

Elections, Fraud, and Election Monitoring in the Shadow of Revolution*

Andrew T. Little

Department of Politics, New York University, New York, USA;
andrew.little@nyu.edu

ABSTRACT

Elections are modeled as a public signal in an incomplete information game of revolution. By changing beliefs about the general level of anti-regime sentiment, elections can make citizens more or less apt to rebel and hence make a successful revolution more or less likely. This effect makes elections valuable to incumbents that are not secure in office as they have more to gain by good results than they have to lose from bad results. Electoral fraud is modeled as a distortion of the public signal, and election monitoring is incorporated as changing the cost of this distortion. In equilibrium, citizens discount the distortion, so the average protest size and probability of revolution are the same as when the incumbent cannot commit fraud. This makes election monitoring valuable to incumbents as it ties their hands and lowers the equilibrium amount of fraud. So, elections may be held that would not occur in the absence of monitoring.

* Many thanks to Avidit Acharya, Scott Ashworth, Georgy Egorov, Alastair Smith, Scott Tyson, participants in the 2010 Institutions in Context Summer School at the University of Tampere, and several anonymous referees for comments and suggestions.

Nearly every country in the world has elections of some kind, though many present no real chance of anyone but the incumbent regime winning. Often this near certainty is further bolstered by large amounts of electoral fraud. Most of these elections — from the blatantly fraudulent to the mostly clean — are to some degree monitored by domestic independent electoral commissions or international observers that are voluntarily invited by the incumbent government. These observations bring up numerous questions. Do elections matter when the winner is not in doubt? Even if an incumbent victory is guaranteed, why risk having a less convincing result than expected? Particularly in the case of noncompetitive elections, why commit massive fraud to make the results even more decisive? Finally, why do incumbent leaders allow for the potential embarrassment of a negative report from domestic or international monitors?

The model presented here seeks to provide insight into these questions by bringing together two lines of thought. First, elections generate public information about the relative strength and popularity of the incumbent leaders, which affects the behavior of the leaders as well as other elites and the general population. Second, in order to prevent exploitation by incumbents that will not give up office voluntarily, citizens must overcome a coordination problem which can be greatly affected by incomplete information.

Combining these ideas, I develop a model where elections generate public information about anti-regime sentiment before citizens decide whether or not to participate in a revolution. Even if there is no chance of a legal turnover of leadership where the incumbent loses and steps down voluntarily, the information generated by the election result affects the potential for irregular means of ousting the incumbent. This adds to growing literatures on revolution as a coordination problem and how the information generated by noncompetitive elections makes them meaningful and potentially valuable to the incumbent.

The first key result from the model is that the public information generation gives an incentive to hold elections for incumbent regimes that may be otherwise reluctant to do so: those that are moderately insecure in office. This follows from the fact that moderately insecure incumbents have the most to gain from an unexpectedly favorable result and less to lose from a bad result. Next, I extend the model to allow the incumbent to commit fraud, modeled as a hidden distortion of the election result. This manipulation will always occur in equilibrium under broad assumptions, but because the citizens are strategic actors they will discount the result, rendering fraud

ineffective. This provides an explanation for the prevalence of election monitoring as a way for leaders to tie their hands from committing inefficient but costly fraud. Still, unless the monitoring is so effective that it renders fraud impossible, fraud will still occur in equilibrium.

This analysis provides some potentially surprising answers to the motivating questions. Noncompetitive elections matter because of the information they generate, and this gives some leaders an incentive to hold elections even if they are costly for exogenous reasons. In particular, the chance of a better-than-expected result gives an incentive to hold elections for incumbents who are somewhat insecure in office — a group that may be unlikely to hold elections for other reasons. The favorable behavior induced by high election results gives an incentive to cheat even in noncompetitive elections. Finally, the ineffectiveness of the distortion in equilibrium gives the incumbent an incentive to make fraud more difficult by inviting domestic and international monitors even without external pressure to do so.

Extant Work

The notion that elections play an informational role is not new: Przeworski (2010) summarizes formulations of this idea dating to Herodotus, see Magaloni (2006) and Blaydes (2011) for more recent work and Gandhi and Lust-Okar (2009) for an overview. However, there have been few attempts to formalize this idea and elucidate its implications. Londregan and Vindigni (2006) show that in an incomplete information bargaining game where conflict may occur, holding an election before bargaining converges the beliefs of the two sides and can create an opening for a peaceful settlement. This mechanism makes elections valuable to incumbents if the information revelation strengthens their bargaining position (Cox, 2009). Little (2011b) provides a general framework for modeling how incumbents hold elections to gather information as well as signal strength (see also Magaloni, 2006; Egorov and Sonin, 2011), which this paper applies to the specific context of revolution and extends by considering the possibility of fraud.

Most of these arguments examine how the incumbent can use the information generated by elections, but other actors may utilize this information as well. In particular, the information conveyed by election results can lead to a wide range of actors — from ordinary citizens to elites — acting more or

less favorably towards the incumbent. This generates incentives to manipulate the results with fraud in order to induce the more favorable behavior.

Attempting to depose an unpopular leader poses a coordination problem. Even if all the citizenry favor a revolt, success requires that a sufficient proportion participate (Granovetter, 1978; Kuran, 1989; Fearon, 2011). So beliefs about whether others are willing to rebel affect individual decisions to participate. The information structure — what citizens know about the other citizens' preferences and beliefs — can play a key role in this environment (Chwe, 1999).

The argument here is that elections affect beliefs about other citizens' levels of anti-regime sentiment.¹ This information can come from more than the vote tally; as Ames (1970) argues about Mexico under the PRI, “Since the winner of a Mexican election is seldom in doubt, and because elections in Mexico usually indicate political skill and regime legitimacy, the number of people turning out on election day may be as interesting as the direction of the vote.”

Beyond the official results and turnout, Tucker (2007) claims reports of electoral fraud served as a coordination device to spur rebellion in the “Colored Revolutions;” I demonstrate here that beliefs about fraud can affect the outcome of a coordination game even when there is a unique equilibrium. A belief that an election result was only high as a result of fraud can lead to collective action short of revolution as well, plausibly accounting for the protests following the Russian Duma elections in late 2011. This example also illustrates that the logic of the model does not only apply to elections where the chief executive office is at stake.

While multiple equilibria are pervasive in canonical complete-information coordination games, these games can have unique equilibria if the actors have private information; such models are often referred to as global games (Morris and Shin, 2003). Global games as defined by Morris and Shin (2003) and closely related models are becoming common in the study of political phenomena such as revolutions (Edmond, 2008; Bueno De Mesquita, 2010a; Shadmehr and Bernhardt, 2011). A natural way to introduce elections into this framework is to treat them as a *public* signal observed by all actors

¹ Fearon (2011) models elections as generating public information about how other citizens voted, though this does not affect beliefs about the incumbent popularity in the same manner as there is no uncertainty about the incumbent type. That is, in Fearon (2011) elections provide information about actions taken by the incumbent, while in this model they provide information about the state of the world.

in addition to their private information. Observing a lower-than-expected incumbent vote share will lead each citizen to increase their belief about other citizens' distaste for the regime and propensity to rebel. Since citizens want to participate in successful revolutions against unpopular leaders, the low election result leads to more participation and a higher probability of successful rebellion. Conversely, a strong election result for the incumbent will lead to less participation.

After analyzing when information generation will make elections valuable to incumbents in this context, I extend the model to consider electoral fraud and election monitoring, topics that have seen little formal analysis. The notion that incumbent leaders, particularly dictators, attempt to manipulate the information available to citizens is also modeled with a global games framework in Edmond (2008) and Egorov and Sonin (2011). A key technical difference between these models and the one here is that Edmond (2008) and Egorov and Sonin (2011) assume the incumbent is fully informed about the parameter representing their strength (θ here) before making their signal-distorting or fraud decision. This generates quite different — though not inconsistent — substantive results that are contrasted in more detail below.²

Other existing models of fraud generally adopt a decision-theoretic approach where incumbents face the sorts of tradeoffs described in the informal literature (see Lehoucq (2003) for an overview): increasing their probability of winning elections by committing fraud against the chance that they will be “caught,” resulting in decreased repressive capacity, legitimacy, or aid (Gandhi and Przeworski, 2009; Magaloni, 2009; Hyde and Marinov, 2008). Fearon (2011) treats citizens as strategic actors in a model where fraud is used to get past a “victory” threshold that is determined by equilibrium behavior, though this relies heavily on equilibrium selection arguments, and, again, the notion that fraud helps win elections.

This general line of reasoning is problematic for two reasons. First, modeling the benefit of fraud as increasing the probability of winning elections is inconsistent with the fact that fraud is pervasive in noncompetitive elections (Simpser, 2011). Second, treating the incumbent as the only strategic actor (or modeling the beliefs of others in a nonstandard fashion) dodges a major

² There are advantages and disadvantages to these assumptions compared to the one in this paper: while citizens perfectly knowing the amount of fraud committed and adjusting is surely not true, neither is the assumption that incumbents are fully informed about their strength nor the equilibrium outcome that incumbents who distort the signal/commit fraud will always be able to commit just enough to survive as in Edmond (2008).

issue in approaching fraud game theoretically: if the defrauded are strategic actors, they must form beliefs about how much is committed. In equilibrium, these beliefs must be correct given the incumbent fraud strategy and the actions of the defrauded must be optimal given this belief; in words, they are generally not fooled.

A rare exception to these pitfalls, Simpson (2011), argues that fraud can serve as a costly signal of repressive capacity. A key difference in this approach is that Simpson treats fraud as an *observed* action, while it is a hidden action in the model below. Some technologies of fraud are observable, such as strong-arming electoral officials to nullify ballots from opposition strongholds, but much fraud is indeed clandestine and hence a poor instrument for this signaling role. That is, while our arguments agree that fraud — particularly in noncompetitive elections — can be fruitfully modeled as an attempt to manipulate information, they differ in what types of fraud are considered; see Little (2011a) for a more detailed comparison of these approaches.³ Simpson (2011) also argues that high levels of fraud may deter coordination if citizens choose to play equilibria with little anti-regime action upon observing high levels of fraud, but the model here shows that in a coordination game with a unique equilibrium the opposite result holds: beliefs that there was a high level of fraud — all else equal — will lead to more anti-regime action.

The model with fraud generates results that are surprising in this context but are standard in “career concerns” models (Holmstrom, 1999). If the citizens adjust for the fraud correctly, as they do in equilibrium, the ability to distort the signal leaves the incumbent no more likely to survive than when they lack the ability to do so. Fraud can be extensive even if it is not only unproductive but also counterproductive. This provides a straightforward answer to the question of why regimes invite election monitors and cheat in front of them (Hyde, 2011). Inviting monitors partially ties the incumbent’s hands and leads to less inefficient and ineffective fraud, but unless this monitoring is perfect there is always an incentive to do some cheating.


This contrasts with accounts of election monitoring that rely on pressure from external actors to hold clean elections (Carothers, 1997; Bjornlund, 2004; Hyde, 2011).⁴ While there is certainly truth to these more

³ Rozenas (2011) treats repression as information manipulation as well. More specifically, repression acts as a mean-preserving rearrangement of posterior beliefs about the incumbent strength, which is more analogous to the role that the decision to hold an election plays here.

internationally focused theories, they seem inconsistent with the fact that monitors are typically invited voluntarily by the incumbent regime. Further, it is not uncommon for governments to ask for monitoring but be turned down. For example, Romania invited United Nations election monitors in 1990 but the UN refused on the grounds that they “did not monitor elections in sovereign countries unless there was a profound threat to regional security” (Bjornlund, 2004, p. 56). Adding a purely domestic and strategic account for the value of election monitoring broadens our understanding of the topic.

The Model

The actors in the model are an incumbent leader I and a continuum of citizens, indexed by j .⁵ The sequence of moves is:



The θ parameter with no subscript represents the general level of anti-regime sentiment, which is never directly observed. In addition to a common prior about θ , all citizens receive a private signal θ_j , corresponding to their personal level of anti-regime sentiment. This signal provides information about θ — citizens rightly believe that when they dislike the regime others tend to feel similarly — and directly affects the citizen utility function. If the incumbent holds an election there is an additional public signal e that provides information about the general level of anti-regime sentiment, though it does not directly affect citizen payoffs. In contrast to Egorov and Sonin (2011), the incumbent has no private information about the anti-regime sentiment, so the citizens make no inferences about θ from the decision to hold an election itself.

⁴ Magaloni (2009) makes a related argument about independent electoral commissions as a means for incumbents to tie their hands, though again with a nonstandard treatment of citizens’ beliefs.

⁵ These actors could correspond to opposition leaders or even members of the incumbent coalition that may rebel (see Boix and Svobik (2008) and Kuhn (2010) for related models that emphasize these interpretations, respectively), but to focus the presentation and be more consistent with models of mass collective action I refer to them generically as citizens.

Each citizen chooses whether to rebel — alternatively referred to as protest — or not. Citizen j 's decision to rebel is denoted $a_j = 1$ for rebelling and $a_j = 0$ for not rebelling. Assume that the rebellion succeeds — which can be interpreted as overthrowing the incumbent or reaching some lesser but still critical mass — if and only if the proportion of citizens rebelling, denoted N , is greater than a fixed and common knowledge threshold T . In general citizens want to participate in successful rebellions but not in failed rebellions, and the incumbent wants the rebellion to fail. Call the portion of the game where the citizens observe their private signals and make their decision the *revolution stage*.

The model would have similar results if the citizens (or incumbent) had payoff functions continuous with respect to the size of protests rather than dependent on whether a “success” threshold is reached. That is, citizens could want to participate in larger protests even if they do not unseat the incumbent, and the incumbent could be hurt by such smaller protests proportional to their size. I will sometimes informally refer to this interpretation, though to be clear the technical results analyze the case where the outcome of concern is only whether the rebellion passes the threshold T .

The incumbent's strategy is the decision to hold an election (notation referring to this contingency will generally have a subscript 1) or not (subscript 0). The citizen strategy has two components: the set of private signals for which they rebel when no election is held, and a function mapping the election result e to the set of private signals for which they will rebel upon seeing that result. The solution concept used is Perfect Bayesian Equilibrium (PBE) with two additional standard restrictions. First, assume that in each revolution stage the citizens use a *cutoff strategy* of the form “rebel if $\theta_j \geq \hat{\theta}$.” That is, citizens rebel if and only if they sufficiently dislike the regime. Second, assume these strategies are symmetric: i.e., all citizens use the same thresholds. Often this is the only form that an equilibrium strategy can take; to ease the presentation I simply assume this form at the outset.

In this model, the requirements for a PBE with symmetric cutoff strategies — stated formally in the Appendix — are (1) the citizens' beliefs about θ are formed by Bayes' rule; (2) when other citizens use the proposed threshold strategy $\hat{\theta}$, rebelling is optimal given this belief if and only if the citizen's anti-regime sentiment θ_j is above $\hat{\theta}$; and (3) the incumbent holds an election if and only if the expected payoff given these cutoff strategies (and their belief about θ , which is the prior) is higher in the lottery over electoral revolution stages than the payoff from the no-election revolution stage.

The incumbent gets a partial payoff normalized to 1 if the rebellion fails and 0 if the rebellion succeeds. In addition, the incumbent pays a cost $k(e) \in \mathbb{R}$ for holding the election, which may depend on the result. This term captures the physical costs of holding an election — campaigning, paying election workers, etc. — as well as any other positive or negative payoffs associated with holding an election. The goal here is not to identify all potential costs and benefits of holding elections, and the k term consolidates the effects not explicitly modeled here.

Since the incumbent makes the electoral decision without knowing e , the relevant cost parameter is the average cost $k \equiv \int k(e)p(e)de$, where $p(e)$ is the probability density function for the election outcomes. I describe this term as a cost, though it may be negative, indicating elections provide an exogenous net benefit to the incumbent. Summarizing, the incumbent expected utility function for not holding (U_0^I) and holding an election (U_1^I) are:

$$\begin{aligned}\mathbb{E}[U_0^I] &= Pr(N_0 < T) \\ \mathbb{E}[U_1^I] &= \int Pr(N_1(e) < T)p(e) de - k\end{aligned}$$

where N_0 is the number of rebelling citizens without an election, and $N_1(e)$ the number of rebelling citizens after an election with result e . To reduce the complexity of some formulas, let $T = 0.5$ — that is, rebellion succeeds if and only if the proportion of rebelling citizens is greater than a half. This sacrifices no meaningful generality as, for the results here, changes in this parameter are redundant with changes in the mean anti-regime sentiment parameter introduced below.

The Citizen Payoffs

Denote the citizen payoffs $u_a^O(\theta_j)$ where $O \in \{S, F\}$ is the outcome of the rebellion (Success, Failure, respectively) and $a \in \{0, 1\}$ corresponds to not rebelling and rebelling, respectively. Each payoff depends on the individual level of anti-regime sentiment θ_j , but not on the potential election result. Assume the following

Assumption 1 *The citizen payoffs have the following properties:*

- (i) $u_1^S(\theta_j) - u_0^S(\theta_j) > u_1^F(\theta_j) - u_0^F(\theta_j)$ for all θ_j

- (ii) The relative value of participating in rebellion conditional on either outcome $(u_1^O(\theta_j) - u_0^O(\theta_j))$, $O \in \{S, F\}$, and $(u_1^S(\theta_j) - u_0^S(\theta_j)) - (u_1^F(\theta_j) - u_0^F(\theta_j))$ are all continuous, twice-differentiable, and increasing in θ_j .
- (iii) There exists a $\underline{\theta}$ and $\bar{\theta}$ such that $u_1^S(\underline{\theta}) - u_0^S(\underline{\theta}) = 0$ and $u_1^F(\bar{\theta}) - u_0^F(\bar{\theta}) = 0$.

The first part and the condition on $(u_1^S(\theta_j) - u_0^S(\theta_j)) - (u_1^F(\theta_j) - u_0^F(\theta_j))$ in the second part capture the coordination dynamics: the relative value of participation is higher when the rebellion succeeds, and this difference is increasing in individual anti-regime sentiment. The rest of the second part implies that regardless of the outcome, the relative value of participation is increasing in distaste for the regime. Part 3 implies that some citizens are willing to rebel or not rebel regardless of the actions of others, and is discussed in more detail below.

Let $q(\cdot)$ be the equilibrium probability of the rebellion succeeding. A citizen receiving signal θ_j will rebel if:

$$q(\cdot)u_1^S(\theta_j) + (1 - q(\cdot))u_1^F(\theta_j) \geq q(\cdot)u_0^S(\theta_j) + (1 - q(\cdot))u_0^F(\theta_j)$$

$$q(\cdot) \geq \frac{u_0^F(\theta_j) - u_1^F(\theta_j)}{(u_1^S(\theta_j) - u_0^S(\theta_j)) + (u_0^F(\theta_j) - u_1^F(\theta_j))} \equiv 1 - h(\theta_j) \quad (1)$$

Working with the $h(\theta_j)$ function — which represents the probability of the regime *surviving* that makes a citizen with anti-regime sentiment θ_j indifferent between rebelling and not — streamlines the algebra.

Equation (1) has three important properties (see the Appendix for a formal statement and proof). First, since $q(\cdot)$ is a probability, if $1 - h(\theta_j) \leq 0$ then a citizen receiving signal θ_j has a dominant strategy to rebel regardless of what other citizens do. Similarly if $1 - h(\theta_j) \geq 1$ the citizen has a dominant strategy to not rebel. Part 3 of Assumption 1 implies that citizens with anti-regime sentiment $\theta_j < \underline{\theta}$ have a dominant strategy to not rebel and those with $\theta_j > \bar{\theta}$ have a dominant strategy to rebel.

Second, for citizens without a dominant strategy ($\theta_j \in (\underline{\theta}, \bar{\theta})$) the RHS of the inequality is decreasing in θ_j by Assumption 1; i.e., citizens with higher anti-regime sentiment require a lower probability of success to rebel. Equivalently, $h(\theta_j)$ is increasing in θ_j .

Finally, given the assumption that citizens are using a strategy of the form “rebel if and only if $\theta_j \geq \hat{\theta}$ ” in each revolution stage, the condition for this

to be an equilibrium is that a citizen observing exactly $\theta_j = \hat{\theta}$ is indifferent between rebelling and not rebelling when other citizens use this strategy. Citizens receiving a higher signal have a higher relative payoff for rebelling *and* will believe the rebellion is more likely to succeed (demonstrated below), and hence rebel. Conversely, citizens with a lower signal do not rebel. This indifference condition is:

$$q(\theta_j = \hat{\theta}, \hat{\theta}) = 1 - h(\hat{\theta}) \quad (2)$$

where the second argument of q refers to the strategy used by other citizens in the relevant revolution stage.

The assumption that some citizens have a dominant strategy to rebel and some have a dominant strategy to not rebel is described in Morris and Shin (2003) as two-sided limit dominance. The assumption that some citizens never rebel should not be too controversial, but assuming some citizens rebel even if they know the revolution will fail is more questionable. This assumption is not unprecedented: (Granovetter, 1978) allows for “rabble-rousers,” and Kuran (1989) has citizens hurt enough by preference falsification to always support their preferred side, both of which are analogous to citizens with a dominant strategy to rebel. In other words, two-sided limit dominance can result if even a small group of citizens get an expressive payoff for rebelling against a regime they strongly dislike independent of the potential for success.

Whether or not these assumptions are more realistic than counterparts where no such signals exist is a difficult if not unanswerable question; I choose the proposed assumptions for two reasons. First, it seems safer to allow for the possibility that in an arbitrarily large group of citizens some are extreme enough to support their preferred side even if this side can not win. Second, by pushing away from the corner solutions of all citizens rebelling or all not rebelling the model becomes simpler to analyze.⁶

The Probability of Rebellion Succeeding

Completing the description of the optimal citizen strategy requires the action for those without a dominant strategy, which requires an assessment of

⁶ Bueno De Mesquita (2010b) and Shadmehr and Bernhardt (2011) discuss how limit dominance assumptions affect the resulting equilibria. The key results here would likely stay unchanged with one-sided limit dominance as long as a stable positive-participation equilibrium exists.

the probability that the rebellion succeeds. This will be affected by the distribution of the general and individual levels of anti-regime sentiment. Assume all actors share a common prior about the systematic anti-regime sentiment θ that is normally distributed with mean μ_0 and standard deviation σ_0 . Let the individual anti-regime sentiment signals be $\theta_j = \theta + \epsilon_j$ where the ϵ_j are independent normal random variables with mean 0 and standard deviation σ_p .

If an election is held, before observing their private signals all citizens observe $e = \theta + \epsilon_e$ where ϵ_e is normally distributed with mean 0 and standard deviation σ_e . The order of the signals does not matter as the citizens observe both signals before taking their action. While the citizen payoffs are not directly affected by the election result, each citizen's belief about the probability that the rebellion will succeed depends on their belief about the anti-regime sentiment of other citizens. Since this belief is affected by the election result, the election result affects the equilibrium strategy.

These assumptions lead to standard results regarding citizens beliefs about θ upon observing their signal and, if held, the election result. Citizens' beliefs after seeing the election result but before the private signal are normally distributed with mean $\mu_1(e) \equiv \frac{\sigma_0^2 e + \sigma_e^2 \mu_0}{\sigma_0^2 + \sigma_e^2}$ and standard deviation $\sigma_1 \equiv \frac{\sigma_0 \sigma_e}{\sqrt{\sigma_0^2 + \sigma_e^2}}$. In words, holding the election has two effects on this belief. First, it affects the *location*, i.e., higher election results lead all citizens to believe the anti-regime sentiment is higher. Second, it increases the *precision* of the belief ($\sigma_1 < \sigma_0$).

Since this posterior belief is normal, the citizen's decision to rebel or not rebel upon observing election result e is exactly the same as it would be in a no-election revolution stage with a different prior. In what follows $\hat{\theta}$ with no arguments refers to a generic or proposed cutoff strategy, and $\hat{\theta}(\mu, \sigma)$ refers to an equilibrium strategy in a revolution stage with a prior on θ with mean μ and standard deviation σ . So, $\hat{\theta}(\mu_0, \sigma_0)$ is a cutoff equilibrium in the no-election revolution stage and $\hat{\theta}(\mu_1(e), \sigma_1)$ is a cutoff equilibrium in a post-election revolution stage with result e .

Referring back to Equation (2), the probability needed to characterize these equilibria is $q(\theta_j = \hat{\theta}, \hat{\theta})$: the probability of rebellion succeeding upon observing signal $\theta_j = \hat{\theta}$ when the other citizens are using cutoff rule $\hat{\theta}$. For any fixed cutoff rule, the proportion of rebelling citizens as a function of the general anti-regime sentiment is continuous and increasing from 0 when θ is arbitrarily low to 1 when θ is arbitrarily high. Since $T \in (0, 1)$, there

exists a θ' such that the rebellion succeeds if and only if $\theta > \theta'$. As a result, after deriving θ' we can write the probability of rebellion succeeding from the perspective of a citizen observing θ_j as $Pr(\theta > \theta' | \theta_j)$.

The details of these calculations add little intuition and are analogous to standard results (e.g., Morris and Shin, 2003), so they are placed in the Appendix. The equilibrium conditions for the no-election revolution stage ($\hat{\theta}(\mu_0, \sigma_0)$) and the election revolution stage ($\hat{\theta}(\mu_1(e), \sigma_1)$) are:

$$\Phi \left(\frac{\left(\frac{\sigma_p^2}{\sigma_0^2 + \sigma_p^2} \right) (\hat{\theta}(\mu_0, \sigma_0) - \mu_0)}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \right) = h(\hat{\theta}(\mu_0, \sigma_0)) \quad (3)$$

$$\Phi \left(\frac{\left(\frac{\sigma_p^2}{\sigma_1^2 + \sigma_p^2} \right) (\hat{\theta}(\mu_1(e), \sigma_1) - \mu_1(e))}{\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}} \right) = h(\hat{\theta}(\mu_1(e), \sigma_1)) \quad (4)$$

where $\Phi(\cdot)$ is the cumulative density function of a standard normal random variable. Both of these equations state that when observing the threshold signal and conjecturing other citizens use this threshold, the probability assigned to the regime surviving (the LHS) is equal to the probability required to make the citizen indifferent between rebelling or not when they have the threshold level of anti-regime sentiment (the RHS).

The assumption that some citizens have dominant strategies not to rebel and some have dominant strategies to rebel ensures that there is at least one $\hat{\theta}$ satisfying these equations and hence at least one equilibrium cutoff strategy. Further, if the prior (or posterior upon observing the election result) is sufficiently diffuse the equilibrium will be unique:

Lemma 1 *The equilibrium citizen strategies in the revolution stage have the following properties:*

- (i) *At least one threshold equilibrium exists, and the number of equilibria is odd.*
- (ii) *If the prior belief or posterior belief after observing the election result about the anti-regime sentiment is sufficiently diffuse, then the threshold equilibrium strategy is unique.*

Proof: See the Appendix. ■

This is a common result in global games and related models. The general intuition is that when the prior is strong the game approaches a complete-information coordination game where the citizens without a dominant strategy always or never rebel. However, when the prior is diffuse, citizens are more uncertain about the signals seen by others, leading to a threshold strategy where those observing the cutoff signal assign a moderate chance of success to the rebellion.

Figure 1 illustrates how the revolution stage plays out for various election results when σ_1 is high enough that there is always a unique

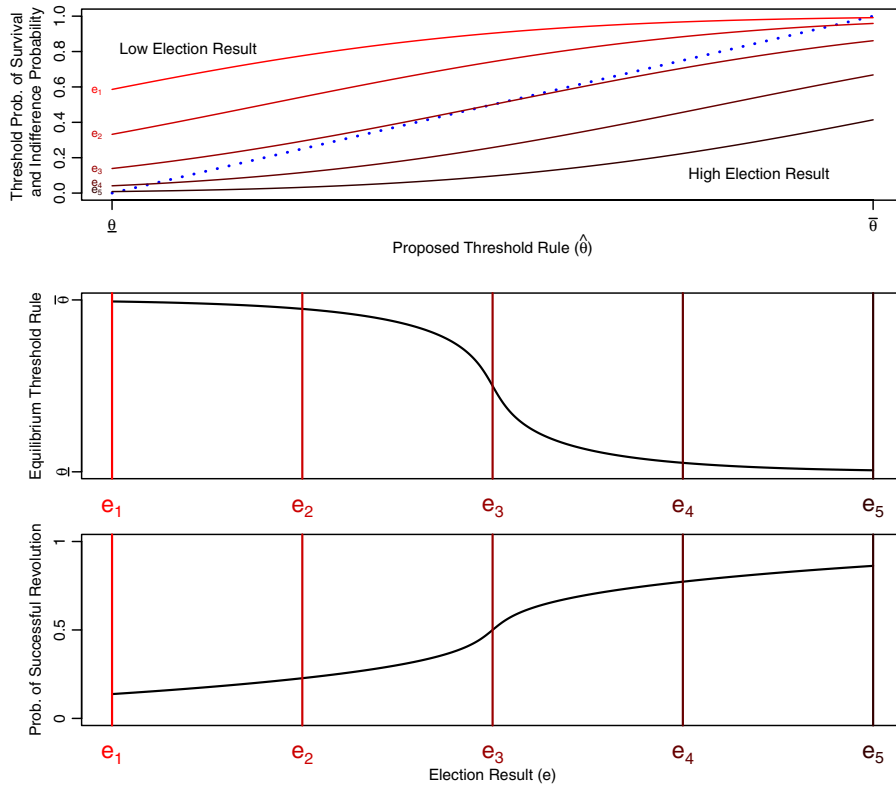


Figure 1. Plot of the indifference condition (top), the resulting equilibrium threshold rule (middle), and the probability of successful rebellion (bottom) as a function of the election result, with $h(\theta) = \theta$, $\sigma_0 = \sigma_e = 0.6$, $\mu_0 = 0.5$ and $e \in \{-0.7, -0.1, 0.5, 1.1, 1.7\}$.

equilibrium — an identical picture could be drawn for the no-election revolution stage with various values of mean anti-regime sentiment (μ_0) taking the place of the election results. The top panel illustrates the equilibrium indifference condition. The dotted line is $h(\hat{\theta})$ and the increasingly dark solid lines are the LHS of Equation (3) for increasing election results ($e_1 < \dots < e_5$). Higher results shift where the LHS attains its maximum slope, but since this is never steeper than the dotted line there is only one intersection and hence the equilibrium is always unique. As we would expect, election results indicating higher anti-regime sentiment (darker solid lines) lead to citizens playing a lower threshold rule (middle panel) and hence a higher probability of successful revolution (bottom panel). The election results plotted range from two standard deviations below to two standard deviations above the average result, indicating that the size of protests and chance of revolution can change dramatically over the range of likely election results. Proposition 1 formalizes this observation:

Proposition 1 *If there is a unique equilibrium in the post-election revolution stage, the probability of revolution is increasing in the election result.*

Proof: See the Appendix. A slightly weaker result holds if there are multiple equilibria. ■

Notably, the incumbent performing well in the election (i.e., low e) does not lead to less protest because citizens inherently care about the result or by bolstering the incumbent's legitimacy — however defined — but because the information generated by a higher election result affects beliefs about how many other citizens will protest.

Most of the remaining intuitive comparative statics common to the global games approach are captured by the h function. Shifting the h function up — equivalent to making the relative value of rebelling higher — lowers the equilibrium threshold, increasing the number of protestors and making successful rebellion more likely. For example, increasing the value of participating in a successful rebellion — perhaps because of high natural resource rents — shifts h up and leads to a higher probability of rebellion. Decreasing the payoff from participating in a failed revolution — perhaps because the incumbent has strong control over a repressive apparatus — shifts h down and makes successful rebellion less likely.

The Electoral Decision — Unique Equilibrium Case

Given the solution to the revolution stages, we can now examine the incumbent's decision to hold or not hold the election, which will depend on whether the revolution stage equilibria are unique. In the interest of space I restrict attention to the unique equilibrium case for the remainder of the paper; see Little (2011b) on interpreting the presence of multiple equilibria in a related game.

Assumption 2 *The prior and electoral signal are sufficiently weak that there is a unique equilibrium in each revolution stage.*

To make the electoral decision, the incumbent needs to know the probability that they will survive the potential rebellion without an election, the probability that they will survive the rebellion for a fixed election result e , and the probability distribution of election outcomes. The first two values follow a similar calculation made by the citizens when deciding to rebel or not. The third comes from a standard result about normal distributions with an uncertain mean, giving expected utilities:

$$\begin{aligned}\mathbb{E}[U_0^I] &= \Phi\left(\frac{\hat{\theta}(\mu_0, \sigma_0) - \mu_0}{\sigma_0}\right) \\ \mathbb{E}[U_1^I] &= \int_{-\infty}^{\infty} \Phi\left(\frac{\hat{\theta}(\mu_1(e), \sigma_1) - \mu_1(e)}{\sigma_1}\right) \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi\left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}}\right) de - k\end{aligned}$$

where $\phi(\cdot)$ is the probability density function of a standard normal random variable. Hence there is an equilibrium where the incumbent holds an election if and only if:

$$\begin{aligned}\Delta_q \equiv \int_{-\infty}^{\infty} \Phi\left(\frac{\hat{\theta}(\mu_1(e), \sigma_1) - \mu_1(e)}{\sigma_1}\right) \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi\left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}}\right) de \\ - \Phi\left(\frac{\hat{\theta}(\mu_0, \sigma_0) - \mu_0}{\sigma_0}\right) \geq k\end{aligned}\tag{5}$$

where Δ_q is the relative value of holding an election due to the change in survival probability.

Figure 2 illustrates the intuition for when the informational effect is a net benefit to the incumbent. The probability of survival as a function of




Figure 2. Probability of survival as a function of election result.

the election result is a backwards s-shaped curve as it must approach 1 for arbitrarily low (i.e., favorable to the incumbent) results and 0 for arbitrarily high (i.e., unfavorable to the incumbent) results. The average result for incumbents likely to stay in power in general (e.g., when μ_0 is low) in the no-election revolution stage is near the top of the s-curve, hence these incumbents have more to lose from a bad election result than they have to gain from a good election result. Conversely, incumbents with a low probability of survival are on the bottom of the s-curve and hence have more to gain from a good result than they can lose from a bad result.

However, for incumbents that are strong or weak enough that they are highly unlikely or likely to lose office regardless of their electoral decision, the informational value of holding an election will be close to zero. So incumbents who are somewhat insecure in office are the ones most likely to hold elections as a result of the mechanism described here, and incumbents who are somewhat secure in office are the ones least likely to hold elections as a result of this mechanism. Proposition 2 summarizes the equilibrium when each revolution stage equilibrium is unique, see the Appendix for a more formal statement and proof⁷:

Proposition 2 *If the prior and electoral signal are sufficiently diffuse, then the citizens have a unique cutoff equilibrium strategy in each revolution stage as defined above, and the equilibrium can be electoral if and only if the informational value of elections is greater than the exogenous cost of elections k .*

⁷ The proof relies on convenient properties of the normal distribution, but likely holds for other distributions. Broadly speaking, the result is driven by the fact that the probability of rebellion succeeding in the unique equilibrium case — as it would be for many distributional assumptions — must be continuously increasing everywhere and approaching 1 for unpopular incumbents and 0 for popular incumbents. For any distribution with these properties and support on the real line this will lead to a roughly s-shaped curve and results analogous to Proposition 2.

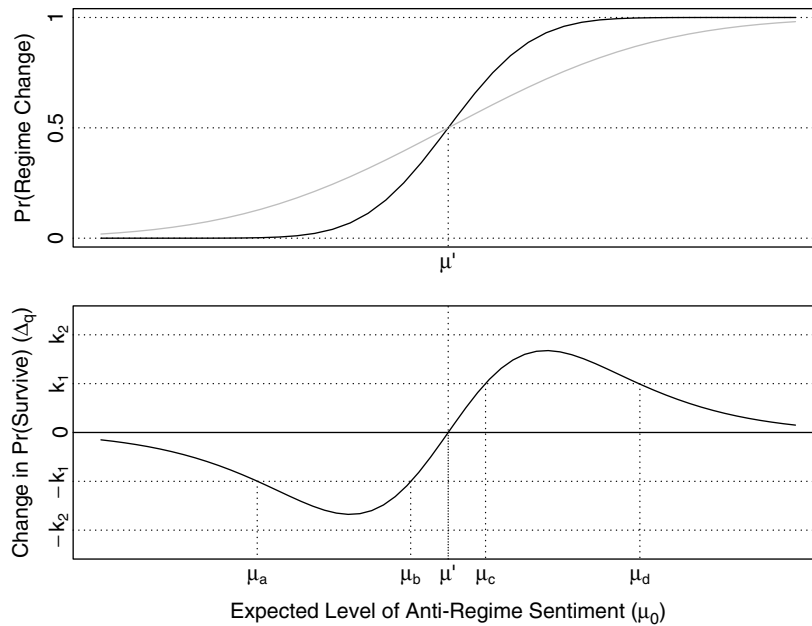


Figure 3. Top panel: the probability of regime change without an election (black) and with an election (grey) as a function of prior mean anti-regime sentiment (μ_0). Bottom panel: change in the probability of survival (Δ_q) when holding an election. Elections are held when the curve is above the horizontal line corresponding to their cost (k). In both panels, $h(\theta) = \theta$, $\sigma_0 = 2$, $\sigma_e = 2$, $\sigma_p = 0.5$. In the bottom panel, $k_1 = 0.125$, $k_2 = 0.25$.

The informational value of elections tends to be positive when the incumbent is weak and negative when the incumbent is strong.

Figure 3 illustrates the electoral decision for $h(\theta) = \theta$. The top panel compares the probability of regime change without an election (black line) and with an election (grey) as a function of the expected anti-regime sentiment. When the regime is likely to survive in general ($\mu_0 < \mu'$), regime change is more likely with an election. When regime change is likely in general ($\mu_0 > \mu'$), incumbents are more likely to survive with an election.

The bottom panel plots the change in the probability of survival (Δ_q) as a function of the expected anti-regime sentiment. The horizontal lines correspond to various values of election costs k ; elections are held when the curve is above k . Hence when $k = k_2$, elections have a high net cost, and

no elections are held. For $k = k_1$, elections are held only for $\mu_0 \in [\mu_c, \mu_d]$. Looking back up at the top panel, this is the range of anti-regime sentiment where the incumbent is moderately insecure in office — the probability of survival is less than $\frac{1}{2}$, but not low enough that they are nearly certain to be removed regardless of the electoral choice.

When $k = -k_1$, elections have a modest net benefit, and elections are held everywhere but $\mu_0 \in [\mu_a, \mu_b]$, which correspond to the incumbents that are moderately secure in office. Finally, when $k = -k_2$, the exogenous benefit to holding elections is high enough that they are always held.

Proposition 2 contrasts interestingly with Egorov and Sonin (2011), who model the decision to hold an election in a similar fashion where the incumbent has (perfect) private information about the anti-regime sentiment. In Egorov and Sonin (2011), it is the *strong* types who tend to hold elections as they are the ones who want to reveal the information generated by elections (see Little (2011b) for a generalization of this argument outside of the global games context). This is not inconsistent with the results presented here. The insecure incumbents identified in Proposition 2 are weak in the sense of the *common knowledge* mean anti-regime sentiment; in Egorov and Sonin (2011) the incumbents who hold elections are those with *private* information indicating they are strong. Thus the incumbents with the greatest incentive to hold elections may be those who are generally perceived to be weak but know that they are stronger than expected, and the leaders most likely to forgo or cancel elections are those perceived to be strong but know they are less popular than commonly thought.

Fraud and Election Monitoring

In the previous section, the election result is an unbiased signal of the incumbent popularity. However, electoral fraud is salient in most cases where the decision to hold an election is in doubt. Fraud plays a large role in many theories of why autocrats or hegemonic parties hold elections (e.g., Magaloni (2006) and Geddes (2006)), so one might think the ability to manipulate the results makes elections appealing to a wider range of incumbents.

In the specification considered below, this intuition will not hold: the ability to commit fraud makes elections *less* valuable to incumbents. Since the citizens know that fraud can be committed, they discount the results of the election when forming their belief about the level of anti-regime sentiment.

Hence in equilibrium, the level of protests and probability of successful revolution in the model with fraud is exactly the same as in a corresponding game with no fraud as described above. The incumbent does not commit less than the expected amount of fraud as this would result in citizens overly discounting the election result and more protest. So, as long as the amount of fraud committed in equilibrium is costly to the incumbent, they are worse off holding an election when fraud is possible than they are in a world where they cannot commit fraud.

Given this result, incumbents inviting election monitors and delegating authority to independent electoral bodies makes sense: in the extreme case where monitoring renders fraud so costly that none is committed in equilibrium, incumbents will be able to get the same probability of maintaining office without incurring any costs. Under certain cost assumptions laid out below, this will hold even if the monitoring is imperfect: while the incumbent will still commit some fraud under the increased monitoring and the monitoring makes committing a fixed amount of fraud more costly, this will lead to a lower equilibrium cost paid and hence a higher incumbent payoff in the electoral revolution stage.

Model with Fraud

Suppose when holding an election the incumbent also chooses an amount of fraud to commit $x \geq 0$, which changes the true election result e to fraudulent electoral result $e_f = e - x = \theta + \epsilon_e - x$. (Recall a lower election result implies less anti-government sentiment.) The citizens do not observe x , but are aware that the incumbent has the ability to manipulate the results. This is admittedly a very stylized conception of fraud, and is intended to represent a wide array of hidden and costly actions the incumbent can take to manipulate the results (see Schendler (2002)).

Fraud entails a cost $c(x)$, where c satisfies $c(0) = 0$, $c'(0) = 0$, $c'(x) \geq 0$, $c''(x) > 0$, and $\lim_{x \rightarrow \infty} c'(x) = \infty$. Loosely speaking, these assumptions imply that falsifying the first vote is free, but more fraud is more costly and exhibits an increasing marginal cost. These standard assumptions can arise if the incumbent chooses the easiest technologies of fraud first, making the marginal cost cheap for low levels of fraud and higher when committing larger amounts of fraud.

The cost term is primarily intended to capture the direct cost of falsifying votes such as paying cronies to commit the falsification and keep quiet

about it, though it could also capture indirect costs such as the potential for international condemnation or reduced aid. The incumbent strategy is now a choice to hold an election or not and the amount of fraud to commit if they hold the election. As above I restrict attention to pure strategies; denote the equilibrium amount of fraud x^* .

As in the game with no fraud, the citizens must form a belief about the level of anti-regime sentiment (θ) based on the prior and election result before observing their private signal, which now must also take into account the expected amount of fraud. For election result e_f and proposed equilibrium amount of fraud x' , the updating is the same as it would be in the no fraud case with election result $e = e_f + x'$, hence the mean of belief analogous to $\mu_1(e)$ is:

$$\mu_f(e_f; x') = \frac{\sigma_0^2(e_f + x') + \sigma_e^2\mu_0}{\sigma_0^2 + \sigma_e^2}$$

and the standard deviation of the belief about θ before seeing the private signal is still σ_1 as defined in the previous section.

Since the incumbent has no private information and is playing a pure strategy, in any equilibrium the citizens must play the optimal strategy given how much fraud is committed, so $x' = x^*$. Thus in equilibrium, $\mu_f(e_f; x^*) = \mu_1(e)$; i.e., the citizens' belief after viewing the fraudulent election result is identical to the belief they have in an analogous revolution stage where fraud is not possible. So, when playing the optimal strategy given this belief, the size of protest and probability of the incumbent surviving after an election are exactly the same as in the game with no fraud.

The final equilibrium condition is that when the citizens use the cutoff rule induced by believing the incumbent commits x^* fraud, this is actually the optimal amount to commit. The payoff for holding an election and committing amount of fraud x when this belief is fixed is:

$$\begin{aligned} \mathbb{E}[U_1^I(x; x' = x^*)] &= \int_{-\infty}^{\infty} \Phi \left(\frac{\hat{\theta}(\mu_f(e_f + x^* - x; x^*), \sigma_1) - \mu_1(e)}{\sigma_1} \right) \\ &\quad \times \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) de - c(x) \end{aligned} \quad (6)$$

If $x > x^*$, the mean belief about the anti-regime sentiment μ_f will be less than the ‘‘correct’’ belief $\mu_1(e)$, and if $x < x^*$ the opposite will hold. That is,

by committing more fraud than expected the incumbent seems stronger than they truly are, and when committing less fraud than expected they seem weaker than they truly are. Both the beliefs that a large amount of fraud was committed and a lower than expected election result will lead to more protest. These effects may have played a role in the Colored Revolutions as well as the less dramatic but still major protests following the Russian Duma election in late 2011.

The key tradeoff for the incumbent is that committing more fraud than expected is costly but leads to a lower probability of successful rebellion. For $x = x^*$ to be in equilibrium, the first order condition obtained by setting the derivative of Equation (6) with respect to x equal to 0 is:

$$c'(x^*) = \left(\int_{-\infty}^{\infty} \frac{1}{\sigma_1 \sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{\hat{\theta}(\mu_1(e), \sigma_1) - \mu_1(e)}{\sigma_1} \right) \phi \left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) de \right) \\ \times \left(\frac{\partial \hat{\theta}(\mu_1(e), \sigma_1)}{\partial \mu_1} \right) \left(\frac{\partial \mu_1}{\partial e} \right)$$

That is, the marginal cost of committing more fraud than expected (the LHS) is equal to the marginal benefit of committing more fraud than expected through decreasing the level of protest (the RHS). The marginal benefit of committing more fraud than expected is independent of the expected level x^* . This results from the fact that the election result only matters through how it affects the citizens' beliefs. Loosely speaking, these beliefs are equally affected by a deviation to falsifying 6% of the votes when the citizens expect 5% as a deviation to falsifying 11% of the votes when the citizens expect 10%.

Given the cost function assumptions, for any positive constant b there will be a unique $x^* > 0$ such that $c'(x^*) = b$, hence there will be a unique x^* such that the first order condition is met when the citizens are playing a best response to x^* . This x^* is the only potential pure strategy equilibrium amount of fraud, and under some additional technical assumptions it will be a global maximizer:

Lemma 2 *Provided c is sufficiently convex, in any electoral equilibrium there is a positive amount of fraud $x^* > 0$.*

Proof: See the Appendix. ■

The fact that this equilibrium has a positive x implies the main result regarding fraud⁸:

Proposition 3 *The incumbent equilibrium payoff in the game without fraud is as high or higher than the payoff in the game with fraud. If the equilibrium is electoral, the payoff in the game with no fraud is strictly higher.*

Proof: This follows from Lemma 2 and the observation above that the probability of surviving the revolution is the same in the game with and without fraud. ■

Election Monitoring

The finding that the incumbent would like to remove their ability to commit fraud provides an explanation for the prevalence of election monitoring. To formalize this intuition, suppose the cost paid to commit amount of fraud x also depends on the level of monitoring m . Before the electoral decision, the incumbent selects monitoring level $m \in [\underline{m}, \bar{m}]$, where $0 < \underline{m} < \bar{m}$. The monitoring decision is observed by citizens, and generates cost function $c(x; m) = mc_0(x)$. That is, higher levels of monitoring raise the cost of committing fraud. Assume that the cost function under the lowest level of monitoring $c(x; \underline{m})$ meets the conditions of Lemma 2, which ensures there is a unique equilibrium level of fraud associated with each level of monitoring.

Denote the equilibrium level of fraud for each monitoring level $x^*(m)$. Since citizens correctly discount the fraud, this decision will not affect the probability of maintaining office. As a result, the optimal level of monitoring is the level that minimizes the equilibrium cost $c(x^*(m); m)$. The model with no fraud can be seen as the limiting case as $m \rightarrow \infty$, and the result below is the continuous analog of the observation that the incumbent prefers the world where they cannot commit fraud.

Lemma 3 *The incumbent payoff equilibrium payoff is strictly increasing in the level of monitoring.*

Proof: See the Appendix. ■

⁸ Bueno De Mesquita (2010a) makes an analogous point when discussing the *ex ante* and *ex post* effectiveness of terrorist violence.

This leads to the main result regarding the game with election monitoring:

Proposition 4 *In the game with fraud and election monitoring:*

- (i) *In any equilibrium where an election is held, the incumbent chooses the highest level of monitoring $m = \bar{m}$.*
- (ii) *For two distinct highest levels of monitoring $\bar{m}_2 > \bar{m}_1$, the set of parameters for which an election is held under the lower level (\bar{m}_1) is a strict subset of the set of parameters for which an election is held under the high level of monitoring (\bar{m}_2).*

The first part follows from Lemma 3. The result is not specific to the global games setup or how monitoring is modeled, see Little (2011a) for an analogous result in a more general setting where monitoring is a noisy signal of the amount of fraud. The second part draws a connection between the electoral decision and the availability of election monitoring. Higher monitoring lead to less utility loss due to fraud, so it makes elections more desirable. While it would be excessively strong to claim that this connection is the prime reason that the number of countries holding elections rose at the same time that election monitoring became available and pervasive, it does imply there might be some unexpected causality in this relationship.

Conclusion

Treating elections primarily as generating a public signal released before some game — in particular, a game of revolution — generates insights into both the conduct of elections and why they are held in the first place. The question of when the informational mechanism makes elections beneficial to the incumbent depends on their security in office without holding an election. Why relatively insecure incumbents who will tend to do less well in elections still hold them may be a particularly hard question, and in this formulation these are precisely the leaders with the most to gain from elections. Empirically testing this proposition poses some difficulties. First, it requires identifying the moderately insecure regimes, which is difficult without using *ex post* information about how long the regime lasted. Data

such as public opinion polls, recent riots, or economic problems could serve as proxies, but could potentially affect the decision to hold elections through other channels.

Elections could be valuable to a wider range of leaders in a related model where they have the opportunity to utilize the information generated by elections as well as the citizens. For example, the incumbent could grant concessions to citizens upon observing poor results to ameliorate the revolutionary pressure. Further, the incumbent could take advantage of good results to consolidate their power. This mechanism would seem to make good results even better and bad results less damaging, making elections overall more valuable.

The results regarding fraud and election monitoring are more straightforward and transparently consistent with the motivating empirical facts. The favorable behavior induced by a convincing election result gives an incentive to commit fraud independent of winning the election by the official tally. The strong incentive to invite election monitoring to reduce (but not eliminate) expectations about how much fraud will occur is consistent with incumbents voluntarily inviting monitors and still cheating in front of them.

The normative implications of the results about election monitoring largely favor the practice. While monitoring does not increase the chances of deposing unpopular incumbents, it does lower the amount of fraud committed in equilibrium, consistent with the empirical findings in Hyde (2007). Presumably many actions associated with fraud such as intimidating opposition voters are costly to more than just the incumbent, so we can loosely say that monitoring may leave all actors better off. Further, as Proposition 4 shows, increasing the level of monitoring available to incumbents increases the set of parameters under which elections are possible.

The results here demonstrate the value of treating elections as generating public information as opposed to assuming they change the nature of the political interaction between incumbent leaders and other actors. The information provided by elections does not only affect incumbent behavior as emphasized in past work, but also *all* political actors from elites to citizens. Applying this general perspective to different information structures and political interactions other than revolution may lead to further insights into the causes and consequences of elections.

Appendix: Complete Derivations and Proofs

The formal definition of the equilibrium concept is:

Definition. A PBE of the baseline model requires: (1) the density of the citizen belief about θ upon observing their private signal without an election $f(\theta|\theta_j)$ and upon observing θ_j and an election result $f(\theta|\theta_j, e)$, (2) the thresholds used by citizens in the revolution stage without an election $\hat{\theta}$ and with an election $\hat{\theta}(e)$, and (3) the incumbent electoral decision $s^I \in \{0, 1\}$. The consistency and rationality conditions are:

$$f(\theta|\theta_j) = \frac{Pr(\theta, \theta_j)}{Pr(\theta_j)} \quad (7)$$

$$f(\theta|\theta_j, e) = \frac{Pr(\theta, \theta_j, e)}{Pr(\theta_j, e)} \quad (8)$$

$$\mathbb{E}[U^C(a = 1|\theta_j, \hat{\theta})] \geq \mathbb{E}[U^C(a = 0|\theta_j, \hat{\theta})] \quad \text{if and only if } \theta_j \geq \hat{\theta} \quad (9)$$

$$\begin{aligned} \forall e : \mathbb{E}[U^C(a = 1|\theta_j, \hat{\theta}(e), e)] \\ \geq \mathbb{E}[U^C(a = 0|\theta_j, \hat{\theta}(e), e)] \quad \text{if and only if } \theta_j \geq \hat{\theta}(e) \end{aligned} \quad (10)$$

$$s^I = 1 \quad \text{if and only if } \mathbb{E}[U_1^I] \geq \mathbb{E}[U_0^I] \quad (11)$$

where $\mathbb{E}[U^C(a|\theta_j, \hat{\theta})]$ denotes the citizen expected utility for taking action a when observing θ_j and other citizens use strategy $\hat{\theta}$, and $\mathbb{E}[U^C(a|\theta_j, \hat{\theta}(e), e)]$ is the analogous utility for the election revolution stages.

For the citizen beliefs, first note that since the incumbent acts first in the model and has no private information about θ : (1) the incumbents belief when making the electoral decision is the prior belief, and (2) the citizens make no inference about θ based on the electoral decision itself. By a standard result for normal priors with a normal signal, the posterior belief about θ upon observing the election result is normally distributed with mean and standard deviation:

$$\mu_1(e) = \frac{\sigma_0^2 e + \sigma_e^2 \mu_0}{\sigma_0^2 + \sigma_e^2} \quad \text{and} \quad \sigma_1 = \frac{\sigma_0 \sigma_e}{\sqrt{\sigma_0^2 + \sigma_e^2}} < \sigma_0$$

Given the same rules of Bayesian updating with normal priors and normal signals, the belief about θ for a citizen receiving signal θ_j is normally distributed with mean $\frac{\theta_j \sigma_0^2 + \mu_0 \sigma_p^2}{\sigma_0^2 + \sigma_p^2}$ and standard deviation $\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}$. Upon observing θ_j and e the belief is normally distributed with mean $\frac{\theta_j \sigma_1^2 + \mu_1(e) \sigma_p^2}{\sigma_1^2 + \sigma_p^2}$ and

standard deviation $\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}$ Or, in the notation above,

$$f(\theta|\theta_j) = \frac{1}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \phi \left(\frac{\theta - \frac{\theta_j \sigma_0^2 + \mu_0 \sigma_p^2}{\sigma_0^2 + \sigma_p^2}}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \right) \text{ and}$$

$$f(\theta|\theta_j, e) = \frac{1}{\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}} \phi \left(\frac{\theta - \frac{\theta_j \sigma_1^2 + \mu_1(e) \sigma_p^2}{\sigma_1^2 + \sigma_p^2}}{\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}} \right)$$

It is optimal to rebel in the revolution stages if and only if $q(\cdot) \geq 1 - h(\theta_j)$, where $q(\cdot)$ is citizen j 's belief that the rebellion will succeed at the relevant information set and $\hat{\theta}$ is the proposed cutoff strategy. By the distribution of the signals, the mass of citizens rebelling when using cutoff strategy $\hat{\theta}$ will be $\Phi\left(\frac{\theta - \hat{\theta}}{\sigma_p}\right)$, which approaches 0 for arbitrarily low θ , 1 for arbitrarily high θ , and is continuous and increasing for all θ . As a result, for any $\hat{\theta}$ there will exist a unique true anti-regime sentiment parameter θ' such that $\Phi\left(\frac{\theta' - \hat{\theta}}{\sigma_p}\right) = T$ and hence the rebellion succeeds if and only if $\theta \geq \theta'$. An analogous argument applies to the election revolution stages, giving:

$$\theta' = \hat{\theta} + \Phi^{-1}(T)\sigma_p \tag{12}$$

$$\theta'(e) = \hat{\theta}(e) + \Phi^{-1}(T)\sigma_p \tag{13}$$

When $T = 0.5$, as assumed in the main text, the terms with $\Phi^{-1}(0.5) = 0$ drop out; the appendix will not make this assumption. So, the citizen's assessment of the probability the rebellion will succeed in the no-election revolution and post-election revolution stages are:

$$q(\theta_j, \hat{\theta}) = Pr(\theta > \theta'|\theta_j) = 1 - \Phi \left(\frac{(\hat{\theta} + \Phi^{-1}(T)\sigma_p) - \left(\frac{\theta_j \sigma_0^2 + \mu_0 \sigma_p^2}{\sigma_0^2 + \sigma_p^2}\right)}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \right) \tag{14}$$

$$q(\theta_j, \hat{\theta}(e), e) = Pr(\theta > \theta'(e)|\theta_j, e)$$

$$= 1 - \Phi \left(\frac{(\hat{\theta}(e) + \Phi^{-1}(T)\sigma_p) - \left(\frac{\theta_j \sigma_1^2 + \mu_1(e) \sigma_p^2}{\sigma_1^2 + \sigma_p^2}\right)}{\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}} \right) \tag{15}$$

Formalizing the statements about the incumbent indifference condition in the main text:

Lemma 4 For $\theta_j < \underline{\theta}$, $h(\theta_j) < 0$. For $\theta_j > \bar{\theta}$, $h(\theta_j) > 1$. For $\theta_j \in [\underline{\theta}, \bar{\theta}]$, h is strictly increasing.

Proof: Part 1 of Assumption 1 ensures $\underline{\theta} < \bar{\theta}$. For $\theta_j < \underline{\theta}$, $u_1^S(\theta_j) - u_0^S(\theta_j) < 0$, and the denominator of h is positive, so $h(\theta_j) < 0$. For $\theta_j > \bar{\theta}$, $u_1^S(\theta_j) - u_0^S(\theta_j) > u_1^F(\theta_j) - u_0^F(\theta_j) > 0$, so the numerator of h is positive and greater than the denominator, hence $h(\theta_j) > 1$. For $\theta_j \in [\underline{\theta}, \bar{\theta}]$, the numerator of $1 - h(\theta_j)$ is $-(u_1^F(\theta_j) - u_0^F(\theta_j))$, which is positive and decreasing, and the denominator is positive and increasing, so $1 - h(\theta_j)$ is decreasing in this range and hence $h(\theta_j)$ is increasing. ■

Combining the results about the citizen beliefs and payoffs, there exists at least one threshold rule $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$ such that $q(\hat{\theta}, \hat{\theta}) = 1 - h(\hat{\theta})$ and citizens with signals $\theta_j < \hat{\theta}$ do not rebel and those observing $\theta_j > \hat{\theta}$ rebel. Combining Equation (14) with the indifference condition where $\theta_j = \hat{\theta}$ gives the final equilibrium condition for the no-election and election revolution stages:

$$\Phi \left(\frac{\left(\frac{\sigma_p^2}{\sigma_0^2 + \sigma_p^2} \right) (\hat{\theta} - \mu_0) + \Phi^{-1}(T)\sigma_p}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \right) = h(\hat{\theta}) \tag{16}$$

$$\Phi \left(\frac{\left(\frac{\sigma_p^2}{\sigma_1^2 + \sigma_p^2} \right) (\hat{\theta}(e) - \mu_1(e)) + \Phi^{-1}(T)\sigma_p}{\frac{\sigma_p \sigma_1}{\sqrt{\sigma_1^2 + \sigma_p^2}}} \right) = h(\hat{\theta}(e)) \tag{17}$$

Lemma 5 *For all other parameters fixed, let $\Theta(\mu, \sigma)$ be the set of threshold rules satisfying Equation (16) with prior mean μ and standard deviation σ . Then $\hat{\theta}(e)$ is an equilibrium cutoff rule in the election revolution stage with election result e if and only if $\hat{\theta}(e) \in \Theta(\mu_1(e), \sigma_1)$.*

Proof: This follows immediately from Equations (16) and (17) and the equilibrium definition. ■

Due to this equivalence, other than for the electoral decision we can analyze on a revolution stage with prior mean μ and standard deviation σ on θ and refer to all the revolution stages with appropriate moments in the prior.

For the uniqueness result, note that the LHS of Equation (16) as a function of the cutoff rule $\hat{\theta}$ is a normal cdf with standard deviation $\frac{\sigma_0 \sqrt{\sigma_0^2 + \sigma_p^2}}{\sigma_p}$, which is monotonically increasing on the interval $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$ and bounded by (0, 1). The RHS is 0 when $\hat{\theta} = \underline{\theta}$, 1 when $\hat{\theta} = \bar{\theta}$, and increasing on this interval. Define:

$$g(\hat{\theta}) \equiv \Phi \left(\frac{\left(\frac{\sigma_p^2}{\sigma_0^2 + \sigma_p^2} \right) (\hat{\theta} - \mu_0) + \Phi^{-1}(T)\sigma_p}{\frac{\sigma_p \sigma_0}{\sqrt{\sigma_0^2 + \sigma_p^2}}} \right) - h(\hat{\theta}) \tag{18}$$

which by the above arguments will be negative at $\underline{\theta}$, positive at $\bar{\theta}$, and continuous by Assumption 1. Therefore g has at least 1 zero between $\underline{\theta}$ and $\bar{\theta}$. By construction, these zeroes are the potential equilibrium thresholds. The same argument holds for the election revolution stages. Formally stating the equilibrium results for the revolution stage:

Lemma 1 *The revolution stage strategies have the following properties:*

- (i) *At least one threshold equilibrium exists, and other than for knife-edge parameters the number of threshold equilibria is odd.*
- (ii) *$\exists \bar{\sigma}$ such that if $\sigma_1 > \bar{\sigma}$, then the threshold equilibrium is unique for all revolution stages.*

Proof: (i) By the continuity of g the first and last zeroes must be at points where $g'(\theta) < 0$ (the “other than for knife-edge parameters” qualifier is for cases where $g'(\theta) = 0$ at the intersection), implying an odd number of intersections.

- (ii) Let $d = \min_{\theta \in (\underline{\theta}, \bar{\theta})} \frac{\partial h}{\partial \theta}$. By Assumption 1, $d > 0$. The maximum slope attained by the LHS of Equation (17) is $\frac{\sigma_p}{\sqrt{2\pi\sigma_1}\sqrt{\sigma_1^2 + \sigma_p^2}}$. This slope is decreasing in σ_1 , so setting the slope equal to d and solving for σ_1 implies that there exists a:

$$\bar{\sigma} \equiv \sqrt{\frac{\sqrt{\sigma_p^4 + 4\frac{\sigma_p^2}{2\pi d^2}} - \sigma_p^2}{2}} > 0$$

such that $\sigma_1 > \bar{\sigma}$ implies $g(\cdot)$ is nonincreasing and hence can only have one zero, guaranteeing a unique equilibrium. Since $\sigma_0 > \sigma_1$, this implies a unique equilibrium in the no-election revolution stage as well. A formal statement of Assumption 2 is $\sigma_1 > \bar{\sigma}$. ■

Proposition 1 follows from implicitly differentiating Equation (18) with σ_1 replacing σ_0 and $\mu_1(e)$ replacing μ_0 . The probability of revolution is increasing because $\frac{\partial \mu_1}{\partial e} > 0$ and $\frac{\partial \hat{\theta}}{\partial \mu} < 0$ at points where $g' < 0$, which must be the case at the first and last intersections as demonstrated by Lemma 1. So if the equilibrium is unique the probability of revolution is increasing in e , and if there are multiple equilibria the lowest and highest possible probabilities of revolution are increasing in e .

The incumbent utility for the no-election revolution stage is:

$$\mathbb{E}[U_0^I] = Pr(\theta < \theta') = \Phi \left(\frac{\hat{\theta}(\mu_0, \sigma_0) + \Phi^{-1}(T)\sigma_p - \mu_0}{\sigma_0} \right)$$

The incumbent expected utility for holding an election comes from integrating over the probability of survival for a fixed election result e times the density of this result. From a standard result about normal distributions with an uncertain and normally distributed mean, the incumbent's density for e is normally distributed with mean μ_0 and standard deviation $\sqrt{\sigma_0^2 + \sigma_e^2}$, so

$$\begin{aligned} \mathbb{E}[U_1^I] &= \int_{-\infty}^{\infty} (Pr(\theta < \theta'(e)|e) - k(e)) \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) de \\ &= \int_{-\infty}^{\infty} \Phi \left(\frac{\hat{\theta}(\mu_1(e), \sigma_1) + \Phi^{-1}(T)\sigma_p - \mu_1(e)}{\sigma_1} \right) \\ &\quad \times \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) de - k \end{aligned}$$

Proposition 2 compares these utilities to determine when elections are held:

Proposition 2 *There exists $\hat{\mu}_{low}, \hat{\mu}_{high}, \hat{\mu}_{low} \leq \hat{\mu}_{high}$ such that:*

- (i) Δ_q is strictly negative for $\mu_0 < \hat{\mu}_{low}$ and strictly positive for $\mu_0 > \hat{\mu}_{high}$
- (ii) Δ_q as a function of μ_0 attains a global minimum for some $\mu_0 < \hat{\mu}_{high}$ and a global maximum for some $\mu_0 > \hat{\mu}_{low}$.

Proof: The general idea of the proof is to demonstrate that as $\mu_0 \rightarrow -\infty$, Δ_q approaches 0 from below, and as $\mu_0 \rightarrow \infty$ Δ_q approaches 0 from above. Some intermediate results:

Lemma 6 *Consider Δ_q as a function of μ_0 . Then:*

- (i) $\Delta_q(\mu_0)$ is continuous in μ_0 .
- (ii) $\lim_{\mu_0 \rightarrow -\infty} \Delta_q = \lim_{\mu_0 \rightarrow \infty} \Delta_q = 0$
- (iii) $\lim_{\mu_0 \rightarrow -\infty} \frac{\partial \hat{\theta}}{\partial \mu_0} = \lim_{\mu_0 \rightarrow \infty} \frac{\partial \hat{\theta}}{\partial \mu_0} = 0$

The first part follows from the continuity of the inputs, the second from the fact that the probability of survival approaches 1 in every revolution

stage when μ_0 is sufficiently small and 0 when μ_0 is sufficiently large. Part iii follows from implicitly differentiating Equation (18).

The derivative of the payoff not holding an election with respect to μ_0 is:

$$\frac{\partial \mathbb{E}[U_0^I]}{\partial \mu_0} = \frac{1}{\sigma_0} \phi \left(\frac{\hat{\theta}(\mu_0, \sigma_0) + \Phi^{-1}(T)\sigma_p - \mu_0}{\sigma_0} \right) \left(\frac{\partial \hat{\theta}}{\partial \mu_0} - 1 \right)$$

When μ_0 is very negative or positive we can ignore the $\frac{\partial \hat{\theta}}{\partial \mu_0}$ term by Lemma 6. Integrating over the noise term ϵ_e , the analogous derivative for the payoff from holding an election is:

$$\begin{aligned} \frac{\partial \mathbb{E}[U_1^I]}{\partial \mu_0} &= \int_{-\infty}^{\infty} \frac{1}{\sigma_1} \phi \left(\frac{\hat{\theta}(\mu_0 + \frac{\sigma_0^2 \epsilon_e}{\sigma_0^2 + \sigma_e^2}, \sigma_1) + \Phi^{-1}(T)\sigma_p - (\mu_0 + \frac{\sigma_0^2 \epsilon_e}{\sigma_0^2 + \sigma_e^2})}{\sigma_1} \right) \\ &\quad \times \left(\left[\frac{\partial}{\partial \mu_0} \hat{\theta}(\mu_0 + \frac{\sigma_0^2 \epsilon_e}{\sigma_0^2 + \sigma_e^2}, \sigma_1) \right] - 1 \right) \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{\epsilon_e}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) d\epsilon_e \end{aligned}$$

When μ_0 is very negative, $\hat{\theta}(\mu, \sigma) \rightarrow \bar{\theta}$ and the terms that dominate are:


$$\begin{aligned} \frac{\partial \mathbb{E}[U_0^I]}{\partial \mu_0} &= -\frac{1}{\sigma_0} \phi \left(\frac{\bar{\theta} + \Phi^{-1}(T)\sigma_p - \mu_0}{\sigma_0} \right) \\ \frac{\partial \mathbb{E}[U_1^I]}{\partial \mu_0} &= -\int_{-\infty}^{\infty} \frac{1}{\sigma_1} \phi \left(\frac{\bar{\theta} + \Phi^{-1}(T)\sigma_p - (\mu_0 + \frac{\sigma_0^2 \epsilon_e}{\sigma_0^2 + \sigma_e^2})}{\sigma_1} \right) \\ &\quad \times \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi \left(\frac{\epsilon_e}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) d\epsilon_e \\ &= -\frac{1}{\sigma_2} \phi \left(\frac{\bar{\theta} + \Phi^{-1}(T)\sigma_p - \mu_0}{\sigma_2} \right) \end{aligned}$$

where $\sigma_2 = \sqrt{(\sigma_e^2 + 1)(\sigma_0^2 + \sigma_e^2)} > \sigma_0$, which follows from the rule used above about a normal pdf with uncertain mean. So for very negative μ_0 , $\frac{\partial \mathbb{E}[U_1^I]}{\partial \mu_0} > \frac{\partial \mathbb{E}[U_0^I]}{\partial \mu_0}$ as these terms compare two tails of normal pdfs with the same mean but the LHS has a higher standard deviation. Reasoning analogously, the same inequality holds as $\mu_0 \rightarrow \infty$. So as $\mu_0 \rightarrow -\infty$, Δ_q approaches 0 from below and for $\mu_0 \rightarrow \infty$, Δ_q approaches 0 from above. By the continuity of Δ_q there must be at least one crossing of zero, call $\hat{\mu}_{\text{low}}$ the first crossing and $\hat{\mu}_{\text{high}}$ the last crossing — these may be equal — proving

part i of the proposition. Combining this with Lemma 6 with the continuity of Δ_q gives part ii. ■

Propositions 3 and 4 (Fraud and Election Monitoring)

The sequence of moves in the game with fraud and election monitoring is:



The strategy of the incumbent is now a monitoring level m , the electoral decision, and the level of fraud when holding an election with monitoring level m . The strategy of the citizens is still their cutoff strategy in the no equilibrium game and their cutoff strategy in the election revolution stage with result e . To be consistent with the equilibrium fraud strategy, denote:

$$\mu_f(e_f; x^*) \equiv \frac{\sigma_0^2(e_f + x^*) + \sigma_e^2\mu_0}{\sigma_0^2 + \sigma_e^2}$$

and standard deviation σ_1 as defined above. Since $e_f = e - x^*$, $\mu_f(e_f; x^*) = \mu_1(e)$ as defined in the baseline model. Hence the citizen strategy must be a cutoff threshold equal to that in an analogous game where fraud is not possible. So for a true election result e , expected amount of fraud x^* , and amount of fraud x , the probability of survival is:

$$\begin{aligned} &Pr(\theta > \theta'(e - x + x^*)|e) \\ &= \Phi\left(\frac{\hat{\theta}(\mu_f(e_f - x + x^*; x^*), \sigma_1) + \Phi^{-1}(T)\sigma_p - \mu_1(e)}{\sigma_1}\right) \end{aligned}$$

Giving an expected payoff for holding an election and committing amount of fraud x when the citizens play the strategy that is a best response to amount of fraud x^* :

$$\begin{aligned} &\mathbb{E}[U_1^I(x; x' = x^*)] \\ &= \int_{-\infty}^{\infty} \Phi\left(\frac{\hat{\theta}(\mu_f(e_f - x + x^*; x^*), \sigma_1) + \Phi^{-1}(T)\sigma_p - \mu_1(e)}{\sigma_1}\right) \\ &\quad \times \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \phi\left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}}\right) de - c(x) \end{aligned}$$

So for $x = x^*$ to be an equilibrium, the first order condition is:

$$c'(x^*) = \left[\int_{-\infty}^{\infty} \frac{1}{\sigma_1} \phi \left(\frac{\hat{\theta}(\mu_1(e), \sigma_1) + \Phi^{-1}(T)\sigma_p - \mu_1(e)}{\sigma_1} \right) \frac{1}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right. \\ \left. \times \phi \left(\frac{e - \mu_0}{\sqrt{\sigma_0^2 + \sigma_e^2}} \right) de \right] \left[\frac{\partial \hat{\theta}(\mu_f, \sigma_1)}{\partial \mu_f} \right] \cdot \left[\frac{\partial \mu_f}{\partial x} \right]$$

The RHS is not a function of x^* and positive since the first term is positive and the two partial derivatives are negative. The cost assumptions on c ensure a unique x^* that satisfies this equation. A sufficient condition for this to be a global maximizer is that for all x :

$$c''(x) \geq \frac{\partial^2 \mathbb{E}[U_1^I(x; x' = x^*)]}{\partial^2 x} \tag{19}$$

Which ensures the objective function is globally concave. The RHS is finite, hence there will always exist cost functions sufficiently convex for x^* to be a global maximizer. This proves Lemma 2. Combining this with the fact that $\mu_f(e_f; x^*) = \mu_1(e)$ proves Proposition 3.

Finally, consider the equilibrium monitoring level. As shown above, the probability of survival is the same under each level of monitoring, so the only difference in the payoff down the game tree under the different levels of monitoring is the cost.

Figure 4 demonstrates the result visually by showing the *derivative* of the cost function for a higher level (m_2) of monitoring and a lower level (m_1).

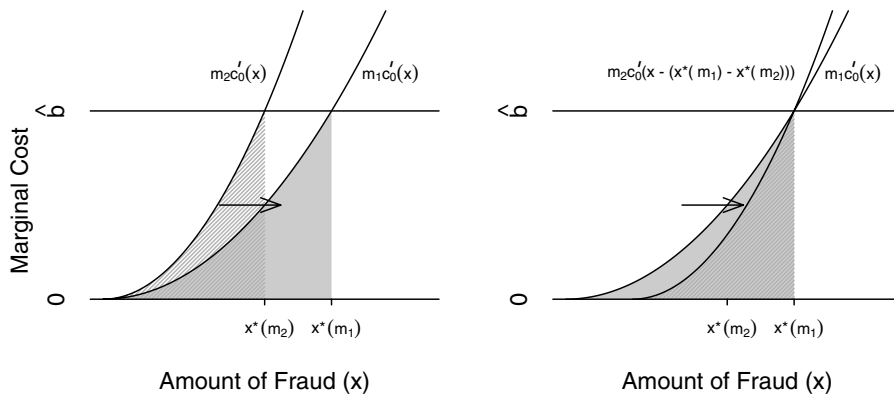


Figure 4. Illustration of the monitoring decision in Proposition 4.

Since $c_0(0) = 0$, the equilibrium cost for each monitoring level is the area under the marginal cost curve up to the equilibrium amount of fraud, which occurs $mc'_0(x) = \hat{b}$. The right panel shows that by translating $m_2c'_0(x)$ by $x^*(m_1) - x^*(m_2)$, the area under this curve is strictly less than the area under $m_1c'_0(x)$. Algebraically:

$$\begin{aligned} c(x^*(m_2), m_2) &= \int_0^{x^*(m_2)} m_2c'_0(x)dx \\ &= \int_{x^*(m_1)-x^*(m_2)}^{x^*(m_1)} m_2c'_0(x - (x^*(m_1) - x^*(m_2)))dx \\ &< \int_0^{x^*(m_1)} m_1c'_0(x)dx = c(x^*(m_1), m_1) \end{aligned}$$

where the inequality comes from the facts that $m_1 < m_2$, c'_0 increasing, and $x^*(m_k) = (c'_0)^{-1}(\hat{b}/m_k)$. As the cost term is strictly decreasing with respect to m , the incumbent selects $m = \bar{m}$, proving part i of Proposition 4. Part ii of Proposition 4 also directly follows from this result: increasing the maximal \bar{m} strictly increases the value of holding an election, making it valuable for a broader range of parameters. ■

References

- Ames, B. 1970. "Bases of Support for Mexico's Dominant Party." *American Political Science Review* 64(1): 153–167.
- Bjornlund, E. C. 2004. *Beyond Free and Fair: Monitoring Elections and Building Democracy*. Woodrow Wilson Center Press.
- Blaydes, L. A. 2011. *Elections and Distributive Politics in Mubarak's Egypt*. Cambridge University Press.
- Boix, C. and M. Svolik. 2008. "The Foundations of Limited Authoritarian Government: Institutions and Power-sharing in Dictatorships." Manuscript.
- Bueno De Mesquita, E. 2010a. "Regime change and Revolutionary Entrepreneurs." *American Political Science Review* 108(3): 446–466.
- Bueno De Mesquita, E. 2010b. "Uniqueness and Multiplicity in Games of Regime Change." Manuscript.
- Carothers, T. 1997. "The Observers Observed." *Journal of Democracy* 8(3).
- Chwe, M. 1999. "Structure and Strategy in Collective Action." *American Journal of Sociology* 105(1): 128–156.
- Cox, G. W. 2009. "Authoritarian Elections and Leadership Succession, 1975–2004." Manuscript.
- Edmond, C. 2008. "Information Manipulation, Coordination and Regime Change." Manuscript.
- Egorov, G. and Konstantin S. 2011. "Incumbency Advantage in Non-Democracies." MPSA 2011 Annual Meeting Paper.
- Fearon, J. 2011. "Self-Enforcing Democracy." *Quarterly Journal of Economics* 126(4): 1661–1708.

- Gandhi, J. and A. Przeworski. 2009. "Holding onto Power by Any Means? The Origin of Competitive Elections." Manuscript.
- Gandhi, J. and E. Lust-Okar. 2009. "Elections Under Authoritarianism." *Annual Review of Political Science* 12: 403–422.
- Geddes, B. 2006. "Why Parties and Elections in Authoritarian Regimes." APSA 2006 Annual Meeting Paper.
- Granovetter, M. 1978. "Threshold Models of Collective Behavior." *The American Journal of Sociology* 83(6): 1420–1443.
- Holmstrom, B. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* 66(1): 169–182.
- Hyde, S. D. 2007. "The Observer Effect in International Politics: Evidence from a Natural Experiment." *World Politics* 60(1): 37–63.
- Hyde, S. D. 2011. "Catch Us if You Can: Election Monitoring and International Norm Creation." *American Journal of Political Science* 55: 201–462.
- Hyde, S. D. and N. Marinov. 2008. "Does Information Facilitate Self-Enforcing Democracy? The role of International Elections Monitoring." Manuscript.
- Kuhn, P. M. 2010. "Fraud, Coordination Problems, and Successful Post-Electoral Protests in Competitive Electoral Autocracies." MPSA 2010 Annual Meeting Paper.
- Kuran, T. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice* 61(1): 41–74.
- Lehoucq, F. 2003. "Electoral Fraud: Causes, Types, and Consequences." *Annual Review of Political Science* 6: 233–256.
- Little, A. T. 2011a. "Fraud and Monitoring in Noncompetitive Elections." APSA 2011 Annual Meeting Paper.
- Little, A. T. 2011b. "An Informational Theory of Elections." Manuscript.
- Londregan, J. and A. Vindigni. 2006. "Voting as a Credible Threat." Manuscript.
- Magaloni, B. 2006. *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*. Cambridge University Press.
- Magaloni, B. 2009. "The Game of Electoral Fraud and the Ousting of Authoritarian Rule." *American Journal of Political Science* 54(3).
- Morris, S. and H. S. Shin. 2003. *Global Games: Theory and Applications*. Number 37 Cambridge University Press.
- Przeworski, A. 2010. "Force and Elections: A General Framework for the Analysis of Political Regimes and their Dynamics." Manuscript.
- Rozenas, A. 2011. "Forcing Consent: Information and Power in Non-Democratic Elections." Manuscript.
- Schandler, A. 2002. "The Menu of Manipulation." *Journal of Democracy* 13(2): 36–50.
- Shadmehr, M. and D. Bernhardt. 2011. "Collective Action with Uncertain Payoffs: Coordination, Public Signals and Punishment Dilemmas." *American Political Science Review* 105: 829–851.
- Simpser, A. 2011. *More than Winning: Why Governments and Parties Manipulate Elections*. Book Manuscript.
- Tucker, J. A. 2007. "Enough! Electoral Fraud, Collective Action Problems, and Post-Communist Colored Revolutions." *Perspectives on Politics* 7(3): 535–551.

