Fraud and Monitoring in Non-competitive Elections

Andrew T. Little

Political Science Research and Methods / *FirstView* Article / September 2014, pp 1 - 21
DOI: 10.1017/psrm.2014.9, Published online: 23 September 2014

 doi:10.1017/psrm.2014.9

Fraud and Monitoring in Non-competitive Elections*

ANDREW T. LITTLE

This article develops a game-theoretic model that reconciles three facts: (1) fraud is pervasive in non-competitive elections, (2) domestic and international monitoring of elections have become nearly universal and (3) incumbent regimes often invite monitoring and still cheat. The incumbent regime commits fraud to manipulate the information generated by a non-competitive election before a political interaction with some audience. *The audience expects fraud, so the incumbent commits fraud because she would appear weak if not doing so. Increasing the visibility of fraud with monitoring is valuable because it lowers the equilibrium level of costly fraud without changing how popular the incumbent appears. The core results hold under multiple extensions, which produce a rich set of comparative static results.*

"Whereas we might be reasonably sure of the result [of the 2012 presidential election], we should not assume that there is little at stake."
Mikhail Khodorkovsky, dissident Russian oligarch (Khodorkovsky 2012)

"Web cameras and transparent ballot boxes help organize an open, transparent and honest election, the kind of election that only Russia has had so far."
Vladamir Churov, head of Russia's Central Elections Commission (*Pano News* 2012)

In March 2012 Vladamir Putin easily re-claimed the presidency of Russia following an eight-year hiatus as prime minister. Putin won over 63 percent of the vote, with the four challengers deemed sufficiently nonthreatening to run earning between 4 percent and 20 percent.[1] According to international monitors from the Organization for Security and Cooperation in Europe (OSCE) and others, this wide margin was partially driven by large amounts of electoral fraud (OSCE 2012).

Leading up to the election, even major opposition figures—such as then-jailed businessman Mikhail Khodorkovsky—acknowledged there was little doubt that Putin would win the 50 percent required to avoid a runoff. Following the election, various domestic monitoring groups estimated the "true Putin vote" to be slightly below or above

* Andrew Little is Assistant Professor, Department of Government, Cornell University, 301 White Hall, Ithaca NY 14853 (andrew.little@cornell.edu). Previous versions of this article were presented at the 2011 Midwest Political Science Association Meeting (with the title "Rational Expectations and Electoral Fraud"), the 2011 Empirical Implications of Theoretical Models Summer Institute at the University of Chicago Harris School, the 2011 American Political Science Association Meeting and the NYU Job Market Paper Workshop. Many thanks to participants in these seminars, as well as Bernd Beber, Tiberiu Dragu, Marko Klašnja, S.P. Harish, James Hollyer, Burt Monroe, Adam Przeworski, Alberto Simpser, Hannah Simpson, Fredrik Sjoberg, Alastair Smith, Mike Tiernay, Dustin Tingley and Josh Tucker for comments and discussion. Supplemental material can be found at http://andrewtlittle.com/research.html

[1] Eleven potential candidates were denied registration, including the Yabloko party nominee Grigory Yavlinskiy.

50 percent (Korolyov 2012).[2] So, depending on exactly how much fraud occurred and how well opposition voters would have coalesced around a single candidate in a second round, Putin almost certainly could have won a runoff with no fraud, and likely could have avoided a runoff with a clean first round. Regardless, Putin would have prevailed easily with substantially less cheating.

On top of the scrutiny brought by international and domestic monitors—many voluntarily invited by the Russian government—Putin's regime spent \$300 million installing webcams in nearly every polling station.[3] In addition to generating less overtly political footage such as canoodling teenagers and makeshift discotheques, the webcams revealed poorly disguised ballot box stuffing and other electoral malfeasance (OSCE 2012).

Many aspects of this case are puzzling. First, why commit so much fraud when Putin would have won a clean—or at least much less fraudulent—election? Second, if fraud is effective for winning elections or for other reasons, and if fraudulent acts "are things that only its victims want publicized" (Lehoucq 2003), why invite monitors and install webcams to make the fraud more visible? Third, why invite the scrutiny of monitors *and* cheat in front of them?

As detailed below, these phenomena occur in many elections beyond the Russian case. The first contribution of this article is to develop a game-theoretic model of a non-competitive election in which the incumbent chooses to commit fraud while also inviting monitoring that makes fraud more visible. The second main contribution is a set of extensions to the baseline model that generates a rich set of empirical predictions that could be used to test this approach in future work. Though the purpose of the model is not to explain webcams and other forms of monitoring in recent Russian elections *per se*, this example will be used to illustrate various aspects of the argument. In contrast to most existing work on these questions, which focuses on international pressure to invite monitors and appear to run clean elections, the incumbent decisions are driven by domestic politics.

The model assumes that non-competitive elections—that is, elections with a trivial probability that the incumbent leader loses and accepts this result—are primarily about information generation and that fraud is an attempt to manipulate this information. The idea that information generation is central to why non-competitive elections are held has become common (for example, Magaloni 2006; Gandhi and Lust-Okar 2009; Blaydes 2011; Little 2011, 2012). If the turnout and/or incumbent vote share generate meaningful information, then the ruling regime has an incentive to distort these outcomes with fraud even when victory is assured. Monitoring—which can include international monitors as well as domestic monitoring groups or "passive" monitoring (Herron 2010) such as installing webcams—is modeled as an imperfect signal of how much fraud was committed.

The actor that the incumbent wants to impress, which I call the *audience*, adjusts the information he draws from the electoral signal based on his conjecture of how much fraud was committed. This renders fraud ineffective in equilibrium. Still, the incumbent cannot commit to choosing no fraud because it is a partially hidden action. So he pays a cost to perpetrate fraud even though the audience is not fooled.

[2] The implied level of fraud according to these groups is close to a more rigorous estimate of the level of fraud in Moscow in the 2011 parliamentary election (Enikolopov *et al.* 2013).

[3] The OSCE requires member states to invite monitors (Kelley 2008). However, this does not always happen in practice: for example, Russian parliamentary and presidential elections in 2007 and 2008 were not observed by the OSCE following a dispute over the size and scope of the mission.

Inviting domestic and international monitoring provides a solution to this problem by credibly lowering the audience's expectation of how much fraud is committed. In equilibrium, the incumbent receives the same average level of support from the audience regardless of the monitoring decision. So in the baseline model, the incumbent invites as much monitoring as possible to tie her hands from committing expensive but ineffective fraud. However, as long as the monitoring signal is not perfectly informative, committing more fraud than expected would lead to higher levels of support. So, the equilibrium level of fraud is strictly positive even though the incumbent chooses as much monitoring as possible. In sum, the motivating empirical facts arise as equilibrium behavior: fraud and taking actions to make it more difficult both occur in a model of a non-competitive election.

The remainder of the analysis extends the model with looser assumptions and other modifications. The central results hold with appropriate caveats (1) in competitive elections, (2) if monitoring imposes direct costs on the incumbent, (3) under an alternative information structure, (4) when there are multiple, biased or strategic monitors and (5) if the incumbent or audience has private information. These extensions also generate comparative static results that could be tested in future empirical work.

EXTANT WORK

The empirical puzzles described above are not limited to recent elections in Russia or even a small number of cases. Fraud is rarely pivotal in determining the winners of elections (Lehoucq 2003), and it is even more common in elections that are not close (Simpser 2013). While a large margin of victory reflects the potentially rampant fraud, it is still puzzling why ruling regimes commit an incremental amount of fraud to raise the total from, as in the recent Russian election, 62 to 63 percent.[4]

Surely the fraud required to get that last percentage point is not about winning. This insight is not new; for example, a seminal article on the types of electoral manipulation in contemporary regimes notes that the greatest offenders hold elections ''without running the risks of democratic uncertainty'' (Schedler 2002). Still, with the exception of some recent and contemporaneous work (for example, Simpser 2013; Gehlbach and Simpser 2013) discussed in more detail below, most of the small (but growing) game-theoretic literature on fraud assumes that it is primarily used to win elections. Tellingly, in one exception to this rule fraud is used to manipulate beliefs in order to win future elections (Simpser 2005). Magaloni (2009) and Fearon (2011) present models in which election monitoring can help incumbents by rendering the election results credible—similar in spirit to the arguments made here—though again in the context of competitive elections.

In addition to an excessive focus on fraud as a tool to pass a victory threshold, existing work tends to ignore the question of how fraud can be effective if the defrauded are strategic actors who form beliefs about the amount of fraud committed and behave accordingly.[5] Existing models with a standard treatment of these beliefs tend to rely on

[4] While fraud is committed by opposition groups and local party officials as well, for clarity I center the argument around national incumbent leaders. While the model here has a top-down approach, it could easily be applied to fraud decisions made at a more local level in order to seem competent to national leadership.

[5] Some models dodge the question of why fraud is effective by either not treating opposition figures and citizens as strategic actors (Gandhi and Przeworski 2009; Hyde 2011), giving the other actors

binary fraud decisions or type spaces: in one, fraud either does not happen in equilibrium or is so rampant that elections are completely uninformative (Fearon 2011); the other assumes that only some types of incumbents have an incentive to commit fraud (Simpser 2005). An appealing aspect of the approach used here is that the incumbent always has an incentive to commit fraud and will always do so, though the amount chosen can be arbitrarily small.

Closer to the approach here, Simpser (2013) argues that blatant manipulation of non-competitive elections is used as a costly signal of repressive capacity.[6] Kuhn (2012) also treats fraud as a hidden action to distort information, though using a model with costly signaling-like dynamics, in the context of a competitive election, and without modeling the choice to invite monitoring.[7] Another informational argument is made in Rozenas (2011) and Gehlbach and Simpser (2013), in which the main effect of fraud is to "jam" the electoral signal (that is, render it less informative) in a manner favorable to the incumbent. While this jamming effect can make fraud useful in some contexts (and could have different implications for inviting monitoring), most fraud has the more direct effect of increasing the votes for the incumbent, which is how I model it here.

While these arguments and mine both center around information, they differ on *how* fraud manipulates information. Most notably, for fraud to serve as an effective costly signal of strength, it must be observed, and jamming models can work whether fraud is directly observed or not. These approaches fruitfully contrast with theories of fraud that ignore blatant manipulation, but a large amount of fraud *is* purposefully clandestine.

The models in this article treat fraud as a *partially hidden* action. The formal structure is most similar to theories of incentives for workers (Prendergast 1999) and, more specifically, "career concerns" models (Holmstrom 1999). This type of model has been applied in political science settings such as political business cycles (Lohmann 1998), politician effort (Ashworth 2005) and terrorist violence (Bueno De Mesquita 2010).

In a previous article I modeled fraud with career concerns dynamics in a game that focuses more on the decision to hold an election and strategic interactions among potential protesters (Little 2012).[8] The baseline model here generates more straightforward and clear versions of some secondary results from that article. This makes the logic of why fraud and monitoring both occur in equilibrium more transparent, and allows for tractable extensions that generate a rich set of empirical predictions and demonstrate that the logic of the model is robust to loosening some of the more problematic assumptions.

(F'note continued)
non-standard beliefs (the incomplete information model of Magaloni 2009) or making these actors aware of fraud but unable to take actions that would induce less fraud (the complete information model of Magaloni 2009).

[6] Simpser (2013) also argues that blatant fraud can be used to discourage citizens from coordinating against the incumbent regime, though others have argued that public signals indicating high levels of fraud can *help* such coordination (Tucker 2007; Little, Tucker and LaGatta 2012).

[7] Simpser (2005) and Kuhn (2012) also have equilibrium with a reversed costly signaling logic in which weak types commit fraud to mimic strong types, who commit no fraud. In both models, fraud is a binary choice, and I am skeptical that this equilibrium would hold if fraud is a continuous choice, as the strong type would always have an incentive to commit a marginal amount of fraud to "separate" from the weak type under a refinement in the spirit of D1.

[8] Egorov and Sonin (2011) treat elections and fraud in a similar fashion. See Little (2012) for a comparison of these models.

Treating fraud in this manner generates new (and sometimes surprising) results about the causes and effects of international election monitoring. By the 1990s, inviting international election monitors was the "best-established, most visible, and often best-funded type of democracy-related assistance" (Carothers 1997). More recently, international monitoring of elections has become nearly universal, particularly (though not exclusively) outside of consolidated democracies (Kelley 2008; Hyde 2011, 2012).

In addition to international monitoring, many forms of domestic monitoring are becoming prevalent. Over half of the countries in the world have meaningful independent electoral commissions (ACE Electoral Knowledge Network 2012). Further, Russia is not the only country to use webcams in polling stations: Azerbaijan did so in nearly 10 percent of polling stations for its 2008 presidential election, and Ukraine followed suit for parliamentary elections in 2012 (Herron 2010; Sjoberg 2012). In an even more extreme form of monitoring, the ballot counting for contested seats in Singapore's Parliament is done in the presence of a representative of the opposition (Au 2011).

If these domestic and international institutions make fraud more difficult, why do incumbent regimes so commonly choose them? Perhaps more puzzling, why cheat in front of voluntarily invited monitors? These questions are not new (see the works cited above), but here I show how both fraud and monitoring can arise in a simple formal model.

Past work on election monitoring focuses on pressure from international actors, both as a benefit of inviting monitoring and as a source of punishment for regimes that are caught cheating (Carothers 1997; Bjornlund 2004; Kelley 2009; Donno 2010; Hyde 2011). The question of why regimes invite monitoring and still cheat is addressed in Kelley (2008) and Hyde (2011, 2012). In general, these explanations assert that regimes gain more legitimacy by inviting monitors and cheating than by not inviting monitors at all, either because refusing monitors is a "self-declaration of cheating"[9] (Kelley 2008) or because regimes have a chance of not being "caught" by the monitors (Hyde 2011).

The argument here provides a domestic and purely instrumental explanation for why monitoring and cheating often go together. This contrasts with the approach of Hyde (2012), who claims election monitoring is a puzzle because, absent international benefits, it "contradicts … states rational self-interest." The central result here suggests that, particularly in the case of non-competitive elections, this is not the case.

An appealing feature of my model is that it does not require assumptions about how the international community interprets the decision to invite monitors, or why (and when) they are willing to punish regimes for fraud. Further, the explanation here is directly linked to a theory of why those who invite monitoring and cheat—which often occurs in non-competitive elections—have an incentive to cheat in the first place. The domestic-centered approach is not intended to argue that the more internationally focused explanations of fraud and monitoring are theoretically or empirically wrong. Rather, the goal is to demonstrate a new way to approach fraud and monitoring that is consistent with some key empirical facts.

THE GENERAL ARGUMENT: FRAUD AS SIGNAL DISTORTING

Before presenting the details, this section describes the general strategic setting that the model is intended to represent. I argue that the results of non-competitive elections matter

[9] The logic of the model is loosely similar to this idea, but differs on why it is valuable to invite monitors to avoid this self-declaration.

because they generate public information that affects how the incumbent interacts with other elites and citizens. For a wide variety of reasons, convincing election results generate higher payoffs to the incumbent from these interactions; the desire to been seen as strong or popular is what is "at stake" in non-competitive elections. This need to appear strong gives the incumbent an incentive to manipulate the results using fraud, potentially well beyond what is necessary to win the election.

To capture this intuition in a simple fashion, the incumbent plays a game with an *audience*, which after the election chooses a level of *support* to give the incumbent. The audience is treated as a unitary actor for convenience, but can also be interpreted as a typical or representative agent from some larger group such as citizens, powerful businessmen, bureaucrats and other members of the regime, or the military. Choosing high support can mean actively supporting the regime in the literal sense (for example, attending pro-regime rallies for citizens, vocally backing the incumbent for elites, or faithfully paying taxes for citizens or elites). Giving support can also be tacit. For example, high support could mean refraining from anti-regime activity such as protesting or attempting to remove the incumbent regime from office by running against them in future elections (or using irregular means). While these notions of support are qualitatively different, what drives the logic of the model is that the incumbent wants the audience to provide high levels of support, and the audience wants to provide high levels of support when the incumbent is strong or popular.[10]

The election result generates a public signal of this incumbent strength. Further, the model assumes this is the *only* role the election plays. This is why I call the model one of a non-competitive election: the election result in itself does not directly affect whether the incumbent remains in office, as would be true in a competitive election.

Fraud is a partially hidden distortion of this information. Technically this could represent *any* hidden action that inflates the election result, such as manipulation of fiscal policy or other aspects of the economy. However, for conceptual clarity and consistency with the emphasis on election monitoring, I focus on the types of manipulation that monitors intend to detect: for example, multiple voting, ballot stuffing or miscounting of votes.

This abstract setting can reflect many political interactions. For example, the audience can be a powerful elite—such as an opposition leader or a high-ranking member of the incumbent coalition—and supporting the incumbent can correspond to accepting lower rents from office or not attempting to remove the incumbent by force. In the Russian case, a natural audience to consider is the powerful businessmen collectively referred to as "oligarchs." In 2000, Putin and the oligarchs agreed to a pact in which the oligarchs would pay taxes and refrain from supporting opposition candidates (Guriev and Rachinsky 2005).[11] Convincing the oligarchs that he has popular support is crucial for Putin to maintain a favorable bargain. More generally, the main assumption for the logic of the model to apply is that the elites are more apt to accept lower rents and not revolt when the incumbent is strong or popular. This article does not seek to identify which particular political actors and interactions are most affected by fraud, which surely varies considerably across cases, but to illustrate a logic of fraud in a transparent and general manner.

[10] See Little (2011) for additional detail and examples of the types of strategic interactions this set-up could capture. Simpser (2013, ch. 4) also contains an extensive list of potential games that could be easily tailored to meet the assumptions used here.

[11] Violating this agreement is what landed Khodorkovsky in jail.

The signal-distorting approach, where fraud is a (partially) *hidden action* used to manipulate *public information*, is not the only way to model fraud with a focus on information. Notably, the baseline is not a costly signaling model (as in Simpser 2013) where fraud is a *observed action* used to affect beliefs about the incumbent's *private information*. I also do not consider the possibility that fraud can be used to render the election less informative (that is, to "jam" the electoral signal, as in Rozenas 2011; Gehlbach and Simpser 2013), which can also be done with observed manipulation. So, the model here is more appropriate for modeling fraudulent actions such as ballot stuffing or making up results, which do not fit well in a costly signaling model in which the "signal" generated by fraud must be observed to be effective.[12] The signal-distorting approach is less appropriate for highly visible strategies such as denying the opposition the ability to campaign effectively, which are more effective at signaling that the incumbent is willing and able to do so.[13]

THE BASELINE MODEL

In the baseline model, there are two actors: an incumbent (*I*; pronoun "she") and an audience (*A*; pronoun "he"). The sequence of moves is:

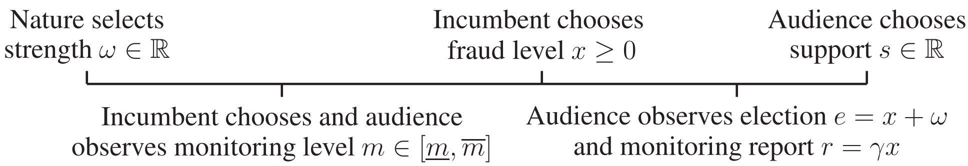

The game begins with Nature selecting the incumbent popularity or strength, denoted ω. This is drawn from common knowledge prior density $f(\omega)$. The parameter is intended to capture genuine popularity among citizens as well as forms of strength that may be reflected in the election result, such as control over a repressive apparatus.

Neither actor has private information about the incumbent popularity. This greatly simplifies the baseline model, and one of the extensions demonstrates that the core results are robust to giving either the incumbent or the audience private information. The only restriction on the prior belief is that $f(\omega) > 0$ for all $\omega \in \mathbb{R}$. So, the equilibrium analysis of the baseline model holds for generally popular or unpopular incumbents, and for when the actors have arbitrarily precise or imprecise information about the incumbent's popularity before the election.

Next, the incumbent chooses a monitoring level m, which the audience observes. Monitoring is treated as a continuous choice, where $\underline{m}$ is the lowest feasible level of monitoring and $\overline{m}$ is the highest possible level. High values of m correspond to inviting more or higher-quality domestic and international observers. Choosing $\underline{m}$ could correspond to inviting no international monitors and choosing election administrators close to the regime. In the baseline model, inviting monitoring imposes no direct costs on the

[12] Hidden fraud does indirectly show up in the election result, but it would be odd to choose clandestine methods to signal strength when observable actions are available.

[13] The hidden/observed distinction is not always clear; e.g., using violence or conspicuously placing secret police in polling booths to intimidate opposition voters is certainly observed by the targets but may not be widely reported. I leave the more subtle distinctions, and the question of how the incumbent chooses among these various technologies, to future research.

incumbent. This applies more plausibly to international monitors who pose little or no financial cost to the regime, and less to expensive types of monitoring such as installing webcams in polling stations, a point addressed in one of the extensions.

The incumbent also selects an amount of fraud to commit $x \geq 0$, which is not *directly* observed by the audience. Since the true level of pro-regime fraud in an election is a function of the actions of many individuals, this is best interpreted as the level of fraud chosen by the top national leadership, holding all other types of fraud constant.[14]

Following the incumbent decisions, the audience observes an election result e that is a function of both the true incumbent popularity and the amount of fraud committed: $e = \omega + x$. That is, the election result will be high when the incumbent is genuinely popular (high ω) or when she commits a high level of fraud (high x). Since $f(\omega) > 0$ for all $\omega \in \mathbb{R}$, the election result can be any real number as well. This proves convenient, as all election results are on the equilibrium path regardless of the incumbent's strategy.

The election result need not correspond to a vote share, but can represent a combination of many observable outcomes from the election: turnout, the number of spoiled ballots or the vote share in particularly important geographic units such as cities with nascent protest movements.[15]

This is admittedly a very reduced-form way to model an election. Typical concerns—such as setting platforms and voters' decision to turn out and who to vote for—are purposefully omitted. All that matters is that the election generates a signal of the incumbent's popularity, and that this signal can be distorted by fraud.

The audience also observers a monitoring report $r = \gamma x$, where $\gamma = 1$ with probability $q(x, m) : \mathbb{R}_+ \times [\underline{m}, \overline{m}] \rightarrow (0, 1)$, and $\gamma = 0$ otherwise. The q function represents the probability that the monitoring report detects the true level of fraud. This is because if $\gamma = 1$ and the incumbent commits fraud, the audience observes a monitoring report greater than zero and directly infers that exactly $x = r$ fraud has been committed. If $\gamma = 0$, the audience will observe $r = 0$ regardless of how much fraud is committed. I often call this a "null" monitoring report. Loosely speaking, this corresponds to a case in which election monitors and other sources do not report any notable malfeasance.

Let q be continuous and differentiable in both arguments, and strictly increasing in level of monitoring $\left(\frac{\partial q}{\partial m} > 0\right)$. This is why m is referred to as the monitoring level: a higher m means there is a higher probability that the true amount of fraud is directly revealed. The probability of detection is either constant or increasing in the amount of fraud $\left(\frac{\partial q}{\partial x} \geq 0\right)$.

As is inevitable in models, numerous aspects of how monitoring and the detection of fraud are formalized are unrealistic or could be done in other defensible ways. To highlight a few issues, most contemporary elections are monitored by multiple observation missions—which could be modeled as strategic actors with influence over how stringent the monitoring is—that release detailed and nuanced reports about various

---

[14] Alternatively, we could that imagine "the incumbent" is a unitary actor representing everyone who contributes to pro-incumbent fraud. Of course, the incumbent relies on a bureaucracy to implement her choice of the fraud level; see Dragu and Polborn (2013) for a model of when bureaucrats are willing to commit illegal acts at the behest of a ruler.

[15] To interpret the election result as something more straightforward like the incumbent vote share, we could let there be an increasing function that maps the abstract election result e to a percentage between 0 and 100. More concretely, if the election result observed is $v = h(e)$, where $h : v \mapsto [0, 1]$ is a strictly increasing and continuous (and hence invertible) function, all the analysis here would go through with $h^{-1}(v)$ replacing e. With some minor modifications to the functional form assumptions, this would not change any of what follows.

aspects of the electoral process. I reduce this process to a unidimensional "report" that either reveals the exact level of fraud or reports that there is no fraud to generate simple and transparent results. As will become clearer when analyzing the model, what matters in the end is that the benefit of committing more fraud than expected decreases in the monitoring level. While there are many potential information structures with which to model election fraud and monitoring that may have different implications, this property could emerge from a wide range of more realistic (but more cumbersome) formulations, some of which are explored in the extensions.

Finally, the audience chooses a support level $s \in \mathbb{R}$. As elaborated below, when making this choice, he takes into account what he observes up to this point—the monitoring level, election result and monitoring report—as well as a conjecture about how much fraud was committed. Crucially, this conjecture will have to be correct in equilibrium.

PAYOFFS

Fraud is costly, and the incumbent wants the audience to provide a high level of support. To formalize this, let the incumbent utility be $u_I(x; s) = b(s) - c(x)$, where b and c are increasing, continuous and twice-differentiable.[16] Further, assume that the benefit function b is strictly increasing and concave ($b' > 0$ and $b''<0$), and that the cost function c is increasing and convex ($c' \geq 0$ and $c'' > 0$) with $c'(0) = 0$ and $\lim_{x\to\infty} c'(x) = \infty$. Loosely speaking, the assumptions about the cost function imply that the cost of falsifying the first ballot is free, but falsifying an additional ballot eventually becomes prohibitively costly.[17]

To capture the intuition that the audience wants to choose high support levels when the incumbent is popular, let his utility function be a quadratic loss function centered around $\omega : u_A = -(s-\omega)^2$.[18] A convenient implication of this utility function is that the audience's best response is to provide the level of support equal to his average belief about the incumbent strength $\mathbb{E}[\omega|\cdot]$. This belief may be conditioned on the election result, monitoring report and the proposed incumbent strategy.

EQUILIBRIUM

I solve for Perfect Bayesian Equilibria with pure strategies (hereafter, equilibria).[19] Given the sequence of moves, write the strategy of the incumbent as the level of monitoring m and the amount of fraud she would commit for each potential monitoring level $x(m)$. A strategy for the audience is a function $s(m,e,r)$ mapping the monitoring level, election result and monitoring report to his level of support.

[16] The continuity and differentiability requirements streamline the technical presentation; qualitatively similar results would hold without these assumptions.

[17] If $c'(0)$ is sufficiently high, or if there is a discontinuity in the cost function at 0, then it would be possible for the incumbent to commit no fraud in equilibrium. If this occurs, the incumbent could still choose $m = \overline{m}$, but is indifferent between any monitoring level that leads to no fraud. See the proof of Proposition 1 in the supplemental information for more technical and substantive analysis of these assumptions.

[18] This functional form is not meant to imply an underlying spatial model, but simply to capture the notion that the audience wants to provide higher support to stronger incumbents. Little (2014) contains a microfoundation for this utility function in a related model.

[19] See the supplemental information for a formal definition of the equilibrium and a discussion of mixed strategies.

The analysis below first finds the equilibrium level of fraud and support—call these $x^*(m)$ and $s^*(m,e,r)$—for a fixed monitoring level m, and then determines the optimal monitoring level (m^*) given how the game plays out following this decision.[20]

When $r > 0$, the audience decision is easy. Since r can only be strictly positive when the report detects the exact amount of fraud (that is, $\gamma = 1$), the audience is at a singleton information set and can directly infer the true incumbent popularity by subtracting the monitoring report from the election result: $\omega = e - r$. So, the audience best response whenever $r > 0$ is to give a support level of $e - r$ regardless of the monitoring level and his conjecture about the level of fraud.

If $r = 0$, the audience does not directly observe the level of fraud, and cannot directly infer the incumbent's popularity. In particular, the audience is in the same information set for any $r = 0$ and ω_1, ω_2 and x_1, x_2 such that $\omega_1 + x_1 = \omega_2 + x_2$. So the audience best response depends on his conjecture about the level of fraud committed; call this $\hat{x}(m)$. In words, the audience best response when $r = 0$ is to subtract the conjectured amount of fraud—which will eventually be pinned down in equilibrium—from the election result. We can write the audience's best response function as:

$$\hat{s}(m, e, r; \hat{x}(m)) = \begin{cases} e - \hat{x}(m) & r = 0 \\ e - r & r > 0. \end{cases} \tag{1}$$

Introducing notation for the audience's expectation of how much fraud will be committed is mostly to clarify the equilibrium derivation, and loosely corresponds to what we would refer to in plain English as expectations of how much fraud will be committed leading up to an election. Still, this mapping is not perfect: in equilibrium, the audience's expectations will be derived purely by strategic reasoning about the incumbent's incentives, and not by factors such as how much fraud the incumbent has committed in past elections.

Now consider the incumbent payoff for choosing an amount of fraud x given this audience best response function. With probability $q(x,m)$ the true amount of fraud is revealed by the monitoring report, and with probability $1 - q(x,m)$ the monitoring report will be null. If the true level of fraud is revealed, the audience directly learns the incumbent popularity and hence the equilibrium level of support is always $e - r = \omega + x - x = \omega$. When the monitoring report does not reveal the amount of fraud, the level of support as a function of the true popularity, conjectured equilibrium level of fraud and actual level of fraud is $e - \hat{x}(m) = \omega + x - \hat{x}(m)$.

Crucially, when the monitoring report is null, the incumbent can get a higher level of support than she would if the audience knew her true popularity by committing more fraud than expected ($x > \hat{x}(m)$). Conversely, the incumbent will appear weaker than she is and garner a lower level of support when committing less fraud than expected ($x < \hat{x}(m)$). That is, for a fixed audience strategy, he *could* be fooled by fraud by an unexpected (that is, off the equilibrium path) level of fraud. However, the standard equilibrium condition that the incumbent and audience play mutual best responses will imply that the audience forms a correct conjecture about the level of fraud.

[20] This follows from the common requirement of a Perfect Bayesian Equilibrium, that player i does not make an inference about Nature's moves by player j's actions if player j has no private information. In fact, the technical requirement here is weaker: all that matters is that the actors play mutual best responses for all m (which is common knowledge), even though there are no proper subgames as Nature moves first.

Since the incumbent is uncertain about her popularity when making the fraud decision, the benefit portion of her expected utility is obtained by integrating over the payoffs derived above with respect to ω. When the monitoring report is correct and the audience always chooses support level ω, the relevant benefit is $\mathbb{E}[b(\omega)] = \int b(\omega)f(\omega)d\omega$.[21] When the monitoring report is null, this term becomes $\mathbb{E}[b(\omega + x - \hat{x}(m))]$, or the expected benefit of committing amount of fraud x when the audience expects $\hat{x}(m)$.

So, the incumbent payoff for choosing fraud level x when the audience expects $\hat{x}(m)$ can be written as:

$$u_I(x, m; \hat{x}(m)) = q(x, m)\mathbb{E}[b(\omega)] + (1 - q(x, m))\mathbb{E}[b(\omega + x - \hat{x}(m))] - c(x)$$

The key trade-off facing the incumbent is that by increasing the amount of fraud, she appears more popular when the fraud is not directly revealed (the second term), but pays more through the cost function c (the third term).[22]

As with any standard equilibrium concept, it must be the case that the incumbent fraud strategy and the audience support level are mutual best responses to each other. In this model, this requires that an equilibrium level of fraud for monitoring level m—again, denote this $x^*(m)$—satisfies two conditions. First, the audience conjecture about the level of fraud must be correct: that is, $\hat{x}(m) = x^*(m)$. Second, given that the audience conjecture is that the incumbent will choose fraud level $x^*(m)$, it is optimal for the incumbent to choose this fraud level. Combining these conditions with the continuity of the primitive functions, for $x^*(m)$ to be an interior equilibrium it must solve:

$$\begin{aligned} 0 &= \frac{\partial}{\partial x}[q(x, m)\mathbb{E}[b(\omega)] + (1 - q(x, m))\mathbb{E}[b(x + \omega - \hat{x}(m))] - c(x)]_{x = x^*(m), \hat{x}(m) = x^*(m)} \\ &= (1 - q(x^*(m), m))\mathbb{E}[b'(\omega)] - c'(x^*(m)) \end{aligned},$$

where $\mathbb{E}[b'(\omega)] = \int b'(\omega)f(\omega)d\omega$ is the expected benefit of committing marginally more fraud than expected *when the monitoring report does not catch the unexpected fraud.* Rearranging gives a more easily interpretable form of the equilibrium condition:

$$\underbrace{c'(x^*(m))}_{\text{Marginal Cost}} = \overbrace{(1 - q(x^*(m), m))}^{\text{Pr(deviation undetected)}} \times \underbrace{\mathbb{E}[b'(\omega)]}_{\text{Marginal Benefit if undetected}} \qquad (2)$$

Equation 2 states that the marginal benefit of committing more fraud than expected (the right-hand side) equals the marginal cost of committing more fraud (the left-hand side). The marginal benefit is strictly positive and non-increasing in $x^*(m)$ because the higher levels of fraud are more likely to be revealed, hence there is a lower chance of the unexpected fraud going undetected. Given the assumptions made about c, the marginal cost is increasing in $x^*(m)$. These two observations (and the assumption about the range of c) ensure that for each monitoring level there is a unique level of fraud $x^*(m)$ such that

[21] An implicit assumption here is that $\mathbb{E}[b(\omega)]$ is finite. Later derivations assume $\mathbb{E}[b'(\omega)]$ and $\mathbb{E}[b''(\omega)]$ are finite as well.

[22] The fact that increasing x may increase the chance that the true level of fraud is revealed through q does not matter on the margin, since in equilibrium there is no difference between the payoff when the report is positive (ω) and zero ($\omega + x - \hat{x}(m)$). More technically, the $\frac{\partial q}{\partial x}$ terms cancel out when deriving equation 2.

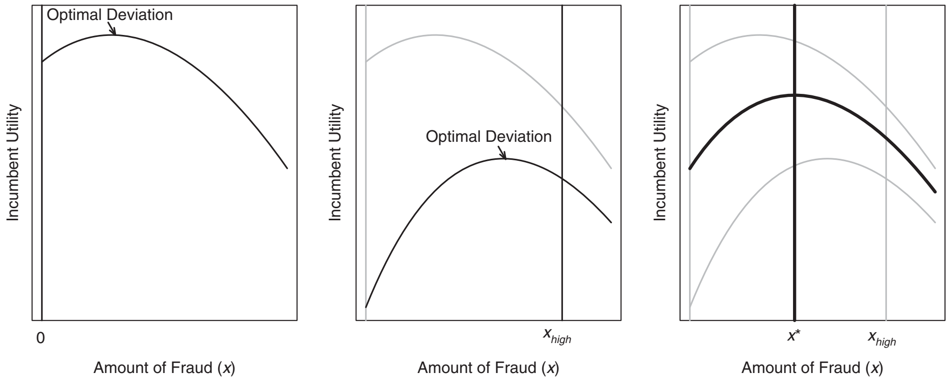

Fig. 1. Incumbent utility as afunction of the amount of fraud for various expected *levels of fraud*
Note: when no fraud is expected (left panel), the incumbent utility is maximized at a higher level, and hence the incumbent would deviate to this level. When proposing an equilibrium with a higher level of fraud (dark curve in middle panel), the incumbent would deviate to a lower amount. When proposing the amount x^* derived from Equation 2, the optimal level of fraud is exactly x^* (right panel).

Equation 2 is met. Given the functional form assumptions for b and c, this solution will be a global maximizer, so:

PROPOSITION 1: For each monitoring level m there exists a unique and strictly positive equilibrium amount of fraud $x^*(m) > 0$ characterized by Equation 2. The equilibrium level of fraud is strictly decreasing in the monitoring level.

PROOF: The existence of an equilibrium $x^*(m)$ for each m is demonstrated above. (See the supplementary information for a proof of uniqueness). The second part follows from implicitly differentiating Equation 2. □

A key characteristic of this part of the equilibrium (that is, all but the monitoring choice) is that the incumbent's fraud does not actually fool the audience. This notion is well described in an open letter from Russian human rights activist Sergei Kovalev to Vladimir Putin in 2008: "You lie, your listeners know this and you know that they don't believe you … Everybody knows everything. The very lie no longer aspires to deceive anyone, from being a means of fooling people it has for some reason turned into an everyday way of life, a customary and obligatory rule for living" (Myagkov, Ordeshook and Shakin 2009, 9).

However, the audience *could* be fooled by an off-the-path deviation—that is, if the incumbent were to choose $x > x^*(m)$ and the monitoring report is uninformative, the audience will give a support level higher than the incumbent strength. This may seem like a technical observation, but it drives both of the central results of the article. An incumbent committing less fraud than expected would appear to be less popular and garner less support from the audience. As illustrated by Figure 1, the incumbent would prefer an equilibrium with no fraud to avoid paying the cost, but this is unattainable; if the audience expects no fraud, the incumbent's optimal response is to commit a strictly positive amount of fraud, as this deviation may not be detected. In other words, the incumbent would prefer to be able to commit to not perpetrating fraud, but absent an ability to make such a commitment, it is always optimal to take fraudulent actions. This is the first main insight from the baseline model.

Inviting high levels of monitoring provides such a commitment device. Since more monitoring makes unexpected fraud more likely to be caught (and hence less effective), it leads to a lower equilibrium level of fraud and hence cost paid. Further, because the monitoring is observed and there is a unique fraud level for each m, the equilibrium level of support is unaffected by the monitoring level. Thus the second main insight from the baseline model is that the incumbent always invites the highest level of monitoring:

PROPOSITION 2: There is a unique pure strategy equilibrium to the baseline model where the incumbent chooses the highest level of monitoring $m^* = \overline{m}$. The equilibrium amount of fraud—which is strictly positive—is given by Equation 2. The audience chooses the support level given by Equation 1 evaluated at $m = \overline{m}$.

PROOF: The equilibrium payoff associated with monitoring level m is $u_I(x^*(m); m) = \mathbb{E}[b(\omega)] - c(x^*(m))$. The change in the incumbent equilibrium utility with respect to m is $\frac{\partial u_I}{\partial m} = -c'(x^*(m))\frac{\partial x^*}{\partial m}$. Since c' is positive and $\frac{\partial x^*}{\partial m} < 0$ from implicitly differentiating Equation 2, $\frac{\partial u_I}{\partial m} > 0$ and hence the incumbent chooses $m^* = \overline{m}$. The remainder of the equilibrium follows from Proposition 1. □

These results present a simple and stark explanation of why incumbent regimes invite domestic and international scrutiny that makes fraud more difficult: under broad assumptions, doing so *always* raises their equilibrium payoff. Despite the inherent international dimension of external monitors, the incumbent decisions are entirely driven by domestic politics, not international pressure or the desire to appear democratic. As long as the monitoring is not perfect, there is still fraud in equilibrium despite this invited scrutiny, which is consistent with the empirical fact that cheating occurs even when monitors are invited. However, if monitoring is extremely strong, as with the ballot counting in Singapore, the amount of fraud will become arbitrarily small.[23]

An alternative way to interpret this result is that monitoring is valuable because it lowers the *expectation* of how much fraud is committed. While Vladamir Churov's boast in the second epigraph about the transparency of Russian elections merits substantial skepticism, it may contain a grain of truth: increasing the transparency of elections can benefit autocrats. The robustness checks will explore some potential reasons why the desire for maximal monitoring is not universal.

One implication of this result is that when monitors are present in an election and detect cheating—even blatant amounts—this does not constitute a failure of the monitors to reduce fraud. Without knowing the counterfactual level of fraud, it is impossible to say whether there would have been more or less without the monitors. While this is ultimately an empirical question,[24] my model presents a theoretical basis for why cheating in monitored elections should not surprise us *and* predicts that fraud would be even more rampant in the absence of monitors.

At first glance, the fact that the availability of monitoring is helpful to leaders holding non-competitive elections (generally autocrats) may seem problematic. Still, the fact that monitoring helps autocrats does not mean it is bad for the audience or for advocates of

[23] Further, if $lim_{x\to 0}\, q(0, \overline{m}) = 0$, the probability of detecting any fraud will approach zero as well.
[24] See, for example, Hyde (2007, 2012) and Herron (2010).

democratic reform. With the utility function here, the audience is equally well off with low or high monitoring. However, if fraud creates negative externalities that harm the audience—which is certainly true if the technologies of fraud used include violence or intimidation—then the audience is better off with a higher level of monitoring. I revisit this point in the discussion of strategic monitors below.

COMPARATIVE STATICS

While the incumbent always chooses the maximal level of monitoring, the model does generate comparative statics on the level of fraud.

PROPOSITION 3: The equilibrium level of fraud is:

(1) decreasing in the maximal level of monitoring
(2) decreasing in the popularity of the incumbent
(3) provided b$'''>0$, increasing in the uncertainty about the incumbent's popularity.

PROOF: Part 1 immediately follows from the first order condition evaluated at $m=\overline{m}$. For Part 2, consider two distributions of popularity such that F_1 first-order stochastic dominates F_2. Then $b''<0$ implies $-b'$ is increasing, so $\mathbb{E}_1[-b'(\omega)]>\mathbb{E}_2[-b'(\omega)]$ and hence $\mathbb{E}_1[b'(\omega)]<\mathbb{E}_2[b'(\omega)]$, where $\mathbb{E}_j$ is the expectation operator with respect to distribution F_j. This implies that the equilibrium level of fraud with distribution F_1 is lower than the equilibrium level of fraud with distribution F_2.

For Part 3, consider two distributions of popularity such that the expectations are equal, but F_1 second-order stochastic dominates F_2 (that is, there is more uncertainty in F_2). Since $b''<0$ and $b'''>0$, $-b'$ is increasing and concave, so $\mathbb{E}_1[-b'(\omega)]>\mathbb{E}_2[-b'(\omega)]$, which implies $\mathbb{E}_1[b'(\omega)]<\mathbb{E}_2[b'(\omega)]$. Hence the incumbent commits more fraud under F_2 □

The intuition behind Parts 2 and 3 is as follows. More popular incumbents care less about marginally increasing the level of support, and hence commit less fraud. Being more uncertain about the incumbent's popularity means that she is more likely to be far less popular than average (giving a higher incentive for fraud) or far more popular than average (giving a lower incentive for fraud). If b' is convex—that is, $b'''>0$—the former effect will be bigger, leading to more fraud on average.

This third derivative is admittedly difficult to interpret—one intuition is that it means the degree to which the marginal returns of more support are diminishing is also decreasing in the level of support. More concretely, this third derivative condition holds for standard increasing and concave functions.[25] Further, it cannot be the case that $b'''<0$ for all s, so the result will hold if the incumbent strength generally lies on the part of b where $b'''>0$.[26]

[25] For example, $b(s)=\log(s)$ defined on $s>0$, $b(s)=s^{\alpha}$ for $\alpha\in(0,1)$ defined on $s\geq 0$, $b(s)=-s^{\alpha}$ for $a<0$ defined on $s\geq 0$, or $b(s)=-\alpha^{-s}$ for $a>1$ defined on $s\in\mathbb{R}$ all have positive third derivatives. So, any shift (i.e., replacing s with $s+k$ in any of the preceding functions) or increasing affine transformation of these utility functions also has a positive third derivative. A counterexample is $b(s)=s-\log(e^{s}+1)$, which satisfies $b'>0$ and $b'<0$, but $b'''<0$ for $s<0$. So, with this benefit function, if the density of f is almost entirely on $\omega<0$, the opposite comparative static would hold.

[26] This is because if b' is decreasing and concave, it must negative for high s. Formally, suppose $b'''<0$ for all s. Take any s_0 and let $d=b''(s_0)<0$. Then $b'''<0$ implies $b''(s)<d$ for all $s>s_0$, so $b'(s)<b'(s_0)+d(s-s_0)$ for $s>s_0$ and $b'(s_0)+d(s-s_0)<0$ for $s>-b'(s_0)/d+s_0$, contradicting $b'(s)>0$.

EXTENSIONS AND EMPIRICAL PREDICTIONS

In this section I describe various modifications and extensions to the baseline model. These serve two purposes. First, they demonstrate that the results are not sensitive to some of the more objectionable simplifying assumptions. Second, adding more parameters and complexity creates new comparative static results, which could be used to test the value of this modeling approach in future work. Formal statements and proofs are available in the supplemental information.

Competitive Elections

The general mechanisms driving the model here need not be absent in competitive elections.[27] For example, suppose that the election result corresponds to a vote share, and that the incumbent wins the election if and only if $e > 0.5$. Under this assumption, fraud may be useful in increasing the probability of winning the election, but this does not preclude the value of manipulating perceptions of popularity conditional on winning the election. In fact, as demonstrated in the supplemental information, if fraud caught by a monitoring report only affects the information learned by the election result (but still counts toward winning the election), the incumbent still invites full monitoring even though this lowers the probability of winning.

If fraud caught by monitors does not count toward winning a competitive election, this will further diminish the incentive for monitoring. However, for the latter to hold (and for elections to be competitive in general) there must be some force compelling the incumbent to follow electoral rules (Little, Tucker and LaGatta, 2012). If this is the case, the incumbent may be more constrained to invite international and domestic monitoring anyways.

Costly Monitoring

The baseline model assumes that the monitoring decision does not directly affect the incumbent's payoff (that is, inviting monitoring is free). This assumption is reasonable in some circumstances, such as international monitors that fund their own missions. However there are exceptions, such as the opening example in which the Russian government spent $300 million on webcams.

To test the robustness of the results to this assumption, suppose the incumbent payoff is $u_I(s, x, m) = b(s) - c(x) - k(m)$, where k is an increasing function that represents the cost of monitoring. Not surprisingly, if the cost k is sufficiently high, the incumbent may invite less than full monitoring.

While assuming that monitoring is costly is not a particularly insightful way to predict why some governments do not always invite as much monitoring as possible, it does allow for comparative static predictions about the equilibrium monitoring level.

PROPOSITION 4: If there is an interior level of monitoring (that is, not $\underline{m}$ or $\overline{m}$), then the equilibrium level of monitoring is:

(1) decreasing in the cost of monitoring
(2) increasing in the cost of fraud

[27] Simpser (2013) also notes that observed manipulation can have valuable "indirect effects" along the lines of the models here as well as "direct effects" on the current election.

(3) decreasing in the popularity of the incumbent
(4) provided $b''' > 0$, increasing in the uncertainty about the incumbent's popularity.

Proof: See the supplemental information.

Unsurprisingly, more costly monitoring—parameterized with a linear scaling of the cost—leads to less monitoring. Second, when the cost of fraud is high (again by a linear scaling), there is a greater incentive to reduce the equilibrium level of fraud with more monitoring.

The third and fourth results are less immediate, but follow from the same logic as Proposition 3. When the incumbent is more popular, she commits less fraud in equilibrium, meaning there is less "demand" to invite costly monitoring. This may explain why Putin did not install webcams until the 2012 election, when he and his party were generally considered to be less popular than in previous elections.

Similarly, when there is a high level of uncertainty about the incumbent's popularity, she commits more fraud, giving a higher incentive to invite monitoring. While the third derivative condition again warrants caution, this result is consistent with an empirical result in Hyde (2011): countries holding their first multiparty elections—presumably a scenario with high levels of uncertainty—are more likely to invite observers.

Combining these results with the comparative statics on the level of fraud makes predictions about the observed relationship between levels of fraud and monitoring. Changes in the cost of monitoring produce an increase in monitoring and a decrease in the level of fraud, which results in a negative correlation between the two in line with the causal effect. However, increases in the level of monitoring induced by a change in the cost of fraud or the beliefs about the incumbent's popularity are accompanied by *more* fraud. This is because anything that increases the marginal benefit of cheating also raises the marginal benefit of inviting more monitoring to reduce (ineffective) fraud. More broadly, this reinforces the point that observing rampant cheating in the presence of international and other kinds of monitoring should not lead us to infer that these activities are counterproductive.

An Alternative Information Structure

As acknowledged above, the sharp distinction between a completely informative and uninformative monitoring report is a major simplification. An alternative way to model the monitoring report is as a normally distributed random variable with mean x—that is, the level of fraud, and some precision (the inverse of the variance), which can be interpreted as the monitoring level.

While more technically cumbersome, this formulation has some advantages relative to the baseline model. It does away with the unrealistic dichotomy between perfectly informative or uninformative monitoring reports. An additional appealing feature is that the monitoring report "matters" in equilibrium. That is, there is variance in the monitoring report in equilibrium, and higher monitoring reports convince the audience that there was more fraud, and lead to less support.[28] This is consistent with recent empirical work demonstrating that negative monitoring reports and other public signals about fraud are associated with post-election protest (Hyde and Marinov 2014; Daxecker 2012; Kuhn 2012).

[28] This requires adding another layer of uncertainty translating the incumbent fraud choice into a "true" amount of fraud. See the supplemental information for details.

If the benefit function b is linear, the results with this information structure are straightforward and qualitatively identical to the baseline model. However, if the benefit function is concave, this reduces the incentive to invite monitoring, since more precise monitoring introduces more uncertainty into the game from the incumbent's perspective. That is, stronger monitoring does not change the average level of support, but it does increase the variance of support that the incumbent receives.

This provides a potential explanation of why monitoring is not universal without assuming it is costly: risk-averse incumbents may prefer the relatively less volatile support levels attained with less monitoring. This strengthens Part 3 of Proposition 4: popular incumbents, who have less to gain from favorable election results/monitoring reports than they have to lose from unfavorable signals, will invite less monitoring.

Multiple, Biased, Strategic Monitors

The baseline model includes key assumptions about the nature of the monitoring report, such as (1) there is one monitoring report, (2) the monitors are not strategic actors and (3) the report is unbiased. Contradictions to these assumptions abound. Contemporary elections are generally observed by many international groups as well as domestic observers. Further, Kelley (2010) argues that "the notion of 'neutral' election observers is a myth" and finds that monitoring groups disagree on their central assessments one-third of the time. Still, loosening these assumptions necessitates only minor caveats to the main results.

The logic of the model is essentially unchanged if there are multiple potential monitoring signals, and if the number of monitors invited is part of the incumbent strategy. Using the information structure from the baseline model, what matters is if *any* monitoring report catches an unexpected deviation. Using the alternative information structure, more monitoring reports simply generate more information about how much fraud there was. So, adding more monitoring reports has the same effect as increasing the monitoring level. If the number of monitors to invite is a strategic choice, the incumbent will invite as many as possible.

Further, little changes in either information structure if there is a *known* bias to the monitoring reports. The incumbent still invites as much monitoring as possible and the audience adjusts for the bias when interpreting the monitoring report, rendering it just as informative as an unbiased report.

Reports with *unknown* bias are more easily modeled in the alternative framework with normally distributed monitoring reports.[29] If the bias is also normally distributed, this leads to a convenient characterization. The audience now subtracts the *average* bias from the monitoring report, and the information he draws from the report with unknown bias is analogous to what he would learn from a less precise monitoring report with known bias. This generates an interesting comparative static result.

PROPOSITION 5. The equilibrium level of fraud is increasing in the uncertainty about the observers' bias.

See the supplemental information for a formal statement and proof. Again, note that it is not the average bias of the monitors that matters, but how uncertain the audience is

[29] The challenge with the baseline model is determining how to refine the audience beliefs about the incumbent's popularity when detecting an off-the-path deviation.

about the bias. This result relies on convenient properties of the normal distribution, but a more general point to take away is that bias in monitoring organizations particularly matters in how it renders their reports less informative.

There are many ways to model the monitors as strategic actors. Doing so can affect both the level of monitoring—as decisions made by monitoring groups may be as important as those made by the regime—and the content of the reports. On the monitoring level, suppose the main goal of the monitors is to reduce the amount of fraud. It immediately follows that they should be eager to perform the strongest level of monitoring, as this will always lead to less fraud. If, for example, we modeled the incumbent decision as an offer to have monitoring at level m followed by the monitors accepting or refusing the invitation (perhaps resulting in the minimal level $\underline{m}$), it is straightforward to see that the incumbent would offer the highest level possible ($\overline{m}$), and the monitors would accept. Equivalently, if the monitors propose m and the incumbent accepts or rejects, the unique equilibrium would still be the monitors proposing (and the incumbent accepting) the highest level of monitoring.

The implications of strategic monitors for the monitoring report itself are potentially complicated. What matters is if some information is conveyed by the monitoring reports. Whether this holds depends on the goals of the monitors. Loosely speaking, monitors that only want to hurt or help the incumbent cannot credibly convey information. For example, some monitoring groups such as the Commonwealth of Independent States—composed of post-Soviet states with largely fraudulent elections—have a reputation of endorsing all elections, which renders their reports uninformative (Hyde 2012, 159). However, if monitors care about their reputation for honesty or have some disincentive to overly distort their findings, they should convey some information in equilibrium, and the results here would remain.

Private Information, Signaling and Blatant Fraud

Recall that the baseline model does not make a costly signaling argument that incumbents blatantly manipulate elections to demonstrate their ability or willingness to do so. Costly signaling arguments are consistent with the general informational approach to fraud advocated here; I focus on the signal-distorting logic for conceptual clarity, as well as the fact that the costly signaling logic is more prevalent in past work (Simpser 2005, 2013; Hyde 2011).

The model can incorporate some elements of a costly signaling argument if the incumbent has private information about some combination of her popularity, the cost to commit fraud or the cost to invite monitoring. The model becomes substantially more complex if such private information is taken into account: the audience observes three different actions and signals (the monitoring level, election result and monitoring report) and forms inferences about both a parameter of uncertainty (the incumbent's popularity) and a partially observed action (the level of fraud). Thus regardless of what type of private information is incorporated, there are many potential equilibria, many of which rely on "strange" off-the-path beliefs, but standard refinements do not apply.

However, when placing restrictions on the off-the-path beliefs, the main results stand. If the incumbent has private information about her popularity, then the less popular types will have a greater incentive to commit fraud. However, it is difficult to construct equilibria in which the types separate on the monitoring level, because the more popular type has an incentive to invite monitoring to make it clearer that they are committing less fraud, and the unpopular type will always want to mimic the popular type. Similarly, if one type can commit fraud more cheaply and does so in equilibrium, the type committing less fraud will want to invite monitoring to demonstrate this fact, so in any proposed equilibrium with separation on monitoring, the low-cost type would want to mimic the high-cost type. That is, absent some force that makes monitoring more costly for one type, there can be no separation on the monitoring choice in equilibrium.[30] Loosening the assumption that the audience has no private information has no effect on the results as long as he is not *perfectly* informed about the incumbent strength; see the supplemental information for details.

If some incumbent types can commit fraud more cheaply and are stronger in the sense that the audience wants to behave more favorably toward them (as in Simpser 2013; Kuhn 2012), then fraud can serve as a signal of repressive capacity as opposed to an attempt to distort the signal about popularity. While this is in some ways a very different model, it generates an analogous result about monitoring. If monitoring more clearly communicates the signal that the incumbent is a strong type, then strong types will want to invite monitoring. As a result, the weak types will mimic the strong types and invite monitors as well.[31]

One issue with the signaling approach—particularly with respect to the fraud decision—is that it is not clear why regimes signal strength with *electoral* repression. Many other repressive technologies are available, both during and far away from elections. An advantage of the signal-distorting approach is that fraud is used to affect signals that are inevitably generated by elections—most notably, the turnout and incumbent vote share.

Still, costly signaling and signal distorting are not necessarily competing theories of fraud. Incumbent regimes need not choose between manipulating electoral results with hidden or blatant methods; it is possible that both types of fraud are present in most (if not all) manipulated elections.

CONCLUSION

As monitoring is simply treated as a way to generate public information or constrain the incumbent's ability to manipulate election results, the model plausibly provides insights into why incumbent leaders build institutions that make elections more free and fair in general. In contrast with past work, the mechanism identified here is independent of international pressure or idiosyncratic preferences for ruling democratically. This insight comes generally from not treating all actions surrounding elections as attempts to increase the probability of winning. When focusing more on the informational aspects of non-competitive elections, the ability to manipulate results at will does not necessarily lead to higher equilibrium payoffs. This mechanism is not enough to push regimes to full democracy, broadly defined, as this must eventually entail the prospect of losing elections and alternation of power. Still, focusing on the informational aspects of elections can help understand the broad range of regimes with non-competitive elections.

[30] As is typical in such signaling games, there is generally a pooling equilibrium in which no types invite monitoring if the audience's off-the-path beliefs are that a type that invites monitoring is unpopular or committing a high level of fraud. While standard refinements do not apply to this more complex setting, such an equilibrium requires that the audience infers that the incumbent is unpopular or committing lots of fraud when making an off-the-path choice to invite monitors, which can be ruled out using a refinement analogous to D1.

[31] A previous version of the article formalizes this claim.

Thus these results speak not only to how we should think about electoral fraud and monitoring, but to the general idea of modeling (non-competitive) elections primarily as information-generating institutions. Expanding the model to include other choices, such as observed manipulation of the result, allowing or not allowing opposition groups to run, and the decision to hold an election in the first place could provide more insights into understanding regime types as the result of a suite of strategic choices.

REFERENCES

ACE Electoral Knowledge Network. 2012. 'Electoral Management Data'. http://aceproject.org/epic-en

Ashworth, Scott. 2005. 'Reputational Dynamics and Political Careers'. *Journal of Law, Economics, and Organization* 21(2):441–66.

Au, Alex. 2011. 'Counting Agent Me'. *Yawning Bread*. Available at http://yawningbread.wordpress.com/2011/05/08/counting-agent-me/

Bjornlund, Eric C. 2004. *Beyond Free and Fair: Monitoring Elections and Building Democracy*. Washington, DC: Woodrow Wilson Center Press.

Blaydes, Lisa A. 2011. *Elections and Distributive Politics in Mubarak's Egypt*. Cambridge University Press.

Bueno De Mesquita, Ethan. 2010. 'Regime Change and Revolutionary Entrepreneurs'. *American Political Science Review* 108(3):446–66.

Carothers, Thomas. 1997. 'The Observers Observed'. *Journal of Democracy* 8(3):17–31.

Daxecker, Ursula E. 2012. 'The Cost of Exposing Cheating: International Election Monitoring, Fraud, and Post-Election Violence in in Africa'. *Journal of Peace Research* 49(4):503–16.

Donno, Daniela. 2010. 'Who is Punished? Regional Intergovernmental Organizations and the Enforcement of Democratic Norms'. *International Organization* 64(4):593–625.

Dragu, Tiberiu, and Mattias Polborn. 2013. 'The Administrative Foundation of the Rule of Law'. *The Journal of Politics* 74(4):1038–50.

Egorov, Georgy, and Konstantin Sonin. 2011. 'Incumbency Advantage in Non-Democracies'. Paper presented at the Annual Meeting of the MPSA, Chicago, IL 31 March–3 April.

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):448–52.

Fearon, James. 2011. 'Self-Enforcing Democracy'. *Quarterly Journal of Economics* 126(4):1661–708.

Gandhi, Jennifer, and Ellen Lust-Okar. 2009. 'Elections Under Authoritarianism'. *Annual Review of Political Science* 12:403–22.

Gandhi, Jennifer, and Adam Przeworski. 2009. 'Holding onto Power by Any Means? The Origin of Competitive Elections'. Unpublished manuscript, New York University.

Gehlbach, Scott, and Alberto Simpser. 2013. 'Electoral Manipulation as Bureaucratic Control'. *American Journal of Political Science*,doi:10.1111/ajps.12122.

Guriev, Sergei, and Andre Rachinsky. 2005. 'The Role of Oligarchs in Russian Capitalism'. *Journal of Economic Perspectives* 19(1):131–50.

Herron, Erik S. 2010. 'The Effect of Passive Observation Methods on Azerbaijan's 2008 Presidential Election and 2009 Referendum'. *Electoral Studies* 29:417–24.

Holmstrom, Bengt. 1999. 'Managerial Incentive Problems: A Dynamic Perspective'. *Review of Economic Studies* 66(1):169–82.

Hyde, Susan D. 2007. 'The Observer Effect in International Politics: Evidence from a Natural Experiment'. *World Politics* 60(1):37–63.

——. 2011. 'Catch Us if You Can: Election Monitoring and International Norm Creation'. *American Journal of Political Science* 55(2):356–69.

——. 2012. *The Pseudo-Democrat's Dilemma: Why Election Observation Became an International Norm*. Ithaca, NY: Cornell University Press.

Hyde, Susan D., and Nikolay Marinov. 2014. 'Information and Self-Enforcing Democracy: the Role of International Election Observation'. *International Organization* 68:329–59.

Kelley, Judith. 2008. 'Assessing the Complex Evolution of Norms: The Rise of International Election Monitoring'. *International Organization* 62:221–55.

——. 2009. 'D-Minus Elections: The Politics and Norms of International Election Observation'. *International Organization* 63:765–87.

——. 2010. 'Election Observers and their Biases'. *Journal of Democracy* 21(3):158–72.

Khodorkovsky, Mikahil. 2012. 'What Is at Stake in the Russian Election?' *International Herald Tribune*, 26 February. Available at http://www.nytimes.com/2012/02/27/opinion/27iht-edkhodorkovsky27.html

Korolyov, Alexei. 2012. 'Russian Monitors: Putin ''Would Still Win'' Despite Fraud'. *RIA Novosti*, 6 March. Available at http://en.rian.ru/russia/20120307/171833877.html

Kuhn, Patrick M. 2012. 'To Protest or Not: The Election Losers' Dilemma'. Unpublished manuscript, University of Rochester.

Lehoucq, Fabrice. 2003. 'Electoral Fraud: Causes, Types, and Consequences'. *Annual Review of Political Science* 6:233–56.

Little, Andrew T. 2014. 'An Informational Theory of Non-competitive Elections'. Unpublished manuscript, New York University.

——. 2012. 'Elections, Fraud, and Election Monitoring in the Shadow of Revolution'. *Quarterly Journal of Political Science* 7(3):249–83.

Little, Andrew T., Joshua Tucker, and Tom LaGatta. 2012. 'Elections, Protest and Alternation of Power'. Paper presented at the 2012 European Political Science Association Meeting, Berlin, Germany, 21–23 June.

Lohmann, Susanne. 1998. 'Rationalizing the Political Business Cycle: A Workhorse Model'. *Economics and Politcs* 10(1):1–17.

Magaloni, Beatriz. 2006. *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*.Cambridge University Press.

——. 2009. 'The Game of Electoral Fraud and the Ousting of Authoritarian Rule'. *American Journal of Political Science* 54(3):751–65.

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

Organization for Security and Co-operation in Europe. 2012. *Russian Federation, Presidential Election – 4 March 2012: Statement of Preliminary Findings and Conclusion*.

Pano News. 2012. 'Observers Detail Flaws in Russian Election'. 5 March. Available at http://newspano.com/news/observers-detail-flaws-in-russian-election

Prendergast, Canice. 1999. 'The Provision of Incentives in Firms'. *Journal of Economic Literature* 1:7–63.

Rozenas, Arturas. 2011. 'Forcing Consent: Information and Power in Non-Democratic Elections'. Unpublished manuscript, Duke University.

Schedler, Andreas. 2002. 'The Menu of Manipulation'. *Journal of Democracy* 13(2):36–50.

Simpser, Alberto. 2005. 'Making Votes Not Count: Strategic Incentives for Electoral Corruption'. PhD Dissertation, Stanford University.

——. 2013. *Why Governments and Parties Manipulate Elections: Theory, Practice, and Implications*. Cambridge University Press.

Sjoberg, Fredrik M. 2012. 'Autocratic Adaptation: The Strategic Use of Transparency and the Persistence of Election Fraud'. *Electoral Studies* 33:233–45.

Tucker, Joshua A. 2007. 'Enough! Electoral Fraud, Collective Action Problems, and Post-Communist Colored Revolutions'. *Perspectives on Politics* 5(3):535–51.