incumbent group)	2.43	2.61	2.71	2.66	2.27	0.97
Expected Turnout (challenger group)	4.10	4.77	5.25	5.43	5.08	3.38
Expected Total turnout	6.53	7.38	7.96	8.09	7.35	4.35
Expected Plurality	-1.67	-1.15	-0.53	0.23	1.20	2.59
Pivot Probability	0.40	0.44	0.47	0.47	0.43	0.26
Probability of Incumbent to win	0.23	0.32	0.42	0.54	0.70	0.91
Total turnout rate	0.24	0.27	0.29	0.30	0.27	0.16

Figure 1: Results of our numerical estimations

As suggested by the results of our numerical estimations, electoral fraud increases turnout rates when it is limited and decreases it when fraud is high. Differently to Vorobyev (2016), where the participation of the Incumbent's supporters decreases with higher probability of ballot rigging, our model predicts that a low level of electoral fraud positively affects participation both of the majority and of the minority party. The intuition behind such a difference depends on the electoral fraud's technology: while in Vorobyev (2016) it is modelled as the probability that the abstainers' vote will be stolen, in our framework it is designed as the number of fake ballots added to the pool, so that

the probability of tied elections increases for low level of fraud. In its turn, the higher the probability of a tied election, the lower the incentive to free ride in the larger group and the higher the incentive to vote for the minority, since they have higher chances of casting a pivotal vote. This statement can easily be verified by looking at the change in pivot probabilities. Strategic motives to vote diminish at a high rate once the fraud goes beyond the level that guarantees the expected victory in binary elections.

3.2 Testable hypothesis

Based on the predictions of the theoretical model we formulated the following hypotheses:

Hypothesis 1. In the Baseline condition a higher turnout is observed for the underdog candidate, as a manifestation of the classic free-riding problem in the larger group.

Hypothesis 2. The relationship between the extent of the fraudulent intervention and the turnout rate has an inverted U-shaped pattern. Limited fraud increases turnout while higher level of fraud decrease it beyond the level observed in the Baseline condition without fraud.

Hypothesis 2a. The turnout patterns described in Hypothesis 2 are observed both in the minority and majority groups.

In addition to these two purely strategic hypotheses derived from the model, we have an additional *behaviorial* hypothesis:

Hypothesis 3. Framing the experiment as an election with vote rigging increases voter turnout with respect to a neutral environment, because of the majority's reaction

to procedural injustice.

4 Experimental design and procedures

We conducted 4 sessions of both the *Neutral Treatment* and the *Framing Treatment* at the experimental laboratory of the University of Milano-Bicocca. Our design is similar to the one used by Levine and Palfrey (2007) but we introduce a fraud mechanism in the electoral system to study its effect on voter turnout.

Our experiment consists of three parts: Baseline, Fraud I and Fraud II.

In part one, participants play the Baseline condition described below for 50 rounds. In part two, the only difference is that a fraud mechanism (Fraud I) is introduced for 25 rounds. Finally, in the third part of the experiment, participants play the last 25 rounds of the experiment where a stronger fraud mechanism (Fraud II) is included in the elections.

While the entire experiment is framed as an election in the *Framing Treatment*, in the *Neutral Treatment* we never refer to elections or voting or ballot rigging, labels are abstract.⁶ In Levine and Palfrey (2007), authors tested for the effect of framing in their voting experiment but with a limited number of observations (two additional sessions with 9 subjects each). Our aim is not only to provide additional and more robust evidence on the possible effect of framing the experiment as an election on voter turnout (as in our Baseline condition) but, more importantly, to investigate the effect of framing the experiment as an election with ballot box stuffing (as in our Fraud I and Fraud II conditions), which is the main objective of our paper. We expect that making subjects aware that the election is characterized by a fraudulent intervention will affect voter turnout in a different way than simply presenting the experiment in neutral terms. The use of abstract terms may eliminate important social, psychological and emotional

⁶For the English version of the instructions (originally in Italian) see the on-line Appendix A.

considerations that are inherent to real world elections characterized by an unfair voting procedure, leading to results that do not reflect what happens in naturally occurring fraudulent elections.

At the beginning of each period, each participant is informed whether he was randomly assigned to group ALPHA or to group BETA⁷ (in the Framing treatment we explicitly refer to Party ALPHA and BETA).⁸ Group A is composed of 9 participants while group B is composed of 18 participants. The two sizes were common knowledge to all participants.⁹ Individuals are asked in each round to choose X (i.e. *to vote* in the Framing Treatment) or Y (*to abstain*). The voting cost is referred to as a Y bonus and is added to a participant's earnings if that participant chooses option Y instead of choosing option X in a round (*election*). Therefore, the voting cost is implemented as an opportunity cost. Bonuses are randomly redrawn in every round, independently for each subject, and subjects are told their own Y bonus before making any decision. We thus assume incomplete information about voters' heterogeneous voting costs. If a participant chooses X, that participant does not receive her Y bonus for that election. Payoffs in each round are determined as follows: if more members of A(B) choose X than members of B(A), then each member of A(B) receives 105 and each member of group B(A) receives 5. In the case of a tie, each member of each group receives 55. In order to test the effect of electoral fraud on individuals' voting behavior, in parts Fraud I and Fraud II everything is the same as in the Baseline conditions, with the only exception that now respectively 2 and 5 fake ballots are automatically added to the minority group.¹⁰

⁷In the following we refer to group A for group ALPHA and to group B for Group BETA.

⁸Random reassignment of subjects to groups A or B was intended to minimize repeated-game effects.

⁹We decided to choose these group sizes for the A and B groups in order to compare experimental data with clear-cut theoretical predictions about voting behavior. When theoretical predictions are not matched, it is less probable that this will happen in a more favourable situation.

¹⁰We considered that the natural field setting in which our results would be most applicable are the ones more likely to involve a decreasing quality of elections, with an electoral system which progressively evolves through a situation of electoral fraud. However, in order to check for the effect of electoral fraud when considering an electoral system characterized by increasing quality, in December 2015 we also ran two more sessions for each treatment where the order of the experimental parts is Fraud II, Fraud I and

Differently to in our theoretical model, we decided not to include a subject playing the role of the Incumbent choosing the optimal level of fraud in the experiment because of three main reasons. First, in the theoretical model, the Incumbent and her preferences are important to understand why the fraud can be limited in size, so that its existence has no impact on our predictions on voter turnout. Second, including such a role in the experiment would have unnecessarily complicated our design, adding an additional level of procedural unfairness. If any, the presence of an Incumbent would strengthen our results. Third, since to the best of our knowledge we are the first to examine the effect of fraudulent elections on voter turnout in a clean and simple way, an Incumbent would have made our analysis closer to the real world but would have also made our results on turnout more difficult to interpret: its inclusion would have brought about new questions such as the Incumbent’s expectations of voter behavior and viceversa, and distributional considerations of the Incumbents’ and her supporters’ payoffs. As a consequence, we decided to leave this dimension to further research.

4.1 Experimental Procedures

The experiment was programmed using zTree (Fischbacher, 2007). Table 2 summarizes our experimental design, which involved four sessions of each of our two treatments. As we have 27 subjects per session, we have collected data from a total of 216 subjects, from October 2014 to January 2015.

All our sessions were conducted at the EELAB experimental laboratory of the University of Milano-Bicocca (Milan, Italy). Subjects were recruited from the undergraduate population of the University of Milano-Bicocca, via the ORSEE software (Greiner, 2004). No subject participated in more than one session of this experiment.

Baseline. Results are presented in the online Appendix B.

Table 2: The experimental design

Treatments	No. of Session	No. of rounds per session	No. of subjects per session	N_A	N_B	Order of the quality of the elections
Framing	4	100	27	9	18	Baseline - Fraud I - Fraud II
Neutral	4	100	27	9	18	Baseline - Fraud I - Fraud II

Once arrived in the laboratory, each participant was randomly assigned to one visually isolated computer terminal. The instructions for the first part were then read aloud and participants could read them on their screen. Individuals were asked to answer a set of control questions on the screen.

Instructions about each of the following parts (Fraud I and Fraud II) were made to appear on the subjects' screen and read aloud only after the completion of the previous instructions.

A test for risk aversion (Holt and Laury, 2002) was implemented as the last task.

At the end of each session, the pay-off of all rounds was added up and the sum was converted into Euros. Before proceeding with payment to the subjects, we asked participants to fill out a demographic form in order to collect information about their age and gender.

The duration of each session was about 60 minutes and the average payment was 14 Euros.

5 Results

Results are summarized from two main perspectives: we first analyse voter turnout at the aggregate level in subsection 5.1, then, in subsection 5.2, we study the impact of fraudulent electoral intervention at the individual level. To account for the fact that participants' choices maybe dependent upon previous periods in the same session, when comparing voting behavior between treatments we report results from a two-sample

Wilcoxon rank-sum (Mann-Whitney) test.¹¹ A more detailed and precise analysis is conducted thereafter, where random effects probit estimations are used to analyse turnout behavior at the individual level.

We initially consider the Baseline condition without any fraud intervention: we focus on framing differences and compare the observed results with our theoretical predictions and with those obtained in previous studies. We then assess the impact of the fraudulent intervention on turnout, in order to estimate to what extent voters are behaving strategically.

5.1 Aggregate level analysis

We start our analysis by considering voter turnout in the Baseline condition. Table 3 shows the aggregate turnout rates for the Baseline condition and compares them with our theoretical equilibria. Figure 2 graphs, for each party and each treatment, the average observed turnout rates throughout the 50 rounds.

Table 3: Turnout rates under the Baseline condition (standard error in parentheses).

	Experimental results		Theoretical equilibrium
	Neutral Treatment	Framing Treatment	
Majority	0.323 (0.010)	0.342 (0.018)	0.228
Minority	0.367 (0.025)	0.317 (0.025)	0.270

One of the central results consistently predicted in strategic costly voting games is that higher turnout is expected among the underdog candidate’s supporters. This is due to the manifestation of the classic free riding incentives in the larger group, which in turn

¹¹When performing the Wilcoxon rank-sum (Mann-Whitney) test we average the data within a session and treat each session as a single observation.

makes voting for the representatives of the smaller group more rewarding. Figure 2 suggests that Hypothesis 1 about the greater participation in the election of the under-dog's group seems to be satisfied for the Neutral treatment but not for the Framing one. In Figure 2 we observe that in the Framing treatment members of the majority group are rather more likely to vote than those in the minority one. Moreover, we observe a slightly more rapid decrease in voting behavior in the Framing treatment rather than in the Neutral one.

Figure 2: Frequency of voting per group in the Neutral and Framing treatment.

By means of a Wilcoxon signed-rank (WSR) for matched pairs we test the null hypothesis being that the frequencies of voting are the same between the minority and the majority group within each treatment. For the purpose of the test, we paired the majority and minority group observed frequency of voting in each session and generated

4 signed differences for each treatment. Our test suggests that the overall average turnout rate (50 rounds) is significantly higher (at the 0.05 level) for the minority group than for the majority group only in the Neutral treatment (Neutral treatment: $p=0.072$, one-sided; Framing treatment: $p=0.136$, WSR test, one-sided). In addition, we perform the same test using data on 10-round averages from the four sessions of each treatment. Our results confirm that the minority group is not more likely to vote than the majority one in the Framing treatment even when considering 10-rounds averages. Differently, in the Neutral treatment, we find a significant difference in voting behavior in rounds 1-10 ($p=0.047$, one-sided), 10-20 ($p=0.034$, one-sided) and 30-40 ($p=0.072$, one-sided).

As a second level of analysis, we investigate whether voter turnout of the minority group is significantly different between treatments. We conducted a Wilcoxon Mann Whitney (MW) test using the four session-level averages for each treatment.

Figure 3: Frequency of voting per group in the Neutral and Framing treatment, ten-round averages.

Figure 3 summarizes turnout using 10-round averages for each of the 4 sessions of each treatment. The average turnout of the minority group in the Neutral treatment is always higher than the average turnout of the minority group in the Framing one, while

we observe the opposite pattern when comparing the majority group. The difference in turnout between the two treatments is constant, in particular for the minority group. Even if a MW test fails to reach significance when comparing all periods together, we reject the null hypothesis of the frequencies of voting between the two treatments coming from the same distribution when considering rounds 10-20 for the minority group ($p=0.042$, MW test, two-sided) and rounds 0-10 for the majority group ($p=0.059$, MW test, two-sided). Results 1 and 2 state our main findings in the Baseline condition, at the aggregate level.

Result 1. *In the Baseline condition the probability of voting is higher than predicted for both the majority and the minority group.*

Support for Result 1 can be found in Table 3. Considerable over-voting can be detected both in the abstract setting and in the framed one. This phenomenon was also observed in Levine and Palfrey (2007), when considering treatments with large electorates (i.e. 27 and 51).

Result 2. *The turnout rate is higher in the minority group than in the majority group only in the Neutral treatment while no evidence of greater support for the underdog candidate is observed in the Framing one.*

Our second result shows that, while the *direction* of our theoretical prediction is in line with what is observed in the Neutral treatment, this is no more true in the Framing one.

Framing the experiment as an election has some behavioral implications. It seems that voters behave less strategically when they found themselves in a voting environment. The underdog's supporters vote less, while the majority group increases its participa-

tion. One possible explanation is the relationship between the electoral frame and the phenomenon of the Incumbency advantage. People may use heuristics when voting:¹² if they like to be associated with the winner in an election, even if this comes at a cost, they will go with the party most capable of winning. In the Framing treatment, being aware of a higher possibility of winning (i.e. being part of the majority party) may make people more likely to vote, whereas when being part of the minority party people do not bother to vote, since the probability of being associated with the winner is lower. In the following, we analyze voters' response to electoral fraud, which is the main objective of our study. Table 4 presents the aggregate turnout rates for both levels of fraudulent intervention.

Table 4: Turnout rates under the Fraud I and Fraud II conditions, in the Neutral and Framing treatment (standard error in parenthesis).

	Fraud I - 2 ballots		Fraud II - 5 ballots	
	Experiment	Theory	Experiment	Theory
Neutral Treatment				
Majority	0.399 (0.012)	0.291	0.377 (0.011)	0.188
Minority	0.444 (0.017)	0.302	0.434 (0.017)	0.108
Framing Treatment				
Majority	0.407 (0.012)	0.291	0.433 (0.013)	0.188
Minority	0.446 (0.019)	0.302	0.450 (0.018)	0.108

Our third result shows that voter turnout in the Fraud I condition is in line with our theoretical predictions. In particular:

¹²Heuristics are not equally used and useful in all (voting) contexts (Fortunato and Stevenson, 2014): heuristic's effectiveness is tied to the electoral context in which it is applied (Bower-bir and Damico, 2013).

Result 3. *Voter turnout increases in Fraud I condition, and this is true both for the minority and majority group and independently of the treatment.*

A low level of fraud increases turnout both in the minority and majority groups, as predicted by theory. In particular, we observe that, with respect to the Baseline condition, voting in the majority party increases to 40.67% and 39.94% in the Neutral and Framing treatment, respectively. Similarly, for the minority group we observe that in 44.56% and 44.44% of the cases people decide to vote. In the following we investigate, by means of a Wilcoxon signed rank test, whether the differences in turnout between the Baseline and Fraud I voting conditions are statistically significant, within the same treatment. Within the Neutral treatment, our findings reject the null hypothesis when using data from all rounds both when considering the minority ($p=0.072$, one-sided) and the majority group ($p=0.034$, one-sided). The same result holds within the Framing treatment ($p=0.034$, one-sided), thus providing evidence of turnout in Fraud I of both groups being significantly higher with respect to the Baseline condition. According to our theoretical model, the minority party should still be more likely to participate in the elections than the majority group. In Table 4 we observe that the turnout rate of the minority group is higher than the majority group in both treatments, even if the difference in total turnout is higher than expected. A set of Wilcoxon signed rank tests which compare turnout between groups within each treatment confirms our theoretical predictions in the Framing treatment, when considering all 25 rounds together ($p=0.072$, one sided). However, even if the difference in voting between the minority and the majority group is consistent with our theoretical predictions, the WSR test does not reach significance in the Neutral treatment ($p=0.137$, one sided).

As a further analysis, we ask whether the impact of the fraudulent intervention is different when considering the Framing or the Neutral treatment. Table 5 reports the

p-value from pairwise MW tests of voting behavior in Fraud I and Fraud II conditions as being the same between treatments, using session-level averages from all 25 rounds ("All") or from 5-round averages. Consider first the Fraud I condition. For all 25 rounds, voting behavior of the minority group is not significantly different when considering the Neutral or the Framing treatment. The same findings are obtained for the majority group, also when considering 5-round averages. With a moderate level of fraud, framing the experiment as an election or in neutral terms does not strongly affect participants' choices.

Table 5: Statistical test of differences in turnout between treatments by groups.

Majority group	Fraud I	Fraud II
All	0.154	0.074*
Rounds 1-5	0.040**	0.281
Rounds 5-10	0.279	0.037**
Rounds 10-15	0.174	0.192
Rounds 15-20	0.124	0.096*
Rounds 20-25	0.055	0.232
Minority group	Fraud I	Fraud II
All	0.386	0.500
Rounds 1-5	0.148	0.190
Rounds 5-10	0.020**	0.385
Rounds 10-15	0.124	0.442
Rounds 15-20	0.123	0.118
Rounds 20-25	0.331	0.121

One sided p value from Wilcoxon Mann Whitney test. Data are averaged at the session level.

Differently, in Fraud II condition we observe that, once again, framing the experiment as an election affects how participants react to an extended fraudulent intervention in favour of the minority group.

Result 4. *In Fraud II condition voter turnout is lower than in Fraud I condition, in accordance with theoretical predictions, only in the Neutral treatment. In the Framing treatment, voters in the majority group increase their participation in the election with respect to a lower level fraud.*

Consistent with our theoretical predictions, we see that, when increasing the fraudulent intervention in favour of the minority group to five ballots, the total voter turnout in the Neutral treatment is lower than when two fake ballots are added. In particular, voter turnout is equal to 41.44 and to 39.59 in Fraud I and Fraud II condition, respectively (WSR test, one-sided, $p=0.034$). Voting behavior follows the opposite pattern when considering the Framing treatment: voter turnout increases in Fraud II condition compared to Fraud I condition, from 41.96 to 43.89. Even if such difference is not statistically significant, according to a Wilcoxon signed rank test where each session is used a single observation, voting behavior observed in our data does not support our qualitative predictions when the experiment is framed as an election.

However, even when the experiment is expressed in neutral terms, participants overvote with respect to our theoretical predictions: while in the Neutral treatment the *direction* of our results is as expected (i.e. voter turnout is lower in Fraud II condition than in Fraud I condition), participation in the elections is still higher than in the Baseline. This tendency for over-participation was also observed in our Baseline condition and in Levine and Palfrey (2007). We further note, as stated in Results 4, that, in the Framing treatment, the majority group participated in voting at a significantly higher rate than in the Fraud I condition. As showed in Table 5, voters in the majority group (the group that is unfairly disadvantaged by ballot rigging) are significantly more likely to vote when the experiment is framed as an election in the Framing treatment than when using abstract terms in the Neutral one (MW test, one sided, $p=0.074$). Our results support the hypothesis that participants' reaction to framed election with fraud is not a purely strategic response. In Figure 4 we provide additional evidence of higher voter turnout of the majority group in the Fraud II condition of the Framing treatment with respect to the Neutral one. Interestingly, we observe that while in the Framing treatment the majority group immediately reacts to increasing fraudulent intervention in the elections, this is not true in the Neutral treatment, where voters do not change

their voting behavior substantially. Moreover, the turnout rate of the minority group develops in the opposite direction in the two treatments: while it decreases over time in the Neutral treatment, voters are more likely to vote when experiencing ballot rigging in the framing one. Such evidence seems to reflect the reaction of the minority group to the increasing turnout of the majority group: since those in the majority are overreacting to the electoral fraudulent intervention, those in the minority also increase their participation in order to win the elections.

Figure 4: Voter turnout across periods in the Baseline, Fraud I and Fraud II parts under the Neutral and Framing treatment.

Light grey dashed line indicates minority group turnout while dark grey solid line stands for average turnout in the majority group. Vertical lines divide, respectively, the Baseline, Fraud I and Fraud II parts of the experiment.

5.2 Individual level analysis

In the previous section we have analysed our experimental data at the aggregate level. Now, an individual level approach will enable us to capture the individual level heterogeneity that characterizes our data. In particular, given that we are dealing with a within subjects design, we have to take into consideration the standard omitted individual-heterogeneity issue. With panel data a random effects probit model tackles that challenge.¹³ Table 6 reports the marginal effects of a series of Probit regressions in which the dependent variable is *Vote*, a dummy variable which takes value one when the participant votes and zero when the participant decides to abstain. Our explanatory variables are the following. *Neutral Treatment* is equal to one for the Neutral treatment and zero otherwise while *FraudI* and *FraudII* are equal to one when participants are playing within the corresponding level of fraudulent intervention and zero otherwise: the omitted dummy variable is the Baseline condition thus the coefficients have to be considered with respect to this category. *Majority* is equal to one if the participant is a member of the majority group and zero otherwise and *Voting cost* is a categorical variable that indicates the randomly assigned voting cost in each round of the experiment. *Experience* takes into consideration the round of the experiment subjects are playing, in order to allow for learning. The binary variables *Vote previous round* and *Close previous round* explicitly consider, respectively, the voting decision and whether competition was close (i.e. one of the two parties won by 1 vote) in the previous round. While the first variable enables us to capture the effect of persistence in voting, besides any strategic evaluations about the participant's party membership and voting cost, the second one evaluates whether the closeness in competition in the previous round makes voters more

¹³We also control for possible observations' dependency within the same session by running a set of probit regressions with clusters at the session level: results are almost unchanged. Similar conclusions are obtained when clusters are based on individuals (since the same individuals participate in many voting decisions over the course of the experiment). See <https://sites.google.com/site/vmaggian/research>.

likely to vote. In all models we include a measure of individual risk aversion.¹⁴

Table 6: Regression table, random effects probit model.

	(1) All	(2) Framing	(3) Neutral	(4) All	(5) All
Vote					
Neutral Tr.	-0.0472 (0.0591)			-0.0342 (0.0627)	0.101 (0.0691)
FraudI	0.513*** (0.0425)	0.738*** (0.0779)	0.491*** (0.0780)	0.626*** (0.0603)	0.630*** (0.0604)
FraudII	0.654*** (0.0624)	0.838*** (0.101)	0.632*** (0.101)	0.749*** (0.0758)	0.752*** (0.0759)
Majority	-0.0929*** (0.0215)	0.118*** (0.0437)	-0.167*** (0.0436)	-0.0225 (0.0308)	0.0804** (0.0377)
Voting cost	-0.0575*** (0.000729)	-0.0560*** (0.00101)	-0.0596*** (0.00106)	-0.0576*** (0.000730)	-0.0577*** (0.000731)
Experience	-0.00673*** (0.000912)	-0.00786*** (0.00128)	-0.00563*** (0.00130)	-0.00675*** (0.000912)	-0.00678*** (0.000912)
Risk Aversion	0.0292*** (0.00576)	0.0469*** (0.00808)	0.0101 (0.00825)	0.0290*** (0.00576)	0.0290*** (0.00576)
Vote previous round	0.302*** (0.0210)	0.288*** (0.0295)	0.312*** (0.0301)	0.301*** (0.0210)	0.301*** (0.0210)
Close previous round	0.105*** (0.0227)	0.131*** (0.0333)	0.0781** (0.0314)	0.103*** (0.0228)	0.104*** (0.0228)
FraudI*Majority		-0.270*** (0.0738)	-0.0371 (0.0747)	-0.156*** (0.0524)	-0.158*** (0.0525)
FraudII*Majority		-0.153** (0.0742)	-0.0845 (0.0743)	-0.119** (0.0525)	-0.119** (0.0525)
Neutral*FraudI				-0.0197 (0.0497)	-0.0203 (0.0498)
Neutral*FraudII				-0.0301 (0.0498)	-0.0306 (0.0498)
Neutral*Majority					-0.204*** (0.0430)
Constant	0.938*** (0.0646)	0.674*** (0.0769)	1.082*** (0.0889)	0.888*** (0.0668)	0.822*** (0.0684)
Constant	-5.090*** (0.569)	-5.999*** (0.946)	-4.652*** (0.772)	-5.094*** (0.569)	-5.082*** (0.568)
Observations	21384	10692	10692	21384	21384
Log lik.	-9941.9	-5022.3	-4889.4	-9936.4	-9925.1
rho	0.00612	0.00247	0.00945	0.00610	0.00617
Chi-squared	6369.7	3181.5	3197.4	6373.5	6375.9
Prob>chi2	0.000	0.000	0.000	0.000	0.000

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

¹⁴In particular, each subject has to make ten choices between two lotteries, one is more remunerative but more risky, while the other is safer but provides the subject with a lower amount of earnings. See Holt and Laury (2002) for more details on the test. We measure risk aversion as the number of safe choices made by the individual, ignoring possible switching from one type of lotteries to another. The analysis reported in this paper changes very little if we instead drop those 42 subjects who switched more than once from B back to A or vice-versa.

First of all, our variables of interest *FraudI* (limited fraudulent intervention with 2 ballots) and *FraudII* (extended intervention with 5 ballots) considerably increase the probability of a subject voting. Using Neutral terms does not have a statistically significant effect per se on voting behavior in Model (1). However, it is possible that framing the experiment as an election with fraud has different effects on the majority and on the minority group, in particular by mobilizing the first to participate in elections in order to fight the electoral procedural unfairness. In model (2) and (3) we separately analyse voting behavior in the Framing and Neutral treatment. As an ulterior confirmation of the analysis made at the aggregate level, we observe that the coefficient of the *Majority* variable is significant in both models but with an opposite sign: while in the Baseline condition of the the Framing treatment the Majority group is more likely to vote than the minority one, the opposite is observed in the Neutral treatment. Moreover, the coefficient of the interaction terms *FraudI*Majority* and *FraudII*Majority* are negative and significant only when considering the Framing treatment, meaning that the difference in voter turnout between the majority and the minority group diminishes in both the Fraud I and Fraud II conditions with respect to the Baseline only in the Framing treatment. Our findings provide additional evidence that Framing the experiment as an election affects voter turnout of the minority and majority groups, in particular when introducing electoral fraud.

In model (5) we further test for the effect of framing. The negative and significant coefficient of the interaction term *Neutral*Majority* shows how much being in the majority rather than in the minority group negatively affects voting behavior in the Neutral treatment with respect to the Framing one. The interaction terms *Neutral*FraudI* and *Neutral*FraudII* are never significant in Model (4): the overall effect of framing is not different when introducing the two types of fraud in our three conditions.

Subjects also respond to changes in the environment, in particular to voting costs and party identity, as shown by the significance of the coefficient of the corresponding vari-

ables in Model (1). Interestingly, voting in the previous round significantly increases the probability of voting only in the Neutral treatment, as shown by the significant coefficient of *Vote previous round* in model (3), while such a carry over effect does not affect voter turnout in the Framing treatment. In model (2) and (3) we further observe that having experienced close electoral competition in a previous round positively affects voter turnout, even if such an effect is higher in the Framing treatment than in the Neutral one, where the probability of voting increases respectively by 15 and 9 percent. Interestingly, in model (2) we further observe that being risk averse significantly increases the probability of voting only in the Framing treatment ($p < 0.10$): the more risk averse a voter is, the more likely he will participate in the election, but only when the environment is sufficiently framed. Finally, the negative coefficient of the *Experience* variable in all models proves that learning occurs, we thus observe convergence of voting behavior in the direction suggested by the equilibrium.

Results presented in Table 6 do not vary when adding control for age, gender and attitudes towards vote to the regression.

5.2.1 Electoral fraud and win probability

Does the Incumbent win more often than the other candidate? Is electoral fraud a good strategy to win elections? In order to answer these questions, in this section we look at the probability of the Incumbent winning the elections depending on the extent of the electoral fraud and on framing.

In Figure 5 we report the probability of winning the election for the minority party across treatments and with respect to the model's predictions. In the Neutral treatment the results of the experiment are closer to our predictions than when considering the Framing treatment. In the Baseline condition the underdog candidate wins by 12% of the total number of elections in the Neutral treatment but only the 7% of elections in the Framing one. While according to a MW test this difference does not reach

significance, it is consistent with our previous results showing that the majority party over-vote when the experiment is framed as an election with respect to a situation where labels are abstract. When adding two fake ballots to the minority group, as expected, we observe an increase in the winning frequency in both the Framing (33%) and the Neutral (32.5%) treatment. Interestingly, a higher level of electoral fraud seems not to have the same effectiveness in increasing the frequency of the minority party winning in the Framing and in the Neutral treatment. Figure 5 shows that in the Framing treatment the minority party is less likely to win elections with respect to the Neutral one: the frequency of victories increases to 81% and to 66% respectively in the Neutral and Framing treatment.¹⁵

Figure 5: Probability to win the elections - Minority party.

¹⁵In the current analysis we consider the frequency of winning the elections by the minority party with respect to the total number of elections, also including the tied elections, by comparing data grouped at the session level. When excluding tied elections from our analysis we observe almost the same results, except for the fact that the difference in the probability of winning the elections in the Fraud II condition between the Framing and Neutral treatment is significant (MW test, two-sided, $p=0.081$).

In table 7, a random effects probit regression provides additional evidence on the role of behavioral reaction to electoral fraud. Our dependent variable is *Majority wins*, which is equal to one when the Majority group wins the elections and zero otherwise.

Table 7: Regression table, random effects probit model.

	(1) All
Majority wins	
Neutral Tr.	-0.187 (0.162)
FraudI	-1.099*** (0.0522)
FraudII	-1.637*** (0.0508)
Neutral*FraudI	0.167*** (0.0548)
Neutral*FraudII	-0.319*** (0.0551)
Experience_FraudI	0.00551** (0.00261)
Experience_FraudII	-0.0282*** (0.00257)
Close previous round	0.0524** (0.0240)
Constant	1.696*** (0.115)
Constant	-3.006*** (0.510)
Observations	21600
Log lik.	-8436.3
rho	0.0472
Chi-squared	6268.4
Prob>chi2	0.000
Standard errors in parentheses	
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$	

We firstly observe that the probability that the majority wins the elections in the Baseline condition is not different when considering the Neutral or the Framing treatment, as shown by the insignificance of the coefficient of the *Neutral Treatment* variable. When labels are framed, the probability of winning the elections for the majority group decreases with the level of fraud: both the coefficients of Fraud I and Fraud II variables are indeed negative and significant. However, the effect of the electoral fraud is shown

to depend on the treatment. The coefficient of the interaction term $Neutral * FraudI$ is positive and significant since the probability of winning with respect to the Baseline condition is more likely to decrease in the Framing treatment than in the Neutral one (respectively, from 95.27 to 70.57 and from 92.16 to 69.1). The coefficient of $Fraud II$ is negative and significant showing that, with respect to the Baseline treatment, adding five fake ballots to the minority party is less effective in reducing the probability of winning the elections for the majority group in the Framing treatment than in the Neutral one. As shown in the previous paragraph the majority group votes more often when the experiment is framed as an election with five fake ballots than when it is framed in neutral terms, counter acting the effectiveness of fraud in winning the election.

6 Conclusion

This paper provides evidence that electoral fraud might increase rather than decrease voter turnout. While a high turnout is usually considered as a signal of a fair election we show, by means of a laboratory experiment, that this is not necessarily the case. Our strategic costly voting model predicts an increase in turnout when the fraudulent intervention is moderate and a sharp drop in participation whenever fraud is pervasive. Results of our laboratory experiment confirm the first part of our theoretical predictions, irrespective of the experimental framing (neutral vs. elections). However, when a more pervasive fraudulent intervention affects the electoral procedure, our theoretical results are not totally matched: depending on whether the experiment is framed as an election or not, a different voting pattern is observed. A slight decrease in participation rates is observed in the Neutral treatment when going from a low to a high level of ballot-box stuffing. Differently, when the experiment is explicitly framed as an election, a higher level of electoral fraud increases turnout among the majority group (compared with a low level of fraud). These differences call for explanations that are beyond pure

mistakes in strategic responses. We believe that such behaviour is an attempt of the majority group, who is negatively affected by the illicit intervention, to react to an unfair procedure, even if such an action is costly.

Our paper shows that voters fail to rationally react to a situation where a number of fake votes are added to the elections: often in the real world people are confronted with situations where the amount of illicit behavior in elections is uncertain, rather than having clear knowledge about the level of the illegitimate intervention. We have now provided a first evidence of voters' behavior in a clear and clean environment. Examining voters' behavior in a situation of uncertainty where the *perception* of fairness plays a role is thus an interesting extension of the paper.

References

- Birch, S., 2010. Perceptions of electoral fairness and voter turnout. *Comparative Political Studies* 43 (12), 1601–1622.
- Bolton, G. E., Brandts, J., Ockenfels, A., 2005. Fair procedures: Evidence from games involving lotteries. *The Economic Journal* 115 (506), 1054–1076.
- Bower-bir, J. S., Damico, N. J., 2013. A tool for all people, but not all occasions: How voting heuristics interact with political knowledge and environment.
- Bratton, M., 2008. Vote buying and violence in nigerian election campaigns. *Electoral Studies* 27, 621–632.
- Butler, M., Marchal, M. A., 2007. Framing effects in political decision making: Evidence from a natural voting experiment. CESifo Working Paper Series 1940.
- Carreras, M., İrepoğlu Yasemin, 2013. Trust in elections, vote buying, and turnout in latin america. *Electoral Studies* 32 (4), 609–619.
- Collier, P., Hoeffler, A., 2009. Democracys achilles heel or, how to win an election without really trying. Working paper series Centre for the Study of African Economies, Department of Economics, Oxford.
- Collier, P., Vicente, P. C., 2012. Violence, bribery, and fraud: the political economy of elections in sub-saharan africa. *Public choice* 153 (1), 117–147.
- Dekel, E., Jackson, M. O., Wolinsky, A., 2008. Vote buying: General elections. *Journal of Political Economy* 116 (2), 351–380.
- Deutsch, M., 1975. Equity, equality, and need: What determines which value will be used as the basis of distributive justice? *Journal of Social Issues* 31 (3).
- Douglas, J. A., 2013. Procedural fairness in election contests. *Indiana Law Journal* 88.

- Druckman, J. N., 2001. On the limits of framing effects: Who can frame? *The Journal of Politics* 63 (4), 1041–1066.
- Druckman, J. N., 2004. Political preference formation: Competition, deliberation and the (ir)relevance of framing effects. *American Political Science Review* 98 (4), 671–686.
- Ellman, M., Wantchekon, L., May 2000. Electoral competition under the threat of political unrest. *The Quarterly Journal of Economics* 115 (2), 499–531.
- Enikolopov, R., Korovkin, V., Petrova, M., Sonin, K., Zakharov, A., 2013. Field experiment estimate of electoral fraud in russian parliamentary elections. *PNAS* 110 (2), 448–452.
- Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Fortunato, D., Stevenson, R. T. ., 2014. Heuristics in context. Manuscript, Rice University.
- Fukumoto, K., Horiuchi, Y., 2011. Making outsiders’ votes count: Detecting electoral fraud through a natural experiment. *American Political Science Review* 105 (3), 586–603.
- Greiner, B., 2004. An online recruitment system for economic experiments. *Forschung und wissenschaftliches Rechnen* 2003. GWDG Bericht 63, 79–93.
- Grosser, J., Kugler, T., Schram, A., 2005. Preference uncertainty, voter participation and electoral efficiency: An experimental study. University of Cologne. Working Paper Series in Economics.
- Grosser, J., Schram, A., 2010. Public opinion polls, voter turnout, and welfare: An experimental study. *American Journal of Political Science* 54 (3), 700–717.

- Holt, C. A., Laury, S. K., 2002. Risk aversion and incentive effects. *American Economic Review* 92 (5), 1644–1655.
- Lehoucq, F. E., 2002. *Stuffing the ballot box: Fraud, electoral reform and democratization in Costa Rica*. Cambridge University Press.
- Levine, D. K., Palfrey, T. R., 2007. The paradox of voter participation? a laboratory study. *American Political Science Review* 101 (1), 143–158.
- McCann, J. A., Dominguez, J. I., 1998. *Mexicans React To Electoral Fraud And Political Corruption: An Assessment Of Public Opinion And Voting Behavior*. Elsevier.
- Minnite, L., Callahan, D., 2003. Securing the vote: ana analysis of election fraud. *Demos* 14.
- Norris, P., 2014. *Why Electoral Integrity Matters*. Cambridge University Press.
- Palfrey, T. R., 2009. Laboratory experiments in political economy. *Annual Review of Political Science* 12, 379–388.
- Palfrey, T. R., Rosenthal, H., 1985. Voter participation and strategic uncertainty. *American Political Science Review* 79 (1), 62–78.
- Schram, A., Sonnemans, J., 1996. Why people vote: Experimental evidence. *Journal of Economic Psychology* 17 (4), 417–442.
- Schuck, A. R., Vliegenthart, R., De Vreese, C. H., 2014. Who’s afraid of conflict? the mobilizing effect of conflict framing in campaign news. *British Journal of Political Science*, 1–18.
- Schuck, A. R. T., Vliegenthart, R., Boomgaarden, H. G., Elenbaas, M., Azrout, R., van Spanje, J., de Vreese, C. H., 2013. Explaining campaign news coverage: How medium, time and context explain variation in the media framing of the 2009 european parliamentary elections. *Journal of Political Marketing* 12 (1), 8–28.

- Simpser, A., 2013. *Why Governments and Parties Manipulate Elections: Theory, Practice, and Implications*. Cambridge University Press.
- Tarrow, S., 1998. *Power in Movement: social movements and contentious politics*. Cambridge University Press.
- Valentino, N. A., Beckmann, M. N., Buhr, T. A., 2001. A spiral of cynicism for some: The contingent effects of campaign news frames on participation and confidence in government. *Political Communication* 18 (4), 347–367.
- Vergne, C., 2009. Turnout in developing countries: The effect of mass media on national voter participation. Center for Studies and Researches on International Development CERDI.
- Vorobyev, D., 2016. Participation in fraudulent elections. *Social choice and welfare* 46 (4), 863–892.
- Wilking, J. R., March 2011. The portability of electoral procedural fairness: Evidence from experimental studies in china and the united states. *Political Behavior* 33 (1), 137–159.
- Wilks-Heeg, S., 2008. Purity of elections in the uk: Causes for concern. Tech. rep., The Joseph Rowntree Reform Trust Ltd.
- Wolak, J., 2014. How campaigns promote the legitimacy of elections. *Electoral Studies* 34, 205–215.