CERGE
Center for Economics Research and Graduate Education
Charles University Prague

Essays on Electoral Fraud

Dmitriy Vorobyev

Dissertation

Prague, December 2013

Dissertation Committee

Libor Dusek (CERGE-EI, Chair)
Levent Celik (CERGE-EI)
Jan Hanousek (CERGE-EI)
Peter Katuscak (CERGE-EI)

Referees

Michael Alvarez (California Institute of Technology)
John Ledyard (California Institute of Technology)

Acknowledgements

I am grateful to my supervisors, Libor Dusek, Levent Celik, Peter Katuscak and Jan Hanousek, for their inspiring guidance and support. I would like to express my gratitude to many people for their time in discussing the work, especially Slava Vinogradov, Jan Zapal and Jung Hun Cho. I am thankful to John Ledyard and Michael Alvarez, who refereed the dissertation, for their helpful comments. I am deeply indebted to Laura Strakova and Deborah Novakova who carefully edited all the chapters. Finally, I would like to thank two institutions, CERGE-EI for enabling me to work in a wonderful environment, and the California Institute of Technology for its generous hospitality during my research stay.

Contents

Introduction

Elections are fundamental to democracy. Over the last decades researchers have extensively studied various aspects of an electoral game such as behavior of voters and candidates or impact of different voting rules on outcomes and efficiency. Though the approaches taken to analyzing elections are incredibly diverse, almost all of them share one common feature: it is assumed that elections are well-functioning mechanisms for converting public preferences into social choice, and that the outcomes of elections are fully determined by votes casted for the candidates. However, in reality fraud has become an integral part of electoral competition in both established democracies and less-than-democratic regimes, despite tremendous efforts exerted by international organizations to ensure transparency in elections. In recent years, both economists and political scientists have started to pay closer attention to elections that lack integrity, but there are still plenty of open questions on all sides of the electoral game in a fraudulent context. In my dissertation I study various aspects of electoral fraud, both theoretically and empirically, to answer some of them.

In the first chapter, I explore how electoral fraud affects voters' participation decisions. For this purpose I analyze a costly voting model of elections, where the incumbent can stuff the ballot box. I find that, in contrast to clean elections where there is usually a unique equilibrium, two stable equilibria may exist: full abstention equilibrium, where the incumbent wins with certainty and which exists only if the incumbent's capability to stuff a ballot box is sufficiently strong; and a more efficient coordination equilibrium, where a substantial share of a challenger's supporters vote and the probability that the incumbent will be defeated is high. Since voters do not take into account positive externality they produce on other voters when deciding to cast their votes, participation in coordination equilibrium is still inefficiently low. Thus, subsidization as well as introducing compulsory voting may improve efficiency. Since the higher capability of the incumbent to stuff a ballot box discourages the participation of his own supporters and can create coordination incentives for the challenger's supporters, higher fraud does not always benefit the incumbent even when it is costless. The model simultaneously explains two empirical observations about fraudulent elections: a positive relationship between fraud and victory margins and a negative effect of fraud on turnout.

All the main findings on electoral fraud in previous literature are derived from the analysis of particular elections in a given country at a given moment, while the issues of fraud dynamics have attracted limited attention. This is surprising, because studying the evolution of fraud seems extremely important from many different perspectives. Comprehensively studying political regimes, designing effective electoral legislation and, especially, assessing the effectiveness of electoral monitoring are much harder to do without an understanding of fraud dynamics. In the second chapter of my dissertation, I study the evolution of electoral fraud over life-times of non-democratic regimes. I present evidence of fraud growth over the lifetime of non-democratic regimes using the example of post-Soviet and Sub-Saharan countries, and provide a theoretical framework to explain the observed tendency. I develop a probabilistic voting model of electoral competition with falsifications where the incumbent faces two types of uncertainty. First, he is uncertain about voters' attitude towards fraud and, second, he does not know for sure his true level of support because of a random component in voters' preferences for the candidates. The model predicts that lower uncertainty about voters' fraud intolerance provides stronger incentives for the incumbent to commit fraud. Over time, with each following election, the incumbent becomes more certain about voters' reaction to fraud as he learns through Bayesian updating and, thus, as the deterrent role of fraud intolerance uncertainty declines, incentives to commit fraud become stronger, providing a growing fraud profile.

One reason electoral fraud suffers from a relative lack of attention in the academic literature is the absence of a reliable measure of fraud. The inability to measure fraud in a consistent way precludes implementation of reliable empirical research on fraudulent elections, which in turn discourages efforts towards a theoretical study of electoral fraud, as it is hard to test any theory in this field. The existing methods of fraud detection are more qualitative than quantitative, often based on the subjective assessment of electoral transparency by observers or other participants of the electoral process, and the results they produce may not always be treated as fully reliable. In the third dissertation chapter, I suggest a simple statistical method for testing for the presence of fraud in the forms of ballot stuffing, multiple voting and vote buying (probably the most widespread fraud techniques), and for estimating the magnitude of fraud, even when very limited official electoral data are available. The method is based on the observation that ballot stuffing, when it takes place in a given precinct, results in an increase both in reported turnout and in the incumbent's vote share. Consequently, such fraudulent precinct moves in turnout distribution towards its right tail . Hence, precincts with relatively low reported turnout are more likely to be clean. Using the information on relatively clean precincts, it is possible to simulate counterfactual data for infected precincts and compare them with the observed data. The extra advantage of an incumbent over the runner-up in precincts with high reported turnout in comparison to precincts with low turnout then implies incidences of ballot stuffing. The method is first piloted on artificial and artificially fraudulent real data, and subsequently applied to test the fairness of the Russian executive elections held between 2000 and 2012, in which transparency and integrity were dubious. Results strongly reject the hypothesis of an absence of ballot stuffing and demonstrate that electoral fraud in Russia has been growing significantly over the last twelve years, providing an additional support to the idea of growing fraud in non-democracies presented in the second dissertation chapter.

Úvod

Volby jsou klíčovým prvkem demokracie. Výzkumníci se v posledních desetiletích intenzivně věnovali různým aspektům volební hry, včetně chování voličů a kandidátů, dopadů rozličných volebních pravidel na výsledky a produktivitu atd. Ačkoli přístupy k analýze voleb jsou neuvěřitelně rozmanité, téměř všechny se však vyznačují jedním společným rysem: předpokládá se, že volby představují dobře fungující mechanismy transformace preferencí veřejnosti ve společenskou volbu a že výsledky voleb jsou zcela určovány hlasy odevzdanými pro kandidáty. Ve skutečnosti se však nedílnou součástí volební soutěže v zavedených demokraciích i v méně demokratických režimech stalo podvodné jednání, a to i přes nesmírné úsilí vynakládané mezinárodními organizacemi k zajištění transparentnosti voleb. Experti na ekonomii a politickou vědu v nedávných letech začali blíže zkoumat volby, které postrádají integritu; přesto však existuje celá řada nezodpovězených otázek na všech stranách volební hry v kontextu nepoctivosti či podvodů. V mé disertační práci se věnujeme zkoumání různých aspektů podvodů ve volbách, a to na rovině teoretické i empirické.

V první kapitole jsme se zaměřili na prozkoumání otázky, jaký vliv má podvod ve volbách na rozhodnutí voličů o účasti. Pro tento účel jsme provedli analýzu nákladného hlasovacího modelu voleb, kdy stávající držitel hlasu může do volební urny přidávat hlasy. Zjistili jsme, že na rozdíl od nezmanipulovaných voleb, kde obvykle existuje jediná rovnováha, mohou existovat dvě stabilní rovnovážné situace: rovnováha dosažená plnou neúčastí, kdy stávající držitel hlasu vyhrává s jistotou, a která existuje pouze v případě, že schopnost stávajícího držitele hlasu naplnit volební urnu je dostatečně silná; a účinnější koordinační rovnováha, kdy hlasuje podstatná část podporovatelů vyzyvatele a pravděpodobnost porážky stávajícího držitele hlasu je vysoká. Vzhledem k tomu, že voliči neberou v úvahu pozitivní externality, násobí svůj vliv na ostatní voliče, aby odevzdali své hlasy; účast na koordinační rovnováze je přesto neefektivně nízká, a proto její produktivitu může zlepšit dotace (sponzorování) nebo zavedení povinné účasti ve volbách. Vzhledem k tomu, že skutečnost, že stávající držitel hlasu má lepší možnost doplnit ve volbách hlasy neexistujících voličů, odrazuje od účasti jeho vlastní podporovatele, a vede k vytváření koordinovaných pobídek na straně podporovatelů jeho vyzyvatele, vyšší míra podvodu není stávajícímu držiteli hlasů vždy ku prospěchu, i s ním nejsou spojeny žádné náklady. Model zárove vysvětluje dvě empirická pozorování týkající se voleb, jejichž součástí byl podvod: pozitivní vztah mezi podvodem a převahou hlasů potřebnou k vítězství (tzv. victory margin) a negativní dopad podvodu na účast ve volbách.

Všechna hlavní zjištění v dřívější literatuře, týkající se falšování voleb a podvodů ve volbách, jsou odvozena z analýzy konkrétních voleb v dané zemi v daném okamžiku, avšak málo pozornosti bylo věnováno otázkám spojeným s dynamikou podvodu (jak to funguje). Což je překvapující, vzhledem k tomu, že studium toho, jak podvod vzniká a vyvíjí se, je mimořádně důležité z mnoha různých pohledů. Komplexní studium politických režimů, návrhy účinné volební legislativy a obzvláš hodnocení účinnosti monitoringu voleb se provádí mnohem obtížněji bez pochopení toho, jaká je dynamika podvodu. V druhé kapitole mé disertační práce sleduji vývoj volebních podvodů během celé doby trvání nedemokratických režimů. Na příkladech post-sovětských zemí a zemí ze subsaharské oblasti předkládám důkazy o tom, že během trvání nedemokratických režimů míra podvodů roste; též přinášíme teoretický rámec pro vysvětlení pozorovaných tendencí. Vyvinuli jsme pravděpodobnostní hlasovací model volební soutěže s falzifikacemi, kdy stávající držitel hlasů čelí dvěma druhům nejistoty. Jednak si není jist postojem voličů k podvodu a za druhé nemá jistotu ohledně své skutečné podpory kvůli faktoru nahodilosti v preferencích voličů vzhledem k daným kandidátům. Tento model předpovídá, že nižší nejistota ohledně netolerance podvodu u voličů vede k vyšší pravděpodobnosti, že stávající držitel hlasů se dopustí podvodu. Postupem doby stávající držitel hlasů s každými následujícími volbami získává větší jistotu ohledně reakce voličů na podvod, protože se poučí pomocí Bayesovské statistiky, a proto jak odrazující úloha netolerantního postoje k podvodu klesá, narůstá inklinace ke spáchání podvodu, v důsledku čehož pozorujeme nárůst profilu volebního podvodu.

Jedním z důvodů, proč se volebním podvodům v akademické literatuře věnuje relativně málo pozornosti, je skutečnost, že neexistuje spolehlivé měřítko podvodu. Neschopnost změřit rozsah podvodu konzistentním způsobem vylučuje provedení spolehlivé empirické studie o podvodných volbách, což dále odrazuje výzkumníky, aby napřeli své úsilí k teoretickému zkoumání podvodu ve volbách, protože na tomto poli je obtížné otestovat jakoukoli teorii. Existující metody zjišování podvodu jsou spíše kvalitativní, než kvantitativní, a často se zakládají na subjektivním zhodnocení transparentnosti voleb pozorovateli či jinými účastníky volebního procesu, přičemž výsledky, k nimž se dospívá, nelze vždy považovat za zcela spolehlivé. Ve třetí kapitole mé disertační práce navrhujeme jednoduchou statistickou metodu testování na přítomnost podvodu ve formě doplování hlasů neexistujících voličů (tzv. ballot stuffing), opakovaného odevzdání jednoho hlasu a nákupu hlasů, které jsou patrně nejrozšířenějšími technikami podvodného jednání u voleb, a také na odhad míry, ve které došlo k podvodu ve formě doplování hlasů neexistujících voličů, které lze uplatnit i v případech, že jsou k dispozici jen velmi omezené oficiální údaje o volbách. Metoda je založena na pozorování, že pokud se v daném volebním okrsku uskuteční podvod ve formě doplování hlasů neexistujících voličů, vede to ke zvýšení hlášeného počtu voličů, kteří se k volbám dostavili, i k nárůstu volebního zisku stávajícího držitele hlasu. V konečném důsledku se tak volební okrsek, v němž došlo k tomuto podvodu, posune na ose účasti hlasujících více směrem doprava. Výsledky z volebních okrsků s hlášenou relativně nízkou účastí ve volbách mají tudíž větší pravděpodobnost, že budou čisté. Použitím informací z relativně čistých volebních okrsků je možné nasimulovat srovnávací údaje pro infikované volební okrsky a porovnat tyto údaje s pozorovanými údaji. Pokud stávající držitel hlasů má mimořádný náskok před vyzyvatelem ve volebních okrscích s hlášenou vysokou účastí ve srovnání s okrsky, kde je hlášena jen nízká účast ve volbách, naznačuje to přítomnost podvodu ve formě doplování hlasů neexistujících voličů. Metoda byla poprvé vyzkoušena na umělých a uměle zkreslených skutečných údajích a následně byla využita ke zkoumání férovosti voleb představitelů exekutivy v Rusku, které se konaly mezi roky 2000 a 2012, a jejichž transparentnost a integrita jsou pochybné. Výsledky jednoznačně nepotvrdily hypotézu o nepřibývání hlasů neexistujících voličů a ukazují, že za posledních dvanáct let v Rusku významným způsobem narůstá podvodné jednání u voleb, což dále podporuje myšlenku prezentovanou v druhé kapitole této disertace.

Chapter 1. Participation in Fraudulent Elections

Abstract

I analyze a costly voting model of elections, in which the incumbent can stuff the ballot box, to investigate how electoral fraud affects the participation decisions of voters. I find that two stable equilibria may exist: first, a full abstention equilibrium, where the incumbent wins with certainty, which exists only if the incumbent's capability to stuff a ballot box is sufficiently strong. Second, a more efficient coordination equilibrium, where a substantial share of a challenger's supporters vote and the probability of the incumbent being defeated is large. Since voters do not take into account positive externality they produce on other voters when deciding to cast their votes, participation in coordination equilibrium is still inefficiently low. Thus, subsidization as well as introducing compulsory voting may improve efficiency. Because the higher capability of the incumbent to stuff a ballot box discourages the participation of his own supporters and creates coordination incentives for the challenger's supporters, higher fraud does not always benefit the incumbent, even when costless. Additionally, the model simultaneously explains two empirical observations about fraudulent elections: a positive relationship between fraud and victory margin and a negative effect of fraud on turnout.

JEL Classification: D72, D73
Keywords: Voting, Fraud, Participation

1.1 Introduction

Participation in elections has been widely studied by economists. Why do voters vote? How do they decide to participate or to abstain? Is the participation level in voluntary elections efficient or are there too many or too few voting? What electoral policies may improve efficiency? Researchers have been addressing these questions for decades. In recent years, both economists and political scientists have begun to pay closer attention to elections that lack integrity, but there are still plenty of open questions on all sides of the electoral game in a fraudulent context, including the impact of fraud on the participation incentives of voters. In this paper I theoretically analyze how the presence of fraud affects voters' participation decisions and thus impacts social welfare. I further explain several puzzling empirical observations about fraudulent elections such as a positive relationship between fraud and victory margins as well as a negative effect of fraud on turnout.

Indeed, if voters anticipate that elections will be tainted by fraud, their decisions must be different. In fact, there is a substantial body of empirical evidence suggesting that voters behave differently in fraudulent elections than they do in 'clean' elections. They are less likely to participate in fraudulent elections, for example. This observation has been verified by McCann and Dominguez (1998), Hiskey and Bowler (2005), and more recently by Simpser (2012) on the example of Mexico, by Birsch (2010), in a study of cross-country electoral survey data from both new and established democracies, and by Landry, Davis and Wang (2010) on the example of elections in China. Though it is well-established that voters reach different decisions in fraudulent elections, the mechanism which leads to these differences has not been investigated in any depth, and the literature on fraudulent elections at best simply assumes that voters have stronger incentives to abstain if they expect fraud.

While the behavior of voters in a fraudulent context has not garnered much attention in the academic literature, the behavior of candidates has been studied fairly well. The literature generally considers fraud to be more than just a means of getting extra votes, and often models it as political pressure or violence. The main question raised by scholars is why and when corrupt incumbents choose to use political pressure in electoral competition. Chatuverdi (2005), for example, studies a competition between parties which can allocate resources between ideological campaigning and political violence. He shows that competing parties use more political pressure if the outcome of elections is more uncertain ex-ante. Once there is a bias towards one candidate in terms of support, competing parties prefer to gain votes through ideological campaigning. Magaloni (2010) focuses on the incentives of a corrupt incumbent to hold clean elections, and shows that under substantial threat of revolt, incumbents prefer to avoid fraud. Recently, Collier and Vicente (2012) have distinguished several illegal strategies used by candidates to affect the outcome of elections, including violence, repression, and electoral manipulations, and show that the choice of strategy adopted depends on the strength of the candidates' support: a weak challenger would prefer to use violence and a weak incumbent would use repression, while an incumbent with strong support would prefer to bribe electoral officials and stuff the ballot box.

Instead of considering fraud as political pressure or violence which affects voters' or candidates' utility, I suggest thinking about it as ballot stuffing. In such a setup voters do not directly suffer from an incumbent's actions but instead anticipate that if they abstain, their votes are likely to be counted in favor of the corrupt incumbent. Such an approach allows me to focus on the effect of fraud on voters' behavior from a purely pivotal perspective.

I therefore model electoral fraud as ballot stuffing, assuming that if a voter does not participate in elections, his unused ballot may be transformed into a vote for the incumbent. Indeed there is a wide range of technologies for rigging elections, and ballot stuffing is just one of them[1]. From the modeling perspective, the variety of fraud technologies which directly influence the reported vote shares of the candidates may be divided into three groups according to the underlying mechanism of lending an advantage to the incumbent. Techniques from the first group transform votes cast for the challenger into votes for the incumbent. This category includes rigging the software for electronic voting machines or designing the ballot so that it consistently leads voters to vote for another candidate than they prefer[2]. The second group utilizes technologies that reduce the number of votes cast for the challenger, such as invalidation and destruction of ballots for wrong candidates. The third group of technologies consists of techniques that directly or indirectly[3] transform unused ballots into additional votes for the incumbent: ballot stuffing, multiple voting, and vote buying, for example.

Though direct evidence is hard to find, the methods of the last group of techniques are likely to be more widespread and account for a larger share of fraudulent activities than the methods from the other two groups, which are much less cost effective and far more limited in adding to the incumbent's advantage. Modeling fraud technology as adding extra votes for the incumbent at the expense of voters who do not participate is, therefore, the most natural way to proceed. In the rest of the paper I thus refer to the fraud technology used in my model as ballot stuffing, though it also accounts for multiple voting, vote buying, and all other fraud techniques that increase the number of votes for the incumbent using, directly or indirectly, the actual ballots of voters who abstained from voting.

I analyze a pivotal costly voting model of elections, in which the incumbent can stuff a ballot box, to investigate how the behavior of voters would change if they know that, were they

[1] See, for example, Lehoucq (2003) for the description of electoral irregularities observed in various elections.

[2] The infamous Florida butterfly ballot in the US 2000 Presidential elections is an example.

[3] I label a technology indirect if it does not literally convert unused ballot papers into votes, but the number of extra votes that the incumbent may get through it is limited by the number of unused ballots, since reported turnout cannot exceed 100%.

to abstain, their votes may be counted in favor of the incumbent. I further investigate what the welfare consequences of this change are, what may be done in order to increase welfare, and how the incumbent would behave if he anticipates voters' response to fraud.

The main findings are the following:

1) Two stable equilbria may exist: an equilibrium with full abstention, where none of the voters vote and the incumbent wins with probability 1, and a coordination equilibrium where a positive share of the challenger's supporters votes, and the probability that the incumbent will lose is high. Full abstention equilibrium exists only when the incumbent's capability to stuff a ballot box is sufficiently large.

2) Coordination equilibrium is likely to deliver higher welfare than full abstention equilibrium. Participation in coordination equilibrium is below an efficient level, since voters do not take into account the positive externality they produce on other voters when making participation decisions. Subsidy and, in some cases, introducing compulsory voting may improve efficiency.

3) Higher capability of the incumbent to stuff the ballot box discourages the participation of the incumbent's supporters, requires stronger coordination among the challenger's supporters, and leads to higher participation of the supporters of the challenger, conditional on coordination being achieved.

4) Higher fraud capability does not always benefit the incumbent, even when costless.

The rest of the paper is organized as follows. In the next section I describe a pivotal private value model of costly voting, where voters decide whether to participate in elections or abstain by comparing their individual specific voting costs with the expected benefit, which involves a probability to cast a decisive vote, i.e. to be pivotal. I then analyze the case where the incumbent can stuff a ballot box perfectly. I show that in addition to the full abstention equilibrium, where none of the voters vote, the incumbent stuffs 100% of votes, and wins with probability one, a relatively more efficient stable coordination equilibrium exists, where a substantial share of the challenger's supporters vote, the number of stuffed ballots is relatively low, and the challenger is likely to win. After characterizing properties of the equilibria, I focus on welfare and show that in a coordination equilibrium voters' participation is inefficiently low, since the voters ignore the externality they produce on other voters when making voting decisions, and thus, subsidizing participation or even introducing compulsory voting may improve welfare. I then generalize the model by allowing fraud to be imperfect. Instead of assuming that the vote of a non-participant is stuffed in favor of the incumbent with certainty, I assume that the incumbent can steal a non-participant's vote with some probability, which can be thought of as the incumbent's fraud capability. This generalization allows me to analyze the whole range of elections, from clean to totally fraudulent, though it makes the analytical solution extremely challenging to obtain. I explore how changes in this probability affect properties of the equilibria, and then study the choice of the incumbent if he is free to choose his fraud capability. In the final section, I discuss how the model fits the empirical evidence regarding fraudulent elections and argue that it can explain several puzzling observations such as a positive relationship between fraud and victory margins and a negative effect of fraud on turnout.

1.2 The Model

Participation in fraudulent elections is analyzed within a pivotal voting framework. Elections are modeled in a way similar to a large body of pivotal costly voting literature, where voters are assumed to make participation decisions based on the probability that their votes can alter the outcome of elections. Costly private value voting models of a similar type have been widely studied by, for example, Palfrey and Rosenthal (1983, 1985), Ledyard (1984), Borgers (2004), and more recently by Krasa and Polborn (2009) as well as Taylor and Yildirim (2010).

1.2.1 Setup

There are N voters ($N \geq 2$) and two candidates to vote for, the incumbent (A) and the challenger (B). Voters have preferences for candidates: B voters support the challenger (B-type) and $N-B$ voters favor the incumbent (A-type). Each voter has an individual specific voting cost c_i drawn from a commonly known distribution F over interval $(0, c_{max}]$ where $c_{max} \leq 1$, independently of his type and other voters. Distribution F is assumed to be continuous with positive density over $(0, c_{max})$ and differentiable cdf. F admits probability density function f which must have no more than one local maximum. If a voter's preferred candidate wins, the voter gains utility 1 if he did not vote, and $1 - c_i$ otherwise. If his favored candidate loses, the voter gains utility 0 if he abstained, and $-c_i$ if he voted. Every individual observes his own cost only. In this model, the supporters of the incumbent and those of the challenger differ ex-ante only in their preferences regarding candidates, while all their other characteristics such as benefits from electing a favored candidate and expected voting costs are the same.

Elections are run under majority rule and, without loss of generality, a tie is resolved in favor of the incumbent. Elections are fraudulent: the incumbent is able to commit fraud through ballot stuffing, meaning that if a voter abstains, his unused ballot may be counted in favor of the incumbent with certain exogenously given probability $\alpha \in [0, 1]$.

Note that if $\alpha = 0$, my model becomes almost identical to the models by Krasa and Polborn (2009) and by Borgers (2004), differing from them only in the assumption on known numbers of supporters for each candidate[4]. Krasa and Polborn (2009) as well as Borgers (2004), in contrast,

[4]Both Borgers (2004) and Krasa and Polborn (2009) assume that a tie is resolved with a toss of a coin, while

assume that the levels of supports of the candidates are ex-ante unknown, and every voter may be of either A-type or B-type with certain commonly known probability[5]. To be fully consistent with both models, I allow for ex-ante unknown levels of support in the generalized version of my model analyzed in Section 3, though it does not have any substantial effect either on the logic of the model or on the results. Thus, my model can be thought of as a generalization of the costly voting models by Krasa and Polborn (2009) as well as by Borgers (2004).

1.2.2 Analysis

I first analyze a simplified model, assuming $\alpha = 1$, wherein the incumbent stuffs the ballot box perfectly: if a voter abstains, his vote is counted in favor of the incumbent with certainty. Though this assumption might seem too strict, it allows for analytical characterization of the properties of equilibria, and understanding of the intuition behind the electoral game. In Section 3 I relax the assumption on perfect fraud and analyze the general model with arbitrary α.

The analysis of voters' behavior in elections with perfect fraud begins from the observation that, conditional on voting, a voter's weakly dominant strategy is to vote for his preferred candidate; thus the analysis focuses on participation decisions only.

Further note that none of the incumbent's supporters have incentives to vote as long as the costs of voting are non-negative. This is because an A-type voter's vote will be counted in favor of the incumbent regardless of whether the voter participates or abstains. Relaxing the assumption on non-negative costs will be discussed further.

Thus, I restrict my attention to the voting behavior of the challenger's supporters. First, note that a B-type voter i decides to vote if and only if his expected benefit exceeds his participation cost:

$$\Pi(p) > c_i. \tag{1.1}$$

$\Pi(p)$ is the voter's probability of being pivotal given that a randomly chosen B-type voter votes with probability p, and, at the same time, expected benefit because the voter's benefit from electing the challenger is 1.

Because the number of votes for the incumbent is at least as large as $N - B$ (number of A-type voters), there is no way the challenger can win elections if the number of his supporters is less than $N - B + 1$. Thus, if $B < N - B + 1$, B-type voters do not have incentives to vote either, and the unique equilibrium is full abstention. The interesting case to analyze is the situation when $B \geq (N + 1)/2$.

I assume that a tie is resolved in favor of the incumbent, but this difference is purely technical and does not crucially affect any result.

[5] In Borgers (2004) the probability that a voter is of a certain type is 0.5, while in Krasa and Polborn (2009) this probability is arbitrary. Thus, Borgers' model is a special case of the model by Krasa and Polborn.

To build the pivotal probability $\Pi(p)$ function, let V be a number of individuals other than i who choose to vote. Thus, V is a random variable that follows a binomial distribution with parameters $B-1$ and p. The probability that V takes a particular value v is then:

$$Prob(V=v) = \binom{B-1}{v} p^v (1-p)^{B-v-1}. \tag{1.2}$$

A B-type voter is pivotal when, in case of his abstention, the numbers of votes for the incumbent and challenger are equal, or when the challenger lacks one vote, which implies that the number of B-type participants must be $N/2$ if N is even, or $(N-1)/2$ if N is odd.

Without loss of generality, assume N is even, since all further calculations can be straightforwardly adjusted for the case when N is odd. Then, the voter is pivotal if and only if $V = N/2$. Note that $N/2+1 \leq B \leq N$ as the number of B-type voters should be larger than the number of A-types and smaller than the total number of voters N. From (1.2) we obtain a probability of being pivotal:

$$\Pi(p) = Prob(V=N/2) = \binom{B-1}{N/2} p^{N/2} (1-p)^{B-N/2-1}. \tag{1.3}$$

It can be seen that this function is non-negative, achieves maximum at $p = \frac{N/2}{B-1}$, equals zero when $p=0$ or $p=1$ whenever $B > N/2+1$. If $B = N/2+1$, then $\Pi(p)$ is strictly increasing in p.

Given pivotal probability function $\Pi(p)$, it is now possible to characterize equilibrium. I search for a within-group symmetric equilibrium where all B-type voters adopt the same voting strategy. Specifically, there must be a common threshold value c^* such that a B-type voter i votes if $c_i \leq c^*$ and abstains otherwise. Thus, c^* should satisfy:

$$\Pi(F(c^*)) = c^*. \tag{1.4}$$

For further analysis the condition can be rewritten as:

$$\Pi(F(c^*)) = F^{-1}(F(c^*)). \tag{1.5}$$

Note that $F(c^*)$ is the expected share of voters with voting costs below c^*, i.e. those who participate in elections. Thus, $F(c^*)$ is the expected turnout of B-type voters. For further analysis I will use short notation $F(c^*) = p^*$. Denoting $F(c) = p$, one can construct a graph in $(p, \Pi(p))$ space.

Figure 1.1: Equilibrium.

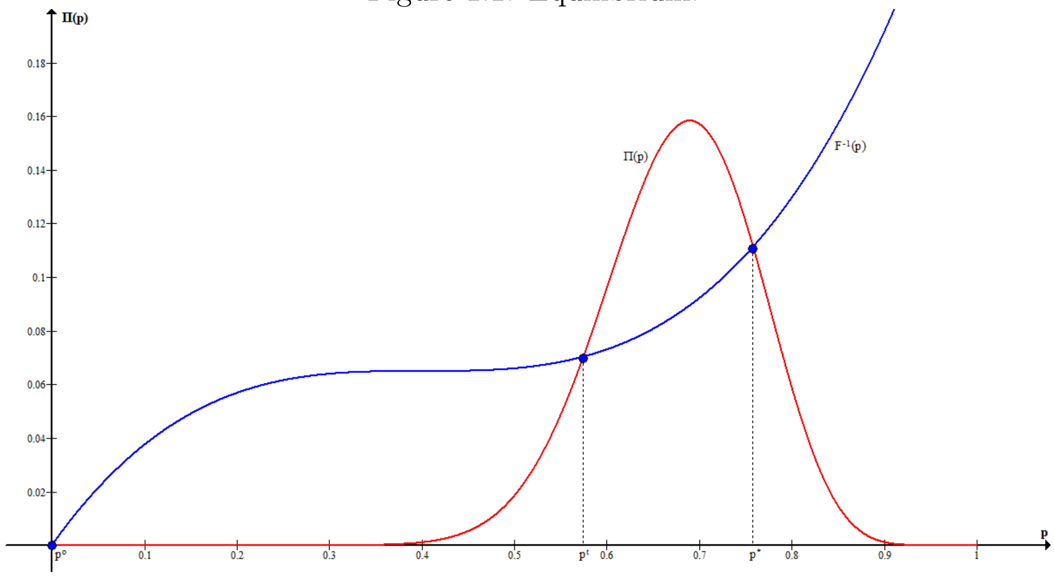

The unimodality of the cost distribution and the assumption on at most one inflection point in this cdf of the distribution together guarantee that there could be up to three points that satisfy equation (1.5). I denote the arguments of these intersections as p^0, p^t and p^*. Note that solution $p^0 = 0$ always exists, while existence of p^t and p^* (which may coincide under certain conditions) depends on the model's parameters. Equilibrium p^0 is an equilibrium with full abstention, while in equilibrium p^* a strictly positive share of B-type voters participates. Note that p^0 and p^* are stable equilibria, while p^t is not: once participation is lower than p^t the model will converge to equilibrium p^0, otherwise - to p^*. Thus, p^t constitutes a participation threshold value, which needs to be enforced in order to achieve stable coordination equilibrium p^*. From now on I focus on stable equilibria only, and show below that coordination equilibrium p^* is likely to be ex-ante more efficient than full abstention equilibrium p^0, and thus solving the collective action problem of achieving p^t turnout level is welfare improving.

As noted above, while equilibrium p^0 always exists, the existence of coordination equilibrium p^* depends on the model's parameters. Necessary and sufficient conditions for the existence of coordination equilibrium p^* can be formulated as follows:

$$\exists p \in (0,1] : \Pi(p) - F^{-1}(p) \geq 0. \tag{1.6}$$

In terms of exogenous parameters, the following condition is sufficient for the existence of coordination equilibrium:

$$\Pi(\frac{N/2}{B-1}) \geq F^{-1}(\frac{N/2}{B-1}). \tag{1.7}$$

Also note that condition (1.7) implies $p^t \leq \frac{N/2}{B-1} \leq p^*$. Conditions under which coordination equilibrium is more likely to exist are summarized in the following proposition.

Proposition 1.1. *Existence.*

1. If coordination equilibrium exists for some N_0, B_0 and cost distribution F_0, then it exists for any $N < N_0$ keeping the ratio of the candidates' support levels fixed;

2. If coordination equilibrium exists for some N_0, B_0 and cost distribution F_0, then it exists for any F which is first-order stochastically dominated by F_0;

3. For any N and B such that $N/2 < B \leq N$ there exists an infinite number of cost distributions F such that coordination equilibrium exists;

Proof: *See the Appendix.*

It is easy to understand the proposition using the graph in Figure 1.1. There are three exogenous components in the model: total number of voters N, distribution of preferences for candidates across voters $N - B$ and B, and cost distribution F. Keeping the ratio of the candidates' support levels fixed, lower N increases function $\Pi(p)$ for all p, since with a lower number of voters every individual is more likely to be pivotal. Thus, with a lower number of voters there is greater likelihood to have coordination equilibrium. The effect of a change in support, holding the population fixed, is unclear. An increase in B would imply that the pivotal probability function achieves its maximum at a lower participation level and that the maximum probability itself is lower. This is not sufficient to make any statement about the likelihood of coordination equilibrium existence without imposing further assumptions on cost distribution. Changes in cost distribution affect the F^{-1} graph only. Whenever the right part (for $p \geq 1/2$) of the inverse cost distribution function shifts down, coordination equilibrium is more likely to exist. From this logic and condition (1.7) it follows that keeping costs sufficiently low is enough to guarantee the existence of coordination equilibrium.

Given that coordination equilibrium exists, the important question is how its properties depend on population size, candidates' support, and voting costs. Coordination equilibrium is characterized by two values: voting rule c^* (which is matched one-to-one to participation level p^*) and participation threshold p^t, which must be enforced in order to guarantee convergence to coordination equilibrium.

Proposition 1.2. *Comparative Statics. If coordination equilibrium exists and sufficient condition (1.7) holds with inequality, then:*

1. Equilibrium participation p^* and threshold value p^t are decreasing in B;

2. p^* is decreasing and p^t is increasing in N, if the support ratio is fixed;

3. If some cost distributions G and F are such that G is first-order stochastically dominated by F, then p^* is higher and p^t is lower for G.

Proof: *See the Appendix.*

Again, Proposition 1.2 is easy to understand with Figure 1.1. An increase in B implies a shift of function Π to the left and down, moving both intersection points between Π and F^{-1} lower. An increase in N results in a shift of Π down, which implies higher p^t and lower p^*. If one cost distribution is first-order stochastically dominated by another, then the graph of the first inverse distribution is lower. Intuitively, a higher share of B-types implies that fewer participants are needed to defeat the incumbent. As a result, fewer individuals will participate and a lower number of participants is sufficient to induce coordination. With a higher number of voters, keeping the support ratio fixed, every voter is less likely to be pivotal, and thus fewer voters will participate in equilibrium. Finally, lower costs straightforwardly result in stronger participation incentives.

The model predicts the existence of two stable equilibria: full abstention equilibrium and coordination equilibrium. In full abstention equilibrium, nobody votes and the incumbent steals all the votes through ballot stuffing, winning with a 100% victory margin. Coordination equilibrium is characterized by a strictly positive participation rate[6] among the B-type voters. Moreover, the turnout is always expected to be higher than 50%. Technically, this result follows from condition (1.7). If the condition is satisfied, then the participation level of B-types must be higher than $\frac{N/2}{B-1}$, which, given that $N/2 + 1 \leq B \leq N$, implies that the expected turnout among the challenger's supporters is more than 50%. Intuitively, the case when fewer than half of B-type voters are supposed to participate cannot be an equilibrium, since this would mean a defeat of the challenger in expectation regardless of the number of A-type voters. Since participation in coordination equilibrium is positive, the number of stuffed ballots is at most $N - \frac{N}{2(B-1)}$, so coordination equilibrium results in ex-ante lower fraud than full abstention equilibrium. Finally, note that in both equilibria, the official reported turnout is always 100% since all the ballots of absentees are converted in favor of the incumbent.

Properties of the full abstention equilibrium, including zero turnout and 100% victory margin, may seem too extreme. A slight modification of the model would generate a less extreme result but keep the logic and all the established properties unchanged. Starting from Riker and Ordeshook (1968), the voting literature often argues that voters' participation in elections is driven not solely by their likelihood to be pivotal, but also by the utility they derive from voting, which can be thought of as a utility from fulfilling a civic duty. Once such a utility is introduced into my model, the results become less extreme. Technically, such a modification means that a

[6] Throughout the paper measures of participation and fraud are in ex-ante terms, since the actual number of participants and thus the actual number of stuffed ballots are random variables, whose exact realizations depend on the realization of voting costs.

voter now compares his voting costs with the expected benefit plus some utility from voting d:

$$\Pi(p) + d > c_i. \tag{1.8}$$

Re-arranging the terms, the last expression can be rewritten as $\Pi(p) > \bar{c}_i$, where $\bar{c}_i = c_i - d$ is a net voting cost. In terms of my model this change means that individual voting costs are now distributed over interval $[c_{min}, c_{max}]$, where $c_{min} < 0$. This change implies that all voters with $\bar{c}_i < 0$, both A-type and B-type, always vote.

Figure 1.2: Equilibrium with negative voting costs.

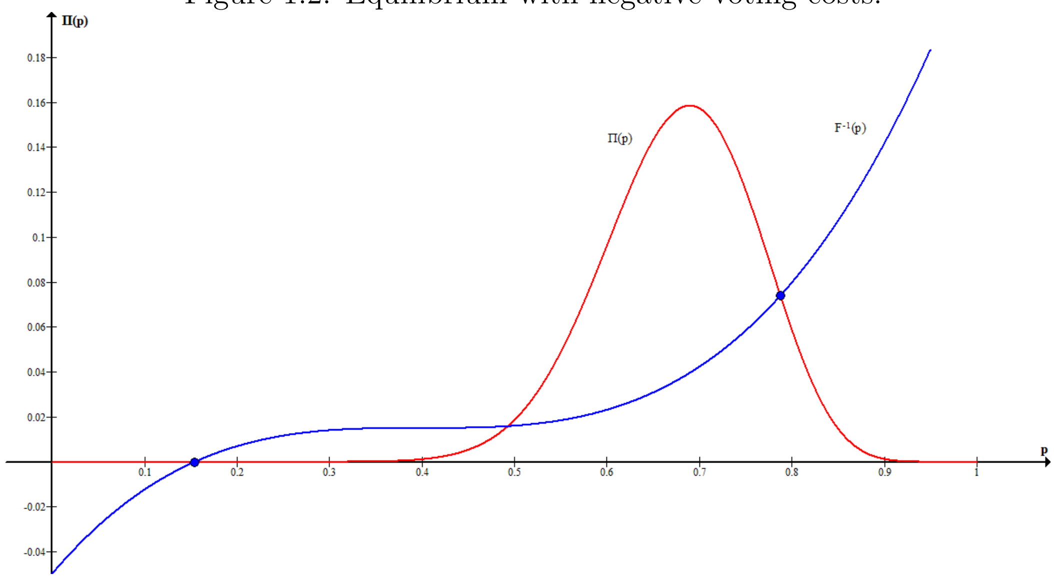

Now, instead of zero participation, the abstention equilibrium is characterized by a strictly positive share of B-type and A-type voters who participate (in expectation) (Figure 1.2). Note that if there is a sufficient number of vote lovers (those who have negative net voting costs) the first equilibrium may disappear: there are enough B-type voters with negative costs to create participation incentives for voters with positive costs and, thus, to induce coordination equilibrium. Since this is the only property of the model which is substantially affected by allowing voting costs to be negative, and all the propositions stated above and below generally stay valid, hereafter I analyze the benchmark model with non-negative costs to avoid unnecessary complications.

Another important characteristic of the equilibria is the probability of the incumbent's victory. Since in full abstention equilibrium, the incumbent stuffs 100% of ballots, he wins with certainty. His winning in coordination equilibrium depends on the actual turnout of B-type voters. Note that the turnout of B-type voters is a random variable. This is because individual voting costs are independent random draws, with unknown exact realization, and thus the exact number of individuals with costs below some particular threshold is also unknown ex-ante. Given some voting coordination equilibrium rule c^*, turnout would follow a binomial distribution with parameters B and $F(c^*)$. Thus, the probability that the incumbent defeats the challenger is the probability that no more than $N/2$ B-type voters cast their votes:

$$w = \sum_{i=0}^{N/2} \binom{B}{i} F(c^*)^i (1 - F(c^*))^{B-i}. \tag{1.9}$$

Since $B > N/2$ the winning probability is less than 1, and it is decreasing in participation of challenger's supporters. Moreover, since in expectation more than $N/2$ B-type voters participate, the probability that the incumbent will win cannot exceed 0.5. Finally, note that the equilibrium participation of the B-type voters is higher than $\frac{N/2}{B-1}$, and thus an upper bound for w can be obtained, which is obviously decreasing in the number of the challenger's supporters and increasing in the number of the incumbent's supporters:

$$w < \frac{1}{(B-1)^B} \sum_{i=0}^{N/2} \binom{B}{i} (N/2)^i (B - N/2 - 1)^{B-i}. \tag{1.10}$$

Given the winning probabilities of the candidates the welfare properties of the equilibria may be investigated.

1.2.3 Welfare

When voting is costly, participation implies a tradeoff between the quality of the aggregation of voters' preferences and participation costs. Higher participation decreases the probability of electing the wrong candidate (preferred by the minority), but at the same time implies higher total costs borne by society. The literature on participation offers different views on whether equilibrium participation in voluntary clean elections is efficient. Krishna and Morgan (2012) analyze a common value model, where voters share the same preferences for candidates but get different signals on the candidates' quality and thus may vote for different candidates; among other things, they show that when voting is costless then a voluntary voting equilibrium is fully efficient as individuals' and social objective functions are the same. Ghosal and Lockwood (2009) develop a model where voters' preferences combine both private and common values (i.e. voters prefer different candidates but also possess heterogeneous information on the common state of the world) and demonstrate that if voters vote according to their private preferences, equilibrium participation is inefficiently high whereas if they vote according to their private information, participation appears to be below the efficient level. In the model of Borgers (2004) as well as of Chakravarty, Kaplan and Myles (2010) a vote cast by a voter produces a negative externality on all other voters by decreasing their pivotal probabilities and thus expected benefits. Since this externality is not taken into account when a voter makes his participation decision, in equilibrium, participation is higher than the welfare-maximizing level. Krasa and Polborn (2009) show that if the levels of the candidates' support are ex-ante different, voting produces a positive externality: by casting a vote, a voter increases the probability that his candidate will win and, if the support levels are not equal, higher participation leads to a larger increase in the welfare of the majority than a decrease in the welfare of the minority. When this effect exceeds the extra participation costs, equilibrium participation is less than the efficient level. My model of fraudulent elections applies a similar logic to that in Krasa and Polborn (2009), though the exact mechanism through which a cast vote affects welfare differs.

First, there are two stable equilibria in this model. To see which of the two equilibria, coordination or full abstention, is more desirable from a social point of view, consider an ex-ante expected welfare evaluated at each of the equilibria. The expected utility of an A-type voter is $1 - w_B$, where w_B is the probability that the challenger will win. Recall that the turnout of B-type voters is a random variable. Given some voting rule $\tilde{c}$ turnout would follow a binomial distribution with parameters B and $F(\tilde{c})$. Thus, the probability that the challenger will defeat the incumbent is the probability that at least $N/2+1$ B-type voters will cast their votes:

$$w_B = \sum_{i=N/2+1}^{B} \binom{B}{i} F(\tilde{c})^i (1 - F(\tilde{c}))^{B-i}. \tag{1.11}$$

Then the expected utility of a B-type voter can be expressed as follows:

$$\int_0^{\tilde{c}} (v_B + \Pi(F(\tilde{c})) - c)\, \mathrm{d}F(c) + \int_{\tilde{c}}^{c_{max}} v_B\, \mathrm{d}F(c) = v_B + \int_0^{\tilde{c}} (\Pi(F(\tilde{c})) - c)\, \mathrm{d}F(c), \tag{1.12}$$

$$v_B = \sum_{i=N/2+1}^{B-1} \binom{B-1}{i} F(\tilde{c})^i (1 - F(\tilde{c}))^{B-i-1}. \tag{1.13}$$

Lemma 1.1. $w_B = v_B + \Pi(F(\tilde{c}))F(\tilde{c})$ *for all* $\tilde{c}$.
Proof: *See Appendix.*

Lemma 1.2. $\frac{\partial w_B}{\partial F(\tilde{c})} = B\Pi(F(\tilde{c}))$ *for all* $\tilde{c}$.
Proof: *See Appendix.*

The second integral in Formula (1.12) is the expected utility a B-type voter would gain if his cost is such that he abstains. Conditional on not casting a vote, the probability that the challenger will win is the probability that out of the other $B-1$ challenger supporters at least

$N/2+1$ participate, which is v_B. If the voter participates, which happens if his cost is below $\tilde{c}$, he incurs a cost but the probability that the challenger will win is now higher. Given that the voter participates, the challenger will win if out of the other $B-1$ voters at least $N/2$ voters participate. The probability of this event equals

$$\sum_{i=N/2}^{B-1} \binom{B-1}{i} F(\tilde{c})^i (1-F(\tilde{c}))^{B-i-1}. \tag{1.14}$$

The latter expression may be rewritten as:

$$\sum_{i=N/2+1}^{B-1} \binom{B-1}{i} F(\tilde{c})^i (1-F(\tilde{c}))^{B-i-1} + \binom{B-1}{N/2} F(\tilde{c})^{N/2} (1-F(\tilde{c}))^{B-N/2-1} = v_B + \Pi(F(\tilde{c})). \tag{1.15}$$

Thus, $\Pi(F(\tilde{c}))$ is the marginal contribution of a B-type voter to the ex-ante probability that the challenger wins, conditional on voting. Then, the voters' expected utility can be expressed as follows:

$$W = (N-B)(1-w_B) + Bv_B + B\int_0^{\tilde{c}} (\Pi(F(\tilde{c})) - c)\,\mathrm{d}F(c). \tag{1.16}$$

Having defined the welfare function, it is possible to compare the efficiency of full abstention equilibrium and coordination equilibrium. First, recall that in this model the voters' utility does not directly depend on the cleanness of the elections, and thus the fact that coordination equilibrium results in much lower fraud than does full abstention equilibrium, which is characterized by 100% ballot stuffing, is irrelevant for the welfare comparison. The only two features that affect voters' welfare are participation costs and the candidates' probabilities of winning. In coordination equilibrium, the probability of choosing candidate B, who is the candidate preferred by a majority, is high, but at the same time there are some participation costs. Clearly, the social welfare gain from the higher probability that the challenger wins is larger if there are more B-type voters in the population. At the same time, according to Proposition 1.2, with the higher share of B-types the expected number of those who choose to vote decreases, implying that total participation expenditures are lower. Thus, intuitively, coordination equilibrium should welfare dominate full abstention equilibrium, at least if the share of B-types is large enough. This result is formalized in the following proposition.

Proposition 1.3. *For any N there exists $B_0 > N/2$ such that for any $B_0 \le B \le N$ coordination equilibrium yields higher expected welfare than does full abstention equilibrium.*
Proof: *See the Appendix.*

Further, to see that coordination equilibrium is still not socially efficient, consider social welfare as a function of some strategy $\tilde{c}$ adopted by all B-type voters:

$$W = (N - B)(1 - w_B) + Bv_B + B \int_0^{\tilde{c}} (\Pi(F(\tilde{c})) - c)\, \mathrm{d}F(c). \tag{1.17}$$

$$W = (N - B)(1 - w_B) + Bv_B + B\Pi(F(\tilde{c}))F(\tilde{c}) - B \int_0^{\tilde{c}} c\, \mathrm{d}F(c). \tag{1.18}$$

Since according to Lemma 1.1 $w_B = v_B + \Pi(F(\tilde{c}))F(\tilde{c})$, the welfare function is then simply

$$W = (N - B)(1 - w_B) + Bw_B - B \int_0^{\tilde{c}} c\, \mathrm{d}F(c). \tag{1.19}$$

Taking the first-order condition with respect to $\tilde{c}$ we obtain:

$$(2B - N)\frac{\partial w_B}{\partial F(\tilde{c})} f(\tilde{c}) - B\tilde{c}f(\tilde{c}) = 0. \tag{1.20}$$

After re-arranging the terms the efficiency condition takes the following form:

$$\frac{2B - N}{B} \frac{\partial w_B}{\partial F(\tilde{c})} = F^{-1}(F(\tilde{c})). \tag{1.21}$$

Let c^o be the optimal voting rule, i.e. the one that satisfies condition (1.21). Recall the equilibrium condition:

$$\Pi(F(c^*)) = F^{-1}(F(c^*)). \tag{1.22}$$

According to Lemma 1.2 $\frac{\partial w_B}{\partial F(\tilde{c})} = B\Pi(F(\tilde{c}))$. Then, since $2B - N \geq 1$ as there are strictly more B-type voters than A-type voters, it must be the case that for all $\tilde{c}$

$$\frac{2B - N}{B} \frac{\partial w_B}{\partial F(\tilde{c})} \geq \Pi(F(\tilde{c})). \tag{1.23}$$

The inequality immediately implies that $c^* \leq c^o$ with equality only in the case when society consists of B-type voters only ($N = B$), and suggests that in coordination equilibrium, participation of B-type voters is below the efficient level. Note that the statements above are valid only when the majority of voters prefer the challenger. When there are more supporters of an incumbent, the only equilibrium is full abstention which is first best, assuming that fraud is costless: the majority candidate wins with certainty at zero participation cost.

1.2.4 Compulsory Voting

Given the established inefficiency, a natural step is to find a way to increase welfare. As in any externality problem, subsidization could be one way to correct for efficiency. Another way is to introduce compulsory voting. A number of studies compare voluntary and compulsory voting from efficiency perspectives. Borgers (2004) establishes that compulsory voting is never welfare improving. Krasa and Polborn (2009) demonstrate that Borgers' result is sensitive to the assumption of equal levels of supports for the candidates, which leads to elimination of the positive externality, and thus means that any increase in participation will decrease welfare. They allow the supports to be ex-ante different and show that under certain conditions compulsory voting may be superior to voluntary voting. Ghosal and Lockwood (2004) show that Borgers' result is also sensitive to the assumption of private values of voters' preferences: once voters' preferences have both private values and common values components, compulsory voting may Pareto dominate voluntary voting. Recently, Krishna and Morgan (2012) compared compulsory and voluntary voting purely in a common values setup and show that voluntary voting welfare dominates compulsory voting when elections are large, regardless of whether voting is costless or costly.

To check whether compulsory voting may improve welfare in elections with fraud, consider the welfare function under voluntary voting evaluated at the equilibrium:

$$W = (N - B)(1 - w_B) + Bw_B - B \int_0^{c^*} c \,\mathrm{d}F(c). \tag{1.24}$$

Now compare this to the welfare function under compulsory voting:

$$W_c = B - N \int_0^{c_{max}} c \,\mathrm{d}F(c). \tag{1.25}$$

Note that $(N - B)(1 - w_B) + Bw_B \leq B$, implying that compulsory voting may be superior to voluntary, though does not have to be since $N \int_0^{c_{max}} c \,\mathrm{d}F(c) \geq B \int_0^{c^*} c \,\mathrm{d}F(c)$. Recall that the participation rate of B-type voters in coordination equilibrium under voluntary voting is inefficiently low. Though compulsory voting results in inefficiently high participation, it still might deliver higher welfare than voluntary voting. The intuition behind this result is that compulsory voting causes the probability of electing the wrong candidate (i.e. the incumbent, who is preferred by the minority) to be zero, but at the same time requires voters to incur large participation costs. Whenever the benefit from the guarantee of choosing the majority candidate exceeds extra participation costs, compulsory voting is preferable. This is more likely to happen when there is a higher share of B-type voters. Clearly, the difference between welfare under compulsory voting and welfare under voluntary voting grows as B increases and approaches N:

$$W_c - W = (2B - N)(1 - w_B) - [N \int_0^{c_{max}} c\, \mathrm{d}F(c) - B \int_0^{c^*} c\, \mathrm{d}F(c)]. \tag{1.26}$$

As a result, compulsory voting may deliver higher welfare than voluntary voting, and the gain from compulsory voting is likely to be larger when there are more supporters of a challenger in the population of voters. This result is very intuitive, since introducing compulsory voting eliminates the possibility of choosing the alternative preferred by the minority, and the gain from the guarantee of electing the 'correct' candidate is higher when he is preferred by a higher number of voters. Indeed, one can always construct a cost distribution such that this gain is outweighed by the costs of voters who would abstain from voting in voluntary elections. If, for example, there are voters with very high costs, which means that cdf of the cost distribution rapidly increases in the tail (to infinity in the extreme case), compulsory voting never improves welfare. However, if the cost distribution is reasonable and not extreme, then compulsory voting might be desirable.

1.3 Generalized Model

1.3.1 Setup

Consider a generalized version of the model presented above. Now, instead of assuming that the vote of a non-participant is stolen with certainty, there is an exogenously given $\alpha \in [0, 1]$ that reflects the probability that the vote is stuffed in favor of the incumbent. Thus, α can be thought of as the incumbent's fraud capability. In addition, in contrast to the perfect fraud model where the levels of the candidates' support were known, in the generalized model only the total number of voters is known, while the exact support levels are uncertain. Instead, there is commonly known probability $\beta \in [0, 1]$ that a voter supports candidate B. The assumption on uncertain support levels does not have a substantial effect on the results derived further, but makes the model fully consistent with the literature on costly voting in clean elections. Once $\alpha = 0$ and $\beta = 0.5$, the generalized model converges to the model of clean elections with ex-ante equal support for candidates analyzed by Borgers (2004). When $\alpha = 0$ and β is arbitrary, the model is very close to the model by Krasa and Polborn (2009)[7]. When $\alpha = 1$ and the numbers of the incumbent's and challenger's supporters are fixed, one has the model of perfect fraud analyzed above. Finally, costs are distributed over the interval $[c_{min}, c_{max}]$ where $c_{min} \geq 0$.

[7] Borgers (2004) and Krasa and Polborn (2009) assume that a tie is resolved with the toss of a coin, while the presented model assumes that the tie is resolved in favor of the incumbent.

1.3.2 Analysis

Consider a B-type voter. Suppose that all other B-type voters adopt voting strategy c_B, i.e. a B-type voter votes if his voting costs are below c_B and abstains otherwise. Similarly, suppose A-type voters adopt strategy c_A. Then the probability that a randomly picked voter votes is $F(c_B)$ and $F(c_A)$ for B-types and A-types respectively.

The probability that there are a incumbent supporters among other $N-1$ voters is

$$P_a^{N-1} = \binom{N-1}{a}(1-\beta)^a\beta^{N-a-1}. \tag{1.27}$$

The probability that k of them participate in elections is

$$P_k^a = \binom{a}{k}F(c_A)^k(1-F(c_A))^{a-k}. \tag{1.28}$$

The probability that m out of another $N-a-1$ B-type voters participate is

$$P_m^{N-a-1} = \binom{N-a-1}{m}F(c_B)^m(1-F(c_B))^{N-a-m-1}. \tag{1.29}$$

A B-type voter is pivotal in two cases. First, if the number of stolen votes is such that the number of votes for each candidate is equal, and second, if the challenger leads by one vote and the voter's ballot is stolen if he abstains. If x votes are stolen, the incumbent gets $x+k$ and the challenger gets m votes. Thus, given a, k and m, a B-type voter is pivotal if and only if $x = m-k$ or $x = m-k-1$. The probability of this event is

$$P_{m-k}^{N-k-m-1} + \alpha P_{m-k-1}^{N-k-m-1}, \tag{1.30}$$

where $P_{m-k}^{N-k-m-1} = \binom{N-k-m-1}{m-k}\alpha^{m-k}(1-\alpha)^{N-2m-1}$.

The probability that a out of $N-1$ voters support the incumbent, k out of these a A-supporters participate, m out of $N-a-1$ challenger supporters participate, and $m-k$ or $m-k+1$ votes out of $N-k-m-1$ non-participants' votes are stolen is then:

$$P_B(a,k,m,c_A,c_B) = P_a^{N-1}P_k^aP_m^{N-a-1}(P_{m-k}^{N-k-m-1} + \alpha P_{m-k-1}^{N-k-m-1}). \tag{1.31}$$

Finally, the probability that a B-type voter is pivotal is a function of voting strategies c_A and c_B adopted by all the A-type and all the B-type voters respectively:

$$\Pi_B(c_A,c_B) = \sum_{a=1}^{N-1}\sum_{k=0}^{a}\sum_{m=k-1}^{B-1} P_B(a,k,m,c_A,c_B). \tag{1.32}$$

Similarly, one can construct a pivotal probability function for an A-type voter:

$$\Pi_A(c_A, c_B) = (1-\alpha)\sum_{a=1}^{N-1}\sum_{k=0}^{a-1}\sum_{m=k-1}^{B} P_A(a,k,m,c_A,c_B). \tag{1.33}$$

where $P_A(a,k,m,c_A,c_B) = P_a^{N-1}P_k^a P_m^{N-a-1}P_{m-k-1}^{N-k-m-1}$.

Note, that there is a $(1-\alpha)$ term in the pivotal probability function for incumbent supporters: an A-type voter will be pivotal only if his vote is not stolen in case of abstention; otherwise his participation decision will not change the outcome.

Symmetric equilibrium is characterized by a pair (c_A, c_B) such that all A-type voters with costs below c_A and all B-type voters with costs below c_B participate, and the others abstain. Equilibrium values of c_A and c_B are the solution for the following system of equations:

$$\begin{aligned} \Pi_A(c_A, c_B) &= c_A \\ \Pi_B(c_A, c_B) &= c_B \end{aligned} \tag{1.34}$$

If one defines a function $L : [c_{min}, c_{max}]^2 \to [c_{min}, c_{max}]^2$ as:

$$L(c_A, c_B) = (max\{min\{\Pi_A(c_A,c_B), c_{max}\}, c_{min}\}, max\{min\{\Pi_B(c_A,c_B), c_{max}\}, c_{min}\}), \tag{1.35}$$

then Brouwer's fixed point theorem will imply the existence of equilibrium.

As in the case of the perfect fraud model, equilibrium voting rules (c_A, c_B) are one-to-one matched to the pair $(F(c_A), F(c_B))$ - equilibrium participation of the incumbent's and the challenger's supporters respectively. Figure 1.3 demonstrates an equilibrium for the following values of parameters: $N = 25, \beta = 0.7, \alpha = 0.3$, costs are distributed uniformly over the interval $[0.01, 0.1]$. The red (darkest) surface is the net expected benefit of a B-type voter as a function of $F(c_A), F(c_B) : \Pi_B(F(c_A), F(c_B)) - F^{-1}(F(c_B))$. Similarly, the blue surface is the net expected benefit of an A-type voter: $\Pi_A(F(c_A), F(c_B)) - F^{-1}(F(c_A))$. The yellow (lightest) surface is a zero-plane. Thus, a point where all three surfaces intersect would be an equilibrium.

Figure 1.3: Equilibrium in the generalized model.

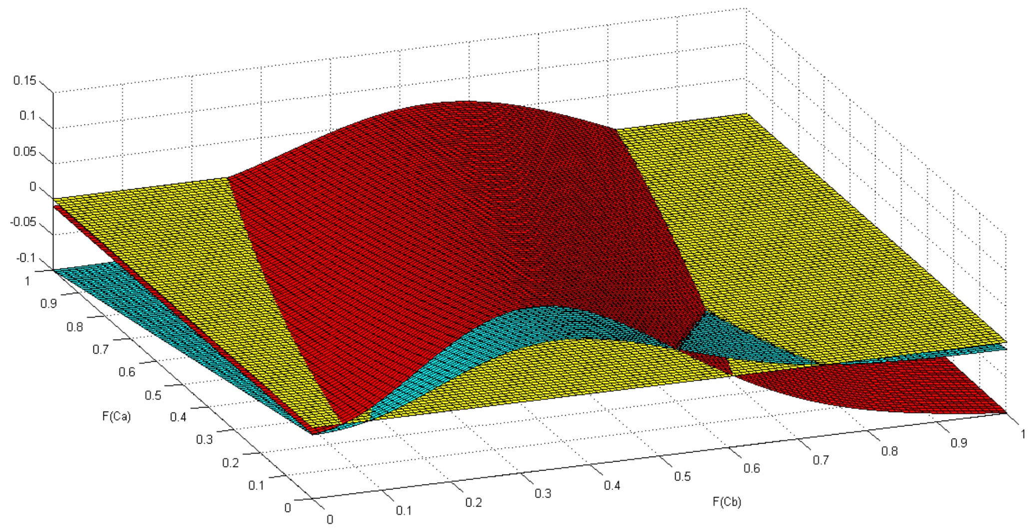

There is one such point on the figure above: $(0.27, 0.63)$. However, there are two more equilibria here. First $(0, 0)$ is an equilibrium (stable) because for both types of voters net expected benefits are negative at this point, and thus no voter has any incentive to vote. Second, $(0, 0.05)$ is an equilibrium (unstable) since the net expected benefit of A-type voters is negative, so they do not vote, while B-types' net expected benefit is exactly 0 at this point.

Obtaining a closed-form solution and even characterizing equilibria for the generalized model is quite challenging. However, numerical simulations provide a number of consistent observations about the equilibria and their properties. As in the case of perfect fraud model, I focus on stable coordination equilibrium. The first question is how changes in fraud capability α affect equilibrium. Consider a simple numerical example with just two voters both supporting challenger with probability 1.

Example 1.1. $N = 2$, $\beta = 1$, costs are distributed uniformly over $[c_{min}, c_{max}]$, where $0 < c_{min} < 1$.

A voter is pivotal if either another voter participates and the vote of the first voter is stolen in case of abstention, or another voter abstains and his vote is not stolen. Thus, equilibrium is given by the following equation:

$$\alpha F(c_B) + (1 - \alpha)(1 - F(c_B)) = (c_{max} - c_{min})F(c_B) + c_{min}.$$

The solution is then $F(c_B^*) = \frac{1-\alpha+c_{min}}{c_{max}-c_{min}-2\alpha+1}$, which is increasing in α.

Note that if the second voter always abstains, the expected benefit of the first voter is $(1 - \alpha)$. Whenever this benefit is less than the minimal possible cost of voting c_{min}, i.e. when

$\alpha > 1 - c_{min}$, the first voter also abstains, and thus $F(c_B^*) = 0$ constitutes another equilibrium.$\square$

In the example above, fraud capability α leads to an increase in the equilibrium participation of B-type voters, conditional on coordination. Simulations of the model for a higher number of voters provide similar results: an increase in α decreases the participation incentives of the incumbent's supporters and increases the participation incentives of the challenger's supporters. To understand how equilibrium participation rates change in response to an increase in α consider Figure 1.4, which displays two-dimensional representation of the equilibrium in the generalized model with the following values of the parameters: $N = 25, \beta = 0.7$, costs are distributed uniformly over the interval $[0.01, 0.1]$. In Figure 1.4 the blue (dark) curve is the solution $c_A(c_B)$ of the first equation of system (1.34) and a red (light) curve is the solution of the second equation. The graph presents the solutions for three distinct values of α: 0.3, 0.5 and 0.7. Equilibrium is determined by the intersection of red and blue curves of the same type which correspond to the same value of α (the blue curve with the highest maximum corresponds to $\alpha = 0.3$, the blue curve with the lowest maximum correspond to $\alpha = 0.7$).

Figure 1.4: Equilibrium for different values of α.

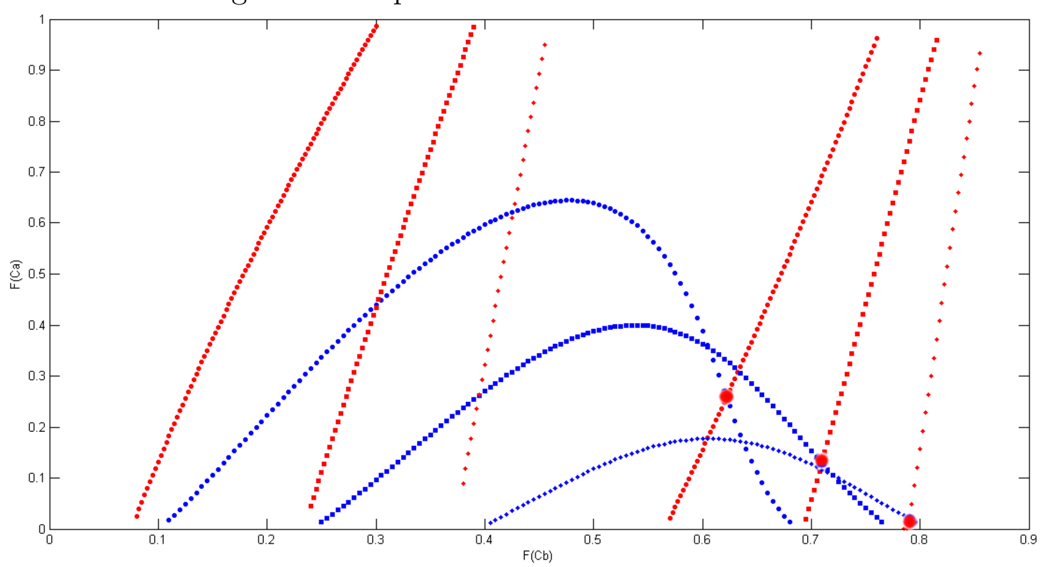

This result is robust to different choices of the model's parameters. Figure 1.5 shows how changes in α affect participation of the incumbent's supporters (blue curve) and the challenger's supporters (red curve) for $N = 25, \beta = 0.7$, costs distributed uniformly over the interval $[0.01, 0.1]$.

Figure 1.5: Coordination equilibrium participation of A-supporters (downward sloping) and B-supporters (upward sloping) as a function of fraud capability

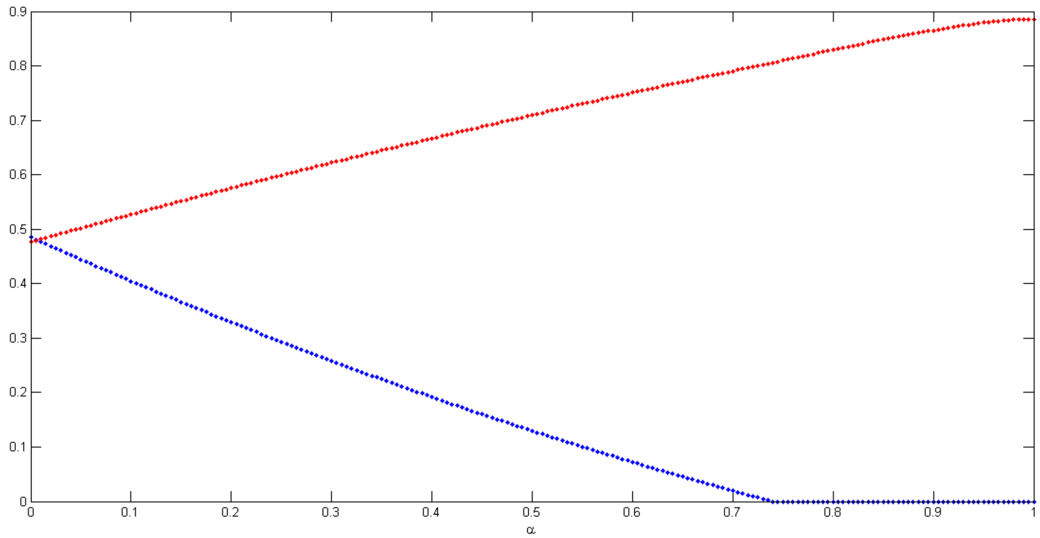

Figure 1.5 shows that the participation incentives of A-type voters decrease with higher α. The intuition is straightforward: the higher the incumbent's fraud capability, the more likely the incumbent will steal an A-type's vote if the voter abstains and avoids cost, the less incentive to participate the voter has. The effect of an increase in α on coordination equilibrium participation of the B-types is the opposite. Higher fraud capability implies that higher participation among B-types is needed to maintain sufficiently high pivotal probabilities. Note that though higher α leads to higher participation of the challenger's supporters in coordination equilibrium, at the same time it requires more people to coordinate in order to achieve this equilibrium.

Further note that when elections are perfectly fraudulent, there are two stable equilibria: full abstention equilibrium and coordination equilibrium. In clean elections, as shown by Borgers (2004) and Krasa and Polborn (2009), there is unique equilibrium. Thus, it must be the case that sufficiently large fraud capability leads to the emergence of a bad full abstention equilibrium. This observation is formalized in Proposition 1.4.

Proposition 1.4. *For given values of N and β, and cost distribution F, there is a unique value α_0 such that for any $\alpha \geq \alpha_0$ full abstention is an equilibrium, and for any $\alpha < \alpha_0$ it is not.*
Proof: *See Appendix.*

The proposition is very intuitive: when fraud capability is low, full abstention cannot be an equilibrium because a single voter has a good chance of influencing the outcome of elections by deviating and participating. But when fraud capability is high, implying that there is a high probability that the vote of a non-participant will be stolen, participating when all the others abstain is unlikely to be profitable as there is a high probability that a sufficient number of votes will be stuffed in favor of the incumbent, and thus deviation from abstention will not change the outcome of elections.

1.3.3 Endogenous Fraud

Suppose now that α is no longer exogenous, but rather the incumbent is free to choose it. Anticipating the response of the voters to any level of fraud capability α, a rational incumbent would choose an α that maximizes his probability of winning the election. Assume for simplicity that fraud is costless. It might seem at first glance that in this case the incumbent should choose maximum possible fraud capability, which is $\alpha = 1$. However an increase in fraud capability has several effects. First, higher α, other things being equal, implies a higher number of stuffed ballots and thus a higher probability that the incumbent wins. But changes in α also affect the participation of voters. Specifically, as stated above, a higher α has a deterrent effect on the participation of the incumbent's supporters and a stimulating effect on participation of the challenger's supporters (conditional on the fact that coordination equilibrium is achieved), which together lead to a decrease in the probability that the incumbent will win. Thus, the resulting effect of an increase in α on the probability that the incumbent wins depends on which of the two effects dominates. To illustrate this intuition, consider an example.

Example 1.2. Suppose there are just three voters, and assume for simplicity that supports are known: one voter is A-type and the other two voters are B-type; costs are distributed uniformly over $[0, 0.5]$. In clean elections, the A-type voter is pivotal if and only if one out of two B-types participates. A B-type voter may be pivotal in two cases, either if another B-type and A-type both abstain or both participate. Thus, pivotal probabilities as functions of voting strategies c_A and c_B are:

$$\Pi_A = 2F(c_B)(1 - F(c_B)),$$
$$\Pi_B = (1 - F(c_A))(1 - F(c_B)) + F(c_A)F(c_B).$$

Then equilibrium is a solution of the following system:

$$4c_B(1 - 2c_B) = c_A,$$
$$(1 - 2c_A)(1 - 2c_B) + 4c_A c_B = c_B.$$

This system has a unique solution, which is approximately $({c_A}^*, {c_B}^*) = (0.21, 0.44)$. The incumbent may win clean elections in two cases: when all three voters abstain, or when the A-type voter participates and at least one B-type voter abstains. Given equilibrium participation strategies $(c_A{}^*, c_B{}^*)$, the probability that the incumbent wins is

$$w_A = (1 - F(c_A{}^*))(1 - F(c_B{}^*))^2 + F(c_A{}^*)(1 - F(c_B{}^*)^2) = 0.103.$$

If elections are perfectly fraudulent, the A-type always abstains, and a B-type is pivotal if and only if another B-type voter participates. Thus, equilibrium is determined by $F(c_B) = c_B$. Clearly, one equilibrium is $c_B{}^* = 0$ which corresponds to full abstention. However, $c_B{}^* = 0.5$, which corresponds to full participation among B-types, is also an equilibrium, as in this case expected benefit $F(0.5) = 1$ always exceeds costs. In full abstention equilibrium the incumbent always wins, while in coordination equilibrium he inevitably loses. Thus, if the incumbent can choose whether to hold clean elections or stuff the ballot box perfectly, clean elections would be preferable if the challenger's supporters coordinate well, even if fraud is costless.$\square$

Figure 1.6 shows the incumbent's winning probability in coordination equilibrium as a function of α for $N = 25, \beta = 0.7$, and voting costs distributed uniformly over the interval $[0.01, 0.1]$. From Figure 1.6 it is clear that that maximum feasible fraud level is not necessarily optimal.

Figure 1.6: The incumbent's winning probability as a function of fraud capability.

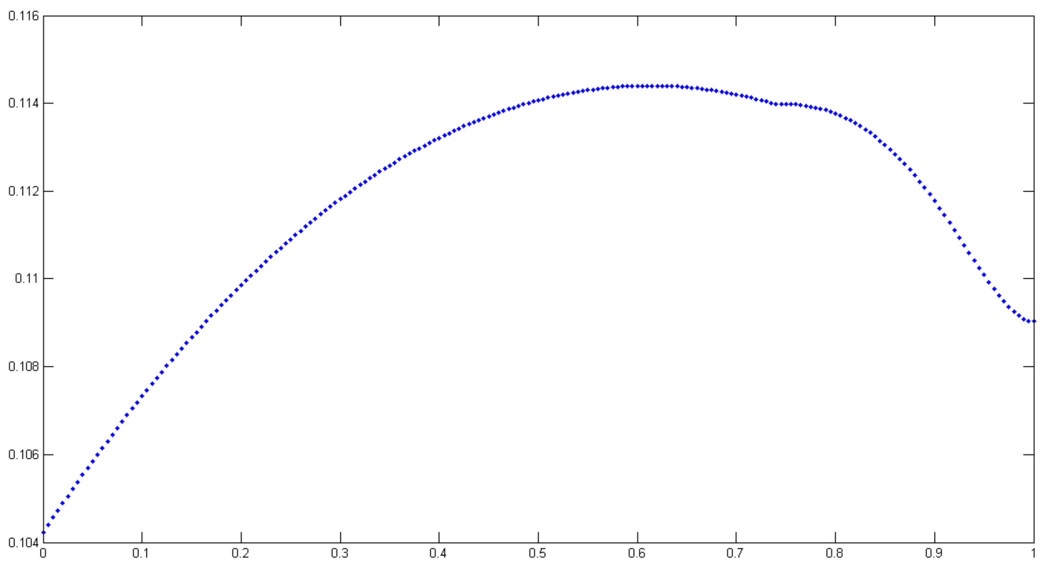

Further, recall Proposition 1.4 which states that a threshold value exists, such that for all α above this value equilibrium with full abstention exists, and for all α below this value it does not exist. From the incumbent's point of view, full abstention equilibrium is the first best as it guarantees the incumbent's victory with certainty. For the parameters used to construct Figure 1.6 the threshold equals $\alpha_0 = 0.22$.

Further, the winning probability achieves its maximum for some α^*. Note that a higher value of α implies that the participation rate of the challenger's supporters in coordination equilibrium is higher, which in turn means that solving the collective action problem is more difficult and thus achieving coordination equilibrium is less likely. Ultimately, when choosing fraud level, the incumbent faces a triple tradeoff: costs of fraud, likelihood of coordination among the challenger's supporters, and winning probability given that coordination is achieved. If the incumbent is strong enough to deter coordination of the opposition's supporters, he would probably choose a low level of fraud close to α_0, while a weak incumbent would prefer a high level fraud close to α^*, which makes coordination harder and provides relatively high chances of winning even if this coordination is achieved.

When fraud capacity α is considered as an endogenous variable which is subject to the incumbent's choice, then given the timing of the model, one may argue that there is a commitment problem if α is easily adjustable. Since, first, the incumbent sets α, then voters make their participation decisions, and only then fraud is realized, the incumbent would always benefit by increasing α from the announced level once voters have made their decisions. If this is the case, then the discussion on optimal α should be thought of as comparative statics with respect to fraud capability rather than the rational incumbent's choice. However, in reality, targeted fraud level is unlikely to be such an easily adjustable variable as it is the result of a comprehensive rigging process which takes place not only on election day, but begins also long before. In this case, though the incumbent does not commit to maintaining the announced level of α, his ability to adjust it at the last moment is very limited.

1.4 Discussion and Empirical Evidence

The presented model generates a number of predictions about voters' and candidates' behavior as well as outcomes in fraudulent elections. It would indeed be valuable to empirically test the validity of these predictions with real electoral data. The problem is that the main results involve the incumbent's fraud capability: its optimal choice, influence on the participation incentives of different groups of voters, and effect on the winning probabilities of the candidates, while the extant, though relatively limited, empirical literature on fraudulent elections primarily studies ex-post realized fraud rather than the ex-ante potential of an incumbent to steal votes. The obvious reason for this is that even ex-post fraud is not an easily measurable variable, so fraud capability seems considerably more challenging to measure or even proxy. Though the main results of the model cannot be easily tested due to this problem, there are still several empirical observations about fraudulent elections that may be used to verify the consistency of the model. Specifically, the model explains two main empirical observations about fraudulent elections: the negative effect of fraud on turnout, and the positive relationship between fraud and victory margin.

A number of survey-based and empirical studies have shown that voters are less likely to participate in elections if they expect fraud. This result was established by McCann and Domingues (1998), who utilized Mexican survey data and found opposition supporters to be more likely to abstain when expectations of fraud are high. The finding was later confirmed by Hiskey and Bowler (2005) who also employed Mexican data to study the impact of procedural fairness on citizens' political engagement. Among other results, they find that individuals who believe that elections are clean are more likely to participate. More recently, Birsch (2010) empirically analyzed cross-country electoral survey data from both new and established democracies and shows that ex-ante fairness of elections is positively related to turnout. She finds that, controlling for a variety of individual- and election-level characteristics, perceived electoral integrity has a strong positive effect on the propensity to vote. Similarly, Landry, Davis and Wang (2010) study local elections in China and conclude that when the race is close, voters perceive elections as fair and are more likely to vote. In contrast to the survey-based research, Simpser (2012) explores Mexican electoral data to assess the relationship between voters' participation incentives and fraud. Using variation in fraud and turnout across Mexican states, and explicitly distinguishing between reported and true turnout, he finds that electoral manipulations discourage voter participation.

The negative relationship between fraud perception and turnout is generally explained by the low incentives of the electorate to participate in costly voting if elections lack competitiveness. Low incentives are assumed to be the result of either direct disutility from participating in corrupt elections (e.g. Simpser, 2008) or from low likelihood of a vote to be pivotal, which in turn comes from a lack of competition (e.g. Birsch, 2010). Though the idea that in fraudulent elections a single voter is less likely to be pivotal, which decreases participation incentives, has been pointed out in the literature, the mechanism behind the relationship between fraud and pivotal probabilities has not been explored in any depth. Indeed, this relationship is not monotonic: fraud can both decrease and increase the competitiveness of elections. It may give a corrupt candidate an overwhelming advantage, but it can also create a chance for an incumbent with low support to make the race competitive. Hence, the effect of fraud on pivotal probabilities is not entirely straightforward and thus the mechanism through which fraud affects participation incentives requires a consistent explanation. This paper provides such an explanation and sheds light on the nature of the relationship between fraud and turnout. In simple terms, the explanation is that the observed correlation is not the result of a direct negative causal effect of fraud on participation, but rather an equilibrium outcome of an electoral game.

The model of elections with ballot stuffing predicts the existence of two stable equilibria. In full abstention equilibrium, none of the voters participate and the incumbent stuffs a lot of ballots. The expectation is that the number of stuffed ballots equals αN if the incumbent's fraud capability is α. In coordination equilibrium, positive shares of both A-type and B-type voters participate, and thus there is less space for ballot stuffing than in full abstention equilibrium. These two equilibria generate the correlation between fraud and turnout observed in real data: higher fraud goes hand-in-hand with lower participation.

This result might seem trivial since the ex-post amount of ballot stuffing is simply a linear function of turnout: the number of stuffed ballots always equals the difference between the total number of voters and the number of participants, multiplied by α. This is true, but the key point here is that the number of participants is not a decreasing function of the number of stuffed ballots or vice versa: the negative correlation between fraud and turnout is generated not by a mechanical linear relationship between turnout and the amount of stuffed ballots, but rather by the negatively related values of turnout and fraud which constitute equilibria. The key point is that the incumbent needs to maintain low participation to guarantee he will win. The only way to ensure this is to commit extensive fraud: if voters fail to coordinate, then their participation incentives are low, given extensive fraud commitment. But extensive fraud is possible only when a lot of voters (all of them in the model) abstain, since ballot stuffing uses residual turnout to transform unused ballots into votes. This situation corresponds to the full abstention equilibrium. Alternatively, if voters coordinate and vote, then there is not much opportunity for ballot stuffing, and the extent of fraud is relatively moderate, which in turn provides sufficient incentives for voters to coordinate. This corresponds to coordination equilibrium. Together, full abstention and coordination equilibria generate a negative fraud-turnout relationship, implying that low (high) turnout is not just a consequence of high (low) fraud, but rather low (high) turnout and high (low) fraud are simultaneously determined equilibrium outcomes.

Similar logic lies behind the explanation for the second empirical observation about fraudulent elections, which relates integrity, victory margin, and fraud excessiveness. Simpser (2008) collects and analyzes a dataset of about 400 executive elections held worldwide between 1990 and 2007, in which multiple candidates were allowed to run, to establish a positive relationship between electoral fraud and victory margin. For each election he uses a variety of available sources ranging from observers' reports and newspaper articles to academic research and poll data to mark the election as either clean or corrupt. The main finding of the analysis is that in fraudulent elections, a high victory margin is observed far more frequently than in clean elections: in about 40% of elections marked as corrupt the victory margin exceeds 40%. This result generally holds for the analysis of alternative datasets on fraudulent elections such as the Database of Political

Institutions by the World Bank (Beck, Clarke, Groff, Keefer, Walsh, 2001) as well as the one collected by Pastor (1999).

From this observation it follows that corrupt politicians often commit excessive electoral fraud. In part, this could arise from the incumbent's uncertainty about the election outcome. If the candidate is risk averse and the costs of fraud are low relative to the stakes of re-election, such uncertainty could provide incentives for excessive fraud. Simpser (2008) explains fraud excessiveness using a two-period voting game where a high victory margin in the first period discourages the participation of opposition supporters in the second period, creating incentives for the incumbent to excessively commit fraud. My model of elections with ballot stuffing suggests an alternative and simpler explanation using a logic similar to that of fraud-turnout correlation. Full abstention equilibrium is characterized by extensive fraud and a 100% victory margin, which is far beyond the level needed to guarantee the incumbent's victory. In contrast, coordination equilibrium implies that elections are relatively clean ex-post and the winning candidate has a reasonable victory margin. Again, the fact that relatively clean elections correspond to a reasonable victory margin, while fraudulent elections are associated with an extremely large margin, comes not from the causal effect of fraud on victory margin, but arises as an equilibrium outcome.

1.5 Conclusion

In this paper I explore the mechanism through which electoral fraud affects the decisions of voters to participate in elections and, thus, social welfare. I analyze a pivotal voter model of elections with costly participation, where the incumbent can stuff the ballot box and voters decide whether to participate in elections or abstain based on a comparison of their subjective probability that their vote will be pivotal with individual-specific participation costs. I show that when a majority of voters support the challenger, two stable equilibria may exist: full abstention equilibrium, where the incumbent wins with certainty and which exists only if the incumbent's capability to stuff a ballot box is sufficiently strong, and a more efficient coordination equilibrium, where a substantial share of a challenger's supporters vote and the probability the incumbent will be defeated is large. I find that participation in the coordination equilibrium is still inefficiently low, since voters do not take into account the positive externality they produce on other voters when participating. Each vote cast by a challenger's supporter increases the probability of the incumbent's defeat, and if the incumbent is supported by a minority of voters, this has a positive effect on the overall welfare of voters. Thus subsidization, and in some cases even the introduction of compulsory voting, may improve efficiency. I then show that higher capability of the incumbent to stuff a ballot box discourages the participation of his own supporters and creates coordination incentives for a challenger's supporters. Hence, fraud does not always benefit the incumbent even when it is costless. Additionally, the model simultaneously explains two empirical observations about fraudulent elections: the positive relationship between fraud and victory margin, and the negative effect of fraud on turnout.

Appendix

Proof of Proposition 1.1. 1. Suppose coordination equilibrium exists for some parameter set $(N_0, B_0, F_0(c))$, and one decreases population size: $N_1 = N_0 - t$. If the support ratio is fixed, then $\frac{B_0}{N_0 - B_0} = \frac{B_1}{N_1 - B_1}$, implying $B_1 = B_0 \frac{N_1}{N_0} = B_0 - \frac{B_0 t}{N_0}$. Note that both $t/2$ and $\frac{B_0 t}{N_0}$ must be integers. Denote:

$$\Pi_0(p) = \binom{B_0 - 1}{N_0/2} p^{N_0/2}(1-p)^{B_0 - N_0/2 - 1},$$

$$\Pi_1(p) = \binom{B_0 - \frac{B_0 t}{N_0} - 1}{N_0/2 - t/2} p^{N_0/2 - t/2}(1-p)^{B_0 - \frac{B_0 t}{N_0} - N_0/2 + t/2 - 1}.$$

After some algebra it can be shown that $\Pi_0(p) \leq \Pi_1(p)$ for all p, which immediately implies the result: if coordination equilibrium exist for parameters $(N_0, B_0, F_0(c))$ then $\exists \bar{p} \in (0, 1]$ such that ${F_0}^{-1}(\bar{p}) \leq \Pi_0(\bar{p}) \leq \Pi_1(\bar{p})$ and thus, equilibrium exists for parameters $(N_1, B_1, F_0(c))$.

2. Suppose coordination equilibrium exists for some cost distribution F_0. Thus, $\exists \bar{p} \in (0, 1]$ such that ${F_0}^{-1}(\bar{p}) \leq \Pi(\bar{p})$. Any F which is first order stochastically dominated by F_0 satisfies $F(c) \geq F_0(c)$ for all c. Hence, $F^{-1}(p) \leq {F_0}^{-1}(p)$ for all $p \in [0, 1]$, including $\bar{p}$. From $F^{-1}(\bar{p}) \leq {F_0}^{-1}(\bar{p}) \leq \Pi(\bar{p})$ existence of equilibrium for cost distribution F follows immediately.

3. Coordination equilibrium exists whenever $\Pi(\frac{N/2}{B-1}) \geq F^{-1}(\frac{N/2}{B-1})$. Because $\Pi(p)$ is continuous, and for any B and N such that $N/2 + 1 \leq B \leq N$ its maximum $\Pi(\frac{N/2}{B-1}) > 0$, there always exists some x such that $\Pi(\frac{N/2}{B-1}) > x > 0$. Let $\bar{x} = sup\{x | 0 \leq x < \Pi(\frac{N/2}{B-1})\}$. Then, any F such that $F^{-1}(\frac{N/2}{B-1}) = \bar{x}$ satisfies $\Pi(\frac{N/2}{B-1}) \geq F^{-1}(\frac{N/2}{B-1})$, which guarantees existence of coordination equilibrium. $\square$

Proof of Proposition 1.2. 1. Let p^* be an equilibrium participation level and p^t an enforcement threshold under some N, B and F, while $\tilde{p}$ is an equilibrium participation level and $\tilde{p}^t$ is an enforcement threshold under N, $B + 1$ and F. Then, p^* is an argument of an intersection point between increasing function $F^{-1}(p)$ and decreasing part of $\Pi(p, B)$, i.e. for $p > \frac{N/2}{B-1}$. Likewise, $\tilde{p}$ is an argument of an intersection point between the same $F^{-1}(p)$ and decreasing part of $\Pi(p, B+1)$, i.e. for $p > \frac{N/2}{B}$. Thus, to prove that $p^* > \tilde{p}$ it is sufficient to show that $\Pi(p, B) > \Pi(p, B+1)$ for $p > \frac{N/2}{B-1}$. It is easy to see that $\Pi(p, B+1) - \Pi(p, B)$ is negative if $p > \frac{N/2}{B}$ and thus for any $p > \frac{N/2}{B-1}$:

$$\Pi(p,B+1)-\Pi(p,B)=\binom{B}{N/2}p^{N/2}(1-p)^{B-N/2}-\binom{B-1}{N/2}p^{N/2}(1-p)^{B-N/2-1},$$

$$\Pi(p,B+1)-\Pi(p,B)=\binom{B-1}{N/2}p^{N/2}(1-p)^{B-N/2-1}[\frac{B}{B-N/2}(1-p)-1].$$

Similarly, p^t is an argument of an intersection point between non-decreasing function $F^{-1}(p)$ and increasing part of $\Pi(p,B)$, i.e. for $p<\frac{N/2}{B-1}$. Likewise, $\tilde{p}^t$ is an argument of an intersection point between the same $F^{-1}(p)$ and increasing part of $\Pi(p,B+1)$, i.e. for $p<\frac{N/2}{B}$. Since $\Pi(p,B)<\Pi(p,B+1)$ for $p<\frac{N/2}{B}$, it follows that $p^t>\tilde{p}^t$.

2. Let population size decrease from some N_0 to $N_1<N_0$ keeping the support ratio $\frac{N_0-B_0}{B_0}$ fixed. Let p^* and p^t be equilibrium participation level and threshold under N_0, while $\tilde{p}$ and $\tilde{p}^t$ are equilibrium participation level and threshold under N_1. Let $N_1=N_0-t$. To keep the support ratio fixed B_1 should be equal to $B_0-\frac{B_0 t}{N_0}$, and $t/2$ as well as $\frac{B_0 t}{N_0}$ must be an integer. Denote

$$\Pi_0(p)=\binom{B_0-1}{N_0/2}p^{N_0/2}(1-p)^{B_0-N_0/2-1},$$

$$\Pi_1(p)=\binom{B_0-\frac{B_0 t}{N_0}-1}{N_0/2-t/2}p^{N_0/2-t/2}(1-p)^{B_0-\frac{B_0 t}{N_0}-N_0/2+t/2-1}.$$

As stated in the proof of Proposition 1, it can be shown that $\Pi(p,B_0)\le\Pi(p,B_1)$ for all p, where $B_1=B_0-\frac{B_0 t}{N_0}$. Because p^* and $\tilde{p}$ are intersections of increasing function $F^{-1}(p)$ with $\Pi_0(p)$ and $\Pi_1(p)$ respectively in their decreasing parts, $p^*<\tilde{p}$. Similarly, since p^t and $\tilde{p}^t$ are intersections of $F^{-1}(p)$ with $\Pi_0(p)$ and $\Pi_1(p)$ respectively in their increasing parts, $p^t>\tilde{p}^t$. Thus, equilibrium participation is decreasing and participation threshold is increasing in population size.

3. Suppose coordination equilibrium exists for some values N, B and cost distribution F: equilibrium participation level is p^* and participation threshold is p^t. Then, according to Proposition 1 coordination equilibrium exists for the same N, B and any cost distribution G which is first-order stochastically dominated by F. Denote participation level under this equilibrium as $\tilde{p}$ and participation threshold as $\tilde{p}^t$. Let us show that $p^*<\tilde{p}$ and $p^t>\tilde{p}^t$.

Assume by contrast that $p^*>\tilde{p}$, and recall that, to be an equilibrium, both p^* and $\tilde{p}$ must be greater than or equal to $\frac{N/2}{B-1}$. Then $F^{-1}(p^*)>G^{-1}(\tilde{p})$. Because p^* satisfies $\Pi(p^*)=F^{-1}(p^*)$

and $\tilde{p}$ satisfies $\Pi(\tilde{p}) = G^{-1}(\tilde{p})$ it must be that $\Pi(p^*) > \Pi(\tilde{p})$. But Π is a decreasing function for any $p \geq \frac{N/2}{B-1}$, implying that $p^* < \tilde{p}$, which contradicts the initial assumption. Thus, $p^* < \tilde{p}$.

Since p^t and $\tilde{p}^t$ are intersections of the increasing part of $\Pi(p)$ with $F^{-1}(p)$ and $G^{-1}(p)$ respectively, and $F^{-1}(p) \geq G^{-1}(p)$ for all p, it immediately follows that $p^t > \tilde{p}^t$. $\square$

Proof of Lemma 1.1. Denote $F(\tilde{c}) = p$ for shorter notation. Then

$$w_B = \sum_{i=N/2+1}^{B} \binom{B}{i} p^i (1-p)^{B-i},$$

$$v_B = \sum_{i=N/2+1}^{B-1} \binom{B-1}{i} p^i (1-p)^{B-i-1}.$$

Consider the first element of w_B and, since $\binom{B}{j} = \binom{B-1}{j-1} + \binom{B-1}{j}$ for all $j < B$, rewrite it as:

$$w_B^1 = \binom{B}{N/2+1} p^{N/2+1}(1-p)^{B-N/2-1} = \binom{B-1}{N/2} p^{N/2+1}(1-p)^{B-N/2-1} + \binom{B-1}{N/2+1} p^{N/2+1}(1-p)^{B-N/2-1}.$$

$$w_B^1 = p\Pi(p) + (1-p)v_B^1.$$

The second element of w_B can be expressed as

$$w_B^2 = \binom{B}{N/2+2} p^{N/2+2}(1-p)^{B-N/2-2} = \binom{B-1}{N/2+1} p^{N/2+2}(1-p)^{B-N/2-2} +$$

$$+ \binom{B-1}{N/2+2} p^{N/2+2}(1-p)^{B-N/2-2}.$$

Which is equivalent to

$$w_B^2 = pv_B^1 + (1-p)v_B^2.$$

Similarly, any j^{th} element of w_B except the first and the last ones ($2 \leq j \leq B - N/2 - 1$) can be expressed as:

$$w_B^j = pv_B^{j-1} + (1-p)v_B^j.$$

The last element of w_B equals

$$w_B^{B-N/2} = p^B = pp^{B-1} = pv_B^{B-N/2-1},$$

where $v_B^{B-N/2-1}$ is the last element of v_B. Summing all the elements of w_B:

$$w_B = \sum_{j=1}^{B-N/2} w_B^j = \sum_{k=1}^{B-N/2-1} v_B^k + p\Pi(p) = v_B + p\Pi(p).\square$$

Proof of Lemma 1.2. Recall the following identity[8]:

$$1 - I_x(a,b) = (1-x)^{a+b-1} \sum_{i=0}^{a-1} \binom{a+b-1}{i} (\frac{x}{1-x})^i,$$

where $I_x(a,b)$ is regularized incomplete beta-function. Denote $F(\tilde{c}) = p$ for shorter notation. Then

$$w_B = \sum_{i=N/2+1}^{B} \binom{B}{i} p^i (1-p)^{B-i},$$

Consider the following regularized incomplete beta-function: $I_p(N/2+1, B-N/2)$. Using the identity above:

$$1 - I_p(N/2+1, B-N/2) = (1-p)^B \sum_{i=0}^{N/2} \binom{B}{i} (\frac{p}{1-p})^i = \sum_{i=0}^{N/2} \binom{B}{i} p^i (1-p)^{B-i} = 1 - w_B,$$

Hence, $w_B = I_p(N/2+1, B-N/2)$.

Also recall Chebyshev's integral:

$$\int x^a (1-x)^b \,\mathrm{d}x = B_x(a+1, b+1),$$

where $B_x(a+1,b+1)$ is incomplete beta-function.

Thus, $\int \Pi(p)\,\mathrm{d}p$ can be expressed as $\binom{B-1}{N/2} B_p(N/2+1, B-N/2)$. By definition of regularized incomplete beta-function:

$$I_p(N/2+1, B-N/2) = \frac{B_p(N/2+1, B-N/2)}{\mathrm{B}(N/2+1, B-N/2)},$$

where $\mathrm{B}(N/2+1, B-N/2)$ is beta-function.

[8]See, for example, Pearson, K., 1968. Tables of Incomplete Beta-Function, Second Edition, Cambridge University Press, page 28.

Since $\mathrm{B}(N/2+1, B-N/2) = \frac{(N/2)!(B-N/2-1)!}{B!}$:

$$B\int \Pi(p)\,\mathrm{d}p = B\binom{B-1}{N/2}B_p(N/2+1, B-N/2) = \frac{B!}{(N/2)!(B-N/2-1)!}B_p(N/2+1, B-N/2) =$$

$$= \frac{B_p(N/2+1, B-N/2)}{\mathrm{B}(N/2+1, B-N/2)} = I_p(N/2+1, B-N/2) = w_B$$

Hence, $B\int \Pi(p)\,\mathrm{d}p = w_B$. To complete the proof it is sufficient to differentiate both parts of the last identity with respect to p. $\square$

Proof of Proposition 1.3. Welfare as a function of some strategy $\tilde{c}$ is expressed as

$$W = (N-B)(1-w_B) + Bv_B + B\int_0^{\tilde{c}} (\Pi(F(\tilde{c})) - c)\,\mathrm{d}F(c).$$

In full abstention equilibrium, where $\tilde{c} = 0$ and both $w_B = 0$ and $v_B = 0$, social welfare is then simply $N - B$. Consider the difference between welfare in coordination equilibrium and welfare in full abstention equilibrium:

$$\Delta W = Bv_B - (N-B)w_B + B\int_0^{c^*} (\Pi(F(c^*)) - c)\,\mathrm{d}F(c).$$

According to Lemma 1.1 $w_B = v_B + \Pi(F(\tilde{c}))F(\tilde{c})$. Then ΔW is simply

$$\Delta W = (2B-N)w_B - B\int_0^{c^*} c\,\mathrm{d}F(c).$$

Since $\int_0^{c^*} c\,\mathrm{d}F(c) < \Pi(F(c^*))$ and $F(c^*) \geq \frac{N/2}{B-1}$, one might obtain a lower bound for ΔW:

$$(2B-N)w_B \geq (2B-N)\sum_{i=N/2+1}^{B}\binom{B}{i}\left(\frac{N/2}{B-1}\right)^i\left(\frac{B-N/2-1}{B-1}\right)^{B-i},$$

$$B\int_0^{c^*} c\,\mathrm{d}F(c) < B\Pi(F(c^*)) \leq B\binom{B-1}{N/2}\left(\frac{N/2}{B-1}\right)^{N/2}\left(\frac{B-N/2-1}{B-1}\right)^{B-N/2-1},$$

$$\Delta W > (2B-N)\sum_{i=N/2+1}^{B}\binom{B}{i}\left(\frac{N/2}{B-1}\right)^i\left(\frac{B-N/2-1}{B-1}\right)^{B-i} -$$

$$-B\binom{B-1}{N/2}\left(\frac{N/2}{B-1}\right)^{N/2}\left(\frac{B-N/2-1}{B-1}\right)^{B-N/2-1}.$$

Note that the lower bound for ΔW is a function of two integers B and N such that $N/2+1 \leq B \leq N$. It can be shown that the lower bound is increasing in B for a fixed value of N. To verify this it is sufficient to go over all possible values of N, and check monotonicity for each value of it and for all $N/2+1 \leq B \leq N$. For this paper monotonicity is checked for all even integers $N \in [2, 1000000]$. Further, the lower bound for ΔW is positive for $B = N$. To see that, it is sufficient to take the first two elements from the sum in the expression for the lower bound evaluated at $B = N$, and observe that for any $N \geq 2$:

$$\frac{N}{(N-1)^N}\left(\binom{N}{N/2+1}(N/2)^{N/2+1}(N/2-1)^{N/2-1} + \binom{N}{N/2+2}(N/2)^{N/2+2}(N/2-1)^{N/2-2}\right) -$$

$$-\frac{N}{(N-1)^{N-1}}\binom{N-1}{N/2}(N/2)^{N/2}(N/2-1)^{N/2-1} =$$

$$\frac{N}{(N-1)^N}\binom{N-1}{N/2}(N/2)^{N/2}(N/2-1)^{N/2-1}\cdot$$

$$\cdot\left(\frac{N}{N/2+1}(N/2) + \frac{N}{(N/2+1)(N/2+2)}(N/2)^2 - (N-1)\right) =$$

$$\frac{N^2}{(N-1)^N}\binom{N-1}{N/2}(N/2)^{N/2}(N/2-1)^{N/2-1}\cdot$$

$$\cdot\left(\frac{N/2}{N/2+1} + \frac{(N/2)^2}{(N/2+1)(N/2+2)} - 1 + 1/N)\right) =$$

$$\frac{N^2}{(N-1)^N}\binom{N-1}{N/2}(N/2)^{N/2}(N/2-1)^{N/2-1}\left(\frac{N}{N/2+2} - 1 + 1/N)\right) > 0.$$

Since the lower bound for ΔW is increasing in B and positive for $B = N$, for any N there exists $B_0 > N/2$ such that for any $B_0 \leq B \leq N$ $\Delta W > 0$. $\square$

Proof of Proposition 1.4. If all the voters abstain, an A-type voter is never pivotal, while a B-type voter is pivotal if and only if none of the non-participants' votes is stolen. Thus, the expected benefit function of a B-type voter at point (c_{min}, c_{min}), which corresponds to full abstention, is $\Pi_B(c_{min}, c_{min}) = (1-\alpha)^{N-1}$, which is strictly decreasing in α. Since at $\alpha = 0$ pivotal function $\Pi_B(c_{min}, c_{min}) = 1 > c_{min}$, and at $\alpha = 1$ $\Pi_B(c_{min}, c_{min}) = 0 < c_{min}$, there exists a unique value of $\alpha = \alpha_0$ such that $\Pi_B(c_{min}, c_{min}) = c_{min}$. For any $\alpha \geq \alpha_0$ $\Pi_B(c_{min}, c_{min}) < c_{min}$, implying that deviation from abstention is never profitable for a B-type voter, and for any $\alpha < \alpha_0$

$\Pi_B(c_{min}, c_{min}) > c_{min}$, implying that deviation is profitable, and thus, full abstention is not an equilibrium. $\square$

References

Beck, T., Clarke, G., Groff, A., Keefer P., Walsh, P., 2001. New Tools and New Tests in Comparative Political Economy: the Database of Political Institutions. *World Bank Economic Review*, 15, 165-176.

Borgers, T., 2004. Costly Voting. *American Economic Review*, 94 (1), 57-66.

Birch, S., 2010. Perceptions of Electoral Fairness and Voter Turnout. *Comparative Political Studies*, 43(12), 1601-1622.

Chaturvedi, A., 2005. Rigging Elections with Violence. *Public Choice*, 125, 189-202.

Chakravarty, S., Kaplan, T., Myles, G., 2010. The Benefits of Costly Voting. *Working Paper*. University of Exeter Business School.

Collier, P., Vicente, P., 2012. Violence, Bribery, and Fraud: The Political Economy of Elections in Sub-Saharan Africa, *Public Choice*, 153 (1-2), 117-147.

Ghosal, S., Lockwood, B., 2009. Costly Voting when Both Information and Preferences Differ: Is Turnout Too High or Too Low? *Social Choice and Welfare*, 33 (1), 25-50.

Hiskey, J., Bowler, S., 2005. Local Context and Democratization in Mexico. *American Journal of Political Science*, 49 (1), 57-71.

Krasa, S., Polborn, M., 2009. Is Mandatory Voting Better than Voluntary Voting? *Games and Economic Behavior*, 66, 275-291.

Krishna, V., Morgan, J., 2012. Voluntary Voting: Costs and Benefits. *Journal of Economic Theory*, forthcoming.

Landry, P., Davis, D., Wang, S., 2010. Elections in Rural China: Competition without Parties. *Comparative Political Studies*, 15, 763-790.

Lehoucq, F., 2003. Electoral Fraud: Causes, Types and Consequences. *Annual Review of Political Science*, 6, 233-256.

Ledyard, J., 1984. The Pure Theory of Large Two-Candidate Elections. *Public Choice*, 44, 7-41.

Magaloni, B., 2010. The Game of Electoral Fraud and the Ousting of Authoritarian Rule. *American Journal of Political Science*, 54 (3), 751-765.

McCann, J., Dominguez, J.,1998. Mexicans React to Political Fraud and Corruption: An Assessment of Public Opinion and Voting Behavior. *Electoral Studies*, 17, 483-503.

Palfrey, T., Rosenthal, H., 1983. A Strategic Calculus of Voting. *Public Choice*, 41(1), 7-53.

Palfrey, T., Rosenthal, H., 1985. Voter Participation and Strategic Uncertainty. *American Political Science Review*, 79 (1), 62-78.

Pastor, R., 1999. The Role of Electoral Administration in Democratic Transitions. *Democratization*, 6, 1-27.

Riker, W., Ordeshook, P., 1968. A Theory of the Calculus of Voting. *American Political Science Review*, 62, 28-42.

Simpser, A., 2008. Cheating Big: On the Logic of Electoral Corruption in Developing Countries. *Working Paper*, University of Chicago.

Simpser, A., 2008. The Measurement of Electoral Corruption: Some Issues and a Cross-National Dataset. *Working Paper*, University of Chicago.

Simpser, A., 2012. Does Electoral Manipulation Discourage Voter Turnout? Evidence from Mexico. *Journal of Politics*, 74 (3), 782-795.

Taylor, C., Yildirim, H., 2010. A Unified Analysis of Rational Voting with Private Values and Group-Specific Costs. *Games and Economic Behavior*, 70 (2), 457-471.

Chapter 2. Growth of Electoral Fraud in Non-Democracies: The Role of Uncertainty

Abstract

Electoral fraud has become an integral part of electoral competition both in established democracies and less-than-democratic regimes. In this paper I study electoral fraud in the non-democratic setting. First, I present evidence of fraud sustainability and growth over the lifetime of non-democratic regimes in post-Soviet and Sub-Saharan countries. Second, I provide a formal model that rationalizes the observed tendency and explains how uncertainty can lead to growing fraud. Specifically, in a probabilistic voting model of electoral competition with falsifications, a corrupt incumbent faces two types of uncertainty: uncertainty about voters' attitude towards fraud and uncertainty about his true support, captured by a purely random component in the voters' utility over candidates. The model predicts that when uncertainty is sufficiently large, higher uncertainty about voters' fraud intolerance provides weaker incentives to commit fraud. Over time the incumbent becomes more certain about voters' reaction to fraud due to learning through Bayesian updating and, thus, as the deterrent role of uncertainty about fraud intolerance declines, the incentives to commit fraud become stronger, providing a growing fraud profile.

JEL Classification: D72, D73, D80
Keywords: Voting, Fraud, Learning

2.1 Introduction

Fair elections are fundamental to democracy. Over the last decades researchers mainly assumed that elections are well-functioning tools for converting public preferences into social choice. However, in reality, cases of manipulating electoral outcomes are quite widespread even in established democracies[9]. In less-than-democratic regimes, strategies to shape electoral results through political pressure and especially electoral fraud[10] are an integral part of political competition. International organizations exert tremendous effort to ensure transparency in elections. However, electoral fraud in non-democracies[11] seems to be not only persistent, but expanding.

All the main findings on electoral fraud are derived from the analysis of particular elections in a given country at a given moment (Lehoucq, 2003), while the issues of fraud dynamics receive limited attention in academic literature. This is surprising, because studying the evolution of fraud seems extremely important from many different perspectives. Comprehensively studying political regimes, designing effective electoral legislation and, especially, assessing the effectiveness of electoral monitoring are much harder to do without an understanding of the dynamics of fraud. To assess the effect of electoral reforms on the integrity of elections or to study the role of international monitoring in improving electoral transparency, one has to understand how the electoral environment changes over time and the sources of those changes. For such purposes, studying fraud dynamics is crucial. This paper is intended to partially fill this gap by studying the role of uncertainty in fraud dynamics.

The contribution of the paper is twofold. Firstly, I discuss evidence suggesting a tendency towards increasing electoral fraud in non-democratic Post-Soviet and Sub-Saharan countries. Secondly, I present a game-theoretical probabilistic voting model with fraud which rationalizes the observed tendency, suggesting evolution of uncertainty about voters' attitude towards fraud as a potential explanation for growing fraud. In particular, a model of electoral competition with falsifications explicitly distinguishes between two types of uncertainty that affect electoral outcomes: aggregate uncertainty about a candidates' true levels of support, captured by a purely random component in voters' utility over candidates, and the incumbent's uncertainty about the degree of fraud voters will tolerate?, represented by his subjective beliefs about the attitude of the voters towards fraud. These two uncertainties prevent the incumbent's learning about voters' fraud tolerance immediately after his first election, providing him with a noisy signal about voters' true attitude towards fraud, which he uses for Bayesian updating of his beliefs.

[9]See the Literature Review section for references.

[10]Following Lehoucq (2003), I define fraud as any illegal act committed with the intent to shape an electoral result.

[11]Hereafter, by non-democracy I mean a country which has been widely criticized for deviating from the principles of democracy although it has formal democratic institutions such as elections. Specifically, I use this term for the countries marked as Not Free or Partially Free in the Freedom House Index of Democracy.

The model predicts that when there is significant initial uncertainty on the incumbent's side, uncertainty about voters' attitude towards fraud negatively affects incentives to commit fraud. Over time this uncertainty decreases due to learning and, thus, the incumbent gains greater incentives to commit fraud.

One can doubt the reliability of uncertainty to explain increasing electoral fraud, suggesting a number of obvious reasons for the observed tendency, such as growing stakes of re-election and decreasing costs of fraud due to learning by doing (see for instance Simpser, 2008). However, costs and stakes determine the level of committed fraud only if there is uncertainty about the outcome of the elections. In this paper I focus on the pure effect of uncertainty and show that it can also provide incentives for increasing fraud. Furthermore, in contrast to conventional wisdom (e.g., Simpser, 2008), I demonstrate that uncertainty does not always increase incentives to commit fraud, and the direction of the effect depends on the nature of the uncertainty.

The paper is organized as follows. In the next section I review the existing literature on electoral manipulations and, particularly, electoral fraud. I then discuss problems of measuring electoral fraud and provide some evidence from Post-Soviet and Sub-Saharan countries, suggesting that electoral fraud has been growing over time. Further, I develop a formal game-theoretic dynamic model of elections with falsifications and show how uncertainty could lead to increasing fraud.

2.2 Literature Review

Classic theories of electoral competition suggest that candidates can influence their chances of being elected only by choosing their policies. Nevertheless, in reality, elections are often associated with a variety of not-always-legal activities that result in redistribution of votes in favor of one or another candidate. A wide stream of both theoretical and empirical literature focuses on different strategies that incumbents can use to influence electoral outcomes. For instance, Glaeser and Schleifer (2005) show how an incumbent can engage in redistributive politics in order to shape the electorate in his favor - the so-called Curley effect. Further, a number of studies analyze political budget cycles when incumbents increase public expenditures or change their composition towards more visible goods in pre-election periods in order to attract votes. This has been empirically documented by Akhmedov and Zhuravskaya (2004), Shi and Svensson (2006) and Guo (2009), for developing countries and, for instance, by Veiga and Veiga (2007) as well as Schneider (2010) for developed ones. Political budget cycles are also widely studied from a theoretical standpoint starting from Rogoff (1990) and continued by, for example, Martinez (2009), who explicitly models how politicians can change their policies to improve their reputations and, thus, increase their chances of re-election.

Electoral fraud could be considered another type of outcome-influencing strategies, widespread under autocracies and dictatorships where fraud tools are more easily available than in pure democracies. Chaturverdi (2005) and, more recently, Collier and Vicente (2010) study pre-election violence as an instrument for shaping electoral results through deterring opposition supporters from voting. To determine when fraud occurs and what should be done to prevent it, Sutter (2003) presents a simple model of rigged elections where the society decides how closely to monitor the elections, suggesting the provision of neutral observers, strengthening the protest option when fraud is detected, and reducing costs of monitoring by, for example, subsidization, as effective fraud prevention tools. However, though all the papers have focused on shaping electoral results through electoral fraud, their purposes, methods of modeling, and underlying assumptions do not allow them to be used as a framework for studying fraud dynamics.

Simpser (2005) makes an exceptional attempt to study electoral fraud in a dynamic setting. He points out that rigged elections are often associated with very high victory margins, implying that incumbents often irrationally commit excessive fraud. Simpser attempts to rationalize such behavior by formalizing the idea that excessive fraud can, first, deter coordination of future opposition and turnout and, second, directly influence the beliefs of opposition supporters that elections will be corrupt and thus discourage their turnout. In a later paper, Simpser (2008) elaborates on this idea and comes up with a model that generates equilibrium with persistent but not growing excessive fraud. Also, the author briefly discusses the potential role of exogenous uncertainty, costs and stakes in his explanation of excessive fraud, but concludes that these features cannot sufficiently explain the high victory margins observed in the data. However, the equilibrium outcome and the latter conclusion are based on strict underlying assumptions, particularly, the assumption that opposition supporters, in the case of the incumbent's victory, gain more utility when they abstain from voting than when they do vote (i.e. only opposition supporters, but not the incumbent's supporters, are discouraged from participating in elections if the incumbent is very likely to win, which is a disputable assumption).

To summarize, the existing literature on electoral fraud cannot satisfactorily explain the observed patterns in the behavior of corrupt incumbents in non-democracies that I have documented in the next section. In particular, questions of sustainability and growth of electoral fraud over the lifetimes of non-democratic regimes as well as reasons and conditions for the occurrence of fraud still call for explanations. This paper presents a model that aims to partially fill the gap, theoretically studying how uncertainty can affect incentives to commit fraud in non-democratic setting and how it can contribute to explaining the increasing fraud profile.

2.3 Dynamics of Electoral Fraud in Non-Democracies

For the analysis of electoral fraud dynamics, two sets of countries with non-democratic regimes are used. When I refer to a regime I mean a period in a country's history when there was either a single leader or several leaders from the same party or family whose ruling methods are considered to be less than democratic. Specifically, I focus on regimes that existed during the last 25 years, between 1988 and 2013, that have formal democratic institutions like direct presidential elections, that were rated as Not Free or Partially Free in the Freedom House Democracy Index (FHDI) for at least 2 years within their lifetimes, and that have survived for at least two terms.

The first country set consists of Post-Soviet countries excluding Lithuania, Latvia and Estonia - democracies that entered the EU in 2003 - and Moldova, where there has been no clear regime since independence and, moreover, no direct presidential elections between 1996 and 2012. Turkmenistan was excluded from the set as a country without direct presidential elections in its whole history. In the end, the first set of countries includes Armenia, Azerbaijan, Belarus, Georgia, Kazakhstan, Kyrgyzstan, Russia, Tajikistan, Ukraine and Uzbekistan. Two of these countries, Georgia and Russia, are characterized by two regimes each. However, the first Russian regime (Yeltsin 1992-2000) is not considered in the further analysis since it is characterized by only one presidential election (1996). Finally, the data set includes 11 non-democratic post-Soviet regimes in 10 countries (See Table 3 of the Appendix for the detailed list).

The second set consists of Sub-Saharan African countries that adopted elections in the early 1990s. From the total of forty nine countries in the region, eight were excluded as fully free and democratic countries (i.e., rated as Free in FHDI[12] in 1991-2013 with the exception of, at most, 1 year: Benin, Botswana, Cape Verde, Mali, Mauritius, Namibia, Sao Tome and Principe and South Africa). An additional two were excluded due to too few multiparty elections (less than two or with a gap of more than 10 years) since 1990 (Burundi, Eritrea). Five more were excluded that did not elect presidents by direct vote of the population (Angola, Ethiopia, Lesotho, Somalia and Swaziland), and three more countries were excluded that have not had a clear regime that has survived for at least two terms between 1990 and 2013 (Comoros, Guinea-Bissau and South Sudan). Finally, the set contains 31 Sub-Saharan countries, among which three (Central African Republic, Kenya and Senegal) are represented by two distinct regimes. This leaves 34 non-democratic African regimes (See Table 4 of the Appendix for the entire list).

To credibly analyze the evolution of electoral fraud, one needs to have some objective measure of fraud. Literature (see, for example. Lehoucq, 2003) suggests several types of sources that can provide valuable information for building such a measure: press, opposition parties' archives with official acquisitions on fraud, complaints submitted to courts, scientific surveys and in-

[12] Freedom House Democracy Index, www.freedomhouse.org/report-types/freedom-world (retrieved 20.03.2013)

terviews with voters, and results of international electoral monitoring. However, all of these sources, except probably the last one, could be biased towards one or another candidate, and, thus, cannot provide fully objective information about fraud. In addition to the fact that comprehensive analysis of such data and collecting relevant information require tremendous effort, the partisan nature of the sources limits their usefulness in measuring fraud. Results of monitoring by electoral observation missions are also limited in their ability to provide useful information on the dynamics of electoral misconduct: public reports mainly contain qualitative rather than quantitative information; monitoring techniques change over time; and the objectivity of the conclusions are often questionable.

Given that electoral fraud is a phenomenon which is hard to measure directly, the only way to assess fraud dynamics is to explore indirect evidence. I analyze several data sources to obtain evidence on the dynamics of electoral fraud in the countries of interest: empirical studies on electoral fraud, variables that might be strongly correlated with fraud such as political and press freedom, and election databases that contain some fraud proxies comparable across different elections. The analysis of all these sources provides a consistent observation: electoral fraud tends to grow over the lifetimes of non-democratic regimes.

I first explore academic empirical literature on fraudulent elections which would be an ideal source of evidence on fraud dynamics. Though empirical literature on electoral fraud dynamics is limited, there are a few papers that try to compare irregularities in consecutive elections in selected countries. Myagkov and Ordeshook (2008) suggest a statistical methodology, based on analysis of the distribution of turnouts over different regions, and apply it to the official data for Russian federal elections between 1993 and 2007 to uncover electoral fraud. They find that, in the mid 1990s, ballot stuffing and some other forms of fraud were mentioned only in a few ethnic Russian regions but then spread to the other regions, both urban and rural, with noticeable acceleration during Putin's administration (2000-2008). The authors stress that once fraud occurs it becomes sustainable: they find that if fraud infected a precinct in Moscow for the first time in the 2004 presidential elections, it is very likely to reoccur there in the 2007 parliamentary elections. An important finding of the paper is that the level of electoral fraud in Russian federal elections has been both sustainable and growing as Putin's regime matures.

Increasing fraud dynamics in Russia after 2000 are also discussed by Treisman (2009), who comprehensively reviews the voting trends in Russia since 1991. In a chapter devoted to electoral manipulations and fraud, by studying a variety of Russian electoral statistics, the author finds that in the early 1990s elections in Russia were almost clean, while since 2000 electoral irregularities have become an integral part of electoral competition. Though manipulations were not serious enough to alter the outcomes, the growth of fraud was marked. These finding are also supported by Mebane and Kalinin (2009), who explore data on Russian federal elections for 2003-2008, looking for deviations from Benford's Law[13] as well as other statistical anomalies that are likely to arise due to fraud. The results show unambiguous growth in electoral fraud in the 2000s: anomalies the methods detect are worse by the end of the period under study than at the beginning[14].

In the third chapter of this dissertation I develop a statistical methodology that allows for detection of fraud in the forms of ballot stuffing, vote buying and multiple voting in official electoral data, and for estimating its magnitude. The methodology is then applied to the elections in Russia held between 2000 and 2012. The results are consistent with the findings of Treisman (2009)and Myagkov and Ordeshook (2008): electoral fraud in Russia has been persistently growing over this period. See section 3.5 of the dissertation for the details.

Using a methodology similar to Myagkov and Ordeshook (2008) and Mebane and Kalinin (2009), Levin et al (2009) analyze electoral data in Venezuela and find some evidence of increasing fraud over time. They analyze the data on two consecutive state-level referenda in 2007 and 2009 and, assuming constant voters' preferences, discover unusual patterns in voting behavior of selected regions that mainly benefit the incumbent. Specifically, the main finding of the paper is that most of the new votes in favor of Chavez in the 2009 referendum came from the regions with large abstention rates in 2007. Though this result cannot serve as strong evidence of fraud, it is more likely to be observed if fraud actually expanded between 2007 and 2009.

The main advantage of the methodology discussed is that it allows the detection of electoral fraud based solely on official election data. However, two main problems have to be mentioned. First, it is difficult (even impossible for the majority of African elections in the 1990s) to obtain such detailed data for all elections of interest. Second, electoral fraud is a comprehensive process, while the method detects mainly ballot stuffing and, thus, may underestimate the magnitude of fraud. Thus, one needs to adjust the evaluation of pure technical fraud by some measure of pre-election activity that directly affects election results. The main part of such activity consists of, for instance, controlling the media and pressuring the opposition (Schedler, 2002, Enikopolov et al, 2009). Hence, indexes of media and political freedom could be used as a proxy for pre-election manipulations.

First, I look at the dynamics of the Freedom House's Media Freedom Index[15], probably the most comprehensive cross country data set on press freedom. The index is an annual survey of media independence, which contains an assessment of the degree of print, broadcast, and internet freedom in practically every country in the world. It provides numerical rankings which reflect the legal environment for the media, political pressures that influence reporting, and economic

[13]Benford's Law states that in large lists of real-life data, digits are distributed in a specific, non-uniform way.
[14]Mebane and Kalinin (2009), page 11.
[15]www.freedomhouse.org/report-types/freedom-press (retrieved 20.03.2013)

factors that affect access to information in a particular country. The index is constructed annually from 1980 to the present, using approximately the same methodology during the entire span, which allows for cross country and time comparisons. I use the index for the countries of interest starting from 1994 (the first year for which the index is available for all the countries of interest) or the year of the first multiparty elections in a country, whichever number is higher, till 2013. The results are summarized in Table 5 and Table 6 of Appendix A. Columns "Press Freedom" contain coefficients and their standard errors in a regression of the index value on time for each regime. Since higher index values correspond to lower press freedom, a positive coefficient implies a decrease of press freedom over time. Generally, the results show that over time politicians in the countries of interest put more pressure on media: 8 post-Soviet regimes of 11 demonstrate a statistically significant positive time trend, while only 2 (Georgia and Tajikistan) shows clear improvement in press freedom. In 15 o of 34 African regimes, press freedom has significantly declined over time, 11 countries demonstrate no significant changes, and only 8 show some improvement.

Further, political freedom could also signal fraud-related pre-election activity, reflecting the transparency of the political environment in a given country. Another popular index, Freedom House Political Rights Index [16], captures the assessment of global political rights at country level. I use the index for 1994-2013 to analyze time trends in political freedom in the countries of interests. The results are summarized in Table 5 and Table 6 of Appendix A. The evolution of political freedom in the countries of interest shows a picture similar to the case of press freedom. Of the Post-Soviet regimes, 7 of 11 demonstrate significant growth of the index over time, which corresponds to a decrease in political freedom, 3 (Azerbaijan and 2 Georgian regimes) have no significant changes and in just 1 (Tajikistan) there is an improvement in political freedom over time. In Africa, for 12 regimes the index grows significantly, for another 15 the trend is not clear, and only 7 show improving political freedom.

The World Bank Database of Political Institutions (Beck et al, 2001) is a popular regularly updated[17] source of data (1975-2009) on political systems and elections around the world. The database contains a dummy variable fraud that captures extra-constitutional electoral irregularities and equals to 0 if elections are considered to be fair and 1 otherwise. Though the variable is just a dummy that does not allow for measuring the magnitude of fraud, and there are some questions regarding how the variable is constructed (for instance, all elections in Russia, Zimbabwe, and in Uzbekistan since 2005 are marked as not fraudulent, though by other measures they were fraudulent), one can gain some inference about the time trends in electoral fraud by

[16] www.freedomhouse.org/report-types/freedom-world, Report "Country Ratings and Status by Region, FIW 1973-2013" (retrieved 20.03, 2013)

[17] Last update of the DPI: April 2010.

looking for a switch from not fraudulent elections to rigged ones within the lifetime of a regime. The results of the analysis of the DPI dataset are summarized in Table 5 and Table 6 of Appendix A: "+" sign means that a switch from not-fraudulent elections to fraudulent during a regime lifetime was observed, "–" sign stands for an opposite switch. Out of 45 regimes from both sets, a switch is observed for 15: in 11 countries there is a switch from not fraudulent to fraudulent, and only in 4 - vice versa.

Next, I examine another popular election data set, National Elections Across Democracy and Autocracy-NELDA (Hyde and Marinov, 2012), which provides detailed information on election events between 1960 and 2006. I analyze the dynamics of several fraud related variables contained in the data set. Specifically, for each election there are three binary answers for the following three questions: "Before elections, are there significant concerns that elections will not be free and fair?", "Is there evidence that the government harassed the opposition?", and "Were there allegations by Western monitors of significant vote-fraud?" As in the case of the analysis of DPI data set, I am looking for switches between answers "no" and "yes" to these questions for elections within each regime of interest. Switch from "no" to "yes" to any of these questions implies a clear decrease in electoral integrity, from "yes" to "no" - an improvement, while the absence of switches is inconclusive. Since the data covers only elections up to 2006, there are just 36 regimes for which at least two observations are available. The results of the analysis are presented in Table 7 of Appendix A. The results show weak evidence of growing fraud. For the first question there are 8 switches from "no" to "yes" (marked with "+" in Table 7), 5 switches from "yes" to "no" (marked with "–") and 3 cases where both switches were observed ("both"). The other 20 regimes did not have switches. For the second question there are 9 "+" marks, 4 "–" marks and 3 regimes with both types of switches. For the third question there are 7 "no"-to-"yes" switches, 4 "yes"-to-"no" switches, 1 regime demonstrated switches of both types, and for 6 regimes data are not available.

Further, evolution of the victory margin[18] could also provide information about fraud. For example, Simpser (2005, 2008) argues that that fraudulent elections are strongly associated with high victory margins: corrupt politicians tend to win elections by large margins. Given this correlation, positive time trends in victory margins may signal growing fraud. The number of presidential elections within the lifetime of a regime varies across countries of interest,with between 2 and 5 elections conducted. As a result, standard errors for time coefficients are relatively large even if the trend is obvious. Table 5 and Table 6 contains the results of the analysis[19]: out of 44 regimes [20], 27 demonstrate clear growth in the victory margin over time.

[18]The victory margin is the difference between the shares of votes cast for the winner and the first runner-up.

[19]absence of standard error means two observations

[20]Several elections in the dataset were boycotted by opposition, and thus victory margins for these elections are irrelevant. One of these elections was held in Cote d'Ivoire, a regime that had only two presidential elections

To summarize, electoral fraud is a difficult-to-measure phenomenon, and there is no objective measure that reliably reflects the magnitude of electoral misconduct. Yet, there is some indirect evidence, and measures of fraud-related phenomena can give an inference of the evolution of electoral fraud over time. The analysis of such indirect evidence in non-democratic regimes in Africa and in Post-Soviet countries suggests that electoral fraud tends to grow over time in the majority of cases. The following section presents a formal model of fraudulent elections to show how uncertainty can contribute to the observed trends.

2.4 The Model

2.4.1 General Setup

This section develops a game-theoretical model of political competition between a corrupt incumbent and a challenger. The incumbent faces a continuum of voters of measure one. Before the elections, the incumbent chooses the level of fraud. The incumbent derives utility only from remaining in office and thus, from his point of view, fraud is just an instrument to manipulate his probability of being re-elected. This eliminates the potential commitment problem typical for non-democracies (Acemoglu and Robinson (2006), Chapter 5) because the incumbent has no incentive to change his policy before the next election campaign, as fraud does not affect the incumbent's utility. The level of chosen fraud corresponds to a unique number $f \in [0, 1]$ - the share of votes that the incumbent can add to his true support. Hereafter, without loss of generality, I assume that the amount of fraud in the model equals the increment in percentage of votes.

All the voters dislike fraud in the same way. Voter i has utility from fraud f

$$V_i(f) = -\beta f, \tag{2.1}$$

where $\beta \in [0, 1]$ is an intolerance parameter that captures voters' attitude towards fraud. Parameter β has a true value, which is, however, unknown to the incumbent. Yet, the incumbent has prior beliefs about its value: $\beta \sim N(\beta_1, \epsilon_1^2)$. Thus, the intolerance parameter β is subject to uncertainty, which I refer to as fraud intolerance uncertainty. The challenger has no option to commit fraud.

The elections are modeled in a modified version of the standard probabilistic voting framework (presented for the first time in Lindbeck and Weibull (1987), later used in Persson and Tabellini (2000) and more recently in, for example, Gregory et al (2011)), where voter i votes

over its lifetime. As a result, there is only 1 data point for this regime for the incumbent, who commits fraud f against the challenger, if

$$V_i(f) + \sigma_i + \delta > 0, \tag{2.2}$$

where σ_i is an individual-specific time-constant preference over the incumbent. Across all the voters, σ_i, which captures a relative ideological bias towards the incumbent, is distributed uniformly over $[-\frac{1}{2\phi}, \frac{1}{2\phi}]$. The distribution is common knowledge.

Uncertainty about voters' preferences, which I hereafter call aggregate electoral uncertainty, is captured by δ, a common for all the voters component, which represents a random preference for the incumbent shared by all voters and which is unknown to the candidates prior to election day. This component is the utility an individual derives from all the incumbent's policies other than fraud relative to all other policies of the challenger. Prior to elections, the value of this component is drawn from a zero-mean normal distribution: $\delta \sim N(0, \psi^2)$, which is known to the incumbent.

If elected, the incumbent gains benefits from remaining in office. The benefits are normalized to 1. The direct costs of fraud are $c(f)$ such that $c(0) = 0$, $c'(0) = 0$, $c'(f) > 0$, $c''(f) > 0$.

Also, $c(1)$ and $c'(1)$ are assumed to be relatively large numbers to guarantee that falsifying 100% of the votes is extremely costly. The incumbent chooses the level of fraud to maximize his expected benefits.

The timing of the game is as follows:

1. The incumbent chooses the level of fraud $f \in [0, 1]$.

2. The voters anticipate f, the elections take place, the results are adjusted by the level of fraud and are announced, and the winner takes office.

3. The payoffs are realized.

Note that correct anticipation of the fraud level by voters is possible only under the assumption of no private information on the side of the incumbent. This means that to anticipate the level of fraud correctly, the voters must know preferences, costs and benefits of fraud as well as the incumbent's beliefs. If one considers this assumption to be too strict, it can be assumed that fraud is fully observable by voters. The latter assumption is not as strict as it seems at first sight because fraud, as discussed above, is a comprehensive process including controlling the media and threatening the opposition, which mainly takes place before elections, and is easily observed by voters. With any of these two assumptions the following analysis is valid.

2.4.2 One Period Analysis

I start with the analysis of a one-period model. For any given level of fraud f, voter i votes for the incumbent if

$$V(f) + \sigma_i + \delta > 0, \tag{2.3}$$

$$-\beta f + \sigma_i + \delta > 0, \tag{2.4}$$

$$\sigma_i > \beta f - \delta. \tag{2.5}$$

Then, the probability that a randomly picked voter votes for the incumbent is

$$P(\sigma_i > \beta f - \delta) = 1 - P(\sigma_i \leq \beta f - \delta) = 1 - \frac{\beta f - \delta + \frac{1}{2\phi}}{1/\phi} = \frac{1}{2} - \phi(\beta f - \delta). \tag{2.6}$$

This is exactly equal to the true share of votes cast for the incumbent for a given realization of δ as there is a continuum of voters of measure 1:

$$\Pi_I = \frac{1}{2} - \phi(\beta f - \delta). \tag{2.7}$$

Note that if elections are clean, i.e., fraud is zero, each candidate will get exactly half of the votes in expectation.

Given the fraud level, the probability that the incumbent wins the elections under the majority rule is then

$$P_w = P(\Pi_I + f \geq \frac{1}{2}) = P\left(\frac{1}{2} - \phi(\beta f - \delta) + f \geq \frac{1}{2}\right) = P\left(\beta - \delta/f - 1/\phi \leq 0\right). \tag{2.8}$$

Denote $X = \beta - \delta/f - 1/\phi$. From the incumben's point of view X is a random variable which, given his prior beliefs about β and distribution of δ, follows $N\left(\beta_1 - \frac{1}{\phi}, \epsilon_1^2 + \frac{\psi^2}{f^2}\right)$.

Given the expected share of votes, the incumbent chooses the level of fraud f such that it maximizes his expected benefit:

$$\max_f P_w(f) - c(f). \tag{2.9}$$

The problem can be rewritten as:

$$\max_f G_X(0) - c(f), \tag{2.10}$$

where $G_X(z) = \frac{1}{\sqrt{2\pi s^2}} \int_{-\infty}^{z} e^{-\frac{(x-\mu)^2}{2s^2}} \mathrm{d}x$, $\mu = \beta_1 - \frac{1}{\phi}$, and $s^2 = \epsilon_1^2 + \frac{\psi^2}{f^2}$.

Proposition 2.1. *Solution f to maximization problem (2.10) satisfies:*

$$-\frac{\psi^2}{f^3}\frac{\mu}{s^2}g_X(0) - c'(f) = 0, \tag{2.11}$$

where $g_X(z) = \frac{1}{\sqrt{2\pi s^2}}e^{-\frac{(z-\mu)^2}{2s^2}}$.
Proof: *See the Appendix.*

Proposition 2.2. *There is unique $f^* \in (0,1)$ that satisfies the first-order condition (2.11) if and only if $\mu = \beta_1 - \frac{1}{\phi} < 0$.*
Proof: *See the Appendix.*

To understand the condition for uniqueness, recall that the incumbent's expected vote share is $E\Pi_I = E(\frac{1}{2} - \phi(\beta f - \delta) + f) = \frac{1}{2} + f(1 - \phi\beta_1)$, which is an increasing function of fraud if and only if $\phi\beta_1 < 1$ or $\beta_1 - 1/\phi < 0$. Otherwise committing fraud does not make sense because it hurts the incumbent in terms of votes. This can happen when people strongly dislike fraud (e.g., β is relatively high), or when there is little heterogeneity among the electorate in terms of ideology (ϕ is high). In the latter case, by committing fraud, which is disliked by everyone, the incumbent loses a relatively large number of his ideological supporters (those with $\sigma_i > 0$) as he is ideologically too close to the challenger who does not commit any fraud. Thus, the condition $\beta_1 - 1/\phi < 0$ guarantees that in expectation committing fraud makes sense, i.e., it provides incumbent with a higher official vote share than he would get if he did not commit fraud.

The second condition (sufficiently large uncertainty) guarantees that the left-hand side of the first-order condition (2.11) intersects the marginal cost function at a point between 0 and 1. It is not binding for any reasonable parameter values, mainly because marginal cost at $f = 1$ is a relatively large number under the assumption that stealing 100% of the votes is extremely costly.

2.4.3 Multi-Period Setup

onsider a multi-period setup where there is a sequence of elections. In the end of period 1, the incumbent observes his vote share:

$$v_1 = \frac{1}{2} - \phi(\beta f_1 - \delta_1) + f_1, \tag{2.12}$$

where f_1 is the first-period fraud. The incumbent knows the exact values of ϕ and f_1, but does not know β and δ_1, and thus cannot decompose the observed value $\beta f_1 - \delta_1$. This value may be rewritten as $m = \beta - \frac{\delta_1}{f_1}$. From the point of view of the incumbent, observing his vote share

(2.12) is equivalent to observing m. Since δ is distributed with zero mean, m can be interpreted as an unbiased signal about true value of β, which is then used by the incumbent for Bayesian updating of his prior beliefs.

Because δ is drawn from normal distribution $N(0, \psi^2)$, signal m is also distributed normally: $m \sim N(\beta, \frac{\psi^2}{f_1^2})$. Given the distribution of the signal and priors $\beta \sim N(\beta_1, \epsilon_1^2)$, the posterior beliefs about β are:

$$\beta|m \sim N\left(\frac{\beta_1\psi^2 + \epsilon_1^2 f_1^2 m}{\psi^2 + \epsilon_1^2 f_1^2}, \frac{\epsilon_1^2\psi^2}{\psi^2 + \epsilon_1^2 f_1^2}\right). \tag{2.13}$$

It is important to note that fraud in the first period affects the beliefs of the incumbent about β in the second period. Specifically, the variance of the beliefs is lower if the first period fraud is higher. Also, note that if there is no fraud in the first period, the beliefs about β do not change: when there is no fraud, there is no way to learn anything about voters' response to it.

Given the updated beliefs, in period 2 the incumbent solves

$$\max_{f_2} G_Y(0) - c(f_2), \tag{2.14}$$

where $Y = \beta - \frac{\delta_2}{f_2} - \frac{1}{\phi}$, $G_Y(z) = \frac{1}{\sqrt{2\pi s_2^2}} \int_{-\infty}^{z} e^{-\frac{(x-\mu_2)^2}{2s_2^2}} \mathrm{d}x$, $\mu_2 = \frac{\beta_1\psi^2+\epsilon_1^2 f_1^2 m}{\psi^2+\epsilon_1^2 f_1^2} - \frac{1}{\phi}$, and $s_2^2 = \frac{\epsilon^2\psi^2}{\psi^2+\epsilon^2 f_1^2} + \frac{\psi^2}{f_2^2}$.

Note that δ_2 is again drawn from the same commonly known distribution $N(0, \psi^2)$ independently from the first-period draw. Further, I will explore the case when δ follows a random walk and show that the results do not substantially differ from the case with independent draws. However, random walk generates an additional effect (to be discussed further) that can contribute to a growing fraud profile and cannot be distinguished from the uncertainty effect. The independent draw assumption eliminates this effect, allowing me to study purely the role of uncertainty in fraud dynamics.

Similar to the first-period case (see Proposition 2.1), the second period first-order condition takes the following form:

$$-\frac{\psi^2}{f_2^3}\frac{\mu_2}{s_2^2} g_Y(0) - c'(f_2) = 0, \tag{2.15}$$

where $g_Y(z) = \frac{1}{\sqrt{\pi s_2^2}} e^{-\frac{(z-\mu_2)^2}{s_2^2}}$.

To understand the conditions under which growing fraud occurs, note that the only components of the maximization problem that are different between two periods are mean and variance of the subjective distribution of β. Recall, that from the incumbent's point of view in period 1 $\beta \sim N(\beta_1, \epsilon_1^2)$, and in period 2, given a signal $\beta|m \sim N(\beta_2, \epsilon_2^2) = N\left(\frac{\beta_1\psi^2+\epsilon_1^2 f_1^2 m}{\psi^2+\epsilon_1^2 f_1^2}, \frac{\epsilon_1^2\psi^2}{\psi^2+\epsilon_1^2 f_1^2}\right)$.

Thus, beliefs are the only thing that determines the difference in optimal choices of fraud in two different periods. For clarity, I separate the effect of changes in beliefs on the mean effect and the variance effect, i.e., changes in optimal fraud between two periods in response to changes in the mean and variance of beliefs, respectively. Further note that variance of the beliefs in the second period is always lower for any $f_1 > 0$: $\epsilon_2^2 = \frac{\epsilon_1^2\psi^2}{\psi^2+\epsilon_1^2 f_1^2} < \frac{\epsilon_1^2}{1+\epsilon_1^2 f_1^2/\psi^2} < \epsilon_1^2$. The following proposition states that this decrease leads to higher optimal level of fraud.

Proposition 2.3. *For some prior beliefs about true value of β distributed according to $N(\beta_0, \epsilon_0^2)$ optimal fraud is decreasing in ϵ_0^2 if $\epsilon_0^2 + \psi^2 > (1/\phi - \beta_0)^{2/3}$.*
Proof: *See the Appendix.*

Proposition 2.3 says that when there is sufficiently large uncertainty, no matter of what type, on the incumbent's side an increase in uncertainty about the level of fraud intolerance leads to lower equilibrium level of fraud. Note that there are two opposite effects of uncertainty on committing fraud. On the one hand, the incumbent is afraid of committing too much fraud when he has doubts about how voters react to it, as the intolerance parameter could easily appear to be high enough to make fraud damaging rather than useful (see Proposition 2.1). On the other hand, higher uncertainty implies that investment in fraud becomes less efficient. According to Proposition 2.3, when uncertainty is relatively high, the first effect dominates. Finally, note that the condition is sufficient, meaning that optimal fraud is decreasing in subjective uncertainty of the incumbent under even less strict circumstances.

Thus, more precise beliefs under sufficiently large uncertainty provides incentives to commit greater fraud, implying that the variance effect pushes the optimal fraud up. The next question is the direction of the mean effect. To answer it, one first needs to know how changes in the mean of beliefs affect the optimal fraud.

Proposition 2.4. *For some prior beliefs about true value of β distributed according to $N(\beta_0, \epsilon_0^2)$ optimal fraud is decreasing in β_0 if $\epsilon^2 + \psi^2 > (1/\phi - \beta_0)^2$.*
Proof: *See the Appendix.*

This result, when optimal fraud is not always decreasing in the expected voter's intolerance, could seem counterintuitive. To understand it note that higher intolerance of fraud should decrease incentives to commit fraud because with higher value of β_0, keeping the variance ϵ^2 fixed, the probability that true β will appear to be high enough to make fraud electorally detrimental to the incumbent instead of beneficial (see Proposition 2.1) is now higher. But according to Proposition 2.3, increased uncertainty about fraud intolerance decreases incentives to commit fraud only when the uncertainty is sufficiently high. Thus, an increase in β_0 induces lower fraud only when uncertainty is relatively high.

The second period mean $\frac{\beta_0\psi^2+\epsilon^2 f_1^2 m_1}{\psi^2+\epsilon^2 f_1^2} > \beta_0$ if $m_1 > \beta_0$, and $\frac{\beta_0\psi^2+\epsilon^2 f_1^2 m_1}{\psi^2+\epsilon^2 f_1^2} < \beta_0$ if $m_1 < \beta_0$. Thus, according to Proposition 2.4, if $m_1 < \beta_0$ and ϵ^2 is high enough, the mean effect pushes the optimal level of fraud up as well as the variance effect, resulting in unambiguously increasing fraud over two periods. If $m_1 > \beta_0$, the updated mean is higher than the prior mean, and in this case the mean and the variance effects affect optimal fraud in opposite ways, and the resulting direction depends on the values of the model parameters and the realized value of the signal. Specifically, the higher the signal, the more likely the mean effect dominates the variance effect, implying a decrease in fraud. Thus, there is a threshold value for signal m_1^* such that if $m_1 > m_1^*$ then the second period optimal fraud is lower than the first one; if $m_1 \leq m_1^*$ then the optimal fraud increases between the two periods and $m_1^* > \beta_0$.

Because m_1 is distributed symmetrically around the true β, realization of the signal is more likely to be below the threshold value m_1^* implying that it is more likely to observe increasing fraud rather than decreasing, if prior beliefs are unbiased ($\beta_0 = \beta$). The likelihood increases if the incumbent overestimates *beta* ($\beta_0 > \beta$).

The analysis could be easily extended to the case of a multiple period game under the assumption of a myopic incumbent. Here, myopia means that the incumbent does not invest in learning by strategically committing excessive fraud. A fully rational incumbent could have incentives to choose relatively more fraud in the first period, bearing some extra risk and extra costs in exchange for learning faster the true value of β. However, the assumption of a fully rational incumbent seems exaggerated, taking into account some features of the real-life electoral environment, where, for example, the length of electoral cycles is rarely less than 4-6 years, which is probably too long to assume commitment of strategic fraud is prevalent.

The crucial thing for the results of the multi period analysis is the conditions stated in Propositions 2.3 and 2.4. The analysis is indeed valid only if the conditions hold over time. Note that both conditions require ϵ^2 and ψ^2 to be sufficiently large, ϵ^2 is decreasing over time due to learning and ψ^2 is constant over time. Thus, eventually the conditions break down. However, the higher ψ^2 and the initial value of ϵ^2, the later the break occurs, allowing the analysis to be valid for sufficiently large number of periods. The following section presents the results of the simulation of the multi-period model.

2.4.4 Simulation of the Multi-Period Model

The multi-period setup assumes T periods. At the beginning of a period t the incumbent solves

$$\max_{f_t} G_H(0) - c(f_t), \tag{2.16}$$

where $H = \beta - \frac{\delta_t}{f_t} - \frac{1}{\phi}$, $G_H(z) = \frac{1}{\sqrt{2\pi s^2}} \int_{-\infty}^{z} e^{-\frac{(x-\mu)^2}{2s^2}} \, \mathrm{d}x$, $\mu = \mu_t(m_{t-1}, f_{t-1} - 1/\phi)$, and

$s^2 = s_t^2(m_{t-1}, f_{t-1} + \frac{\psi^2}{f_t^2})$, $_t$ and s_t^2 are mean and variance of the incumbent's beliefs about β at period t, which are both functions of the previous period fraud and signal.

To analyze the dynamics of fraud I simulate the multi-period model with a sequence of 6 periods. I repeat the sequence of elections 30 times, starting from the same parameter values. In each period of a sequence I solve the incumbent's problem given a quadratic cost function, resolve uncertainty by randomly drawing a value of δ_t from the specified distribution, and update the incumbent's beliefs. If the incumbent has lost elections, the game is over, otherwise the next period starts.

The benchmark model parameters are as follows: $\beta = 0.15$, $\epsilon = 0.05$, $\psi = 0.001$, $\phi = 6$. Prior beliefs are unbiased. The choice of the parameter values is not a result of calibration due to the obvious reasons discussed above. Instead, the parameters are chosen in a way that guarantees reasonable relationships between them. First, β and ϕ are set such that committing fraud makes sense (see Proposition 2.2). Second, ϵ is chosen such that it guarantees reasonable uncertainty of the value of β, allowing the incumbent to assign a relatively high probability to an outcome where the voters are fraud intolerant and high fraud becomes electorally detrimental to the incumbent (see Proposition 2.2). Finally, variance ψ of the purely random component δ is chosen such that it is relatively lower than ϵ, guaranteeing that a toss of a coin does not decide too much in the model and the incumbent does not lose very often due to bad luck.

To show that the exact values of the parameters do not exclusively determine the model predictions, I simulate the model for another two parameter sets, chosen in the same way as described above, in addition to the benchmark one. As a result, the model is simulated for three distinct sets:

$$(\beta, \epsilon, \psi, \phi) = (0.15, 0.05, 0.001, 6), (0.5, 0.4, 0.01, 1.5), (0.8, 0.25, 0.025, 1). \tag{2.17}$$

In the following figures each kinked line represents optimal fraud as a function of time.

Figure 2.1: Optimal fraud for different parameter sets (2.17).

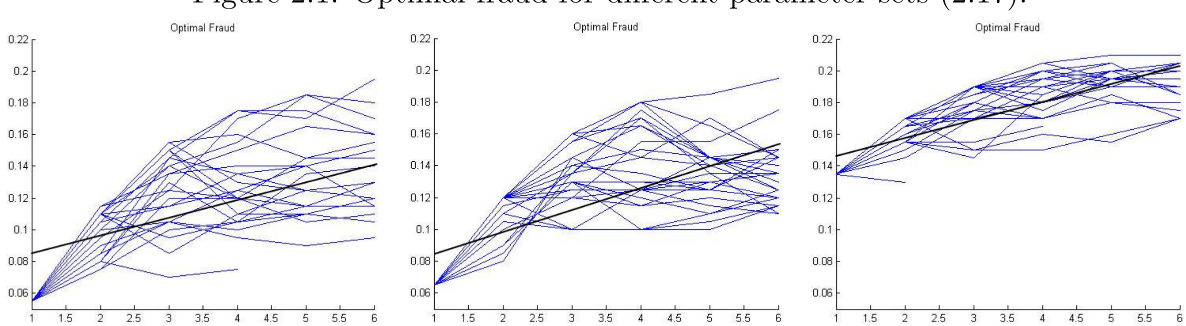

To document the significance of the growing trend I simulate the model with a sequence of

6 periods 1000 times for the parameter sets used above and run a simple regression of optimal fraud on time. The results show that in the model, fraud on average grows by about 1.1-1.5 percentage points every period.

Variable	Set 1	Set 2	Set 3
Time	0.01531***	0.01330***	0.01165***
	(0.00019)	(0.00018)	(0.00013)
Const	0.06175***	0.07278***	0.13356***
	(0.00072)	(0.00070)	(0.00049)
N	1000	1000	1000

Table 2.1: Time trend of optimal fraud for the benchmark model

One of the essential model components is the aggregate electoral uncertainty captured by parameter δ. The benchmark dynamic model assumes that every period δ is independently drawn from the same normal distribution. One can argue that non-partisan preferences for candidates could be time dependent. Finally, I allow for this by making δ follow a random walk instead of being independently drawn: $\delta_t \sim N(\delta_{t-1}, \psi^2)$. Again, the results of the time regressions demonstrate the significance of fraud growth.

Figure 2.2: Optimal fraud in the model with random walk for different parameter sets (2.17).

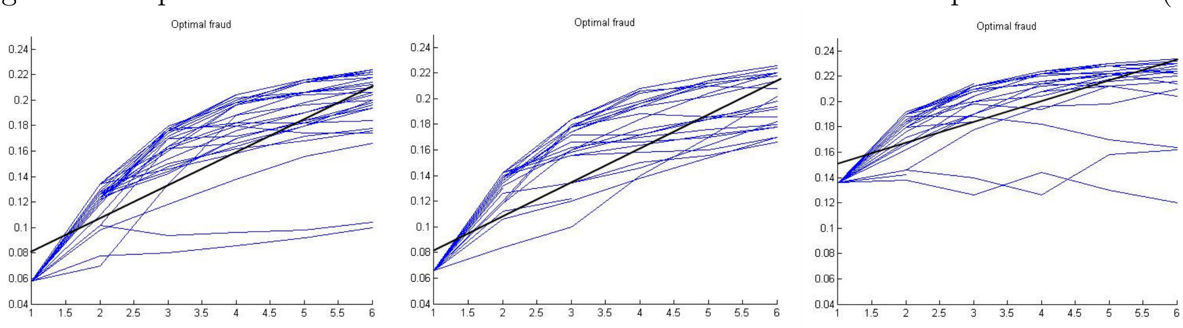

Variable	Set 1	Set 2	Set 3
Time	0.02632***	0.02521***	0.01611***
	(0.00023)	(0.00022)	(0.00016)
Const	0.05486***	0.06517***	0.13651***
	(0.00085)	(0.00086)	(0.00061)
N	1000	1000	1000

Table 2.2: Time trend of optimal fraud for the model with random walk

Indeed, random walk for aggregate uncertainty seems to be more realistic than independent draws. However, the growing fraud profile in this case could be a result of two effects: in addition to the effect of learning about fraud tolerance, there is also an effect of aggregate uncertainty. In contrast to the benchmark case where the aggregate uncertainty was constant over time, it is growing due to the random walk process. As a result, incentives to commit higher fraud over time are increased not only by more precise beliefs about β, but also by higher aggregate uncertainty about electoral support. Thus, for the benchmark case the assumption on independent draws for the aggregate uncertainty component allows me to distinguish two uncertainty time effects and, thus, seems more plausible for the purposes of the paper.

To summarize, the model simulation results in a consistent growing fraud profile which is robust to parameter choice and underlying law of motion for aggregate electoral uncertainty. When uncertainty is sufficiently large, incentives to commit fraud increase when an incumbent's uncertainty about β decreases. Over time an incumbent's beliefs about β become more precise due to learning and, thus, the deterrent role of uncertainty about fraud intolerance declines, implying that the incentives to commit fraud become stronger, leading to a growing profile.

An important observation is that fraud generally grows at a decreasing rate, which is a result of fast learning that takes place mainly in early periods. To clearly illustrate the speed of information gathering, in Figure 2.3, I represent the evolution of uncertainty for the three parameters sets already used above. One can notice that the standard deviation of subjective beliefs rapidly decreases in the few first periods.

Figure 2.3: Optimal fraud in the model with random walk for different parameter sets (2.17).

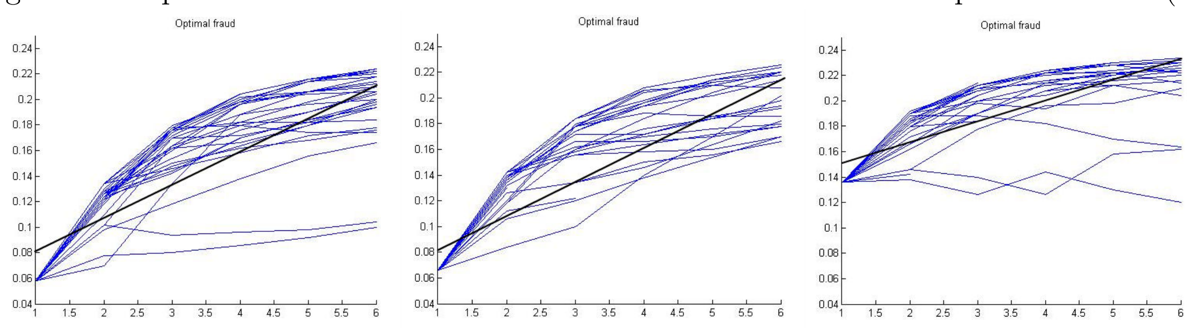

Finally, it is worth stressing that the model puts aside the effects of the growing stakes of re-election and decreasing costs of fraud because of learning-by-doing, focusing purely on the role of uncertainty in fraud dynamics. Obviously, being introduced into the model, these components would just magnify the results, making the model predictions even stronger.

2.5 Conclusion

This paper consists of two parts. The first part explores different available information on electoral fraud in post-Soviet and Sub-Saharan countries. All the explored sources, including the academic literature, electoral data and data on political freedom, provide consistent, though indirect, evidence of growing fraud: electoral manipulations tend to grow over the lifetime of a non-democratic regime.

The second part provides a theoretical model of electoral competition with falsifications designed in the traditional probabilistic voting framework, which specifically studies the effect of uncertainty on the incumbent's incentives to commit fraud. The model explicitly distinguishes between two types of uncertainty that affect electoral outcomes: aggregate electoral uncertainty, captured by a purely random component in voters' utility over candidates, and subjective uncertainty about voters' fraud intolerance, represented by the variance of an incumbent's beliefs about the fraud intolerance component in voters' utility. The most important findings of the model are as follows.

First, when uncertainty (no matter of what type) is sufficiently large, incentives to commit fraud increase when an incumbent's uncertainty about fraud intolerance decreases. Second, because in the multi-period setup an incumbent's uncertainty decreases as a result of learning through Bayesian updating, increasing fraud is more likely to be observed. Third, optimal fraud demonstrates an increasing profile at a decreasing rate. This is explained by fast learning that mainly takes place in early periods: the incumbent quickly absorbs information about the true value of the intolerance component in voters' utility function. The predictions are robust to functional forms of the model components as well as choice of parameter values.

Appendix A

Country	Regime Years	Elections	Leaders
Armenia	1998-	1998, 2003, 2008	Kocharyan, Sargsyan
Azerbaijan	1993-	1993, 1998, 2003, 2008	H. Aliyev, I. Aliyev
Belarus	1994-	1994, 2001, 2006	Lukashenko
Georgia 1	1995-2003	1995, 2000	Shevardnadze
Georgia 2	2004-	2004, 2007	Saakashvili
Kazakhstan	1991-	1991, 1999, 2005	Nazarbaev
Kyrgyzstan	1990-2005	1991, 1995, 2000	Akayev
Russia	2000-	2000, 2004, 2008, 2012	Putin, Medvedev
Tajikistan	1992-	1995, 1999, 2006,	Rahmon
Ukraine	1994-2004	1994, 1999	Kuchma
Uzbekistan	1990-	1991, 2000, 2007	Kerimov

Table 2.3: Regimes. Post-Soviet Countries.

Country	Regime Years	Elections	Leaders
Burkina Faso	1987-	1991, 1998, 2005, 2010	Compaore
Cameroon	1982-	1992, 1997*, 2004, 2011	Biya
CAR 1	1993-2003	1993, 1999	Patasse
CAR 2	2003-	2005, 2011	Bozize
Chad	1990-	1996, 2001, 2006, 2011	Deby
Congo (Brazzaville)	1997-	2002, 2009	Sassou-Nguesso
Cote d'Ivoire	1960-1999	1990, 1995*	Houphouet-Boigny, Bedie
Djibouti	1977-	1993, 1999, 2005*, 2011	Aptidon, Guelleh
DR Congo (Kinshasa)	1997-	2006, 2010	Kabila, Kabange
Equatorial Guinea	1987-	1992, 2006, 2009	Mbasogo
Gabon	1987-	1991, 1998, 2005, 2010	O. Bongo, A. Bongo
the Gambia	1994-	1996, 2001, 2006, 2011	Jammeh
Ghana	1982-2000	1992, 1996, 2000	Rawlings
Guinea	1984-2008	1993, 1998, 2003	Conte
Kenya 1	1978-2002	1992, 1997, 2002	Moi, Kenyatta
Kenya 2	2002-	2002, 2007	Kibaki
Liberia	2005-	2005, 2011	Johnson-Sirleaf
Madagascar	2001-2009	2001, 2006	Ravalomanana
Malawi	1994-	1994, 1999, 2004, 2009	Muluzi, Mutharika
Mauritania	1984-2005	1992, 1997*, 2003	Taya
Mozambique	1986-	1994, 1999, 2004, 2009	Chissano, Guebuza
Niger	1999-2010	1999, 2004	Tandja
Nigeria	1987-	1999, 2003, 2007, 2011	Obasanjo, Yar'Adua
Rwanda	1994-	2003, 2010	Bizimungu, Kagame
Senegal 1	1960-2000	1983, 1988, 1993, 2000	Diouf
Senegal 2	2000-2012	2000, 2007, 2012	Wade
Seychelles	1977-	1993, 1998, 2001, 2006, 2011	Rene, Michel
Sierra-Leone	2007-	2007, 2012	Koroma
Sudan	1989-	1996, 2000, 2010	al-Bashir
Tanzania	1960-	1995, 2000, 2005, 2010	Nyerere, Mwinyi, Mkapa, Kikwete
Togo	1993-	1993*, 1998, 2003, 2005, 2010	Eyadema, Gnassingbe
Uganda	1986-	1996, 2001, 2006, 2011	Museveni
Zambia	1991-2011	1991, 1996, 2001, 2006 2008, 2011	Chiluba, Mwananwasa, Banda
Zimbabwe	1987-	1990, 1996*, 2002, 2008	Mugabe

*Elections were boycotted by opposition parties.

Table 2.4: Regimes. African Countries

Country	Political Rights	Press Freedom	Victory Margin	DPI
Armenia	0.161 (0.026)***	0.767 (0.092)***	7.6 (2.34)	+
Azerbaijan	0	0.495 (0.092)***	10.7 (7.16)	+
Belarus	0.124 (0.018)***	1.263 (0.192)***	24.7 (4.50)	no
Georgia 1	0.050 (0.066)	-2.351 (0.564)***	8.2	no
Georgia 2	0.083 (0.065)	-0.067 (0.560)	-66.3	no
Kazakhstan	0.023 (0.009)**	1.167 (0.044)***	12.3 (1.81)	+
Kyrgyzstan	0.158 (0.040)***	1.790 (0.260)***	5.6	no
Russia	0.099 (0.022)***	2.088 (0.151)***	6.3 (6.88)	no
Tajikistan	-0.033 (0.013)**	-1.232 (0.225)***	-20.9	no
Ukraine	0.145 (0.031)***	3.200 (0.306)***	11.4	+
Uzbekistan	0.022 (0.009)**	1.021 (0.124)***	-0.3	no

Table 2.5: Fraud Proxies Dynamics. Post-Soviet Countries

Country	Political Rights	Press Freedom	Victory Margin	DPI
Burkina Faso	0.006 (0.016)	0.189 (0.043)***	-4.8 (0.28)	+
Cameroon	-0.035 (0.015)**	-0.726 (0.126)***	31.7 (10.31)	+
CAR 1	0.309 (0.088)***	-0.752 (0.661)	16.6	no
CAR 2	0	0.321 (0.216)	13.7	no
Chad	0.081 (0.014)***	0.293 (0.110)**	13.7 (3.12)	–
Congo (Brazzaville)	0.082 (0.038)*	0.091 (0.162)	-15.5	no
Cote d'Ivoire	0	2.800 (0.901)**		+
Djibouti	0.023 (0.027)	0.917 (0.181)***	11.3 (1.15)	no
DR Congo (Kinshasa)	0.172 (0.062)**	0.429 (0.152)**	0.5	n/a
Equatorial Guinea	-0.053 (0.025)*	1.009 (0.085)***	-2.8 (0.06)	no
Gabon	0.083 (0.012)***	1.304 (0.156)***	-1.2 (12.54)	no
the Gambia	-0.103 (0.051)*	1.185 (0.193)***	12.3 (2.85)	–
Ghana	-0.383 (0.053)***	-0.075 (0.038)	-20.9 (6.18)	–
Guinea	0.022 (0.012)*	0.209 (0.231)	29.2 (17.20)	no
Kenya 1	0.127 (0.092)	2.103 (0.444)***	-20.6 (11.14)	+
Kenya 2	0.145 (0.031)***	-0.900 (0.250)***	-28.6	no
Liberia	0	-1.000 (0.244)***	62.6	no
Madagascar	0.366 (0.072)***	2.883 (0.506)***	27.5	no
Malawi	0.082 (0.029)***	1.153 (0.180)***	6.6 (5.36)	no
Mauritania	-0.099 (0.031)***	-1.000 (0.263)***	18.6	no
Mozambique	0.053 (0.013)***	-0.425 (0.173)**	4.7 (5.07)	no
Niger	-0.021 (0.073)	0.647 (0.438)	11.2	no
Nigeria	-0.114 (0.126)	-0.800 (0.400)	2.5 (6.31)	no
Rwanda	0 (0)	-0.050 (0.153)	-3.4	no
Senegal 1	0.039 (0.022)*	-2.702 (1.012)**	-43.4	–
Senegal 2	0.042 (0.044)	2.157 (0.229)***	-24.3 (27.88)	no
Seychelles	0	0.596 (0.138)***	-8.24 (8.10)	n/a
Sierra-Leone	-0.143 (0.082)	-2.300 (0.277)***	12.1	no
Sudan	0	-0.704 (0.127)***	-30.4	n/a
Tanzania	-0.104 (0.015)***	0.037 (0.056)	1.9 (8.98)	+
Togo	-0.066 (0.018)***	-0.040 (0.174)	2.5 (1.12)	no
Uganda	0.029 (0.030)	1.338 (0.111)***	-5.2 (6.12)	+
Zambia	-0.004 (0.036)	0.254 (0.169)	-14.6 (5.75)	+
Zimbabwe	0.089 (0.018)***	2.107 (0.338)***	5.1 (32.88)	no

Table 2.6: Fraud Proxies Dynamics. African Countries

Country	Question 1	Question 2	Question 3
Armenia	+	no	no
Azerbaijan	no	no	no
Belarus	+	+	+
Burkina Faso	–	–	no
Cameroon	no	–	–
CAR 1	–	+	no
Chad	+	no	no
Cote d'Ivoire	no	no	n/a
Djibouti	no	both	n/a
Equatorial Guinea	no	no	–
Gabon	–	no	n/a
Gambia	–	both	n/a
Georgia 1	+	no	+
Ghana	both	no	no
Guinea	no	no	n/a
Kazakhstan	no	no	no
Kenya 1	no	no	–
Kyrgyz Republic	no	+	+
Madagascar	–	+	–
Malawi	+	–	no
Mauritania	no	+	no
Mozambique	no	no	no
Niger	+	no	no
Nigeria	no	no	no
Russia	+	+	+
Senegal 1	no	no	no
Sudan	no	+	no
Tajikistan	no	+	no
Tanzania	both	no	no
Togo	no	–	both
Uganda	no	no	+
Ukraine	no	+	+
Uzbekistan	+	no	n/a
Zambia	no	no	no
Zambia	both	no	no
Zimbabwe	no	both	+

Table 2.7: NELDA Database

Appendix B

Proof of Proposition 2.1. The first-order condition to the maximization problem (2.10):

$$\frac{dG_X(0)}{df} = c'(f),$$

$$\frac{dG_X(0)}{df} = \frac{dG_X(0)}{ds^2}\frac{ds^2}{df} = -\frac{2\psi^2}{f^3}\frac{dG_X(0)}{ds^2},$$

$$\frac{dG_X(0)}{ds^2} = \frac{1}{2\sqrt{2\pi}s^3}\int_{-\infty}^{0} e^{-\frac{(x-\mu)^2}{2s^2}}\,\mathrm{d}x + \frac{1}{2\sqrt{2\pi s^2}}\int_{-\infty}^{0}\frac{d}{ds}\left(e^{-\frac{(x-\mu)^2}{2s^2}}\right)\mathrm{d}x = -\frac{1}{2s^2}G_X(0) + I, \quad (2.18)$$

$$I = \frac{1}{2\sqrt{2\pi s^2}}\int_{-\infty}^{0}\frac{d}{ds}\left(e^{-\frac{(x-\mu)^2}{2s^2}}\right)\mathrm{d}x = \frac{1}{2\sqrt{2\pi s^2}s^2}\int_{-\infty}^{0}\frac{(x-\mu)^2}{2s^2}e^{-\frac{(x-\mu)^2}{2s^2}}\,\mathrm{d}x.$$

Using a substitution $y = \frac{x-\mu}{\sqrt{2s^2}}$:

$$I = \frac{1}{\sqrt{\pi}s^2}\int_{-\infty}^{-\mu/\sqrt{2s^2}} y^2 e^{-y^2}\,\mathrm{d}y. \quad (2.19)$$

Integrating by parts with substitution $U = z$ and $dV = y^2e^{-y^2}dz$:

$$\int_{-\infty}^{a} y^2e^{-y^2}\,\mathrm{d}y = -\frac{1}{2}ae^{-a^2} + \frac{1}{2}\int_{-\infty}^{a} e^{-y^2}\,\mathrm{d}y.$$

Applying this result to the integral (2.19):

$$I = \frac{1}{\sqrt{\pi}s^2}\int_{-\infty}^{-\mu/\sqrt{2s^2}} y^2e^{-y^2}\,\mathrm{d}y = \frac{1}{\sqrt{\pi}s^2}\left(\frac{1}{2}\frac{\mu}{\sqrt{2s^2}}e^{-\frac{\mu^2}{2s^2}} + \frac{1}{2}\int_{-\infty}^{\mu/\sqrt{2s^2}} e^{-y^2}\,\mathrm{d}y\right).$$

Going back to original notation:

$$I = \frac{1}{2s^2}\left(\frac{\mu}{\sqrt{2\pi s^2}}e^{-\frac{\mu^2}{2s^2}} + G_X(0)\right) = \frac{1}{2s^2}\left(\mu g_X(0) + G_X(0)\right).$$

Plugging the last expression to (2.18):

$$\frac{dG_X(0)}{ds^2} = -\frac{1}{2s^2}G_X(0) + I = -\frac{1}{2s^2}G_X(0) + \frac{1}{2s^2}\left(\mu g_X(0) + G_X(0)\right) = \frac{\mu}{2s^2}g_X(0).$$

Combining all the results:

$$\frac{dG_X(0)}{df} = -\frac{2\psi^2}{f^3}\frac{dG_X(0)}{ds^2} = -\frac{2\psi^2}{f^3}\frac{\mu}{2s^2}g_X(0) = -\frac{\psi^2}{f^3}\frac{\mu}{s^2}g_X(0).\square$$

Proof of Proposition 2.2. The incumbent's expected share of votes $\Pi_I = \frac{1}{2} + (1 - \phi\beta_1)f$ is increasing in f if and only if $\beta_1 - 1/\phi < 0$. Otherwise, it is always optimal to choose zero fraud. The left hand side of the first-order condition (2.11) $L(f) = -\frac{2\psi^2}{f^3}\frac{2s^2}{g_X(0)}$ is a strictly decreasing function of fraud whenever the first part of the proposition is satisfied. Because marginal cost $c'(f)$ is an increasing function of fraud, and $L(0) = \frac{1/\phi - \beta_1}{\psi\sqrt{2\pi}} > 0 = c'(0)$, there is a unique intersection between $L(f)$ and $c'(f)$. To show that the intersection point is between 0 and 1, it is sufficient to show that $L(1) < c'(1)$:

$$L(1) = \frac{\psi^2(1/\phi - \beta_1)}{\epsilon_1^2 + \psi^2}\frac{1}{sqrt2\pi(\epsilon_1^2 + \psi^2)}e^{-\frac{1/(\phi-\beta_1)^2}{\epsilon_1^2+\psi^2}} < \frac{1/\phi - \beta_1}{\sqrt{2\pi(\epsilon_1^2 + \psi^2)}}.$$

Thus, to have optimal fraud less than 1, it is sufficient to have $\frac{1/\phi - \beta_1}{\sqrt{2\pi(\epsilon_1^2+\psi^2)}} < c'(1)$. $\square$

Proof of Proposition 2.3. Consider the first-order condition (2.11) and denote $L = -\frac{2\psi^2}{f^3}\frac{2s^2}{g_X(0)} - c'(f)$. Denote the solution for the first-order condition (2.11) as f^* and use the implicit function theorem:

$$\frac{\partial f^*}{\partial \epsilon_0^2} = -\frac{\partial L/\partial \epsilon_0^2}{\partial L/\partial f^*}.$$

Since $\mu < 0$

$$\frac{\partial L}{\partial f^*} = \frac{\psi^2 \mu g_X(0)}{({f^*}^2\epsilon_0^2 + \psi^2)^2}\left(\frac{\psi^2\mu^2}{f^2\epsilon_0^2 + \psi^2} + 3\epsilon_0^2\right) - c''(f) < 0.$$

Further

$$\frac{\partial L}{\partial \epsilon_0^2} = -\frac{\psi^2 \mu g_X(0) f^*}{2({f^*}^2\epsilon_0^2 + \psi^2)^2}\left(\frac{\mu^2}{s^2} - 3\right).$$

The latter expression is negative if and only if $\frac{\mu}{s^2} < 3$. Because $\frac{\mu^2}{s^2} = \frac{\mu^2}{\epsilon_0^2 + \psi^2/{f^*}^2} < \frac{\mu^2}{\epsilon_0^2+\psi^2}$ to guarantee $\frac{\partial L}{\partial \epsilon_0^2} < 0$ it is sufficient to have $\frac{\mu^2}{\epsilon_0^2+\psi^2} < 3$ or $\epsilon_0^2 + \psi^2 > \mu^2/3$, which in turns guarantee $\frac{\partial f^*}{\partial \epsilon_0^2} < 0$. $\square$

Proof of Proposition 2.4. Consider the first-order condition (2.11) and denote $L = -\frac{2\psi^2}{f^3}\frac{2s^2}{g_X(0)} - c'(f)$. Denote the solution for the first-order condition (2.11) as f^* and use the implicit function theorem:

$$\frac{\partial f^*}{\partial \epsilon_0^2} = -\frac{\partial L/\partial \epsilon_0^2}{\partial L/\partial f^*}.$$

Since $\mu < 0$

$$\frac{\partial L}{\partial f^*} = \frac{\psi^2 \mu g_X(0)}{({f^*}^2 \epsilon_0^2 + \psi^2)^2} \left(\frac{\psi^2 \mu^2}{f^2 \epsilon_0^2 + \psi^2} + 3\epsilon_0^2 \right) - c''(f) < 0.$$

Thus, to prove the proposition it is enough to show that $\frac{\partial L}{\partial \mu} < 0$.

$$\frac{\partial L}{\partial \mu} = -\frac{\psi g_X(0)}{f^*({f^*}^2 \epsilon_0^2 + \psi^2)} \left(1 - \frac{\mu^2}{s^2} \right).$$

The latter expression is negative whenever $1 - \frac{\mu^2}{s^2} < 0$ or $\epsilon_0^2 + \psi^2 > \mu^2$.$\square$

References

Acemoglu, D., Robinson, J., 2006. Economic Origins of Dictatorship and Democracy. *New York: Cambridge University Press.*

Akhmedov, A., Zhuravskaya, E., 2004. Opportunistic Political Cycles: Test in a Young Democracy Setting. *The Quarterly Journal of Economics*, November, 119(4), 1301-1338.

Beck, T., Clarke, G., Groff, A., Keefer P., Walsh, P., 2001. New Tools and New Tests in Comparative Political Economy: the Database of Political Institutions. *World Bank Economic Review*, 15, 165-176.

Chaturvedi, A., 2005. Rigging Elections with Violence. *Public Choice*, 125, 189-202.

Collier, P., Vicente, P., 2012. Violence, Bribery, and Fraud: The Political Economy of Elections in Sub-Saharan Africa. *Public Choice*, 153 (1-2), 117-147.

Enikopolov, R., Petrova, M., Zhuravskaya, E., 2009. Media and Political Persuasion: Evidence from Russia. *American Economic Review*, 101(7), 3253-3285.

Glaeser, E.L., Schleifer, A., 2005. The Curley Effect: The Economics of Shaping the Electorate. *The Journal of Law, Economics, & Organization*, 21 (1), 1-19.

Gregory, P., Schroder, P., Sonin, K., 2011. Rational dictators and the killing of innocents: Data from Stalin's archives. *Journal of Comparative Economics*, 39 (1), 34-42.

Guo, G., 2009. China's Local Political Budget Cycles.*American Journal of Political Science*, 53, 621-632.

Hyde, S., Marinov, N., 2012. Which Elections Can Be Lost? *Political Analysis*, 20(2), 191-2010.

Lehoucq, F., 2003. Electoral Fraud: Causes, Types and Consequences. *Annual Review of Political Science*, 6, 233-256.

Levin, I.,Cohn, G., Ordeshook, P., Alvarez, R., 2009. Detecting Voter Fraud in an Electronic Voting Context: An Analysis of the Unlimited Reelection Vote in Venezuela. *Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Online Proceedings.

Lindbeck, A., Weibull, J., 1987. Balanced-Budget Redistribution as the Outcome of Political Competition. *Public Choice*, 52, 273-297.

Martinez, L., 2009. A Theory of Political Cycles. *Journal of Economic Theory*, 144, 1166-1186.

Mebane W., Kalinin, K., 2009. Comparative Election Fraud Detection. *Presented at the Annual Meeting of the APSA.*

Myagkov, M., Ordershook, P., 2008. Russian Elections: An Oxymoron of Democracy. *VTP Working Paper.*

Persson T., Tabellini, G., 2000. Political Economics: Explaining Economic Policy. *MIT Press.*

Rogoff, K., 1990. Equilibrium Political Budget Cycles. *American Economic Review*, 80, 21-36.

Schedler, A., 2002. Elections without Democracy: The Menu of Manipulation. *Journal of Democracy*, 13(2), 36-50.

Schneider, C., 2010. Fighting with One Hand Tied behind the Back: Political Budget Cycles in the West German States. *Public Choice*, 142, 125-150.

Shi, M., Svensson, J., 2006. Political Budget Cycles: Do They Differ Across Countries and Why? *Journal of Public Economics*, 90, 1367-1389.

Simpser, A., 2005. Making Votes Not Count: Strategic Incentives for Electoral Corruption. *Mimeo*, University of Chicago.

Simpser, A., 2008. Cheating Big: On the Logic of Electoral Corruption in Developing Countries. *Mimeo*, University of Chicago.

Sutter, D., 2003. Detecting and Correcting Electoral Fraud. *Eastern Economic Journal*, 29(3), 433-451.

Treisman, D., 2009. Elections in Russia, 1991-2008. *Working Paper*, Moscow State University - Higher School of Economics.

Veiga, L., Veiga, F., 2007. Political Business Cycles at The Municipal Level. *Public Choice*, 131, 45-64.

Chapter 3. Towards Detecting and Measuring Ballot Stuffing

Abstract

This paper proposes a method for detecting electoral fraud in the form of ballot stuffing. As ballot stuffing increases both turnout and the incumbent's vote share in precincts where it occurs, precincts with low reported turnout are more likely to be clean. Information on clean precincts is used to simulate counterfactual data for infected precincts, which are then compared to the observed data. The method is applied to the 2006 Finnish presidential elections. The test fails to reject the hypothesis of no ballot stuffing for the original presumably clean data, but detects artificially imputed fraud and provides a correct estimate of its magnitude. The same method implies that in the presidential elections in Russia held between 2000 and 2012, ballot stuffing was a significant issue, and the number of ballots stuffed in favor of the incumbents had been persistently growing over the period. Regional-level analysis suggests that this is a result of both increasing magnitude of fraud and expansion of electoral falsification across the regions of Russia.

JEL Classification: D72, D73
Keywords: Elections, Fraud Detection

3.1 Introduction

Despite its importance, electoral fraud suffers from a relative lack of attention in the academic literature. Probably the main reason for this is the absence of a reliable measure of fraud. Indeed, not only measuring but even detecting fraud is problematic. The existing methods of fraud detection are more qualitative than quantitative, often based on the subjective assessment of electoral transparency and fairness by observers or other participants of the electoral process, and the results they produce may not always be treated as fully reliable. The few attempts to rigorously analyze electoral data for the presence of fraud have usually required a large amount of data, which handicaps efforts to measure fraud, proxy it, or even detect it with a reasonable degree of confidence. It further precludes implementing reliable empirical research, which in turn discourages efforts towards a theoretical study of the nature and consequences of electoral fraud.

This paper proposes a statistical mechanism for testing the fairness of elections when the available data are limited. The methodology enables elections to be tested for the presence of electoral fraud in the form of ballot stuffing using official detailed electoral data. The mechanism is applied to test the fairness of the Russian presidential elections between 2000 and 2012, whose transparency and integrity are often considered to be in doubt, and to obtain an estimate of the magnitude of ballot stuffing in Russia.

The paper is organized as follows. The next section discusses existing approaches to detecting fraud. Section 3.3 presents a methodology that enables testing of the fairness of elections based on official electoral results. In Section 3.4, the methodology is applied to several datasets. First, I create artificial electoral data, show that the test fails to reject the null hypothesis of no ballot stuffing, then impute fraud of about 2% and test the data once again, resulting in a strong rejection of the null hypothesis. Second, I perform the same exercise for data on the 2006 presidential elections in Finland. The test cannot reject the hypothesis of fair elections for the original data, but rejects the hypothesis, once 2% fraud is imputed. Third, I apply the test to data on the Russian presidential elections 2000, 2004, 2008 and 2012 at both country and regional levels. The hypothesis of no ballot stuffing is strongly rejected in all cases, and the test implies that the number of stolen votes in favor of incumbents had been growing persistently between 2000 and 2012.

3.2 Literature Review

Detecting fraud in data is not new. The general idea underpinning most fraud detection statistical techniques is tracing unusual patterns in the observed data that might be explained by fraud. Such techniques have been successfully used to uncover fraud in a variety of domains, from sports betting (e.g., Wolfers, 2006)) and education (e.g., Jacob and Levitt, 2003) to online auctions (e.g., Pandit et al., 2007) and banking (e.g., Quah and Sriganesh, 2008). A number of recent papers review fraud detection techniques for specific fields such as telecommunications (Becker et al., 2010), health care (Li et al., 2008) and finance (Sudjianto et al., 2010). Though the specific design of fraud detection techniques does depend on the nature of the data and type of expected fraud, all fraud detection methods share enough features to be divided into two main groups: supervised and unsupervised. Methods of the supervised type assume that there are two data samples available for the analysis: the one which is affected by fraud, and the one which is not. In this case, labeling a new data set as clean or fraudulent is essentially a comparison with benchmark samples. When such samples are not available, the unsupervised methods are applied. They do not use benchmark samples and instead look for outliers in an observable sample. Due to the nature of data on elections and frequently changing electoral environments, electoral fraud detection methods have to be of an unsupervised type.

Attempts to detect fraud in the electoral process used to be rare and unsystematic. Lehoucq (2003), in his comprehensive review of studies on electoral fraud, mentions a number of papers that look for traces of electoral fraud in elections in Argentina, Peru, Colombia, England, Ireland, Germany, Spain, Mexico and some Asian countries[21]. The majority of these studies detect fraud using descriptive evidence such as surveys, interviews and documents; none uses statistical methods. Even though such qualitative approaches can provide insight into the presence of electoral fraud in given elections, they require tremendous effort to collect relevant data and may yield results with limited application and replicability.

Due to the limitations of qualitative approaches, researchers have started to pay attention to the statistical analysis of electoral data with the aim of detecting electoral fraud. The largest and most rapidly growing approach to electoral fraud detection is digit analysis, which analyzes digit patterns in electoral data to identify anomalies that may appear due to fraud.

Beber and Scacco (2008) suggest a methodology based on the idea that people are bad random number generators: if elections are fair, the distribution of insignificant digits (e.g., digits at the third decimal place and further) in electoral outcomes (i.e. data on turnouts and vote shares) must be close to uniform, but if there are manual changes in outcomes there must be biases in generating digits. The idea is supported by a statistical comparison of outcomes from Swedish and Nigerian elections. However, such a method is limited to detecting manipulations with electoral returns; it is unlikely to produce a result if electoral outcomes are shaped in a more sophisticated way than manually changing digits in election protocols; it does not provide any estimate of the magnitude of fraud.

In contrast to Beber and Scacco (2008), a number of recent papers have analyzed the first significant digits in official electoral data to find deviations from Benford's Law (Benford, 1938).

[21]For the references see Lehoucq (2003).

The law states that the first digits in real data are distributed in a specific non-uniform way. Deviations from the law are found by Roukema (2009) in data on the last Iranian presidential elections, Cantu and Saiegh (2010) in Argentinean elections, Pericchi and Torres (2004) in Venezuela, Mebane (2006, 2010) in the US and Mexico, Mebane and Kalinin (2009) in Russia, and by Breunig and Goerres (2011) in the Bundestag elections in Germany. Despite their merits, these methods of digit analysis are subject to criticism, which casts doubt on their relevance for detecting fraud in electoral data (e.g., Brady, 2005). Recently Deckert et al. (2011) have shown that deviation from Benford's Law can arise in electoral data regardless of whether the elections are rigged or fair, and that the methods essentially do not differ from a random draw in their ability to mark elections as clean or fraudulent.

A number of authors have suggested alternative methods for discovering fraud in official electoral data. Myagkov and Ordeshook (2008) study the fairness of Russian federal elections between 1993-2007, by examining a variety of patterns in Russian electoral data such as turnout distributions across regions and precincts, and vote flows between different elections. They conclude that ballot stuffing as well as other fraud techniques of the 1990s were frequently used in a few Russian ethnic republics, and then spread to other regions of the country. They apply similar techniques to several elections in Russia and Ukraine looking for the presence of fraud (see Myagkov et al, 2009, for a detailed discussion).

Electoral fraud in Russia is also discussed by Treisman (2009), who reviews the trends in voting in Russia since 1991. In a chapter devoted to electoral manipulations and fraud, by studying a variety of Russian electoral statistics in a way similar to Myagkov and Ordeshook (2008) and Mebane and Kalinin (2009), the author finds that in the early 1990s, the elections in Russia were nearly clean, whereas since 2000 electoral irregularities have become an integral part of electoral competition. Although Treisman's approaches are reasonable and capable of producing reliable conclusions, they are mainly based on a visual analysis and comparison of electoral data.

More recently, Levin et al (2009) have elaborated on the approaches by Myagkov and Ordeshook (2008) and explore electoral data in Venezuela for the presence of fraud. They analyze the data on two consecutive state-level referenda in 2007 and 2009 and, assuming time-constant voters' preferences, discover unusual patterns in the voting behavior of selected regions that mainly benefit the incumbent. Specifically, most of the new votes in favor of Chavez in the 2009 referendum came from regions with large abstention in 2007. To obtain this result, Levin et al. explore three types of indicators. First, they perform digit analysis of the electoral outcomes. Second, they study the flow of votes between the two elections by estimating the proportion of vote share in the first elections that "flows" to each alternative in the second elections in order to see whether there is a noticeable increase in support of one of the alternatives in regions with a substantial increase in turnout. Finally, they look closely at the relationship between the time changes in turnout and share of votes cast for the incumbent. As in Myagkov and Ordeshook (2008), Levin et al's analysis is primarily based on a comprehensive investigation and description of data patterns rather than statistical testing.

In short, even though electoral fraud appears to be very widespread, existing means of detecting fraud are primarily descriptive and qualitative and highly dependent on the nature of available data. Though statistical studies of electoral data with focus on fraud do exist, they are mainly focused on exploring unusual patterns in data and require a tremendous amount of information. Thus, there is still a need for a rigorous method to detect fraud and measure its extent, especially for cases when available data are restricted. This paper attempts to make progress towards designing such a method.

3.3 Fraud Detecting Methodology

The suggested approach is based on the observation that ballot stuffing increases both turnout and votes cast for the corrupt candidate (hereafter I assume that fraud is implemented in favor of the incumbent). If elections are subject to ballot stuffing, which takes place not in all but in a selected number of precincts, this observation immediately implies that a precinct with lower reported turnout is more likely to be clean. The general idea behind the methodology described below is to use the information from such low turnout precincts to simulate counterfactual data for high turnout and likely fraudulent precincts, and test for systematic difference between counterfactual and observed data.

The suggested procedure allows for testing for the presence of ballot stuffing even when very limited data are available. Suppose that the only data available for the analysis are precinct level turnouts and candidates' vote shares. If elections are not fraudulent, and there is a certain degree of homogeneity between electoral precincts, the distribution of turnout across precincts should be close to bell-shaped. Clearly, if electoral districts are similar in terms of characteristics that might determine turnout, or alternatively if turnout is weakly affected by characteristics in which the districts differ, then distribution of turnout should be approximately normal (see, for example, Myagkov and Ordeshook (2008) or Levin et al. (2009) for more detailed discussion).

In turn, ballot stuffing, when it takes place in a given precinct, results in an increase both in reported turnout and in the incumbent's vote share. Consequently, such a fraudulent precinct moves in turnout distribution towards its right tail. As a result, distribution of turnout in fraudulent elections will be skewed to the left and have a thicker right tail. Furthermore, in a precinct in the right distribution tail, i.e. those with high reported turnout, an incumbent will have an advantage in comparison to the other precincts. The idea behind the suggested methodology is to check whether the incumbent has such an advantage and whether it could be considered natural.

Suppose the following statistic is computed from the available data:

$$s = \frac{V_I/V_R|t \geq t^*}{V_I/V_R|t < t^*}. \tag{3.1}$$

The nominator of the statistic is the ratio of the incumbent's and the runner-up's shares in the right tail of the turnout distribution (i.e., in those precincts where the turnout is above some threshold t^*). The denominator is the same ratio, but computed over the left distribution tail. I use this statistic to test the null hypothesis that elections are fair. Under the null hypothesis, the statistic should be close to one if there is no objective systematic relationship between turnout and voting in favor of one or another candidate, meaning that if elections are fair, the ratio of the incumbent's and the runner-up's shares should not systematically differ in the precincts with high and low turnout. But if there is ballot stuffing in favor of the incumbent, the incumbent's share in the precincts with high turnout will be relatively higher, meaning that the statistic will be greater than one (less than one, if ballot stuffing is in favor of the challenger).

Indeed, there could be an objective correlation between the turnout and the vote shares for the candidates if the supporters of one candidate are more politically active than the supporters of the others. In this case this correlation will be present over the whole dataset, including the left tail. In other words, if such a natural relationship between turnout and vote shares exists, it can be estimated using left tail clean data only, and fraud, if it exists, will make this correlation higher in the right tail. The procedure described below is designed to test not for the presence of correlation between turnout and incumbent's vote share in the right tail data, but rather for the presence of extra correlation in comparison to the left tail.

Using the statistic, I test the null hypothesis that there is no ballot stuffing. To conduct the test, one needs to know only a distribution of the statistic under H0. To obtain such distribution the following procedure is proposed.

First, recall that ballot box stuffing increases both turnout and the incumbent's share of votes. This means that a fraudulent precinct moves to the right tail of the turnout distribution. This in turn means that in the case of rigged elections, the left tail of the observed turnout distribution contains a lower number of fraudulent precincts than the right one. Moreover, the higher the scale of fraud, the further to the right a fraudulent precinct moves, implying that a larger tail of the distribution remains clean. Second, recall that if there is no ballot stuffing and precincts are in some sense homogenous, turnout distribution across precincts should be approximately normal. Assuming some particular shape of the true turnout distribution (for example, normal), I choose its parameters such that the distribution fits the left-tail data, i.e.

those with turnout below some threshold value t^* (I discuss the choice of the threshold below). Next, I estimate the relationships between turnout and vote shares in the clean left tail by simply regressing vote shares on turnout:

$$V_{Ii} = \alpha + \beta t_i + \epsilon_i. \tag{3.2}$$

Note that the purpose of these regressions is not to establish a causal effect of turnout on vote shares, but rather to find a correlation and then extrapolate it on the simulated right tail.

Then I repeat the following simulation multiple times. At each simulation step I generate a new turnout distribution across precincts $\bar{t_i}$ as a random draw from the fitted normal distribution and predict vote shares. To make predicted vote shares consistent with clean left tail data, I first maintain the same relationships between the vote shares and turnout as in the observed clean left tail, and, second, introduce additional noise into predicted vote shares such that their variances evaluated over the left tail are the same as the variances of observed left tail vote shares. Specifically, vote shares for the incumbent and challenger (runner-up) are predicted as

$$\bar{V_{Ii}} = \tilde{\alpha} + \tilde{\beta}\bar{t_i} + u_i. \tag{3.3}$$

where $\tilde{\alpha}$ and $\tilde{\beta}$ are random draws from normal distributions with means α and β, and standard deviations equal to corresponding standard errors from regression 3.2. The latter means that when predicting the vote share, I do not just use coefficients obtained from regression 3.2, but allow them to vary across simulations according to the precision of the estimation. If a predicted vote share exceeds 1, it is equalized to 1. Errors u_i are drawn from a zero mean normal distribution. Variance of this distribution is chosen such that

$$Var(\bar{V_{Ii}})|t < t^* = Var(V_{Ii})|t < t^*. \tag{3.4}$$

By allowing coefficients to vary and making variance of predicted vote shares to be the same as variance of actual vote shares, I guarantee that the simulated right tail data are consistent with observed relatively clean left tail data.

Once the vote shares are predicted, a statistic (3.1) can be computed. Repeating this simulation multiple times and computing the statistic on each step, one can obtain a distribution of the statistic under the null hypothesis that elections are fair, and tabulate critical values. Having computed the statistic for observed data, one can now test the hypothesis. Note that if the value of the statistic appears to be above the right tail critical value, it implies that ballot stuffing took place in favor of the incumbent. In contrast, a statistic below the left tail critical value signals ballot stuffing in favor of the challenger. Given critical and actual values of the statistic it is now easy to obtain an estimate of ballot stuffing magnitude by calculating the vote share that incumbent has to obtain in the right tail precincts in order to equalize the observed value of the statistic to its critical value of desired confidence level. The difference between actual vote share and counterfactual vote share calculated in this way would give an estimate of the magnitude of ballot stuffing.

When there are more than two candidates in elections the procedure is slightly different. Because turnout can be related differently to vote shares of each candidate and there are several candidates, a challenger's vote share cannot be predicted by simply subtracting the incumbent's vote share from one. Instead, his vote share should be obtained in a similar way as incumbent's one (formulas (3.2) and (3.3)).

The ballot stuffing detection procedure is based on a number of implicit assumptions. First, I assume that ballot stuffing occurs in a small number of precincts. Suppose instead that ballot stuffing of relatively the same magnitude would occur in all precincts. This means that turnout and share of votes cast for the incumbent increase in all precincts. Thus, there will be no systematic difference between left and right tail data due to fraud, which is needed for identification. Second, the fraud should be of reasonable magnitude in that it should result in a noticeable increase in turnout to move the precinct to the right tail of the turnout distribution. Together these two assumptions say that fraud, in order to be detected, should move the precinct where it occurred to the right tail of the distribution.

The methodology described above explicitly distinguishes between the left tail and right tails of turnout distribution by using a turnout threshold t^*. Ideally, t^* should be chosen such that all precincts with turnout below t^* are clean and the lowest reported turnout among the fraudulent precinct is slightly above t^*. In practice such a choice is challenging. More likely, there will still be some fraudulent precincts even in the left turnout distribution tail, but less than in the right one, meaning that fraud detection is still possible though the fraud magnitude will be underestimated in this case. On the one hand, the low value of the threshold allows for the capture of small-scale fraud and fraud in low turnout precincts, since they are more likely to appear above the threshold. On the other hand, low t^* will not allow detection of even large-scale fraud if it appears in a very small number of precincts, as the contribution of the fraudulent precincts in the statistic will be relatively small due to a large number of clean precincts. Also, low t^* will result in a small number of data points in the left tail, which are used for estimation of the natural relationship between turnout and vote shares, which will decrease the power of the test. On the other hand, a high threshold value would make it easier to reject the null hypothesis of no ballot stuffing if there is large-scale fraud, but could fail to reject the hypothesis when fraud is balanced. Thus, the choice of threshold generally depends on the data as well as some prior information about the nature and the extent of fraud.

One way to endogenize the choice of the turnout threshold is to analyze the values of coefficient on turnout from regression (3.2) for different thresholds. Clearly, going over different threshold values from low to high, at some point the left tail data would include fraudulent precincts. As a result, coefficient β from regression (3.2) will start growing if fraud is in favor of an incumbent and decreasing if fraud is in favor of a challenger. Thus, the value of turnout threshold at which the coefficient of turnout starts growing (decreasing) would be a natural choice for t^*. The value of t^* specified in this way can itself send signals about the nature of fraud. If t^* appears to be high, that would mean that in each fraudulent region the magnitude of ballot stuffing was huge as a substantial share of tainted precincts ended up in the very right tail of the reported turnout distribution. Alternatively, reasonably low t^* means that fraud in a given precinct was not extreme, though the total number of fraudulent regions could still be substantial. Indeed, there could be cases when sharp changes in the value of β are not observed at all (for instance when elections are clean). In this case, the only way to define t^* is to make some reasonable, yet ad-hoc choice, for instance, some number between 0.5 and 0.8.

It is important to notice that the particular choice of threshold value can affect the corresponding estimate of ballot stuffing magnitude. Recall that the estimate described above is the difference between the actual number of votes for the incumbent and the number of votes that incumbent should have received in order not to reject the hypothesis of no ballot stuffing at the desired confidence level. Since a higher value of threshold effectively means that a higher fraction of data is considered clean, and thus a lower share of data is considered potentially fraudulent, the fraud estimate will generally be a decreasing function of threshold. Thus, this estimate should be thought of as a lower bound of ballot stuffing rather than its measure. Though, generally, direct comparison of such estimates across different elections would not be entirely correct, under certain circumstances it might be still useful for getting an idea aboutthe relative extent of ballot stuffing. See Section 3.5 for further discussion.

Noticeable growth or decline of coefficient β starting from some turnout threshold value itself signals the existence of fraud as in the absence of ballot stuffing the coefficient should not change sharply. Thus, testing for the broken trend in β as a function of t^* could be another, simpler way to test for the presence of ballot stuffing. Alternatively, because the suggested methodology is based on the observation that turnout in clean elections should follow approximately a normal distribution, one can simply test for the distribution symmetry. However, the suggested methodology has a number of advantages over these alternatives. First, the method would indicate the direction of ballot stuffing if it exists. Depending on whether the observed value of the statistic falls to the left or right tail of its distribution under the null hypothesis, one can always say whether ballot stuffing is in favor of the incumbent or the challenger. Second and most important, in contrast to the symmetry and broken trend tests, the suggested method provides some information on the magnitude of ballot stuffing.

Finally, it is important to note that ballot stuffing is just one technique for rigging elections, while the whole range of potential techniques is wide (Lehoucq, 2003). As a result, the suggested methodology tests for the presence of and provides an estimate of just one particular fraudulent activity, which however is very widespread and popular, and which accounts for a substantial share of voting fraud due to its obvious cost effectiveness. Moreover, technically the suggested methodology is intended to detect any activity that leads to simultaneous increase in reported turnout and vote share of a candidate. Ballot stuffing is not the only rigging technique that leads to this. Such activities as multiple voting and vote buying also result in increase of turnout and an incumbent's vote share in a regions where they occur, and thus the suggested methodology is fully appropriate for tracing out their steps in electoral data.

3.4 Testing Fairness of Elections

In this section I apply the described methodology of detecting ballot stuffing to several distinct datasets. I first generate artificial clean electoral data and then impute fraud into them. I apply the test to the original clean data and then to the fraudulent data to show that the test raises a red flag in case of fraudulent data only. Then, I apply the methodology to real data from the 2006 Finnish presidential elections. As the integrity of these elections was never subject to debate, I consider them an example of presumably clean elections and show that the test cannot reject a null hypothesis of no ballot stuffing for this data, but it does reject the null hypothesis once fraud is artificially imputed.

3.4.1 Artificial Data

First, I show that the method is capable of detecting electoral fraud of a reasonable magnitude in artificial data. For this purpose, I create a dataset that consists of turnout and candidates' vote shares. Specifically, I generate 1000 observations for turnout t that follow a normal distribution with 0.5 mean and 0.1 standard deviation. Each observation represents data for a precinct. Then I generate vote shares for the incumbent allowing for a natural correlation between vote shares and turnout as well as noise drawn from normal distribution $N(0, 0.05)$:

$$V_{Ii} = 0.05t_i + \epsilon_i. \tag{3.5}$$

I then apply the methodology described above to the simulated data. First, I need to choose a threshold value of turnout to define left and right tails. There is no clear trend break in the coefficient on turnout from regression (3.1) as a function of turnout percentile where the threshold value is evaluated, and the variation in the coefficient is not substantial (See Figure

3.2). I choose threshold value at the 61st percentile, which implies that the left tail contains precincts with turnout below 0.556.

Then I estimate the relationship between the incumbent's vote share in the left tail of the turnout distribution (i.e., in precincts where turnout is less than 0.556) by running a regression of V_I on t. Next, I choose the parameters of the normal distribution such that it fits the left tail of the observed turnout distribution as precisely as possible. Having the regression coefficients, their standard errors and turnout distribution parameters, I predict incumbent vote share in the right tail of the turnout distribution (i.e., in precincts where $t \geq 0.556$) allowing for variation in coefficients (coefficients for prediction are randomly drawn from distributions consistent with the estimated coefficient means and standard errors) and noise. Noise is added so that variances of the incumbent's predicted and original vote shares are the same for precincts with $t < 0.556$.

Once the right tail data are constructed, the test statistic is calculated. Then I repeat the procedure of the right tail prediction 5000 times, calculate the statistic on each simulation, get the distribution of the statistic under the null hypothesis of no ballot stuffing, tabulate critical values and compare them to the value of the statistic from the observed data. The value of the statistic from the observed data is 1.059, while the 90% critical value is 1.093 and the 10% critical value is 0.973. Thus, the test cannot reject the null hypothesis.

Figure 3.1: True turnout distribution and distribution of the statistic under the null hypothesis for the original data.

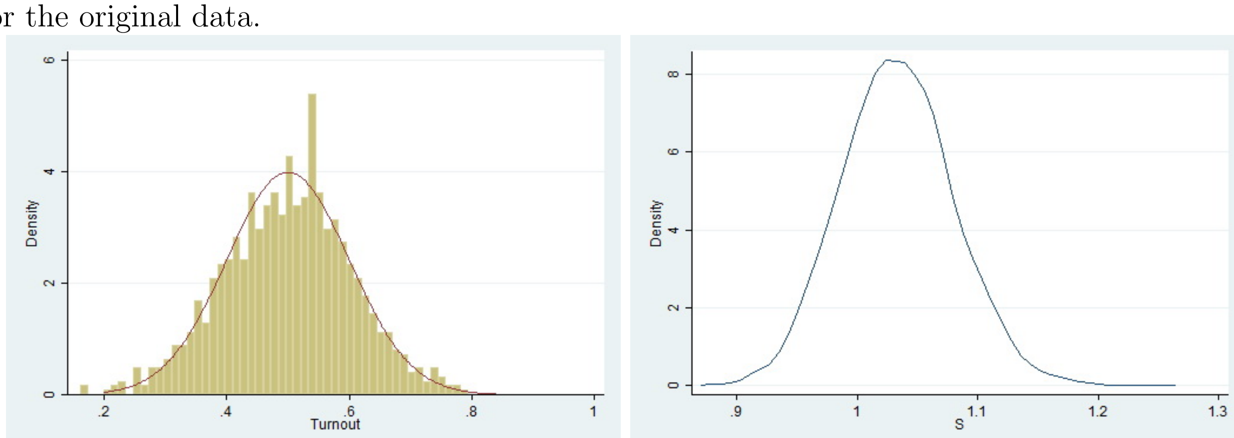

Then, I impute fraud in the data. I randomly choose 150 precincts (15%). In each of them I give the incumbent additional votes: in every spoiled region, the incumbent receives an additional number of votes f_i proportional to the size of the district E_i. Then, if we denote $\theta = f_i/E_i$ where E is the size of the district, after fraud turnout $\hat{t_i}$ and incumbent's vote share $\hat{V_{Ii}}$ can be expressed in terms of before fraud turnout t_i and the vote share V_{Ii} as follows:

$$\hat{t_i} = t_i + \theta. \tag{3.6}$$

$$\hat{V_{Ii}} = \frac{V_{Ii}t_i + \theta}{t_i + \theta}. \tag{3.7}$$

Indeed, the higher θ, the easier it is for the test to reject the hypothesis of no ballot stuffing. Thus, I choose the smallest value of θ such that the null hypothesis is rejected at the 99% confidence level. To guarantee 99% confidence rejection, θ should approximately be 0.18, which gives the incumbent about extra 1.9% of fraudulent votes on an aggregate level measured as the difference between his before and after fraud vote shares. Figure 3.3 shows after fraud distributions of turnout. One can see that fraud results in a thicker right distribution tail.

Once fraudulent data are generated I apply the detecting procedure described in Section 3.3. For this test threshold value is chosen at 55.5th percentile, where according to Figure 3.2 the coefficient on turnout starts to grow persistently.

Figure 3.2: Coefficient on turnout as a function of threshold value for clean artificial data (dashed line) and data with imputed fraud (solid line).

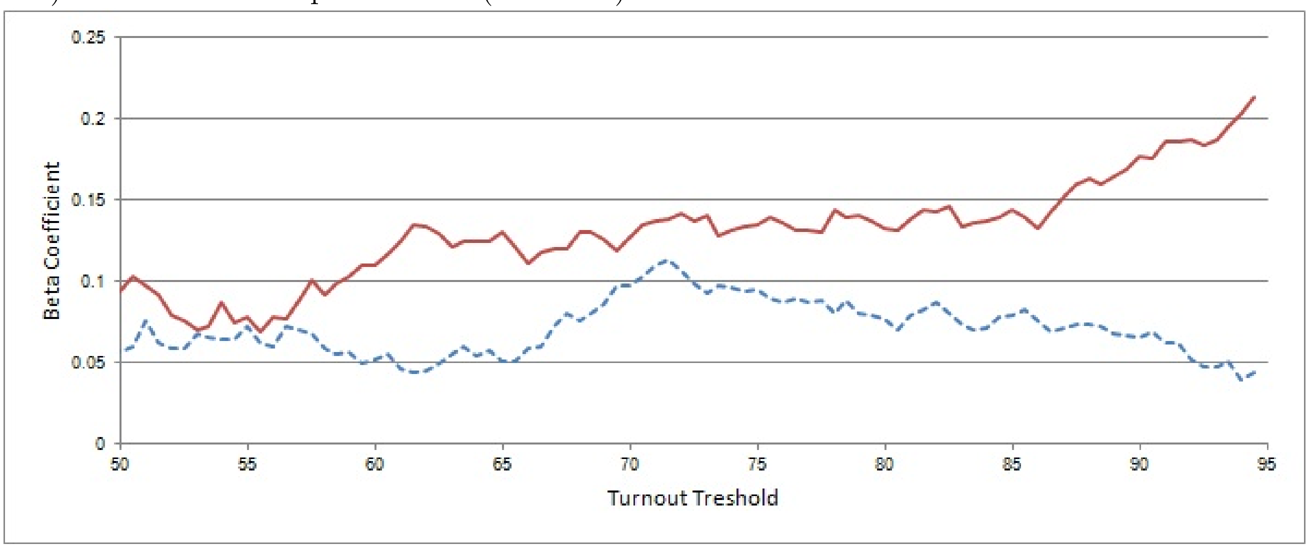

With $\theta = 0.18$, the value of the statistic is 1.215 and the 99% critical value is 1.185. Figure 3.3 presents the distribution of the statistic under the null hypothesis of no ballot stuffing. One can see that this distribution is not exactly the same as the one in Figure 3.1. This is because the distribution on Figure 3.3 is obtained using the after fraud data, and, as was discussed above, the left tail of turnout distribution can still contain fraudulent precincts, which fully account for the observed difference.

Figure 3.3: True turnout distribution and distribution of the statistic under the null hypothesis for fraudulent data.

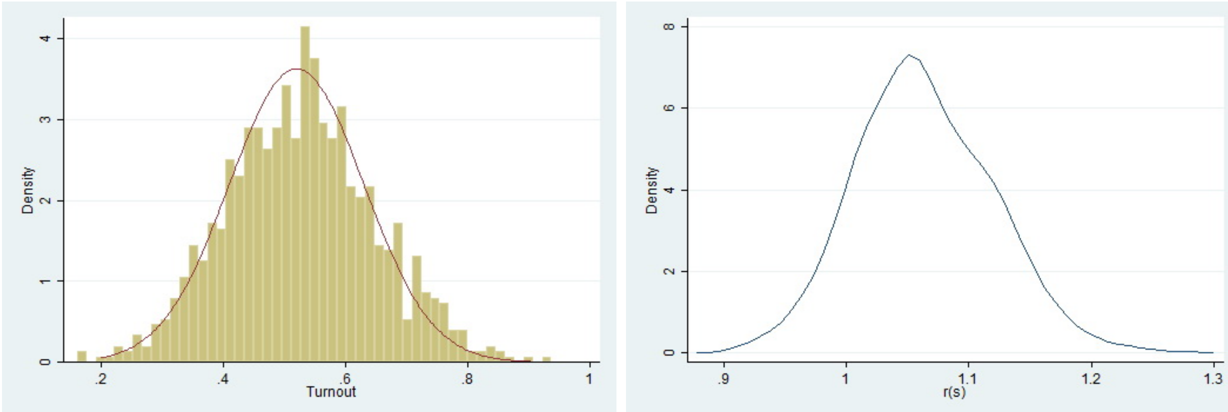

The exercise shows that the suggested methodology does not reject the hypothesis of fair elections even if there is a natural positive correlation between turnout and share of votes for the incumbent, but successfully detects with 99% confidence less than 2% fraud.

3.4.2 Finnish 2006 Presidential Elections

In this section the methodology is applied to real data that came from a presumably clean first round of the 2006 presidential elections in Finland. These elections were chosen as an example of direct executive elections whose integrity can be hardly put in doubt[22]

Another reason why Finnish presidential elections were chosen to test the ballot stuffing detection methodology is that these elections were in some sense close to the Russian presidential elections, which are extensively analyzed in the next section. Though it is hard to believe that Russian and Finnish elections are truly comparable in any dimension, this is probably the best match one could do: the dates of the elections were not too far apart (I analyze Russian presidential elections of 2000, 2004, 2008 and 2012), the electoral systems in both countries are close, and the importance of elections is in some sense similar: directly elected presidents of Russia and Finland both have an executive power in contrast to the majority of European countries. In fact, only few countries in Europe have a directly elected president as an executive

[22]OSCE Office for Democratic Institutions and Human Rights (ODIHR), the largest and probably the most experienced organization that deploys elections observation missions in Europe, was requested to observe the Finnish Parliamentary elections of 2007 held one year after the presidential race. In their report OSCE analysts recommended that no OSCE/ODIHR election observation or assessment activity shall be undertaken in connection with the 18 March 2007 parliamentary elections. A tradition of democratic elections in Finland is accompanied by a commensurate level of public trust. All interlocutors expressed their overall confidence in the electoral process, and no immediate issues were brought to the attention of the Needs Assessment Mission that would necessitate OSCE/ODIHR involvement. Republic of Finland. Parliamentary Elections 18 March 2007, OSCE/ODIHR Needs Assessment Mission Report. Page 4. Available at www.osce.org/odihr/elections/finland/24126 (retrieved 01.10.2012).

(Armenia, Azerbaijan, Belarus, Bulgaria, Cyprus, Finland, France, Georgia, Lithuania, Poland, Romania, Slovakia, Ukraine) and Finland seems to be the best choice from this sample if one would like to have an example of a country which has as many as possible similarities with Russia in terms of electoral environment and, what is the most important for this paper, the highest confidence in electoral transparency and integrity.

For the analysis I use the data from the first round of the elections. The main reason for such a choice is again an intention to make Finnish elections as comparable as possible to the Russian elections analyzed later. Because since 1996, the second round in Russian presidential elections has never been held due to one of the candidates winning in the first round, only first round Russian data are available for the analysis. Thus, it necessary to also analyze first round data for Finland, as the ballot stuffing detection procedure for more than two candidate elections slightly differs from the one applied to the artificial data in the previous section. As discussed in Section 3.2, in such case the method requires analysis of the correlation between turnout and vote shares not only for the incumbent but also for the challenger.

The dataset consists of 461 municipality-level (lowest available level) data observations which came from the Finnish public authority Statistics Finland[23]. To perform the analysis I first need to choose turnout threshold value. Following the approach suggested in Section 3.3, I draw the coefficient on turnout from regression (3.2) as a function of threshold value. It can be seen from Figure 3.5 that there is no clear break in the trend, which would suggest a choice of the threshold. So, I choose two different threshold values at the 56.5th percentile where there is small growth of the coefficient and at the 88th percentile where one can see a small decline in the graph.

In both cases, the test does not reject the hypothesis of no ballot stuffing: with turnout threshold value chosen at the 56.5th percentile the statistic is 0.685, while the 90% and the 10% critical values are 0.801 and 0.467 respectively, and for the 88th percentile threshold the value of statistic is 0.601, while the 90% and the 10% critical values are 0.733 and 0.422 respectively. Note that a critical value below 1 suggests that natural correlation between turnout and votes for the incumbent is lower than between turnout and votes for runner-up.

[23]www.tilastokeskus.fi

Figure 3.4: Turnout distribution and distribution of the statistic under H0 for the original data and 56.5 threshold.

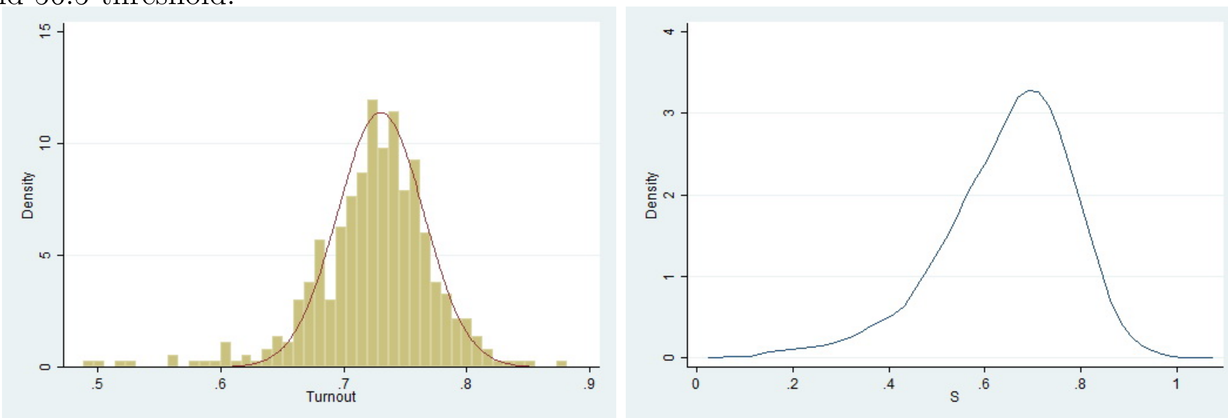

Then, as in the previous case, I randomly choose 69 precincts (15%), artificially spoil them by adding additional 20% of votes in favor of the incumbent which gives him approximately extra 2.5% of votes, and conduct the test once again. Now the threshold value is chosen at the 81th percentile. With these 2.5% fraudulent data, the test rejects the null hypothesis of fair elections at the 99% confidence level: statistic 1.018, 99% critical value 0.928.

Figure 3.5: Coefficient on turnout as a function of threshold value for clean Finnish data (dashed line) and data with imputed fraud (solid line).

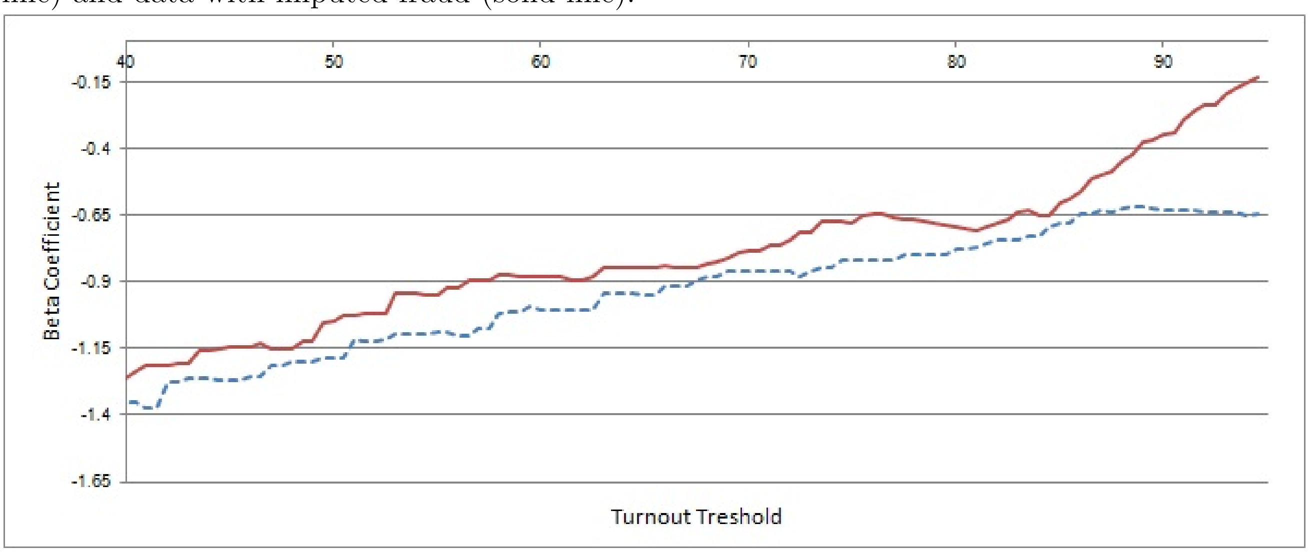

Figure 3.6: After fraud turnout distribution and the distribution of the statistic under H0 for fraudulent data.

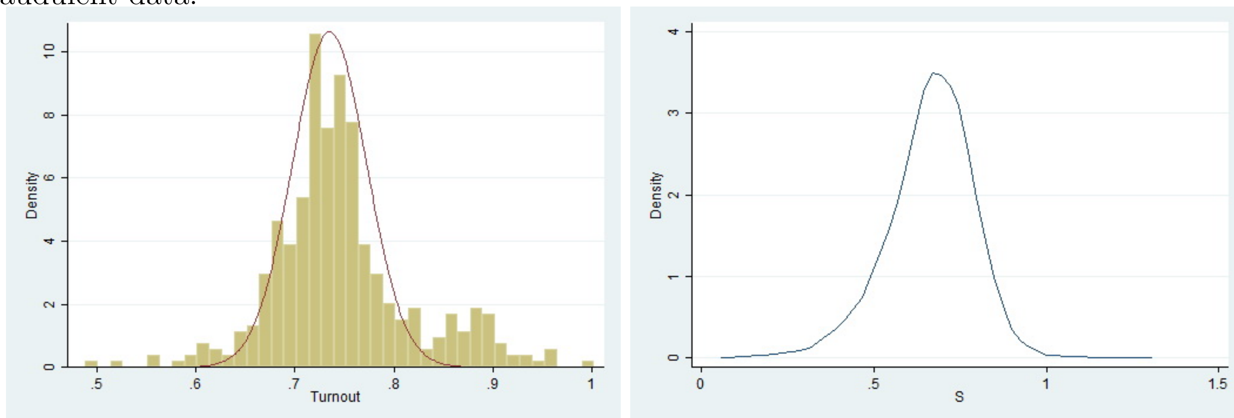

Since fraud was artificially imputed in the data, it is easy to test the ability of methodology to provide an estimate of the ballot stuffing magnitude or, as discussed in Section 3.3, an estimate of the lower bound of ballot stuffing. Recall that the estimate is obtained as the difference between actual number of votes for the incumbent and the number of votes that incumbent should have received just in order not to reject the hypothesis of no ballot stuffing at the desired confidence level.

The imputed 20% of votes in 69 precincts corresponds in this exercise to approximately 100 200 extra votes for the incumbent. These votes give the incumbent a 24.74% reported victory margin, while without fraud the incumbent would win at 22.24% margin. With turnout threshold value at 81th percentile, the methodology provides the following estimates for various confidence levels:

Confidence Level	Counterfactual Victory Margin	Stuffed Votes	Underestimated Fraud
90	22.77%	79 500	20 700 (20.6%)
95	23.22%	62 000	38 200 (38.1%)
99	24.14%	24 600	75 600 (75.4%)

Table 3.1: Estimates of Ballot Stuffing Magnitude

According to Table 3.1 with 99% confidence the number of ballots stuffed in favor of the incumbent is at least 24, 600, and victory margin should not exceed 24.14%, while the observed margin is 24.74%. For 90% confidence level these numbers are 79, 500 and 22.77% respectively, which is quite close to the actual number of stuffed ballots (100 200) and the true margin (22.24%). Given that the methodology performs well on artificial and artificially fraudulent data in terms of both testing data for presence of ballot stuffing and estimating ballot stuffing magnitude, the next natural step is to apply it to real presumably fraudulent data.

3.5 Russian Presidential Elections

In this section I apply the methodology to the Russian presidential elections of 2000, 2004, 2008 and 2012. The fairness and transparency of Russian elections are often questioned, and evidence of electoral misconduct regularly appear in academic research (see, for example, Treisman (2009), Myagkov and Ordeshook (2008), Sakwa (2005)), reports of international observers[24], press, etc.

In all the four election the incumbents won with an overwhelming advantage in the first round by receiving more than 50% of the votes. The officially reported results of the elections are summarized in the following table.

Year	Turnout	Winner's Votes	Runner-Up's Votes	Victory Margin
2000	68.64%	39 740 467 (52.99%)	21 928 468 (29.24%)	23.75%
2004	64.38%	49 565 238 (71.31%)	9 513 313 (13.69%)	57.62%
2008	69.70%	52 530 712 (70.28%)	13 243 550 (17.72%)	52.56%
2012	65.34%	45 602 075 (63.60%)	12 318 353 (17.18%)	46.42%

Table 3.2: Official Results of the Russian Presidential Elections of 2000-2012.

Using the described fraud detection methodology I analyze these four consecutive presidential elections in Russia in order to test for the presence of ballot stuffing and to obtain some comparable measures of its magnitude.

3.5.1 Country-Level Analysis

First, I follow the same approach as in the previous sections to analyze polling station level data obtained from the central elections commission of Russia, that contain information about the number of registered voters, turnout and votes cast for candidates at each polling station in the 2000, 2004, 2008 and 2012 elections. The data contain approximately 95,000 observations for each elections.

The results of the analysis are presented in Table 3.3 and provide clear evidence of persistent growth in ballot stuffing in Russian elections between 2000 and 2012. The estimates of ballot stuffing are provided for the 95% and 99% confidence levels. The percentages in the columns "Ballot Stuffing" are the differences between officially reported incumbent's victory margin and counterfactual incumbent's victory margin, corrected for ballot stuffing of the corresponding level of confidence. One may notice that the difference between the 95% and 99% confidence estimates are very small, while in the example of Finland (see Table 3.1) it was relatively large. The main reason for this is the size of the data used for the analysis: large Russian data allow for more precise analysis. Finally, it is important to note that though the estimates of the absolute

[24]See, for example, reports of OSCE on various federal-level elections in Russia: www.osce.org/odihr/elections/russia (retrieved 20.03.2013)

numbers of stuffed ballots are substantial, ballot stuffing was not pivotal for determining the outcome of any of these four elections: without these ballots the incumbent would still win with an overwhelming advantage in the first round. Even though the suggested methodology underestimates the true magnitude of fraud (see Section 3.3 for the discussion), underestimated fraud is unlikely to be large enough to alter any of the outcomes.

Year	Threshold	Statistics	95%		99%	
			Critical Value	Ballot Stuffing	Critical Value	Ballot Stuffing
2000	44.5	1.084	1.048	0.58% (559 600)	1.055	0.47% (455 500)
2004	73.5	2.170	0.945	3.98% (5 927 700)	0.948	3.96% (5 355 600)
2008	71.0	2.379	1.221	4.35% (6 236 200)	1.226	4.33% (6 204 800)
2012	40.0	1.525	0.937	9.14% (10 258 700)	0.941	9.08% (10 295 200)

Table 3.3: Russian presidential elections 2000-2012.

Since I want to not only test the data for the presence of fraud but also credibly compare its magnitude across these four elections, I perform an additional analysis in which I use the same threshold value for all four tests - 60th percentile value of the turnout distribution. As discussed in Section 3.3, the particular choice of threshold value affects the estimate of ballot stuffing magnitude, and thus comparison of the estimates across different elections obtained with different thresholds may not be fully correct. In fact, this estimate should be thought of as a lower bound of the magnitude of ballot stuffing, and thus it may still provide some information about the relative fairness of different elections. Given that the threshold value is chosen at the level of 60th percentile for all four elections, a higher estimate of the lower bound would signal, though imperfectly, higher magnitude of ballot stuffing. The results of the analysis are summarized in Table 3.4.

Year	Statistics	95%		99%	
		Critical Value	Ballot Stuffing	Critical Value	Ballot Stuffing
2000	1.184	1.011	1.50% (1 433 600)	1.014	1.47% (1 408 900)
2004	1.614	0.947	4.28% (6 343 800)	0.951	4.26% (6 311 800)
2008	2.128	1.189	5.89% (8 199 100)	1.194	5.86% (8 156 900)
2012	1.983	1.019	7.34% (8 562 500)	1.022	7.31% (8 538 300)

Table 3.4: Russian presidential elections 2000-2012. Common threshold.

The results of the analysis with common threshold are fully consistent with the results presented in Table 3.3 though the numbers differ slightly: the magnitude of ballot stuffing has been growing persistently between 2000 and 2012. Again, the estimated ballot stuffing was not pivotal: even without it the incumbent would have won the first round with an overwhelming advantage in all four elections.

3.5.2 Regional-Level Analysis

As discussed in Section 3.3, one of the crucial assumptions that underlies the described fraud detection approach is a certain degree of homogeneity between electoral precincts. When there are systematic differences in voters' behavior across precincts, the method might produce not fully correct conclusions. The case of Russia is an example of such situation. Russian regions are very different in various aspects (economic and social conditions, demography, cultural and historical peculiarities, etc) which might affect voters' behavior. If the ballot stuffing detection methodology is applied to a dataset that contains electoral information on precincts with very distinct voting patterns, then the results and conclusions regarding fairness of the elections and especially estimates of the magnitude of fraud might be biased.

One way to deal with this issue is to split the data set on more homogenous subsets, apply the method to each subset separately, obtain fraud estimates for all the subsets and then aggregate them. In the case of Russia, splitting the country-level dataset on subsets by regions seems to be the most natural approach. Indeed, dividing data on a large number of relatively homogenous subsets requires original dataset to be sufficiently large, and such detailed data are not always available. Fortunately, the Russian central election commission provides such data: the lowest level available datasets (polling station level) contain about 95,000 observations for each elections, and for the regions the number of observations varies between 700 and 3000, with several exceptions for very small regions, which is sufficient to perform the analysis.

I perform the analysis separately for each region for the elections of 2000, 2004 and 2008. In 2008 there were 83 regions in Russia. Complete analysis for all three elections was possible for 62 out of 83 regions. For several small regions there were too few observations (e.g. Chukotskiy and Yamalo-Nenetskiy autonomous discticts, the republic of Yakutiya) and in some regions fraud was so extensive that the implementation of the methodology was not possible (see details below). As in the previous section, in the analysis, I use the same turnout threshold at 60th percentile of the turnout distribution in order to make fraud estimates comparable across both different elections and different regions. The results of the regional analysis are summarized in Tables 3.8-3.14 of the Appendix. The regions are organized by federal districts[25] of Russia. For each of the three elections the tables contain the incumbent's reported margin of victory (VM) and ballot stuffing estimates (fraud) both in percentages and by absolute number of stuffed ballots. Recall once again that ballot stuffing estimates are lower bounds of actual fraud and may substantially underestimate the real level of falsifications (see Section 3.3 for the discussion).

The elections of 2000 were relatively clean: fraud is detected in 13 regions and its magnitude is usually moderate with an average of about 0.8%, measured as the difference between actual

[25] Federal districts are 8 geography-based groupings of federal subjects (regions) of Russia for the convenience of operation and governing.

and 99% confidence fraud-corrected victory margins of the incumbent. In the 2004 elections, fraud is detected in 22 regions with an average of 1.32%. Finally, in 2008, fraud presents in 45 regions from the sample with an average of 2.29%. The 2008 elections are the most fraudulent in terms of both number of fraudulent regions and estimated fraud magnitude. The highest level of ballot stuffing in 2008 is found in the republics of Tatarstan and Chuvashiya, Belgorodskaya, Voronezhskaya, Moskovskaya, Orlovskaya, Penzenskaya, Rostovskaya, Samarskaya and Tymenskaya oblasts, Primorskiy kray, as well as at the city of Moscow. It is important to notice that there are almost no exits from the pool of fraudulent regions over time: once fraud appears in a region it almost certainly remains there. This finding is consistent with Myagkov and Ordeshook (2008), who argue that ballot stuffing and some other forms of fraud in the mid-1990s presented only in a few ethnic Russian regions but then spread to the other regions with noticeable acceleration during the 2000s. There are just 12 regions out of 62 in the sample where fraud is not detected at 99% confidence in all three elections: Astrakhanskaya, Kirovskaya, Kurganskaya, Leningradskaya, Lipetskaya, Magadanskaya, Tambovskaya and Tulskaya oblasts, the republics of Hakasiya and Komi as well as Yugra autonomous district.

Table 3.5 contains the results of the regional-level analysis aggregated at the level of federal districts. The table demonstrates that in the Russian elections held in the 2000s, the most severe fraud took place in southern and Caucasian regions, central Russia (mainly due to the capital) as well as Volga regions, among which there are a number of ethnic republics (Bashkortostan, Chuvashiya, Mariy El, Mordoviya, Tatarstan, Udmurtiya). In contrast, Siberia, Ural and Northwestern regions are the cleanest in all the elections. Note that the numbers in Table 3.5 are aggregates of numbers from the regions for which the analysis was possible. The analysis was not possible for some very small regions and regions with very extensive fraud. Since regions where such extensive fraud took place are mainly from the South (all Caucasian republics, Krasnodarskiy kray) and the Volga area (Bashkortostan, Mordovia), the aggregated results substantially underestimates the amount of fraud in Volga, Southern and Northern Caucasian federal districts in comparison to the other districts.

As discussed above, the suggested methodology is based on a number of assumptions, one of most crucial of which is that fraud must be implemented in a small number of precincts. If the majority of precincts in a given region are fraudulent, then the analysis is not able to detect fraud, as it treats some share (60% in case of this particular analysis) of data as clean. Some Russian regions are an excellent example of cases in which the number of fraudulent precincts is so large that the suggested methodology fails to detect falsification at all. Figures 3.7 and 3.8 depict turnout distributions in the republic of Bashkortostan and Saratovskaya oblast. Both regions shows extremely skewed distributions with extremely high mean in all three elections, which signals extensive falsifications. The suggested fraud detection methodology cannot be

Federal District	2000		2004		2008	
	VM (%)	Fraud (%)	VM (%)	Fraud (%)	VM (%)	Fraud (%)
Central	19.46	0.18	52.27	0.91	48.02	3.08
Ural	30.35	0	61.85	0.15	53.16	1.51
Volga	24.29	0.86	55.42	1.34	49.69	2.98
Northwestern	44.11	0.12	61.65	0.11	51.62	1.48
Southern and North Caucasian	20.52	0.86	50.17	3.18	51.12	2.36
Far Eastern	13.85	0.49	47.37	1.62	44.42	1.81
Siberian	8.58	0.04	50.07	0.46	44.75	0.61

Table 3.5: Russian presidential elections 2000-2008. Regional-level analysis. Federal districts aggregates.

applied to such data. A similar picture can be observed in the republics of Adygeya, Chechnya, Dagestan, Ingushetiya, Kabardino-Balkariya, Kalmykiya, Karachevo-Cherkessiya, Mordoviya, North Osetiya and Tyva.

Figure 3.7: Distribution of turnout across polling stations in the republic of Bashkortostan in the elections of 2000, 2004 and 2008.

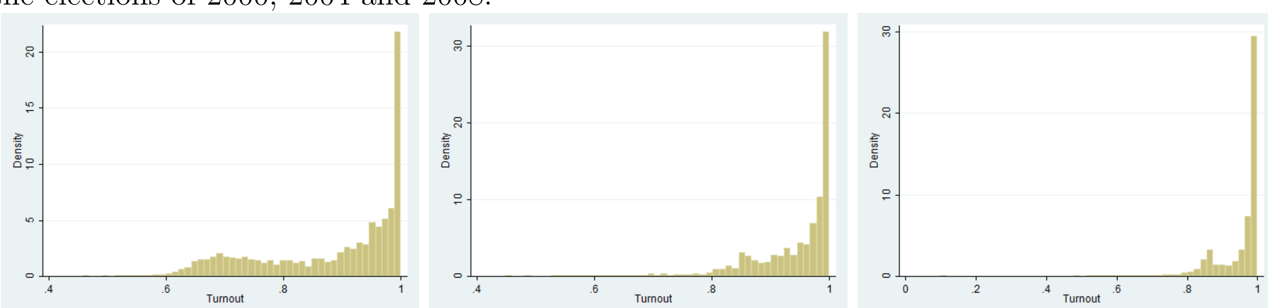

Figure 3.8: Distribution of turnout across polling stations in Saratovskaya oblast in the elections of 2000, 2004 and 2008.

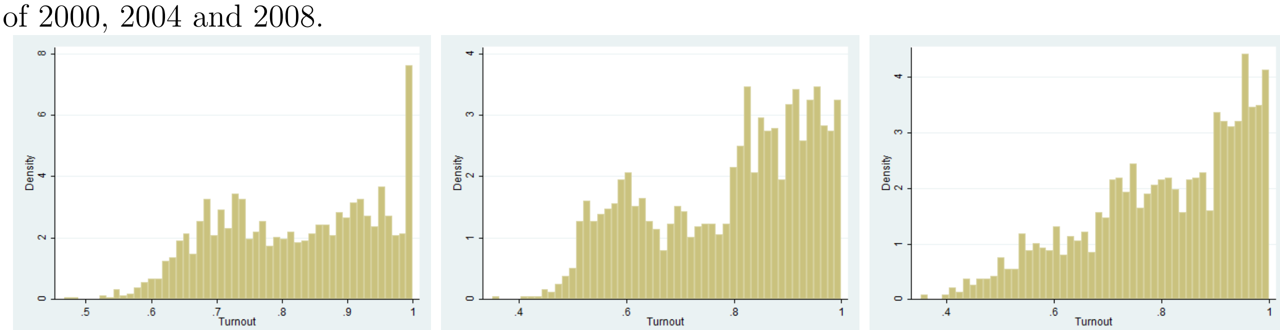

In several regions, such extremely skewed distributions are observed in 2008 or/and 2008 only, while in 2000 the turnout distributions have reasonable shape, and the data allows performance of the analysis. Such a situation can be observed in Krasnodarskiy kray (see Figure 3.9) as well as in the republics of Buryatiya and Mariy El.

Figure 3.9: Distribution of turnout across polling stations in Krasnodarkiy kray in the elections of 2000, 2004 and 2008.

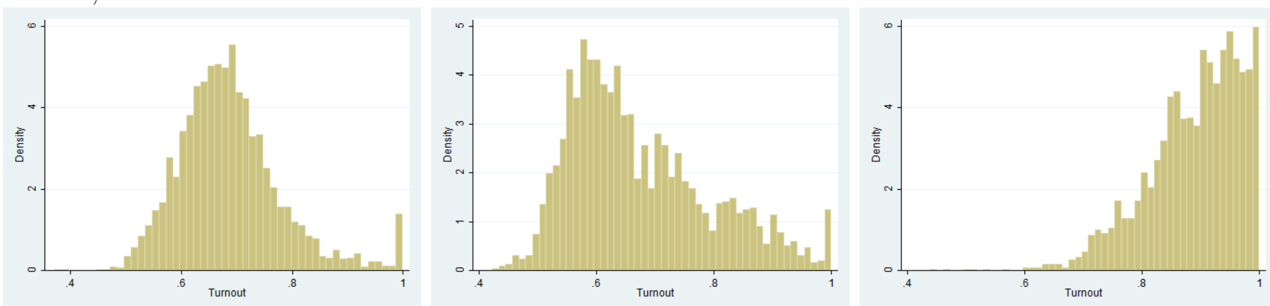

One way to deal with such a situation is to apply the method in a usual way to the first elections, and then to use the same underlying turnout distribution for all subsequent elections. Indeed, the results produced by such an approach may not be fully precise, but can still provide some idea about the extent of ballot stuffing and its dynamics over the period. The estimates are presented in Table 3.6. For each of the three elections, the tables contain the incumbent's reported victory margin (VM) and ballot stuffing (fraud) estimates both in percentages and in absolute numbers of stuffed ballots. Again, the results are consistent with the general trend: fraud in these regions has been growing since 2000, both in absolute numbers of stuffed ballots and extra victory margin (except Bashkortostan). One may notice that in comparison to other regions (see Tables 3.8-3.14 of the Appendix), fraud estimates are not as large as can be expected given extreme turnout distributions. The reason for such a result is that in the analysis I again use the 60th turnout percentile threshold, while if fraud is very extensive a lower value of the threshold should be used (see Section 3.3 for the discussion). Table 3.7 contains fraud estimates obtained using the same analysis but with 50th turnout percentile value as the threshold. The estimates are much higher in all three cases, making the regions among the most fraudulent in Russia.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Bashkortostan	32.17	4.47 (142 100)	87.82	1.25 (238 800)	80.08	1.84 (242 700)
Krasnodarskiy kray	14.10	0	48.10	0	58.24	1.32 (95 000)
Saratovskaya	29.63	0.92 (19 753)	53.76	1.43 (43 200)	59.37	3.06 (102 600)

Table 3.6: Russian presidential elections 2000-2008. Regions with extreme turnout distributions. High turnout threshold.

To summarize, the main findings of the regional-level analysis are the following. First, ballot stuffing had been growing substantially between 2000 and 2008. This is a result of both increasing

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Bashkortostan	32.17	6.75 (203 700)	87.82	1.76 (321 500)	80.08	2.93 (347 400)
Krasnodarskiy kray	14.10	0	48.10	0	58.24	2.15 (142 800)
Saratovskaya	29.63	1.85 (39 000)	53.76	2.58 (76 000)	59.37	4.32 (167 900)

Table 3.7: Russian presidential elections 2000-2008. Regions with extreme turnout distributions. Low turnout threshold.

magnitude of fraud and the expansion of electoral falsification across the regions of Russia. This finding could serve as an additional argument and motivation for the idea presented in the second chapter of this dissertation, where I argue that fraud has a tendency to grow over the lifetime of non-democratic regimes. Second, ballot stuffing is very persistent: once it appears in a region, it remains present in subsequent elections. Third, in some regions of Russia, primarily ethnic republic and southern regions, ballot stuffing is so extensive that the fraud detection methodology suggested in the paper cannot be applied. For the regions where analysis is possible, the most severe fraud is detected again in ethnic republics, southern regions and the capital.

3.6 Conclusion

This paper suggests a simple statistical method to test for the presence of ballot stuffing using official detailed electoral data. The method is based on the observation that ballot stuffing increases both turnout and the incumbent's vote share in precincts where it occurs. Hence, precincts with relatively low reported turnout are more likely to be clean. Using the information on relatively clean precincts, it is possible to simulate counterfactual data for spoiled precincts and to compare them with the observed data.

The method is first piloted on artificial data and artificially fraudulent real data, and subsequently applied to test the fairness of the Russian executive elections in 2000, 2004, 2008 and 2012, whose transparency and integrity are dubious. Results strongly reject the hypothesis of no ballot stuffing in all four elections, while the estimates of the magnitude of ballot stuffing suggest that fraud has been persistently growing over time. However in none of the elections was ballot stuffing sufficiently large to alter the outcome. Finally, regional-level analysis of Russian electoral data shows that the most severe ballot stuffing takes place primarily in ethnic republics and southern regions of Russia.

Appendix

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Belgorodoskaya	7.76	0.33 (3000)	27.2	2.37 (26265)	47.33	4.12 (67153)
Bryanskaya	-2.93	0	40.29	0	34.72	3.04 (26818)
Ivanovskaya	22.87	0.33 (2716)	50.84	0.14 (1368)	44.77	1.08 (8902)
Kaluzhskaya	17.18	0	55.43	0	43.72	1.43 (13320)
Kostromskaya	33.45	0	50.68	0	39.71	0
Kurskaya	10.3	0.69 (5487)	44.16	0	42.44	0
Lipetskaya	-6.55	0	42.37	0	44.05	0
Moscow City	27.05	0	61.21	1.87 (191213)	55.05	4.62 (428352)
Moskovskaya	19.97	0.30 (13141)	60.47	0.74 (58525)	52.41	3.78 (277217)
Orlovskaya	1.22	0	37.63	0	43.61	3.26 (27362)
Ryazanskaya	12.30	0	59.52	0	36.59	1.38 (12708)
Smolenskaya	17.70	0	44.03	0	34.73	2.25 (18053)
Tambovskaya	6.78	0	39.31	0	53.18	0
Tverskaya	29.82	0	55.20	0	48.39	1.64 (23661)
Tulskaya	11.45	0	47.03	0		0
Vladimirskaya	22.43	1.46 (15372)	53.37	0.43 (6393)	42.11	1.28 (15769)
Voronezhskaya	24.79	0	43.33	2.05 (41977)	43.69	6.46 (137818)
Yaroslavskaya	43.14	0.44 (6021)	58.64	0.11 (1634)	42.93	1.80 (19516)

Table 3.8: Russian presidential elections 2000-2008. Central Federal District.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Chelyabinskaya	16.54	0	56.01	0	45.41	1.51 (49701)
Kurganskaya	11.98	0	47.78	0	44.29	0
Sverdlovskaya	46.65	0	68.56	0	55.80	0.91 (41771)
Tyumenskaya	25.35	0	62.37	1.14 (22129)	69.45	3.94 (98036)
Yugra	39.53	0	67.40	0	52.2	0

Table 3.9: Russian presidential elections 2000-2008. Ural Federal District.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Chuvashiya	1.4	0	45.59	0.77 (8867)	43.91	3.91 (45689)
Kirovskaya	31.33	0	45.91	0	62.13	0
Nizhegorodskaya	21.00	0	48.66	0	37.91	0.72 (21099)
Orenburgskaya	2.73	0	34.22	0	34.50	0.21 (3364)
Penzenskaya	11.00	0	44.40	0.53 (7012)	52.36	5.68 (88866)
Permskiy	40.83	0	62.66	0	50.60	1.24 (28646)
Samarskaya	12.45	0	44.16	0	41.45	4.22 (98080)
Tatarstan	48.81	3.49 (137965)	75.98	3.26 (272004)	66.30	4.90 (299358)
Udmurtiya	36.18	0	66.62	0.05 (995)	54.38	1.75 (27212)
Uljanovskaya	9.18	0	46.65	0	45.59	3.30 (37777)

Table 3.10: Russian presidential elections 2000-2008. Volga Federal District.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Arkhangelskaya	39.54	1.21 (14047)	68	0.24 (4696)	48.21	0.54 (6394)
Kaliningradskaya	36.69	0	58.07	0	38.85	3.15 (21427)
Kareliya	47.23	0	64.02	0	49.97	1.42 (8653)
Komi	38.14	0	61.94	0	56.95	0
Leningradskaya	47.64	0	66.92	0	52.26	0
Murmanskaya	50.52	0.21 (2163)	59.68	0	47.03	1.73 (14601)
Novgorodskaya	43.44	0	57.84	0	45.62	0.92 (5290)
Pskovskaya	36.91	0	54.16	0	49.83	4.10 (30199)
St Petersburg City	45.42	0	67.73	0.28 (17873)	55.49	2.22 (108752)
Vologodskaya	47.91	0	64.27	0	52.62	2.57 (7740)

Table 3.11: Russian presidential elections 2000-2008. Northwestern Federal District.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Astrakhanskaya	34.32	0	47.8	0	57.99	0
Volgogradskaya	19.53	0	41.07	0	38.02	0.20 (3935)
Rostovskaya	20.79	2.05 (59462)	58.27	6.15 (304304)	61.67	3.64 (199793)
Stavropolskiy	15.73	0	43.22	0	41.63	1.89 (39977)

Table 3.12: Russian presidential elections 2000-2008. Southern and North Caucasian Federal Districts.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Amurskaya	15.89	0	46.13	0	43.85	1.32 (10380)
Magadanskaya	39.50	0	59.68	0	42.92	0
Primorskiy	4.15	0.42 (14432)	42.25	2.10 (34428)	44.04	4.28 (68709)
Sakhalinskaya	15.94	0	54.14	0.96 (4529)	42.26	0
Khabarovskiy	21.49	0	51.5	2.32 (33482)	45.99	0

Table 3.13: Russian presidential elections 2000-2008. Far Eastern Federal District.

Region	2000		2004		2008	
	VM	Fraud	VM	Fraud	VM	Fraud
Altayskiy	4.09	0	49.05	0	37.1	0.9 (15901)
Buryatiya	1.84	0	49.5	1.47 (12655)	52.37	0.64 (6449)
Chitinskaya	13.65	0	57.71	0	48.47	1.84 (19633)
Irkutskaya	17.07	0	45.19	0.3 (5160)	39.42	1.94 (37491)
Kemerovskaya	10.08	0	60.2	1.73 (60544)	62.05	0
Krasnoyarskiy	15.54	0.23 (3730)	49.21	0	41.85	0.54 (11869)
Novosibirskaya	1.62	0	41.51	0	37.34	0.63 (13196)
Omskaya	-4.74	0.09 (974)	49.62	0	41.18	0.12 (2578)
Tomskaya	27.18	0	53.80	0.02 (140)		0
Hakasiya	6.04	0	43.06	0	37.69	0

Table 3.14: Russian presidential elections 2000-2008. Siberian Federal District.

References

Alvarez, M., Hall, T., Hyde, S., 2008. Election Fraud: Detecting and Deterring Electoral Manipulation. *Brookings Institution Press, Washington, D.C.*

Benford, F., 1938. The Law of Anomalous Numbers. *Proceedings of the American Philosophical Society*, 78, 4, 551572.

Berber, B., Scacco, A., 2008. What the Numbers Say: a Digit Based Test for Election Fraud Using New Data from Nigeria. *Presented at the Annual Meeting of the APSA.*.

Becker, R., Volinsky, C., Wilks, A., 2010. Fraud Detection in Telecommunications: History and Lessons Learned. *Technometrics*, 52(1), 2033.

Brady, H., 2005. Comments on Benford's Law and the Venezuelan Election. *Mimeo.*

Breuning, C., Goerres, A., 2011. Searching for Electoral Irregularities in an Established Democracy: Applying Benford's Law Tests to Bundestag Elections in Unified Germany. *Political Analysis*, 19(4), 409-433.

Cantu, F., Saiegh, S., 2011. Fraudulent Democracy? An Analysis of Argentina's Infamous Decade Using Supervised Machine Learning. *Political Analysis*, 19(4), 409-433.

Decker, J., Myagkov, M., Ordeshook, P., 2011. Benford's Law and the Detection of Election Fraud. *Political Analysis*, 19(3), 245-268.

Jacob, B., Levitt, S., 2003. Rotten Apples: An Investigation of The Prevalence and Predictors of Teacher Cheating. *The Quarterly Journal of Economics*, August, 843-877

Lehoucq, F., 2003. Electoral Fraud: Causes, Types and Consequences. *Annual Review of Political Science*, 6, 233-256.

Levin, I.,Cohn, G., Ordeshook, P., Alvarez, R., 2009. Detecting Voter Fraud in an Electronic Voting Context: An Analysis of the Unlimited Reelection Vote in Venezuela. *Electronic Voting Technology Workshop / Workshop on Trustworthy Elections*, Online Proceedings.

Li, J., Huang, K., Jin, J., Shi, J., 2008. A Survey on Statistical Methods for Health Care Fraud Detection. *Health Care Management Science*, 11(3), 275-287.

Lukinova, E., Myagkov, M., Ordeshook, P., 2011. Ukraine 2010: Were Tymoshenko's Cries of Fraud Anything More Than Smoke? *Post-Soviet Affairs*, 27(1), 37-63.

Mebane W., 2006. Election Forensics: Vote Counts and Benford's Law. *Working Paper*, University of Michigan.

Mebane W., Kalinin, K., 2009. Comparative Election Fraud Detection. *Presented at the Annual Meeting of the APSA.*

Myagkov, M., Ordershook, P., 2008. Russian Elections: An Oxymoron of Democracy. *VTP Working Paper.*

Myagkov, M., Ordeshook, P., Shakin, D., 2009. The Forensics of Election Fraud: Russia and Ukraine. *Cambridge University Press.*

Pandit, S., Chau, D., Wang, S., Faloutsos, C., 2007. NetProbe: A Fast and Scalable System for Fraud Detection in Online Auction Networks. *Proceedings of the 16th international conference on WWW.*

Pericchi, L., Torres, D., 2004. La Ley de Newcomb-Benford y sus aplicaciones al Referendum Revocatorio en Venezuela. *Reporte Tecnico No-Definitivo*, 2a.

Quash, J., Sriganesh, M., 2008. Real-Time Credit Card Fraud Detection Using Computational Intelligence. *Expert Systems with Applications*, 35(4), 17211732.

Roukema, B., 2009. Benfords Law Anomalies in the 2009 Iranian Presidential Election. *Mimeo.*

Sakwa, R., 2005. The 2003-2004 Russian Elections and Prospects for Democracy. *Europe-Asia Studies*, 57(2), 369-398.

Sudjianto, A.,Nair, S., Yuan, M., Zhang, A., Kern, D., Cela-Daz, F., 2010. Statistical Methods for Fighting Financial Crimes. *Technometrics*, 52(1), 5-19.

Treisman, D., 2009. Elections in Russia, 1991-2008. *Working Paper*, Moscow State University - Higher School of Economics.

Wolfers, J., 2006. Point Shaving: Corruption in NCAA Basketball. *American Economic Review*, 96(2), 279-283.