Deliver the Vote!

Micromotives and Macrobehavior in Electoral Fraud

Ashlea Rundlett and Milan W. Svolik*

Abstract

Most election fraud is not conducted centrally by incumbents but rather locally by a multitude of political operatives. How does an incumbent ensure that his agents deliver fraud when needed and as much as is needed? We address this and related puzzles in the political organization of election fraud by studying the perverse consequences of incentive conflicts between incumbents and their local agents. These incentive conflicts result in a herd dynamic among the agents that tends to either oversupply or undersupply fraud, rarely delivering the amount of fraud that would be optimal from the incumbent's point of view. Our analysis of the political organization of election fraud explains why fraud sometimes fails to deliver victories, why popular incumbents often preside over seemingly unnecessary fraud, and it predicts that the extent of fraud should be increasing in both the incumbent's genuine support and reported results across precincts. A statistical analysis of anomalies in precinct-level results from the 2011-12 Russian legislative and presidential elections supports our key claims.

*Department of Political Science, University of Illinois at Urbana-Champaign. Emails: rundlet2@illinois.edu and msvolik@illinois.edu.

1 Introduction

"You may have won the election, but I won the count!" was Anastasio Somoza's rebuke to an opponent who accused him of rigging an election.[1] A burgeoning literature depicts the many ways in which incumbents attempt to "win the count" and conducts increasingly sophisticated analyses of their detection and deterrence. Yet while "winning the count" may be directed and facilitated from above, its execution is primarily local.[2] The most frequent forms of election fraud – ballot box stuffing, multiple voting, voter intimidation, or the falsification of counts – are ultimately executed at the level of individual polling stations, not by the incumbent but rather a machinery that typically consists of hundreds of political operatives, party members, and state employees.[3]

In spite of rich descriptive accounts of such local-level fraud in qualitative and historical literature, most formal and analytical research on election fraud treats its political organization and execution as unproblematic. The incumbent's machinery of manipulation is assumed to operate as a unified political actor, under the incumbent's perfect political control. This approach leaves us with a number of puzzles: How does the incumbent ensure that his local agents deliver fraud when needed and exactly as much as is needed? What motivates local agents to engage in fraud when doing so may result in criminal prosecution and conviction? Why does locally-conducted electoral fraud succeed in delivering a victory in some elections but fail in others?

In this paper, we study these puzzles and demonstrate that incentive conflicts in the political organization and execution of election fraud have far-reaching implications for its conduct, success, and empirical fingerprints. Two related but distinct incentive conflicts

[1] Richard Gott, "The Spanish left settles down for the long haul," *The Guardian*, June 17, 1977.

[2] On the conduct and forms of election fraud, see Birch (2011), Lehoucq (2003), Schedler (2012), and Simpser (2013).

[3] On the local execution of fraud, see Cantú (Forthcoming), Benton (2013), Larreguy, Olea, and Querubin (2014), Lehoucq (2003), Martinez-Bravo (2014), and Simpser (2013).

critically shape the political organization of election fraud: the principal-agent problem between an incumbent and his local agents, and the collective action problem among the agents. At the heart of the principal-agent problem is a conflict of interest between the incumbent and his agents about when to engage in fraud and how much of it to conduct. The incumbent's preferences were eloquently summarized in John F. Kennedy's facetious response to questions about the role of his farther's wealth in his political success: "I have just received the following wire from my generous daddy. It says, 'Dear Jack, don't buy a single vote more than is necessary. I'll be damned if I'm going to pay for a landslide!' "[4] That is, even those incumbents who are willing to engage in fraud if it is needed for a victory want to avoid unnecessary fraud that will only raise suspicions. Most agents, meanwhile, prefer to conduct fraud when it carries the least risk – when the incumbent's victory is assured and the agents' actions are unlikely to be investigated. Agents are most reluctant to engage in fraud when the incumbent's victory is in doubt and they worry about being prosecuted if the challenger wins the election. Thus agents are least willing to engage in fraud precisely when the incumbent needs it most!

This principal-agent problem between the incumbent and his agents is compounded by a collective action problem among the agents. It is most pronounced when the incumbent narrowly trails the challenger. In these scenarios, the incumbent's agents understand that, if only enough of them engaged in fraud, they could secure the incumbent's victory. At the same time, however, each agent's doubts about other agents' actions lead her to question the prudence of her own engagement in fraud. Hence even when fraud could assure the incumbent's victory, the agents' fear of its ultimate failure may turn it into a self-fulfilling prophecy.

In order to rigorously examine the interplay between principal-agent and collective

[4] Gridiron dinner in Washington, D.C., 15 March 1958; see Sabato (2013, 46).

action problems in the political organization and execution of election fraud, we develop a formal model with two key, novel features. First, the incumbent does not engage in fraud directly but instead depends on the illicit collaboration of a large number of local agents who must be motivated by the promise of a reward. This is a departure from existing formal research where the incumbent's machinery of fraud is assumed to act as a unitary actor (Egorov and Sonin 2012; Fearon 2011; Gandhi and Przeworski 2011; Little 2012; Rozenas 2013; Simpser 2013; Svolik and Chernykh 2013). A key aspect of this departure concerns the contingency of the agents' rewards – as well as their punishment – on the incumbent's political survival: Each agent understands that she will obtain the promised reward only if the incumbent is reelected and may face prosecution if she engages in fraud but the incumbent is ultimately defeated.[5]

The second key feature of our formalization concerns the limited information available to the incumbent and his local agents when they are deciding on whether to engage in fraud. The difficulties that incumbents in hybrid regimes face when gauging their genuine popularity have been highlighted in research on electoral manipulation and democratization (Gehlbach and Simpser 2011; Little 2013; Miller 2013; Rozenas 2013) and parallel classic accounts of incentives for "preference falsification" under authoritarianism (Kuran 1991; Wintrobe 1998). The novel feature of our model is in how the structure of this information paucity is tailored to the context of electoral manipulation: While both the incumbent and his local agents have only imperfect information about the incumbent's genuine popularity, each local agent's information is much more precise than the incumbent's but at the same time confined to her own precinct.

The chief macro-political consequence of these two novel features is a herd dynamic

[5]Such politically contingent rewards play a key role in Gehlbach and Simpser's (2011) analysis of bureaucratic incentives under authoritarianism, Smith, de Mesquita, and LaGatta's (2013) study of group-level incentives in partisan electoral competition, and Dragu and Polborn's (2013) analysis of administrative constraints on the arbitrary exercise of government power.

among the agents that tends to result in either overwhelming victories for the incumbent or, less often, his resounding defeats. We obtain this prediction by a natural application of the global game approach to the analysis of collective action problems (Carlsson and van Damme 1993; Morris and Shin 2003).[6] Its key advantage is to transform a setting that would otherwise suffer from a multiplicity of equilibria with contradictory predictions into one with a unique, tractable, and politically intuitive equilibrium. In our setting, the agents' incentives result in a unique tipping-point equilibrium according to which an agent engages in fraud only if her local, private perception of the incumbent's popularity exceeds a threshold. The intuition is as follows: The higher the incumbent's genuine popularity in an agent's precinct, the more popular she infers the incumbent is nationwide; consequently, she anticipates that fewer agents need to engage in fraud in order to secure the incumbent's victory, which in turn lowers her own risk of engaging in fraud. Thus while never observed perfectly by either the incumbent or the agents, the incumbent's genuine nationwide popularity ends up playing a central role in coordinating the agents' attempts to resolve their collective action problem.

The perverse consequence of such individual-level incentives is a herd dynamic at the aggregate level: Jointly, agents will tend to either oversupply or undersupply fraud, rarely delivering the amount of fraud that would be optimal from the incumbent's point of view. At one extreme, when the incumbent is unpopular and needs fraud most, agents will tend to underdeliver it; at the other extreme, agents will deliver excess fraud when not needed at all. Put simply, the incumbent cannot order 51% of the vote and expect to get precisely that. The aggregate amount of fraud will approximate the desired level only when he narrowly trails the challenger. It is only in such elections that fraud will be both politically

[6]See Boix and Svolik (2013), Bueno de Mesquita (2010), Edmond (2013), Little (2012), and Shadmehr and Bernhardt (2011) for applications of global games and related techniques to the analysis of collective action problems in protests, repression, regime change, and authoritarian power-sharing.

decisive for the incumbent's victory and successful in securing it.

Our analysis of microincentives in the political organization of election fraud improves our understanding of the resulting macrobehavior in a number of ways. First, the equilibrium dynamic that we just described helps us understand the puzzling, often contradictory accounts of incumbents who enjoy genuine popularity and at the same time engage in seemingly unnecessary fraud. In a seminal analysis of the Institutional Revolutionary Party's (PRI) demise in Mexico, Magaloni (2006) observes that, while certainly present, fraud only served to embellish the already impressive popularity that the PRI enjoyed before the 1980s (see also Greene 2007; Simpser 2013). Similarly, students of contemporary Russia are puzzled by the embarrassingly obvious fingerprints of fraud in elections that the United Russia party, and especially Vladimir Putin and Dmitry Medvedev, could have won cleanly.[7] According to a leading explanation for these perplexing outcomes, inflated margins of victory serve to signal the incumbent's invincibility and thus deter potential challengers or defectors (Magaloni 2006; Simpser 2013). Our analysis suggests an alternative mechanism: Rather than an intentional strategy, overwhelming incumbent victories may be the unintended byproduct of the principal-agent and collective action problems in the political organization of election fraud: Because individual local agents are most willing to conduct fraud when it carries the least risk – when the incumbent is genuinely popular – we should not be surprised to observe genuine popularity go hand in hand with excessive fraud at the aggregate level.

This logic also helps us understand why local-level fraud, even when encouraged by the incumbent, sometimes fails to secure his re-election. Because fraud is by definition illegal, the incumbent's capacity to motivate agents to engage in fraud on his behalf is limited to

[7] See Myagkov, Ordeshook, and Shakin (2009, 178), Simpser (2013, 170), Treisman (2011, especially pp. 97, 103, and 348), and Daniel Treisman, "What Keeps the Kremlin Up All Night," *The St. Petersburg Times*, February 19, 2008.

the tacit promise of a reward upon his reelection. Such politically contingent inducements, however, are least effective precisely when the incumbent needs the agents' collaboration most – when he lacks genuine popularity. Our analysis of the ensuing collective action problem highlights how individual agents' worries about the incumbent's eventual defeat reverberate among them and, if sufficiently pronounced, multiply into an avalanche of defections from the incumbent. Thus even though any single agent's decision to refrain from fraud may be inconsequential at the national level, a minor decline in the incumbent's popularity may instigate a large aggregate shift from nearly to universal participation in fraud to an almost complete abstention. Agents' fears of the incumbent's defeat become a self-fulfilling prophecy.

This reasoning suggests a mechanism of fraud deterrence that has not been explored by existing theoretical research but is implicit in recent empirical research. Extant analytical treatments of election fraud focus on the threat of a post-election protest or violence as the chief deterrent against fraud (Fearon 2011; Little 2012; Przeworski 2011; Svolik and Chernykh 2013; Tucker 2007). By contrast, our arguments highlight that a major reason for the failure of local-level fraud may be the incumbent's inability to muster the machinery of fraud in the face of his declining popularity.[8] This focus on the incentives faced by local-level agents parallels empirical research on election monitoring, the deterrent effect of which is also hypothesized to occur at the level of individual polling stations (Hyde 2008; Ichino and Schündeln 2012). Our results, however, suggest that the direct effect of local-level deterrents – whereby they raise the risk of engaging in fraud for individual agents who are being monitored – may not be the most consequential one.[9] Rather, the

[8]In the controversial 1988 Mexican presidential election, for instance, the ruling PRI had to resort to top-level manipulation after local-level fraud proved insufficient after years of the party's declining popularity. See Ginger Thompson, "Ex-President in Mexico Casts New Light on Rigged 1988 Election," *The New York Times*, March 9, 2004, p.A-10; and Castañeda (2000, 231-239).

[9]After all, only a small fraction of polling stations is visited by election observers during any single election; see e.g. Hyde (2011).

primary consequence of monitoring may be to heighten the collective action dilemma among all agents, even those who are not being monitored: When monitoring occurs, all agents anticipate that much greater efforts must be exerted at non-monitored polling stations in order to secure the incumbent's victory. To our knowledge, such indirect, systemic consequences of local fraud deterrents have not yet been examined either empirically or theoretically.

An improved understanding of the microincentives faced by the agents who ultimately execute fraud also helps us anticipate its empirical fingerprints. The prevailing approach to fraud detection focuses on the identification of statistical anomalies in voting or turnout but is often less explicit about the political process that generates them.[10] Our model clarifies that anomalies indicative of fraud may be the unintended consequence of incentive conflicts in the political organization of fraud and predicts a specific pattern that such anomalies should follow: Their occurrence across precincts should not be uniform but rather increasing in both the incumbent's genuine popularity and his vote share. Yet at the same time, such anomalies alone do not imply that the incumbent stole an election that would have otherwise been won by the challenger. In fact, the fingerprints of fraud may be most pronounced precisely when fraud is not politically decisive.

We find empirical support for our arguments when we examine the pattern of fraud in the 2011 legislative and 2012 presidential elections in Russia. We confirm a finding from earlier analyses of these elections (Gehlbach 2012; Klimek et al. 2012; Kobak, Shpilkin, and Pshenichnikov 2012), according which one form of electoral manipulation involved the rounding of the incumbent Vladimir Putin's and United Russia Party's vote shares to some higher multiple of five by the regime's local operatives.

[10] See Ahlquist, Mayery, and Jackman (2013), Alvarez, Hall, and Hyde (2008), Beber and Scacco (2012), Cantú (Forthcoming), Cantú and Saiegh (2011), Cho and Gaines (2007), Hyde and Marinov (2009), Ichino and Schüdeln (2012), Mebane and Kalinin (2009), Myagkov, Ordeshook, and Shakin (2009), and Sjoberg (2013).

Crucially, we also identify a previously unnoticed pattern that is anticipated by our model: The extent of such anomalies is increasing in the incumbent's precinct-level vote share. In order to quantify the extent of fraud, we develop a measure of the ruggedness in the distribution of Putin's and United Russia's results based on kernel density estimation techniques. Using two different benchmarks, we find that the distribution of Putin's and United Russia's results is too rugged at results corresponding to multiples of five to occur by chance and this ruggedness is indeed increasing in their precinct-level vote share, as anticipated by our theoretical arguments. Our approach additionally reveals that, instead of ballot stuffing, fraud in these elections was most likely executed by the stealing of votes from the leading opposition candidate Gennady Zyuganov and his Communist Party.

2 The Model

Consider the following electoral manipulation game between an incumbent and his agents.[11] Each agent i operates in one among a continuum of precincts of equal size and decides whether to engage in fraud on behalf of the incumbent at the time of the election. We denote agent i's engagement or not in fraud by $a_i = \{f, n\}$, respectively. The incumbent, however, does not observe whether any agent engaged in fraud; he only observes the *precinct-level election result* R_i – his share of the vote in agent i's precinct. Before the election, therefore, the incumbent promises each agent a reward (a higher salary, promotion, perks) commensurate with the election result in her precinct. More precisely, each agent obtains the payoff wR_i after the incumbent's victory, where $w \geq 0$ and we refer to is as the *reward factor*.

The precinct-level election result R_i depends on the incumbent's genuine *precinct-level*

[11] Proofs of all technical results as well as the discussion of alternative parameterizations of our information structure can be found in the supplementary appendix.

popularity S_i and whether the agent engaged in fraud on behalf of the incumbent, $a_i = \{f, n\}$, as follows:[12]

$$R_i = \begin{cases} S_i + F & \text{if } a_i = f; \\ S_i & \text{if } a_i = n. \end{cases}$$

Above, the parameter F, $0 < F < \overline{F}$, denotes the share of the precinct-level election result due to the agent's fraud and we interpret it as a measure of the precinct agents' *fraud capacity*.[13]

Crucially, the agents obtain the reward wR_i only if the incumbent is re-elected. By contrast, if the incumbent loses, each agent's payoff depends on whether she engaged in fraud at the time of the election. If she did not engage in fraud, she obtains the payoff 0. If she did engage in fraud, the agent obtains the payoff $-cF$, where $c > 0$ stands for the political *cost* of fraud. It reflects the potential investigation of allegations of fraud in the agent's precinct after the challenger takes office, possibly resulting in the agent's criminal prosecution and conviction.[14] Thus we are effectively assuming that the likelihood of a conviction or the severity of punishment are increasing in F, the share of the precinct-level election result due to the agent's fraud.[15]

[12] This binary-action model has the same implications as a more general one, in which each agent chooses the amount of fraud from the interval $f \in (0, F)$. Given the payoffs in Figure 1, an agent who optimally engages in a positive amount of fraud does so at its maximum feasible level, effectively choosing either $f = 0$ or $f = F$.

[13] This "additive" assumption about fraud production is only one – and an intentionally simple one – among several plausible ways of formalizing fraud production. Letting $R_i = (1+F)S_i$ or $R_i = \frac{(1+F)S_i}{(1+F)S_i+(1-S_i)}$ results in qualitatively identical insights but less transparent algebra. In the supplementary appendix, we derive $\overline{F}$, the maximum admissible value of F that is implied by our informational assumptions and the global game framework (see below); $\overline{F} = \frac{1}{2}(1 - 4\epsilon)$ for $\epsilon < \frac{1}{4}$.

[14] Two wisecracks from Chicago machine politics illustrate this two-sided contingency of the agents' payoffs: i) "The guiding principle of Chicago politics was that victors got patronage spoils; losers had to find honest work." ii) "When a Cook County election is stolen, it stays stolen." See Grossman, Guy, "Old way still wins the votes," *Chicago Tribune*, 26 October 1994, and "It's No News Here," *Chicago Tribune*, 23 November 1968, p. 16, respectively.

[15] We may also interpret c less politically, as the agent's costly effort. In that case, $-cF$ should also be included in the top-left cell of the payoff table in Figure 1.

		Election result	
		$R \geq \frac{1}{2}$	$R < \frac{1}{2}$
Agent i's action	$a_i = f$	$w(S_i + F)$	$-cF$
	$a_i = n$	wS_i	0

Figure 1: Agent i's payoffs as a function of her fraud decision a_i and the election result R

The agents' payoffs are summarized in Figure 1. We see that each agent's incentive to engage in fraud depends on her expectation about the *national-level election result* R. If the agent expects the incumbent to win, $R \geq \frac{1}{2}$, then she prefers to conduct fraud since $w(S_i + F) > wS_i$. If, on the other hand, the agent expects the incumbent to lose, $R < \frac{1}{2}$, then she prefers to play fair since $-cF < 0$.

The overall election result R depends on both the incumbent's genuine popularity and the actions of the agents. More specifically, the incumbent's *popularity* at the time of the election θ, $0 < \theta < 1$, corresponds to the fraction of the electorate that actually voted for the incumbent. We assume that θ is not perfectly known by any of the players but is commonly believed to be uniformly distributed on the unit interval $(0, 1)$. Instead, each agent privately observes the fraction S_i of her precinct that voted for the incumbent, which is correlated with θ in the following way: S_i is uniformly distributed on the interval $(\theta - \epsilon, \theta + \epsilon)$, $0 < \epsilon < \overline{\epsilon}$.[16] We think of ϵ as "small" and interpret it as a measure of heterogeneity in the incumbent's support across precincts.[17] Thus when each agent decides whether to engage in fraud on behalf of the incumbent, she has only imperfect information about the incumbent's genuine, national-level popularity.[18]

[16]The possibility of $S_i < 0$ and $S_i > 1$ when θ is within an ϵ distance of the boundaries 0 and 1, respectively, is irrelevant for the strategic analysis that follows (and could be avoided by letting S_i be uniformly distributed on the intervals $(0, \theta + \epsilon)$ and $(\theta - \epsilon, 1)$ when θ is within an ϵ distance of 0 and 1, respectively.) The advantage of an information structure based on the Uniform distribution is the availability of closed form solutions for key quantities in our analysis. In section 2.2 and the supplementary appendix, we present results based an alternative information structure according which θ and S_i follow the (appropriately transformed) Normal distribution.

[17]In the supplementary appendix, we derive $\overline{\epsilon}$, the maximum admissible value of ϵ that is implied by our informational assumptions, $\overline{\epsilon} = \frac{1}{4}(1 - 2F)$ for $F < \frac{1}{2}$.

[18]Note that while S_i is informative about θ, the agents lack a common knowledge of θ for an arbitrarily

Because we model the incumbent's agents as atomless players on a continuum of precincts, any single agent's decision to engage in fraud on behalf of the incumbent will be inconsequential at the national level. Jointly, however, the agents' actions affect the election result as the overall election result R amounts to

$$E[R] = \int_{\theta-\epsilon}^{\theta+\epsilon} \frac{1}{2\epsilon} \left(S_i + \mathbf{1}_{\{a_i=f\}} F\right) \, dS_i = \theta + \phi F. \tag{1}$$

Above, $E[R]$ denotes the expected value of R, $\mathbf{1}_{\{a_i=f\}}$ is an indicator function that equals 1 if agent i engaged in fraud and 0 otherwise, and ϕ is the fraction of agents that engaged in fraud. We think of our assumption of atomless agents along a continuum of equally-sized precincts as capturing a country with a large number of precincts that are small relative to the country as a whole.[19]

The national-level election outcome R thus depends on the precinct agents' fraud capacity F, the fraction of agents engaging in fraud ϕ, and the incumbent's election-day popularity θ. If $\theta \geq \frac{1}{2}$, then $R \geq \frac{1}{2}$ and the incumbent wins the election regardless of the agents' actions. If $\theta < \frac{1}{2} - F$, on the other hand, then the incumbent will be defeated even if all agents conducted fraud on his behalf, $R < \frac{1}{2}$. Only when $\frac{1}{2} - F \leq \theta < \frac{1}{2}$ does the election outcome depend on the fraction of agents ϕ engaging in fraud. In turn, if the agents were able to observe θ perfectly, they would all refrain from fraud when $\theta < \frac{1}{2} - F$, engage in fraud if $\theta \geq \frac{1}{2}$, and condition their actions on the actions of others when $\frac{1}{2} - F \leq \theta < \frac{1}{2}$. In the latter, politically most interesting case, all agents engaging in fraud and refraining from fraud both constitute a Nash equilibrium.

small ϵ. For a discussion of this feature of global games, see Morris and Shin (2003).

[19]This approximation works well since the national-level election result in such a country is effectively the mean of precinct-level results, $R = \frac{1}{N}\sum_{i=1}^{N} R_i$. By the law of large numbers, this mean converges in probability to (1), $R = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N}\left(S_i + \mathbf{1}_{\{a_i=f\}}F\right) = \lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N} S_i + F \lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\mathbf{1}_{\{a_i=f\}} = \theta + \phi F$. Our results extend to a setting with a finite number of agents; c.f. Morris and Shin (2003, Appendix B).

This indeterminacy disappears in our setting where agents do not directly observe the incumbent's national-level popularity θ. While each agent's precinct-level result S_i is only an imperfect signal of θ, some values of S_i allow him to perfectly infer the outcome of the election. More specifically, our model for the distribution of S_i implies that if $S_i < \frac{1}{2} - F - \epsilon$, then $\theta < \frac{1}{2} - F$ and thus $R < \frac{1}{2}$. On the other hand, if $S_i \geq \frac{1}{2} + \epsilon$, then $\theta \geq \frac{1}{2}$ and thus $R \geq \frac{1}{2}$. For these values of S_i, therefore, each agent optimally refrains from and engages in fraud, respectively. But when $\frac{1}{2} - F - \epsilon \leq S_i < \frac{1}{2} + \epsilon$, agent i's optimal action depends on her inference about other agents' precinct-level results and actions.

Consider therefore Bayesian Nash equilibria in threshold strategies $\sigma(S_i)$ for agents with precinct-level results on the interval $\frac{1}{2} - F - \epsilon \leq S_i < \frac{1}{2} + \epsilon$. According to these strategies, agent i engages in fraud if and only if the incumbent's popularity in her precinct is at least some threshold value S^*,

$$\sigma(S_i) = \begin{cases} \text{engage fraud, } a_i = f, & \text{if } S_i \geq S^*; \\ \text{do nothing, } a_i = n, & \text{if } S_i < S^*. \end{cases}$$

We will refer to S^* as the agents' *fraud threshold.*

When an agent whose precinct-level result is S_i engages in fraud, she expects the payoff

$$\Pr\left[R \geq \frac{1}{2} \mid S_i\right] w(S_i + F) - \Pr\left[R < \frac{1}{2} \mid S_i\right] cF\,. \tag{2}$$

Meanwhile, the agent's expected payoff from doing nothing is

$$\Pr\left[R \geq \frac{1}{2} \mid S_i\right] wS_i\,. \tag{3}$$

According to the threshold strategy $\sigma(S_i)$, the threshold agent in whose precinct the incumbent's popularity is $S_i = S^*$ must be indifferent between engaging in fraud and doing nothing. Letting $S_i = S^*$ and equating (2) to (3), we see that the following indifference condition holds for the threshold agent:

$$\Pr\left[R < \frac{1}{2} \mid S_i = S^*\right] = \frac{w}{c+w}\,. \tag{4}$$

The indifference condition in (4) highlights the central role that each agent's expectation about the outcome of the election plays in her decision to engage in fraud. The smaller the reward factor w, the stronger must be the threshold agent's expectation that the incumbent will win. This will occur when the fraction of agents that engage in fraud ϕ satisfies the following *majority condition*

$$R \geq \frac{1}{2} \quad \text{or equivalently} \quad \theta + \phi F \geq \frac{1}{2}\,.$$

We may therefore refer to the value of ϕ at which the incumbent wins by a bare majority as the *majority threshold* ϕ^*,

$$\phi^* = \frac{\frac{1}{2} - \theta}{F}\,.$$

According to our assumptions about the distribution of S_i and the threshold strategy $\sigma(S_i)$, the fraction of agents that engage in fraud in equilibrium is

$$\phi = \frac{(\theta + \epsilon) - S^*}{2\epsilon}\,.$$

In turn, the majority threshold implies that the threshold agent's belief that the incumbent will lose the election is

$$\begin{aligned}
\Pr\left[R < \frac{1}{2} \mid S_i = S^*\right] &= \Pr\left[\phi < \phi^* \mid S_i = S^*\right] \\
&= \Pr\left[\frac{(\theta+\epsilon) - S^*}{2\epsilon} < \phi^*\right] \\
&= \Pr\left[\theta < S^* + 2\epsilon\phi^* - \epsilon\right]. \qquad (5)
\end{aligned}$$

Substituting the majority threshold ϕ^* into (5), we obtain

$$\Pr\left[R < \frac{1}{2} \mid S_i = S^*\right] = \Pr\left[\theta < \frac{FS^* - \epsilon S^* + \epsilon}{F + 2\epsilon}\right].$$

Given that θ and S_i are uniformly distributed on the intervals $(0, 1)$ and $(\theta - \epsilon, \theta + \epsilon)$, respectively, the threshold agent believes that the incumbent's popularity θ is uniformly distributed on the interval $(S^* - \epsilon, S^* + \epsilon)$.[20] In turn,

$$\begin{aligned}
\Pr\left[R < \frac{1}{2} \mid S_i = S^*\right] &= \frac{\frac{FS^* - \epsilon S^* + \epsilon}{F + 2\epsilon} - (S^* - \epsilon)}{2\epsilon} \\
&= \frac{\frac{1}{2} - S^* + \epsilon}{F + 2\epsilon}. \qquad (6)
\end{aligned}$$

Substituting (6) into the indifference condition in (4), we see that the agents' fraud threshold must be

$$S^* = \frac{1}{2} - F\frac{w}{c + w} + \epsilon\frac{c - w}{c + w}.$$

Jointly, the fraud threshold S^* and majority threshold ϕ^* imply the existence of a *popularity threshold* θ^* such that, in equilibrium, the incumbent loses the election if $\theta < \theta^*$ and wins the election if $\theta \geq \theta^*$. That is, when the incumbent's genuine popularity is

[20]This holds for the posterior density of $\theta|S_i$ for the range of S_i under consideration, $\frac{1}{2} - F - \epsilon \leq S_i < \frac{1}{2} + \epsilon$. The posterior density of $\theta|S_i$ within a 2ϵ distance of the boundaries 0 and 1 is different; see the supplementary appendix for details.

exactly at the popularity threshold θ^*, he wins by a bare majority,

$$\theta^* = \frac{1}{2} - \phi^* F \quad \text{and} \quad \phi^* = \frac{(\theta^* + \epsilon) - S^*}{2\epsilon}\,.$$

Substituting S^* above and solving for θ^* and ϕ^*, we see that

$$\theta^* = \frac{1}{2} - F\frac{w}{c+w} \quad \text{and} \quad \phi^* = \frac{w}{c+w}\,.$$

Finally, consider the incumbent's optimal choice of the reward factor w in light of the thresholds S^*, θ^*, and ϕ^*. Letting $b > 0$ be the incumbent's payoff from winning the election and normalizing his payoff from losing to 0, his expected payoff becomes

$$b(1-\theta^*) - wE[R] = b(1-\theta^*) - w\left(E[\theta] + \phi F\right) = b(1-\theta^*) - w\left(\frac{1}{2} + \phi F\right)\,. \qquad (7)$$

In (7), we used that fact that $E[\theta] = \frac{1}{2}$ and the equilibrium probability of the incumbent's victory is $1 - \theta^*$, given our assumption that, before the election, θ is commonly believed to be uniformly distributed on the unit interval $(0, 1)$. Treating θ^* and ϕ as functions of w and maximizing (7) with respect to w, we obtain

$$w^* = \sqrt{\frac{cF[cF + 2\epsilon(b+c)]}{F^2 + 2\epsilon F + \epsilon}} - c\,,$$

which is positive as long as $b > \frac{c}{2F}$. Intuitively, the optimal reward factor w is increasing in the incumbent's payoff from winning the election b. When $b \leq \frac{c}{2F}$, meanwhile, the incumbent does not value victory enough to be willing to pay for any fraud and optimally chooses $w^* = 0$.

Proposition 1 summarizes our results so far.

Proposition 1 (Collective Action and Election Fraud)**.** *In the unique Bayesian Nash equilibrium,*

i. the incumbent chooses the reward factor w^**;*

ii. agent i *engages in fraud if* $S_i \geq S^*$ *and does nothing otherwise;*

iii. the incumbent wins the election if $\theta \geq \theta^*$ *and is defeated otherwise;*

iv. the fraction of agents that engage in fraud when the incumbent barely wins the election is ϕ^**;*

where

$$w^* = \sqrt{\frac{cF[cF + 2\epsilon(b+c)]}{F^2 + 2\epsilon F + \epsilon}} - c \ \textit{if}\ b > \frac{c}{2F}, \ \textit{and}\ w^* = 0 \ \textit{otherwise},$$

and

$$S^* = \theta^* + \epsilon\frac{c - w^*}{c + w^*}, \quad \theta^* = \frac{1}{2} - F\phi^*, \quad \textit{and} \quad \phi^* = \frac{w^*}{c + w^*}.$$

Proof. Follows from the text. See the supplementary appendix for the derivation of w^* and the upper bounds on F and ϵ. □

2.1 Comparative Statics and Political Implications

In order to highlight the political implications of our results so far, consider an illustration based on the parameters $c = 1$, $F = \frac{2}{10}$, $\epsilon = \frac{1}{10}$, and $b = 70$, which yield $S^* = 0.3$, $\theta^* = 0.35$, $\phi^* = 0.75$, and $w^* = 3$. That is, in equilibrium, agents engage in fraud only if the incumbent's popularity in their precinct is greater than 30% and fraud secures the incumbent's victory only if his national-level popularity is greater than 35%, or equivalently, when at least three-fourths of agents participate in fraud. Figure 2 employs these values to plot the equilibrium election result R^* as a function of the incumbent's genuine popularity θ.

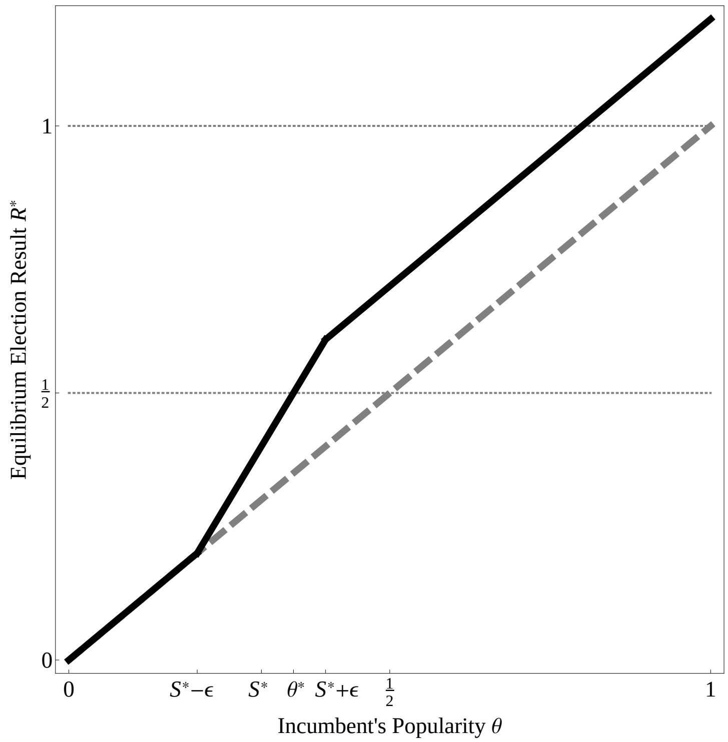

Figure 2: The effect of the incumbent's genuine, national-level popularity θ on the equilibrium election result R^* (solid black line)

We see in Figure 2 that the agents' equilibrium behavior can be partitioned into four qualitatively distinct intervals over θ. At very low levels of the incumbent's genuine popularity, $0 < \theta < S^* - \epsilon$, no agent observes a precinct-level popularity high enough to warrant engaging in fraud and all agents correctly anticipate the incumbent's defeat. When $S^* - \epsilon \leq \theta < \theta^*$, *strategic miscoordination* occurs: if enough agents engaged in fraud, they could ensure the incumbent's victory for some values of θ on this interval (since $\frac{1}{2} - F = 0.3$), but because the incumbent's popularity is too low, an insufficient number of agents ends up engaging in fraud. By contrast, when $\theta^* \leq \theta < \frac{1}{2}$, successful coordination occurs: enough agents engage in fraud to secure an *undeserved victory* for the incumbent. If we take the share of the election result that is due to fraud as a measure of such undeservedness, then the incumbent's victory is most undeserved when $\theta = S^* + \epsilon$ and all

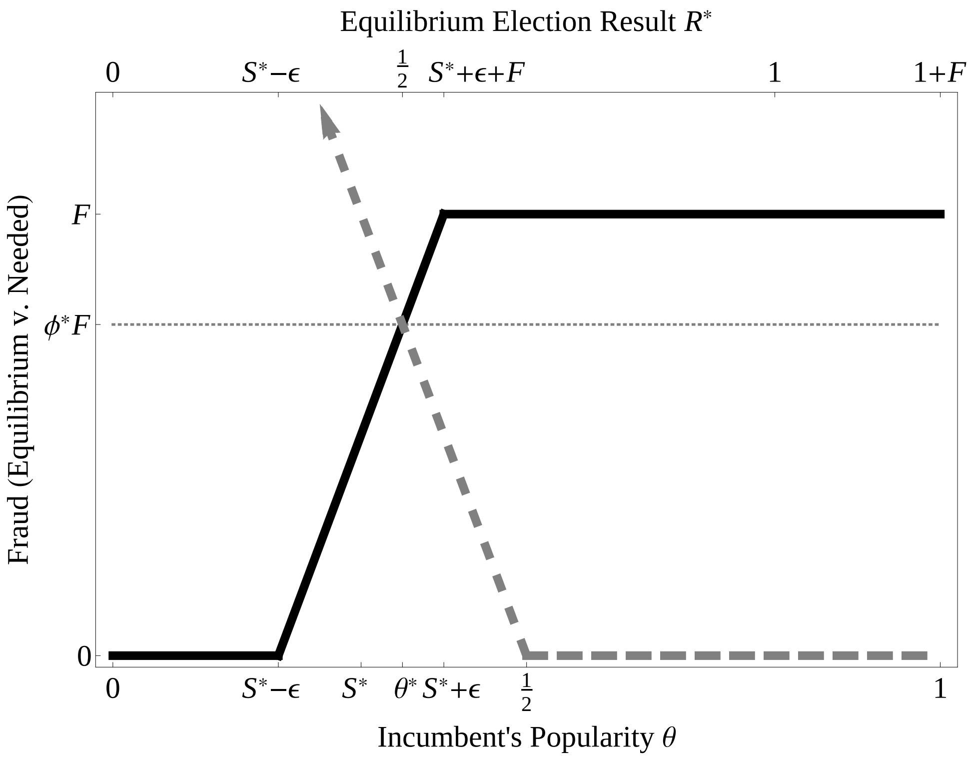

Figure 3: The association between the incumbent's genuine, national-level popularity θ (bottom axis), the equilibrium election result R^* (top axis), the equilibrium level of fraud F^* (solid black line), and the level of fraud needed for a victory $\hat{F}$ (dashed gray line)

agents engage in fraud. Finally, when $\frac{1}{2} < \theta < 1$, fraud occurs but it is *unnecessary*: the incumbent is popular enough to win without fraud. In fact, when $\theta \geq \frac{1}{2} + \epsilon$ all agents are aware that the incumbent will prevail without their complicity; they nonetheless engage in fraud because it boosts the election result in their precincts and thus leads to a higher reward. A perverse consequence of these incentives are national-level election results exceeding 100% at high values of θ.[21]

This conflict between the incumbent's needs and the agents' equilibrium behavior is illustrated in Figure 3. The bottom axis denotes the incumbent's national-level popularity θ; the top axis denotes the equilibrium election result R^*; the dashed gray line plots the

[21] This effect would disappear, i.e. the plot of R^* would flatten and converge to the 45 degree line at high levels of θ, if we assumed that, in addition to the strategic cost of fraud, agents face a direct cost of fraud that is increasing in the amount of fraud that they engage in.

level of fraud $\hat{F}$ that the incumbent needs for a victory; and the solid black line plots the level of fraud F^* that occurs in equilibrium. When $\theta < \theta^*$, the incumbent needs more fraud for a victory than the agents deliver, $F^* < \hat{F}$. When $\theta \geq \theta^*$, on the other hand, the agents engage in a level of fraud that is unnecessary from the incumbent's point of view, $F^* > \hat{F}$, collectively delivering up to 20% in excess of what the incumbent needs. Only at exactly θ^* are the incumbent's needs and the agents' equilibrium behavior in balance: This is when the incumbent needs the national level of fraud to add up to 15% ($\phi^* F = 0.15$) and just the right fraction of agents – three-fourths ($\phi^* = 0.75$) – delivers it.

Precinct-level results thus play a central role in forming agents' beliefs about the incumbent's national-level popularity and, in turn, the degree of their coordination needed for the incumbent's victory. But the precise values of the thresholds S^*, θ^*, and ϕ^* are also shaped by the parameters F, c, ϵ, b, and the equilibrium reward factor w^*.

Consider the popularity threshold θ^*, which is central not only strategically but also normatively: As θ^* declines, the interval $\frac{1}{2} - \theta^*$ along which the incumbent secures an undeserved victory becomes larger. An increase in the fraud capacity F and the reward factor w^* lowers θ^* and thus expands the range of θs along which the incumbent prevails in spite of being opposed by a majority of the electorate. That is, the greater the amount of fraud that each agent can produce within her precinct, the lower the demands on the agents' coordination. Meanwhile, the greater the agents' compensation for precinct-level results, the greater the risk that each agent is willing to take when engaging in fraud. The opposite holds for the cost parameter c. At borderline values of F, w^*, and c ($F \to 0$, $w^* \to 0$, or $c \to \infty$), no fraud occurs in equilibrium when it is actually needed by the incumbent as $\theta^* = \frac{1}{2}$, $S^* = \frac{1}{2} + \epsilon$, and $\phi^* = 0$. Intuitively, no agent is willing to risk prosecution when there is no way to inflate the incumbent's vote share, when there is no personal benefit from doing so, or when the cost of failure is extreme.[22]

Crucially, while a greater reward factor w^* lowers the fraud and popularity thresholds S^* and θ^*, it does not eliminate the collective action problem among the agents.[23] Observe that as w^* tends to infinity, $\lim_{w\to\infty} \phi^* = 1$, and in turn, $\lim_{w\to\infty} S^* = \frac{1}{2} - F - \epsilon$ and $\lim_{w\to\infty} \theta^* = \frac{1}{2} - F$, but

$$S^* > \frac{1}{2} - F - \epsilon \quad \text{and} \quad \theta^* > \frac{1}{2} - F \quad \text{for any} \quad w^* > 0.$$

That is, even when the agents' compensation is arbitrarily large, there will be values of the incumbent's popularity at which the agents could deliver the incumbent's victory by conducting fraud but will fail to do so out of the fear that an insufficient fraction among them will engage in fraud.

2.2 An Alternative Information Structure: The Normal Model

The key advantage of the uniform information structure that we have employed so far is the availability of closed form solutions for the fraud, popularity, and majority thresholds, and in turn, the ease with which their political implications can be studied. While simplifying our analysis, the uniform model effectively assumes that the incumbent is ignorant of his likely national-level popularity when setting the reward factor w.

In order the examine how the incumbent's anticipated popularity shapes his choice of w, we now develop an alternative information structure – the Normal model. Specifically, we denote by θ' the probit-transformed version of the incumbent's popularity θ and assume

[22] These comparative statics for F and c (as well as for ϵ and b), which we discuss in the supplementary appendix, continue to hold after accounting for their indirect influence on θ^* via their effect on the incumbent's choice of w^*.

[23] The converse holds for the cost parameter c: while a greater cost of engaging in fraud raises the fraud and popularity thresholds S^* and θ^*, it cannot entirely prevent fraud from succeeding in equilibrium; $S^* < \frac{1}{2} + \epsilon$ and $\theta^* < \frac{1}{2}$ for any $c > 0$.

it follows the Normal distribution with mean $\theta_0^{'}$ and variance σ_0^2, $\theta' \sim \mathcal{N}(\theta_0^{'}, \sigma_0^2)$. Analogously, we let $S_i^{'}$ denote the probit-transformed version of the agents' signals S_i and assume it follows the Normal density with the mean θ' and variance σ^2, $S_i^{'} \sim \mathcal{N}(\theta', \sigma^2)$.[24] We transform θ' and $S_i^{'}$, whose support is on $(-\infty, \infty)$ onto $(0, 1)$, the natural interval θ and S_i, via the probit link, $\theta' = \Phi^{-1}(\theta)$ and $S_i^{'} = \Phi^{-1}(S_i)$. The key advantages of this information structure are that it allows for prior beliefs about θ of an arbitrary mean and precision, and that it maintains a positive amount of uncertainty about θ and S_i along the entire support $(0, 1)$.

The Normal model implies that the incumbent is genuinely supported by a majority of the electorate, $\theta \geq \frac{1}{2}$, when $\theta' \geq 0$. The fraction of agents ϕ that engage in fraud in equilibrium corresponds to one minus the cumulative distribution function of the $\mathcal{N}(\theta^{*'}, \sigma^2)$ density evaluated at $S^{*'}$ and, after observing the incumbent's (probit-transformed) popularity in her precinct $S_i^{'}$, agent i believes that the incumbent's national-level (probit-transformed) popularity θ' follows the Normal density with the mean $\frac{\sigma_0^2 S_i^{'} + \sigma^2 \theta_0^{'}}{\sigma_0^2 + \sigma^2}$ and the variance $\frac{\sigma_0^2 \sigma^2}{\sigma_0^2 + \sigma^2}$.[25]

Unlike in the case of the uniform parameterization, the threshold agent's indifference condition for the Normal model – which we state explicitly in the supplementary appendix – does not have a closed form solution. But jointly with the majority condition it results in a unique set of equilibrium fraud, popularity, and majority thresholds, which can be obtained numerically.[26] The parameter values c=1, $F = \frac{2}{10}$, $b = 100$, $\sigma_0^2 = \frac{1}{4}$, $\sigma^2 = \frac{1}{16}$, and $\theta_0 = \frac{1}{2}$, for instance, imply $w^* = 4.76$, $\theta^* = 0.32$, $S^* = 0.22$, and $\phi^* = 0.90$. That is, agents

[24]Paralleling our earlier interpretation of ϵ, we think of the variance σ^2 of $S_i^{'}$ as a metric of heterogeneity in the incumbent's support across precincts.

[25]This is the standard Bayesian inference for the Normal distribution according to which the posterior mean of $\theta'|S_i^{'}$ is a weighted average of the prior mean $\theta_0^{'}$ and the precinct-level signal $S_i^{'}$ (with the weights in proportion of the prior variance σ_0^2 to the signal variance σ^2); see e.g. Bernardo and Smith (1994, 439). We plot illustrations of the posterior density $\theta|S_i$ in the supplementary appendix.

[26]Uniqueness in the Normal model obtains as long as the signal S_i is sufficiently precise relative to the prior belief about θ; see the supplementary appendix for a precise statement.

engage in fraud only if the incumbent's popularity in their precinct is greater than 22% and fraud secures the incumbent's victory if his national-level popularity is greater than 32%, or equivalently, when at least 90% of agents participate in fraud. This is illustrated in the left panel in Figure 4, which plots the effect of the incumbent's actual national-level popularity θ on the equilibrium and needed levels of fraud as a solid black line (this is the analogue of Figure 3 from our earlier discussion.) We see that a key conclusion from our earlier analysis remains unchanged: in equilibrium, fraud is either over- or under-supplied due to the interplay of the collective action and principle-agent problems. The key difference between the uniform and the normal models is that positive levels of fraud occur at all values of θ in the latter. That is, even when θ is close to 0 (or 1), there is a small measure of agents who observe precinct-level signals implying that the incumbent is overwhelmingly popular (or unpopular.)

The right panel in Figure 4 illustrates the relationship between the incumbent's pre-election belief about his popularity and his equilibrium choice of the reward factor w^* at $\sigma_0^2 = \frac{1}{4}, 1, 2$ (black, dark gray, and light gray lines, respectively.) While equilibrium comparative statics of this relationship are too mathematically complex to be tractable, simulations suggest that the incumbent views his choice w as "insurance" against electoral defeat. That is, given only vague prior information about his eventual popular support θ, the incumbent is willing to pay the agents to conduct fraud even when he expects to prevail, hedging against the odds that he is being too optimistic. As the precision of his prior information about θ improves, however, the incumbent significantly reduces the reward factor at high and low levels of his expected popularity: When he expects an overwhelming victory, compensating the agents would only result in unnecessary fraud and thus a waste of resources; when he expects an overwhelming defeat, a consequential reward factor would have to be so high as to render fraud too expensive. Crucially and as

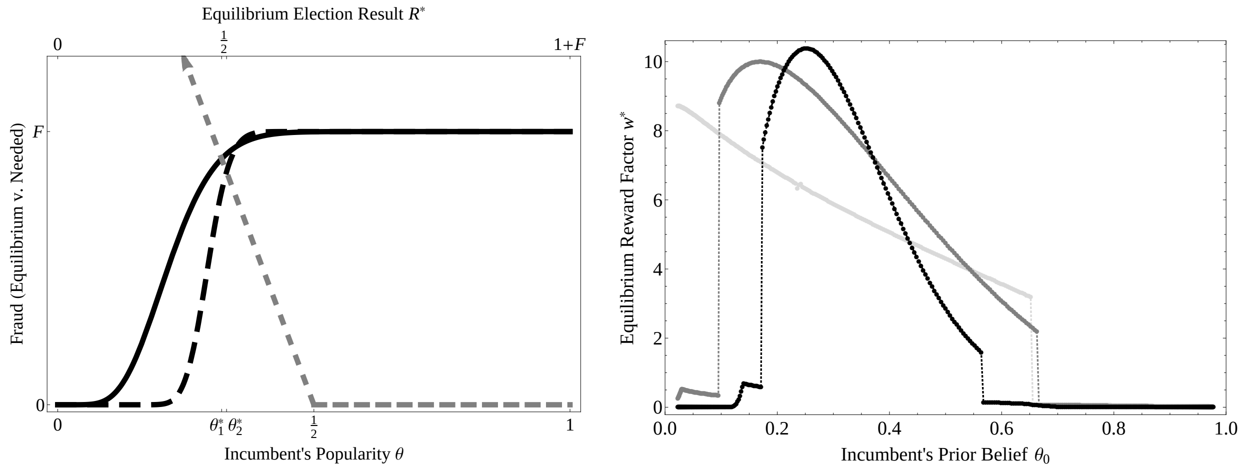

Figure 4: The effect of the incumbent's actual national-level popularity θ on the equilibrium v. needed level of fraud (left); the effect of the incumbent's prior belief θ_0 (about θ) on the equilibrium reward factor w^* (right)

illustrated by Figure 4, for most plausible values σ_0^2, the incumbent optimally chooses nontrivial levels of compensation even when he expects to prevail. Thus in general, fraud serves both those incumbents who know they cannot win a clean election and those who can but prefer to insure against an unlikely defeat.

Finally, the Normal model highlights particularly well the contrast between the rigidity of the equilibrium outcome when the incumbent's popularity is above θ^* and the resounding defeats that occur as the incumbent's popularity crosses below θ^*. The dashed black line in the left panel in Figure 4 plots the effect of θ on the level of fraud when agents receive highly precise signals, $\sigma^2 = \frac{1}{100}$. We see that when each agent has nearly perfect information about the incumbent's national-level popularity θ, shifts in the agents' perception of the incumbent's popularity result in herd-like coordination: On the one hand, virtually all agents conduct fraud on behalf of the incumbent regardless of the actual value of θ as long as $\theta > \theta^*$; on the other hand, the minor shift in the incumbent's popularity from just above to just below θ^* results in his defeat by a margin of about F% as virtually all agents change their behavior from conducting fraud to refraining from it.

3 Empirical Analysis

Our theoretical analysis leads to a number of empirical predictions. First, the conflict between the incumbent and his agents over when and how much fraud to conduct results in either the undersupply or oversupply of fraud from the incumbent's point of view. Second, a key feature of the herd dynamic that occurs among the incumbent's agents is that small shifts in their perception of the incumbent's viability may trigger large aggregate shifts in the amount of fraud conducted, and this effect will be strongest around the threshold θ^*. Jointly, these predictions anticipate a pattern of elections that are won or lost by large margins correlated with the incumbent's popularity. In these elections, the incumbent's defeat may be instigated by only a minor decline in his genuine popularity, and while electoral fraud may be widespread, it will be politically decisive in only a fraction of them. Finally, our analysis predicts that the extent of fraud – and thus its fingerprints – should be increasing in both the incumbent's genuine support as well as his vote share across precincts. Hence we should expect a positive association between the incumbent's popularity and the extent of fraud across both multiple elections and individual precincts in a single election.

As a first step toward the empirical assessment our arguments, we focus on the last of these predictions: That the extent of fraud across individual precincts should be increasing in the incumbent's share of the vote. Unlike the incumbent's genuine support across precincts, this quantity is readily observable in most elections. In order to measure the extent of fraud, we take advantage of a particular form of fraud that arguably took place during the 2011 legislative and 2012 presidential elections in Russia: the rounding of the incumbent Vladimir Putin's/United Russia Party's precinct-level vote share to some higher multiple of five by the regime's local operatives (Gehlbach 2012; Klimek et al. 2012;

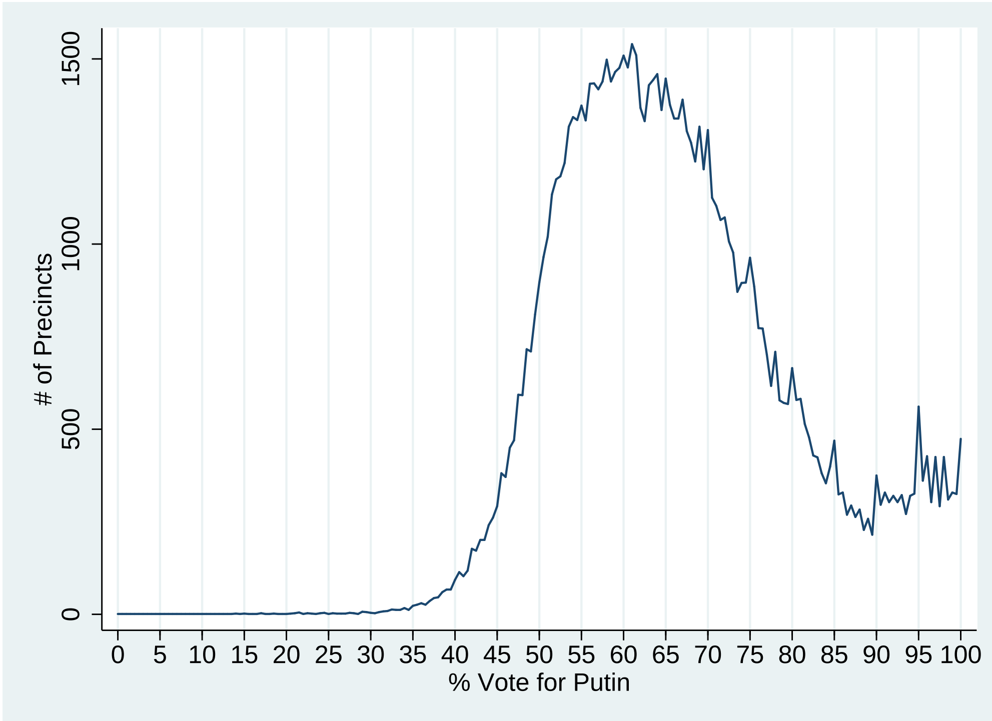

Figure 5: The distribution of Vladimir Putin's precinct-level vote share in the 2012 presidential election

Kobak, Shpilkin, and Pshenichnikov 2012). While this form of fraud was most likely only one among several forms of electoral manipulation during these elections, its execution at the level of individual precincts provides an opportunity to assess whether the extent of fraud was indeed increasing in the incumbent's precinct-level vote share – as our theoretical model anticipates.

Several features of the 2011 legislative and 2012 presidential elections in Russia match the key elements of our model: Extant empirical research shows that fraud did indeed occur in these elections (Enikolopov et al. 2013); that it was executed locally by operatives within the state bureaucracy, the public sector, and the United Russia party (Frye, Reuter, and Szakonyi 2012); and that the regime had only rough information about its genuine support due to significant misreporting in public opinion surveys (Kalinin

2013). Due to space constraints, we focus below on the 2012 presidential election; our analysis of the 2011 parliamentary election – which provides even stronger support for our arguments – can be found in the supplementary appendix.

As a preliminary step, we establish that multiples of five are indeed over-represented in Vladimir Putin's precinct-level vote shares. Figure 5 plots the distribution of Putin's vote share in the 2011 presidential election across more than 90,000 precincts. In spite of the large number precincts, we see a suspicious lack of smoothness due to spikes that mostly coincide with multiples of five, especially in the range from 55% to 100%. In order to examine the distribution of digits in precinct-level vote shares more formally, we round each candidate's vote share to the nearest multiple of 0.5, extract the unit and the first decimal place digits (e.g. both 76.481 and 46.532 become 6.5), and pool them into the twenty resulting digit pairs.[27] The distribution of these pooled digit pairs in Vladimir Putin's precinct-level results is displayed in Figure 6.

Consistent with our discussion above, precinct-level vote shares that end in either 0.0 or 5.0 are over-represented for Putin, and crucially, only for Putin.[28] This is confirmed by a series of likelihood ratio independence tests. Assuming that neighboring digits should be distributed approximately uniformly, we compute the G^2 statistic (Agresti 1996, 36) for the frequencies of 0.0 and 5.0 and the two digit pairs to their left and right. These are the digit pairs $\{9.0, 9.5, 0.0, 0.5, 1.0\}$ and $\{4.0, 4.5, 5.0, 5.5, 6.0\}$, respectively. The G^2 statistics (74.8 and 40.0 with $df = 4$) strongly suggest that these digit frequencies are not uniform (both p-values $= 0$). Once we exclude the digit pairs 0.0 and 5.0, however, the remaining digit

[27] Our results do not depend on the extent of rounding; rounding to units or one decimal place leads to identical conclusions. We dropped all precincts with fewer than 50 voters in order to exclude special-category precincts (hospitals, military units) and to eliminate small-N effects on precinct-level vote shares.

[28] Vote shares that end in 0.0 tend to be over-represented for the three minor candidates (Prokhorov, Zhirinovsky, and Mironov) but only because there are many precincts in which these candidates receive close to 0% of the vote. After excluding precincts where any candidate received less than 1% of the vote, precinct-level vote shares that end in either 0.0 or 5.0 remain over-represented for Putin but not for any other candidate.

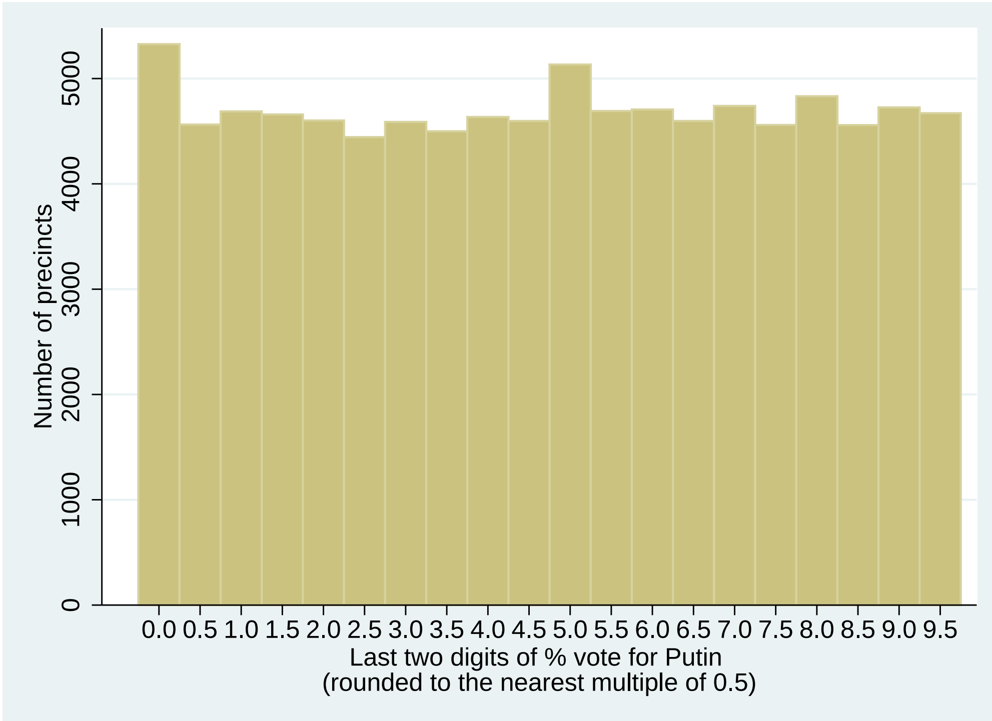

Figure 6: The distribution of the pooled unit and the first decimal place digits in Putin's precinct-level vote share (after rounding to the nearest multiple of 0.5)

frequencies are consistent with uniformity ($G^2 = 3.2$ and 1.7 with $df = 3$ implying p-values of 0.37 and 0.65, respectively). Thus the lack of uniformity in Putin's digit distribution is overwhelmingly due to the over-representation of the multiples of 5 among his precinct-level vote shares.[29]

[29]This is implied by the partitioning property of chi-squared statistics (Agresti 1996, 39-40); an analysis based on standardized residuals implies the same conclusion.

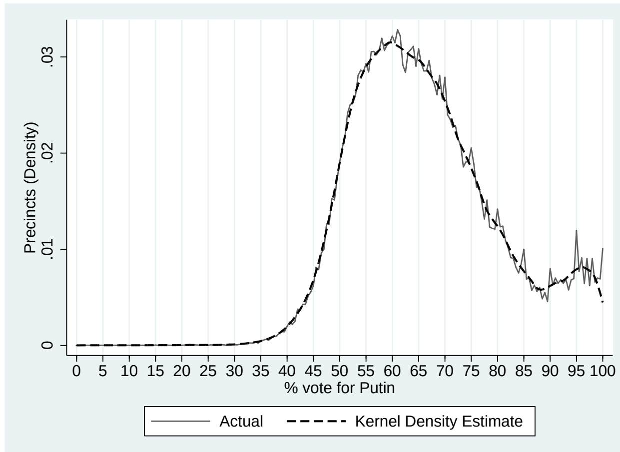

(a) Putin

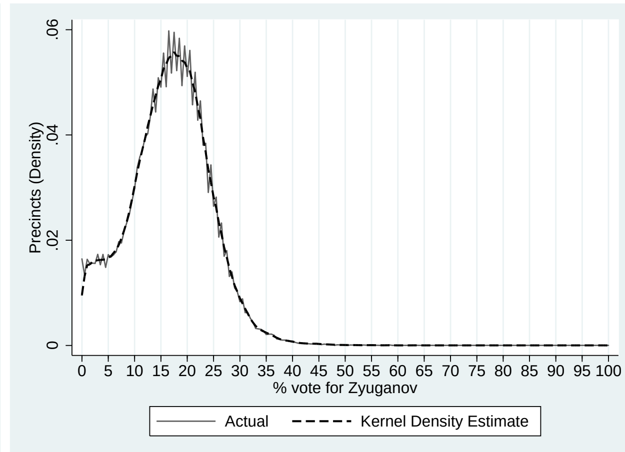

(b) Zyuganov

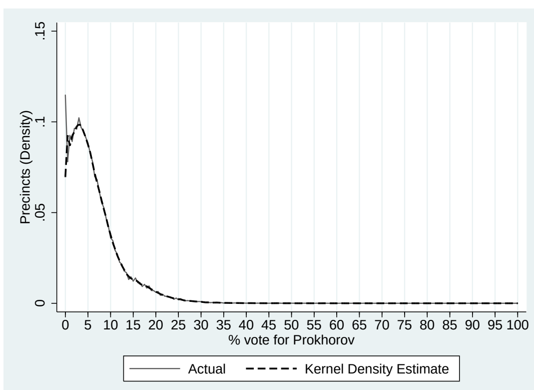

(c) Prokhorov

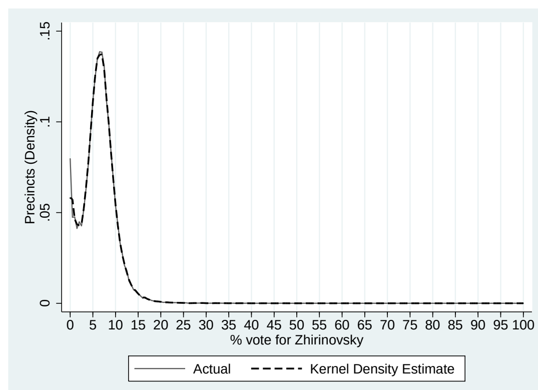

(d) Zhirinovsky

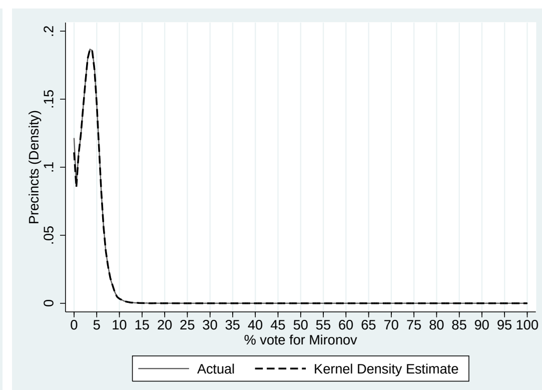

(e) Mironov

Figure 7: The distribution (gray solid line) and kernel density estimate (black dashed line) of each candidate's precinct-level results

We now turn to our main empirical test and examine whether the over-representation of the multiples of 5 for Putin is indeed increasing in his precinct-level vote share as predicted by our model. In order to do so, we first develop a measure of the ruggedness in the distribution of Putin's precinct-level results: We take the difference between the empirical distribution of a candidate's precinct-level results and its optimal kernel density estimate. The former is based on a histogram with 0.5 bin width; the latter employs the (optimal) Epanechnikov kernel and the optimal bandwidth estimate (Cameron and Trivedi 2005, Chapter 9).[30] Figure 7 plots the kernel density estimate of each candidate's precinct-level results by a black dashed line along with their actual empirical distribution (gray solid line). We see a nearly perfect overlap between the distribution of precinct-level vote shares and the corresponding kernel density estimates for the three minor candidates (Prokhorov, Zhirinovsky, and Mironov), some ruggedness unrelated to multiples of 5 for Zyuganov, and significant ruggedness correlated with multiples of 5 for Putin.

Figure 7 further highlights that the ruggedness in the distribution of Putin's precinct-level results is not only substantial but also increasing in his precinct-level vote share as our theoretical framework anticipated. In order to quantify how anomalous this ruggedness is and to evaluate the strength of its association with Putin's precinct-level vote share, we employ two standards. The first is empirical: We judge the ruggedness of Putin's precinct-level results by the standard of his four competitors. Specifically, we calculate the difference between the empirical distribution of Putin's competitors' precinct-level results and their kernel density estimate, pool these residuals, and use their 95th and 99th percentiles as our first benchmark for judging how anomalous the ruggedness of Putin's precinct-level results is.[31] Figure 8 plots these residuals separately for Putin (diamonds),

[30]The optimal bandwidth estimate minimizes the mean integrated squared error based on a Gaussian kernel. In the supplementary appendix, we confirm the robustness of these results by using both twice and half the optimal bandwidth estimate (as recommended in Cameron and Trivedi 2005, 304).

[31]In order for these residuals to be comparable across candidates, we divide them by the height of the

Zyuganov (squares), and the remaining three minor candidates (Prokhorov, Zhirinovsky, and Mironov). We see that with the exception of a few poorly fitting corner values for the minor candidates, all residuals above the 95th and 99th percentile benchmarks belong to either Putin or Zyuganov (see below for our accunt of Zyuganov's residuals). Crucially, only Putin's residuals coincide with the multiples of five and are significantly increasing in his vote share.[32]

We arrive at identical conclusions when we employ an alternative, theoretical benchmark for evaluating the ruggedness of Putin's precinct-level results. We compute the 95% and 99% asymptotic confidence intervals for the kernel density estimate of Putin's results and treat the observations that lie outside these confidence intervals as anomalously rugged. Just as before, Putin's and Zyuganov's residuals are significantly larger than those of the remaining candidates and only Putin's residuals are increasing in his precinct-level vote share.[33] Hence judging by two different standards – that of Putin's competitors and asymptotic confidence intervals – the distribution of Putin's results is indeed suspiciously rugged at multiples of five and this ruggedness is increasing in his precinct-level vote share as anticipated by our theoretical arguments.

Our analysis of the ruggedness in the distribution of precinct-level results in the 2012 Russian presidential election also sheds light on how fraud was conducted. Throughout our discussion above, Zyuganov was the only candidate other than Putin with a significant amount of ruggedness in his precinct-level results. Yet crucially, this ruggedness did not coincide with multiples of five and was not increasing in his vote share. This observation suggests that rather than stuffing ballots in order to round Putin's precinct-level vote share

corresponding candidate's KDE. Each residual then measures the difference between a candidate's empirical distribution of precinct-level results and its KDE relative to its own height.

[32]When we regress Putin's residuals on his vote share, the vote share coefficient is positive and statistically significant at the 0.01 significance level; the regression coefficient on the other candidates' vote share is statistically significant but negative.

[33]Due to space constraints, these plots are presented in the supplementary appendix.

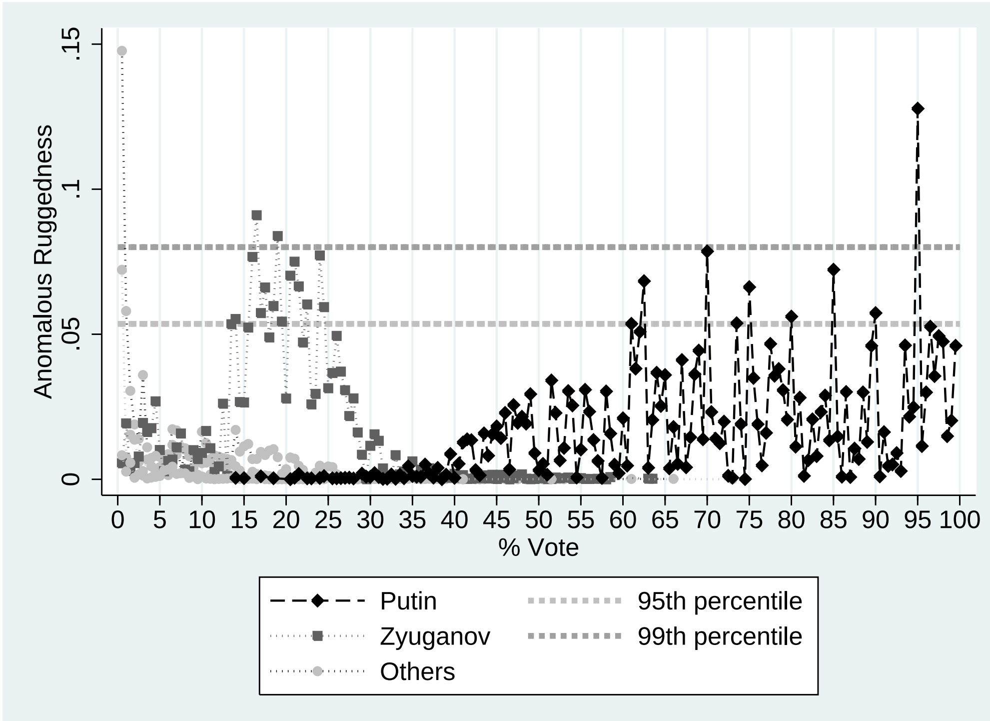

Figure 8: The difference between the empirical distribution of each candidate's precinct-level results and its kernel density estimate

to some higher multiple of five, Putin's local operatives may have been instead stealing votes from Zyuganov. This makes logistical sense: Zyuganov was the only major opposition candidate in this election and hence the only candidate with a number of votes large enough in most individual precincts that could be transferred to Putin's column in order to round his vote share to some higher multiple of five. In order to evaluate this hypothesis, we add Putin's and Zyuganov's precinct-level votes and examine the ruggedness in the resulting distribution of vote shares. As Figure 9 reveals, the significant ruggedness in the two candidates' individual vote-share distributions now disappears – supporting our hypothesis of vote-stealing from Zyuganov.[34]

[34]This is confirmed by a formal analysis of the KDE residuals as above; see the supplementary appendix. Vote-stealing from Zyuganov also explains why the digit methods of Beber and Scacco (2012), Cho and Gaines (2007), and Mebane and Kalinin (2009) do not detect anomalies in precinct-level vote *totals* (rather than vote *shares*) for Putin/United Russia: the number of votes needed to round the incumbent's vote share to some higher multiple of 5 was not fabricated but rather computed by local operatives.

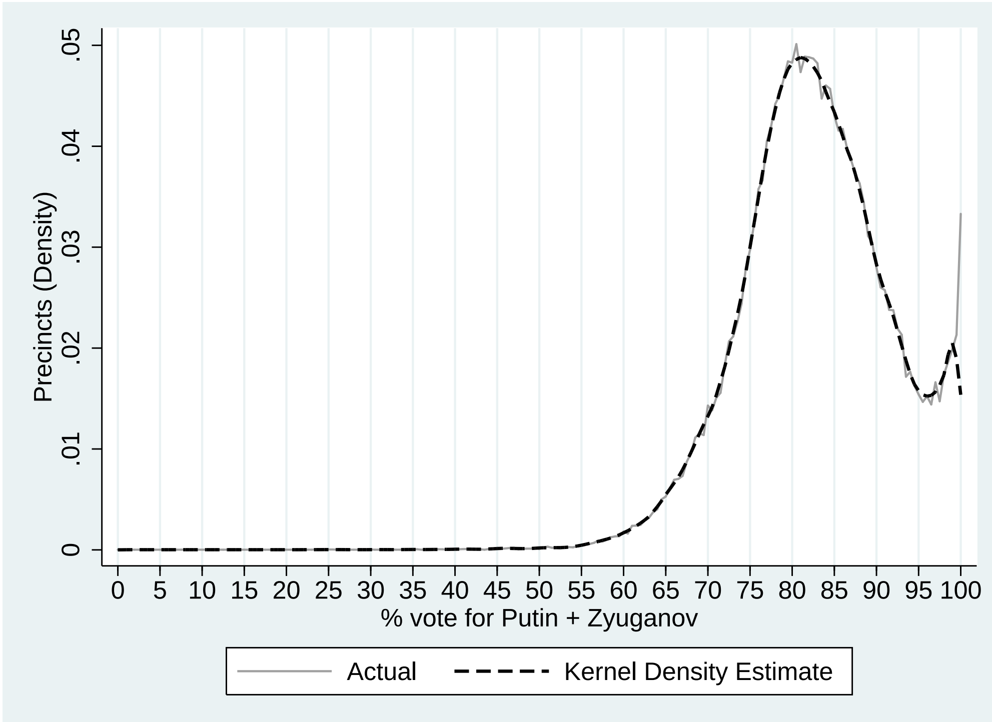

Figure 9: The distribution (gray solid line) and kernel density estimate (black dashed line) of the sum of Putin's and Zyuganov's precinct-level vote share

To summarize, our analysis of the 2011-2012 Russian legislative and presidential elections (see the supplementary appendix for the former) supports a key prediction from our theoretical model: The extent of fraud in these elections is indeed increasing in the incumbent Vladimir Putin's and United Russia's precinct-level vote share and, crucially, in only their vote share. We took advantage of a particular type of fraud that took place during these elections – the rounding of the regime candidate's and party's vote share to some higher multiple of five. Consistent with the findings from earlier analyses of these elections, we found that precinct-level vote shares corresponding to multiples of five are indeed over-represented in Putin's and United Russia's results but not those of other candidates or parties. By analyzing the ruggedness in the distribution of precinct-level results of all competing candidates and parties, we found that the votes necessary for this fraud were most likely stolen from the leading opposition candidate Gennady Zyuganov and his Communist Party.

4 Conclusion

Why do incumbents who could arguably win a clean election engage in fraud? Our analysis of the principle-agent and collective action problems in the political organization of electoral fraud suggests one answer. Because most fraud is local, and thus ultimately perpetrated by a large number of operatives, incumbents have only imperfect control over whether and how much fraud will be conducted. Each agent understands that the difference between her promotion and prosecution rests on whether her involvement in fraud will result in the incumbent's eventual victory or defeat. In turn, the agents' perception of the incumbent's (lack of) popularity and the ensuing burden on their collective complicity plays a crucial role. The aggregate result is too much fraud for incumbents who need it the least and too little fraud for those who need it the most. The seeming invincibility of some incumbents and the suprising fragility of others are thus two sides of the same political logic.

The political dynamic that emerges out of the interplay of the principle-agent and collective action problems that we examine naturally extends to other settings, most directly to patronage politics and repression. Just as in the present setting, the rewards and punishment of patronage brokers and repressive agents are contingent upon the success or survival of their principals. In patronage politics, for example, a broker's effort on behalf of a candidate should be driven by her expectation of the candidate's ultimate victory and hence the likelihood that her effort will not be wasted. Meanwhile in repression, the willingness of a repressive agent to engage in legally questionably suppression of opposition should be driven by his expectation of the regime's survival and hence his immunity from prosecution. Just as in the case of fraud, candidates and leaders who want to motivate their operatives to engage in electioneering or repression have only limited control over whether and how much of it will be conducted. Our analysis of election fraud thus suggests that the interplay of the principle-agent and collective action problems in the political organization of clientelism and repression should have similar, potentially undesirable aggregate consequences.

References

Agresti, Alan. 1996. *An introduction to categorical data analysis*. New York: Wiley.

Ahlquist, John S., Kenneth R. Mayery, and Simon Jackman. 2013. "Alien Abduction and Voter Impersonation in the 2012 US General Election: Evidence from a survey list experiment." Unpublished Manuscript, University of Wisconsin and Stanford University.

Alvarez, R. Michael, Thad E. Hall, and Susan D. Hyde, eds. 2008. *Election Fraud: Detecting and Deterring Electoral Manipulation*. Washington, D.C.: Brookings Institution Press.

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

Benton, Allyson Lucinda. 2013. "The Territorial Logic of Electoral Manipulation in an Electoral Authoritarian Regime: The Case of Mexico." Unpublished Manuscript, CIDE.

Bernardo, Joseé M., and Adrian F.M. Smith. 1994. *Bayesian theory*. Chichester, New York: Wiley.

Birch, Sarah. 2011. *Electoral malpractice*. Oxford: Oxford University Press.

Boix, Carles, and Milan W. Svolik. 2013. "The Foundations of Limited Authoritarian Government: Institutions and Power-sharing in Dictatorships." *Journal of Politics* 75(2): 300–

316.

Bueno de Mesquita, Ethan. 2010. “Regime change and revolutionary entrepreneurs.” *American Political Science Review* 104(3): 446–466.

Cameron, A. Colin, and Pravin K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.

Cantú, Francisco. Forthcoming. “Identifying Irregularities in Mexican Local Elections.” *American Journal of Political Science* .

Cantú, Francisco, and Sebastián M. Saiegh. 2011. “Fraudulent Democracy? An Analysis of Argentina’s Infamous Decade Using Supervised Machine Learning.” *Political Analysis* 19(4): 409–433.

Carlsson, Hans, and Eric van Damme. 1993. “Global Games and Equilibrium Selection.” *Econometrica* 61(5): 989–1018.

Castañeda, Jorge G. 2000. *Perpetuating power: How Mexican presidents were chosen*. New York: New Press.

Cho, Wendy K. Tam, and Brian J. Gaines. 2007. “Breaking the (Benford) Law: Statistical Fraud Detection in Campaign Finance.” *The American Statistician* 61(3): 218–223.

Dragu, Tiberiu, and Mattias Polborn. 2013. “The Administrative Foundation of the Rule of Law.” *The Journal of Politics* 75(4): 1038–1050.

Edmond, Chris. 2013. “Information manipulation, coordination, and regime change.” *Review of Economic Studies* 80(4): 1422–1458.

Egorov, Georgy, and Konstantin Sonin. 2012. “Incumbency Advantages in Non-Democracies.” Unpublished Manuscript, Northwestern University and New Economic School.

Enikolopov, Ruben, Vasily Korovkin, Maria Petrova, Konstantin Sonin, and Alexei Zakharov. 2013. “Field experiment estimate of electoral fraud in Russian parliamentary elections." *Proceedings of the National Academy of Sciences* 110(2): 448–452.

Fearon, James D. 2011. "Self-enforcing Democracy." *The Quarterly Journal of Economics* 126(4): 1661–1708.

Frye, Timothy, Ora John Reuter, and David Szakonyi. 2012. "Political Machines at Work: Voter Mobilization and Electoral Subversion in the Workplace." Unpublished Manuscript, Columbia University and University of Wisconsin.

Gandhi, Jennifer, and Adam Przeworski. 2011. "Holding onto Power by Any Means? The Origins of Competitive Elections." Unpublished Manuscript, Emory University and New York University.

Gehlbach, Scott. 2012. "Electoral Fraud in Russia: Report from the Russian Blogosphere." *The Monkey Cage* .

Gehlbach, Scott, and Alberto Simpser. 2011. "Electoral Manipulation as Bureaucratic Control." Unpublished Manuscript, University of Wisconsin and University of Chicago.

Greene, Kenneth F. 2007. *Why Dominant Parties Lose: Mexico's Democratization in Comparative Perspective.* New York: Cambridge University Press.

Hyde, Susan D. 2008. "The Observer Effect in International Politics: Evidence from a Natural Experiment." *World Politics* 60(1): 37–63.

Hyde, Susan D. 2011. *The Pseudo-Democrat's Dilemma: Why Election Observation Became an International Norm.* Ithaca, N.Y.: Cornell University Press.

Hyde, Susan D., and Nikolay Marinov. 2009. "Does Information Facilitate Self-Enforcing Democracy? The Role of International Election Monitoring." Unpublished Manuscript, Yale University.

Ichino, Nahomi, and Matthias Schündeln. 2012. "Deterring or Displacing Electoral Irregularities? Spillover Effects of Observers in a Randomized Field Experiment in Ghana." *Journal of Politics* 74(1): 292–307.

Kalinin, Kirill. 2013. "Unifying the Concepts of Electoral Fraud and Preference Falsi cation: Case of Russia." Unpublished Manuscript, University of Michigan.

Klimek, Peter, Yuri Yegorov, Rudolf Hanel, and Stefan Thurner. 2012. "Statistical detection of systematic election irregularities." *Proceedings of the National Academy of Sciences* 109(41): 16469–16473.

Kobak, Dmitry, Sergey Shpilkin, and Maxim S. Pshenichnikov. 2012. "Statistical anomalies in 2011-2012 Russian elections revealed by 2D correlation analysis." Unpublished Manuscript, Imperial College London.

Kuran, Timur. 1991. "Now Out of Never: The Element of Surprise in the East European Revolution of 1989." *World Politics* 44: 7–48.

Larreguy, Horacio A., Cesar E. Montiel Olea, and Pablo Querubin. 2014. "The Role of Labor Unions as Political Machines: Evidence from the Case of the Mexican Teachers' Union." Unpublished Manuscript, Harvard University and New York University.

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

Little, Andrew T. 2012. "Elections, Fraud, and Election Monitoring in the Shadow of Revolution." *Quarterly Journal of Political Science* 7(3): 249–283.

Little, Andrew T. 2013. "An Informational Theory of Noncompetitive Elections." Unpublished Manuscript, New York University.

Magaloni, Beatriz. 2006. *Voting for autocracy: Hegemonic party survival and its demise in Mexico*. New York: Cambridge University Press.

Martinez-Bravo, Monica. 2014. "The Role of Local Officials in New Democracies: Evidence from Indonesia." *American Economic Review* 104(4): 1244–87.

Mebane, Walter R., and Kirill Kalinin. 2009. "Comparative Election Fraud Detection." Unpublished Manuscript, University of Michigan.

Miller, Michael K. 2013. "Elections, Information, and Policy Responsiveness in Autocratic Regimes." Unpublished Manuscript, Australian National University.

Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Applications." In *Advances in Economics and Econometrics*, ed. Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky. New York: Cambridge University Press pp. 56–114.

Myagkov, Mikhail, Peter C. Ordeshook, and Dimitri Shakin. 2009. *The Forensics of Election Fraud*. Cambridge: Cambridge University Press.

Przeworski, Adam. 2011. "Force and Elections." Unpublished Manuscript, New York University.

Rozenas, Arturas. 2013. "Office Insecurity and the Informational Foundations of Non-Competitive Elections." Unpublished Manuscript, New York University.

Sabato, Larry. 2013. *The Kennedy half-century: the Presidency, assassination, and lasting legacy of John F. Kennedy*. New York, NY: Bloomsbury.

Schedler, Andreas. 2012. "The Politics of Uncertainty: Sustaining and Subverting Electoral Authoritarianism." Unpublished Manuscript, CIDE.

Shadmehr, Mehdi, and Dan Bernhardt. 2011. "Collective Action with Uncertain Payoffs: Coordination, Public Signals, and Punishment Dilemmas." *American Political Science Review* 105(4): 829–851.

Simpser, Alberto. 2013. *Why Governments and Parties Manipulate Elections: Theory, Practice, and Implication*. New York: Cambridge University Press.

Sjoberg, Fredrik M. 2013. "Making Voters Count: Evidence from Field Experiments about the Efficacy of Domestic Election Observation." Unpublished Manuscript, Harriman Institute Working Paper, Columbia University.

Smith, Alastair, Bruce Bueno de Mesquita, and Tom LaGatta. 2013. "Group Incentives and Rational Voting." Unpublished Manuscript, New York University.

Svolik, Milan, and Svitlana Chernykh. 2013. "Third-Party Actors and the Success of Democracy: How Electoral Commissions, Courts, and Observers Shape Incentives for Election Manipulation and Post-Election Protest." Unpublished Manuscript, University of Illinois and Oxford University.

Treisman, Daniel. 2011. *The return: Russia's journey from Gorbachev to Medvedev.* New York: Free Press.

Tucker, Joshua A. 2007. "Enough! Electoral Fraud, Collective Action Problems, and Post-Communist Colored Revolutions." *Perspectives on Politics* 5(4): 535–551.

Wintrobe, Ronald. 1998. *The political economy of dictatorship.* Cambridge: Cambridge University Press.