Electoral Manipulation as Bureaucratic Control

Scott Gehlbach University of Wisconsin–Madison
Alberto Simpser University of Chicago

Bureaucratic compliance is often crucial for political survival, yet eliciting that compliance in weakly institutionalized environments requires that political principals convince agents that their hold on power is secure. We provide a formal model to show that electoral manipulation can help to solve this agency problem. By influencing beliefs about a ruler's hold on power, manipulation can encourage a bureaucrat to work on behalf of the ruler when he would not otherwise do so. This result holds under various common technologies of electoral manipulation. Manipulation is more likely when the bureaucrat is dependent on the ruler for his career and when the probability is high that even generally unsupportive citizens would reward bureaucratic effort. The relationship between the ruler's expected popularity and the likelihood of manipulation, in turn, depends on the technology of manipulation.

On September 9, 2001, Alexander Lukashenko won reelection as president of Belarus with 75.7 % of the vote. Although pre-election polls indicated that Lukashenko could have won a free and fair election, various independent observers documented widespread abuse of the electoral process that combined to produce a landslide victory.[1] Why did Lukashenko steal an election that he could have won without fraud, or at least with less fraud? Contemporaneous accounts suggest that Lukashenko resorted to electoral manipulation in an effort to shore up support within the bureaucracy, a key group without whose support it is impossible to govern in Belarus. As explained by Belarusian newspaper editor Alexander Tomkovich, only a large victory could ensure Lukashenko the continued cooperation of state officials: "He probably could have won without fraud or intimidation, albeit with a smaller margin. But politically, he could not afford such a victory.... A landslide first-round victory is a must—one can never know if the bureaucracy would change its mind in between the two rounds."[2]

Mohamed Morsi's fall from power in Egypt in mid-2013 illustrates what might have happened to Lukashenko had he not secured the allegiance of the bureaucracy. After Morsi took power in 2012 with a narrow victory in Egypt's first democratic election for president, various elements of the Egyptian state apparently conspired to withhold essential public services, including power and policing, in a bid to undermine support for the Morsi presidency. As noted in the *New York Times*, "Despite coming to power through the freest elections in Egyptian history, Mr. Morsi was unable to extend his authority over the sprawling state apparatus."[3] The mass protests that preceded Morsi's ouster in July 2013 were motivated in part by dissatisfaction with the performance of the Egyptian bureaucracy, broadly defined.

Scott Gehlbach is Professor and Associate Chair, Department of Political Science, University of Wisconsin–Madison, 110 North Hall, 1050 Bascom Mall, Madison, WI 53706-1316 (gehlbach@polisci.wisc.edu). Alberto Simpser is Assistant Professor, Department of Political Science, University of Chicago, 5828 S. University Avenue, Chicago, IL 60637 (asimpser@gmail.com).

For helpful comments, we thank Bruce Bueno de Mesquita, Ethan Bueno de Mesquita, Andrew Little, Monika Nalepa, Milan Svolik, four anonymous reviewers, and seminar and conference participants at Cornell University, the Higher School of Economics (Moscow), the Juan March Institute, the University of Illinois, the EITM Alumni Conference, the Political Economy in the Chicago Area Conference, and the annual meetings of the American Political Science Association and the Midwest Political Science Association. Work on this article was completed when Gehlbach was a Senior Research Fellow at the International Center for the Study of Institutions and Development at the Higher School of Economics (Moscow).

[1] See, for example, the report of the Office for Democratic Institutions and Human Rights of the Organization for Security and Cooperation in Europe, available at http://www.osce.org/odihr/elections/belarus/14459.

[2] "It's Lonely at the Top in Belarus," *Moscow Times*, September 11, 2001.

[3] "Sudden Improvements in Egypt Suggest a Campaign to Undermine Morsi," *New York Times*, July 10, 2013.

American Journal of Political Science, Vol. 59, No. 1, January 2015, Pp. 212–224

 DOI: 10.1111/ajps.12122

These cases illustrate an essential feature of politics in many countries: Bureaucratic compliance is often crucial for political survival, yet securing that compliance requires that political principals convince agents that their hold on power is secure. Way, for example, observes that in contemporary electoral authoritarian regimes, bureaucrats in the military, regional governments, and the media "have been less likely to follow orders to favor the incumbent if they thought that he or she faces serious challenges," and that such bureaucrats avoid taking actions "that might offend an opponent who might subsequently gain power" (2006, 173).

We argue that electoral manipulation can help rulers solve this problem of bureaucratic compliance. In particular, by influencing bureaucrats' beliefs about the ruler's hold on power, manipulation can encourage bureaucrats to work on behalf of the ruler when they would not otherwise do so. We provide a formal model that illustrates the role that electoral manipulation can play in bureaucratic control and identifies conditions under which an incumbent ruler might choose to exercise this option.

We assume a representative citizen with unknown support for the ruler, whose ultimate survival depends on an action taken by a bureaucrat following an election that has no direct payoff relevance for either the ruler or bureaucrat. The role of electoral manipulation is to influence the bureaucrat's perception of the citizen's support, and thus responsiveness to bureaucratic effort, by interfering with the electoral expression of that support (i.e., by "jamming" the signal provided by voting). Our baseline model assumes that the act of manipulation itself is observed, but that the effect of manipulation on the election outcome is uncertain. Together, these assumptions imply that the bureaucrat may be persuaded by electoral manipulation to take an action that is not in his personal best interest: to cooperate with the ruler even when the citizen may not be responsive to bureaucratic effort.

A key assumption of our model is that the bureaucrat is dependent on the ruler for his continued employment, future promotion, or other career concerns. Politicized bureaucratic employment characterizes many weakly institutionalized environments—see, for example, Reuter and Robertson (2012) on Russia and Martinez-Bravo (2014) on Indonesia—with potential consequences for the behavior of policy makers.[4] In our setting, the rationale for manipulation disappears as dependence on the ruler does.

[4]Vlaicu and Whalley (2013), for example, demonstrate that policy makers who serve at the will of elected officials may behave more like politicians than like bureaucrats, relative to the benchmark established in models such as Alesina and Tabellini (2007, 2008).

Our baseline model produces the following predictions. Electoral manipulation is more likely when the ruler's expected popularity is neither too high nor too low, when electoral manipulation is moderately effective, and when the likelihood is high that even a generally unsupportive citizen would reward bureaucratic effort. We demonstrate these relationships in a simple version of our model in which the representative citizen is assumed to vote sincerely for or against the ruler, thus focusing on the central strategic interaction between the ruler and bureaucrat. We then show that sincere voting emerges naturally in the full version of our model, where a strategic citizen has (unknown) preferences over the production of some good that depends on both the ruler's survival and the bureaucrat's effort.

We subsequently show that our central result—that electoral manipulation functions as a tool of bureaucratic control—holds under various alternative technologies of electoral manipulation. We consider manipulation technologies with arbitrarily high levels of secrecy (in contrast to the baseline model, where manipulation is fully observable), and we permit the ruler to choose from among manipulation technologies with different effectiveness. We also study an extension of our model in which the ruler observes his true level of support before deciding whether to manipulate, as in the phenomenon of "election-night" fraud. Electoral manipulation occurs in equilibrium (for some parameter values) under all these alternative formulations, but in the second and third extensions, manipulation is unambiguously more likely when the ruler is more popular. The relationship between the ruler's expected popularity and the likelihood of manipulation is thus sensitive to the technology of manipulation.

Our analysis makes several contributions to the literature. First, it provides an explanation for the observation that many rulers work much harder to manipulate than one might expect if their goal was only winning the immediate election. Our model can thus be read as a special case of Simpser's (2005, 2013) argument that electoral manipulation can be a tool for strategically influencing the information available to societal actors, and thus is about more than winning. This informational perspective on electoral manipulation has been widely adopted in the literature, including by Little (2012, 2013), who shows that electoral manipulation is inefficient from the incumbent's perspective, as manipulation does not increase the incumbent's survival probability in equilibrium. The inefficiency in Little's work is not present in our model, a consequence of our assumption that the mapping from manipulation to electoral outcome is uncertain, so that the bureaucrat may not be able to back out the citizen's type from (observed) manipulation and

the election outcome. Other work in the formal literature on electoral manipulation includes Svolik (2012), who like us examines the impact of third-party actors on incentives for electoral manipulation, though his focus is quite different, examining "suppliers" of information (i.e., electoral commissions, courts, and observers) rather than "consumers" (i.e., the bureaucracy). Rozenas (2013) shows that "repression," which very closely approximates manipulation in our model, is increasing in the ruler's expected popularity, whereas Egorov and Sonin (2012) demonstrate that electoral fraud may be nonmonotonic in the dictator's expected popularity; we find that the relationship between these two variables depends on the technology of electoral manipulation.

Second, and related, our analysis holds relevance for the study of electoral authoritarianism. Authoritarian regimes that regularly hold elections have become increasingly common in recent decades (Howard and Roessler 2006; Hyde and Marinov 2012; Levitsky and Way 2002, 2010; Schedler 2006), part of a more general phenomenon of democratic-seeming institutions in authoritarian settings (e.g., Boix and Svolik 2013; Gandhi 2008; Gandhi and Przeworski 2007; Gehlbach and Keefer 2012; Lust-Okar 2006). Our analysis shows that electoral manipulation in such regimes can be used to shore up the ruler's or ruling party's power by increasing its control over subordinates.

Third, our article contributes to a vast literature on the political control of the bureaucracy. One branch of this literature addresses optimal bureaucratic control when the principal is unlikely to remain in power (e.g., de Figueiredo 2002; Horn 1995; McCubbins, Noll, and Weingast 1987, 1989; Moe 1990; Rothenberg 1994), but most such work takes perceptions of the principal's hold on power as given.[5] We show that electoral manipulation can be used to influence perceptions of the ruler's hold on power and thus encourage bureaucratic compliance.

A Simple Model

Environment

Consider an extensive game of incomplete information with two strategic players: an incumbent ruler (r) and a bureaucrat (b). In addition, assume a nonstrategic, representative citizen of type $t \in \{p, o\}$, where p indicates a "proponent" and o an "opponent" of the ruler. (We model the citizen as a strategic actor in the full model presented in the following section.) The citizen's type is unobserved by the ruler and bureaucrat, who share a common prior belief that $\Pr(t = p) = q$. The bureaucrat acts as the ruler's agent, choosing whether to exert effort. The ruler's survival probability, in turn, depends on both bureaucratic effort and the citizen's type. Information about the latter may be revealed by an election that the ruler can attempt to manipulate (we assume that the impact of manipulation on the election outcome is uncertain) but that does not directly determine the ruler's survival, as discussed below.

Formally, we assume the following information structure and timing of events:

1. Nature determines the citizen's type $t \in \{p, o\}$, which is observed by neither the ruler nor the bureaucrat.
2. The ruler chooses whether to manipulate the election, $m \in \{0, 1\}$. The bureaucrat observes the manipulation choice m.
3. The citizen (nonstrategically) votes against or for the ruler, $v(t) \in \{0, 1\}$, where $v(o) = 0$ and $v(p) = 1$ (i.e., the opponent opposes and the proponent supports the ruler). The election outcome $\hat{v}(m, v) \in \{0, 1\}$ (which does not directly determine the ruler's survival) is jointly determined by the ruler's and citizen's actions, according to a process described below. The bureaucrat observes the election outcome $\hat{v}$ but not the citizen's action v.
4. The bureaucrat chooses an effort level $e \in \{0, 1\}$.
5. The ruler survives or not, $\rho \in \{0, 1\}$, according to a process described below.

A key assumption of the model, common in the literature (Fearon 2011; Little 2012; Rozenas 2013; Simpser 2013; Svolik 2012), is that the election outcome may provide information that influences subsequent actions but does not directly determine the ruler's survival. One interpretation of this assumption, consistent with the survival process described below, is that the ruler can only be removed by force, but the willingness of the citizen to use force is correlated with her willingness to vote against the incumbent in the election. Alternatively, one can think of the ruler as sufficiently popular to initially survive without the support of the representative citizen, but that with some probability, the representative citizen is pivotal in a subsequent election. A related interpretation is that the election is to a comparatively weak institution (e.g., to a parliament in a super-presidential system), so that the primary function of the election is to transmit information about voter preferences prior to a subsequent election to a comparatively

[5] For an exception, see Egorov and Sonin (2009).

strong institution (Meirowitz and Tucker 2007). For simplicity, we use the term *election* in what follows to refer to the process described in the third step above, as opposed to the subsequent retention or replacement of the ruler.

Another key assumption is that electoral manipulation is observable. This can be motivated on empirical grounds. Some tactics of manipulation, such as bribing or intimidating voters, are by their nature visible to many citizens, including agents of the ruler. Others, such as stuffing ballot boxes or tampering with voter lists, are often utilized blatantly, even though in principle they could be kept secret. Still other tactics of manipulation, such as padding voter lists or tampering with the counting of the votes, can be, and often are, detected and publicized by domestic or international monitors. Further below, we consider the possibility that manipulation is only partially observable.

Nonetheless, even when electoral manipulation is visible, the extent to which such effort might have affected vote totals is often unclear. We capture this consideration by assuming that the mapping from manipulation to election outcome is uncertain, as follows. If $m = 0$, then $\hat{v} = v$, that is, if there is no manipulation, then the election outcome faithfully reflects the citizen's choice (i.e., type). In contrast, if $m = 1$, then $\hat{v} = v$ if $v = 1$, whereas if $v = 0$, $\hat{v} = 1$ with probability h and $\hat{v} = 0$ with probability $1 - h$, where the parameter $h \in (0, 1)$. Thus, h is the probability that manipulation is successful.[6] As we will see, this uncertainty implies that manipulation can be effective in confusing the bureaucrat about the citizen's type, even though the act of manipulation is observed.[7]

The further assumption that the ruler cannot choose the effectiveness of manipulation (i.e., that h is exogenous) reflects potential agency problems among the various actors tasked with carrying out manipulation.[8] Thus, for example, the level of manipulation may be inefficiently large, from the perspective of the ruler, when such agents compete to deliver strong election results (e.g., Myagkov, Ordeshook, and Shakin 2009, 7). That said, many of our key results hold if we assume that the ruler is able to choose a particular h, as we discuss below.

[6] Rozenas (2013) employs a similar formulation in his analysis of electoral "repression."

[7] On Bayesian persuasion, see Kamenica and Gentzkow (2011). It is easily verified that our setting is a special case of Kamenica and Gentzkow's model and that our assumptions satisfy their conditions: The bureaucrat exerts no effort in the absence of additional information, and the bureaucrat's effort choice is constant in a neighborhood around the prior belief of the citizen's type.

[8] Implicitly, we assume that the bureaucrat in our model has responsibilities different from electoral manipulation itself. A natural but nontrivial extension of our model would incorporate agents whose responsibility it is to manipulate elections.

We model the process that determines the ruler's survival, captured by the indicator $\rho \in (0, 1)$, as follows. We assume a simple complementarity between the citizen's type and the bureaucrat's effort choice:

$$\begin{aligned} \Pr(\rho = 1 \mid e, t) &= e \text{ if } t = p, \\ &= \beta e \text{ if } t = o, \end{aligned} \tag{1}$$

where $\beta \in (0, 1)$. Bureaucratic effort is a necessary and sufficient condition for the ruler to remain in power if the citizen is a proponent, whereas if the citizen is an opponent, the ruler may lose power even when the bureaucrat exerts effort. Intuitively, the ruler relies on the bureaucracy to remain in power (i.e., in the last stage of the game, not in the nonbinding election), but the degree to which bureaucratic effort translates into survival depends on support for the ruler, as captured by the citizen's type. We provide a microfoundation for Equation (1) further below when we model the citizen as a strategic actor.

To close the model, we specify preferences for the actors. We assume that the ruler has von Neumann-Morgenstern preferences over lotteries over terminal histories represented by the utility function

$$u_r = \rho\zeta - m\kappa. \tag{2}$$

Thus, the ruler receives an exogenous rent $\zeta > 0$ from survival, and he pays a direct cost κ of manipulation.

The bureaucrat, in turn, has preferences over lotteries over terminal histories represented by

$$u_b = \rho w - e\eta, \tag{3}$$

where the parameters $w, \eta > 0$. A straightforward interpretation of this equation is that the bureaucrat's employment, and thus wage, are tied to the ruler's survival. Other relationships are easily captured by the same preferences, including promotion or other career concerns that depend on the ruler's survival.[9]

To focus on the interesting case, we restrict attention to a particular region of the parameter space. First, we assume $\beta w < \eta < w$, which ensures that the bureaucrat's problem is nondegenerate: There are beliefs about the citizen's type that justify effort and no effort, respectively. Second, we assume $q < \bar{q} \equiv \frac{\eta - \beta w}{w - \beta w}$, which implies that, in equilibrium, the bureaucrat prefers not to exert effort unless the election is sufficiently informative about the citizen's type. The quantity $\bar{q}$ is bounded by 0 and 1, given the assumption $\beta w < \eta < w$.

[9] In a working paper version of this article, we alternatively assume that the ruler can name an outcome-contingent wage that is paid if and only if the ruler survives.

Analysis

Our solution concept is perfect Bayesian equilibrium; we restrict attention to pure-strategy equilibria. We begin with the bureaucrat's effort decision e. Let $\hat{q}$ represent the bureaucrat's posterior belief that the citizen is a proponent, having observed both the ruler's manipulation decision and the election result (i.e., $\hat{q} = \Pr(t = p \mid m, \hat{v})$). Given the survival process in Equation (1) and the bureaucrat's preferences expressed by Equation (3), the bureaucrat's expected payoff from exerting effort is

$$\hat{q}w + (1 - \hat{q})\,\beta w - \eta. \tag{4}$$

With probability $\hat{q}$, the citizen is a proponent, in which case effort secures the ruler's survival and the bureaucrat's "wage" w, whereas with probability $1 - \hat{q}$, the citizen is an opponent, implying that the ruler survives with probability β; in either case, the bureaucrat bears a cost of effort η. In contrast, the payoff from choosing $e = 0$ is simply zero, as the ruler is replaced with certainty. The bureaucrat therefore prefers to exert effort if and only if

$$\hat{q} \geq \tilde{q} = \frac{\eta - \beta w}{w - \beta w}. \tag{5}$$

Now consider the ruler's decision to manipulate. If the ruler does not manipulate, then given the assumption of sincere voting, the bureaucrat updates his belief to $\hat{q} = 1$ after observing $\hat{v} = 1$, thus justifying $e = 1$. In contrast, the bureaucrat believes he is facing an opponent with certainty when $\hat{v} = 0$, thus inducing $e = 0$. The ruler's expected payoff from choosing $m = 0$ is therefore $q\zeta$: He survives and receives the rent from holding office ζ if and only if the citizen is a proponent, which occurs with probability q.

What if the ruler manipulates? As with the case of no manipulation, the election result is $\hat{v} = 0$ if and only if the citizen is an opponent, so that the bureaucrat chooses $e = 0$ upon observing $\hat{v} = 0$. In contrast to the case of no manipulation, however, the bureaucrat's posterior belief upon observing $\hat{v} = 1$ is

$$\hat{q} = \tilde{q} \equiv \frac{q}{q + (1 - q)\,h}, \tag{6}$$

given that manipulation is successful with probability h. If $\tilde{q} \geq \bar{q}$ (see Equation 5), then it is a best response for the bureaucrat to exert effort if and only if he observes $\hat{v} = 1$. In contrast, if $\tilde{q} < \bar{q}$, then the bureaucrat prefers to choose $e = 0$, regardless of the election outcome, thus forcing the ruler from office with certainty. A necessary condition for the ruler to choose $m = 1$ is therefore $\tilde{q} \geq \bar{q}$, as otherwise the ruler could profitably deviate to $m = 0$ and save the cost of manipulation.

Focus on the case where $\tilde{q} \geq \bar{q}$. By manipulating, the ruler induces the bureaucrat to exert effort not only when the citizen is a proponent, but also (with probability h) when she is an opponent. Given the survival process in Equation (1) and the ruler's preferences expressed by Equation (2), the ruler's expected payoff from manipulation is therefore

$$[q + (1 - q)\,h\beta]\,\zeta - \kappa. \tag{7}$$

The ruler survives with certainty if the citizen is a proponent (which occurs with probability q) and with probability $h\beta$ if the citizen is an opponent, and he bears the cost of manipulation κ.

For it to be a best response for the ruler to choose $m = 1$, Expression (7) must be greater than $q\zeta$, the expected payoff from choosing $m = 0$, that is,

$$\kappa \leq (1 - q)\,h\beta\zeta. \tag{8}$$

Together, this condition and $\tilde{q} \geq \bar{q}$ constitute necessary and sufficient conditions for existence of an equilibrium in which the ruler manipulates. An equilibrium with no manipulation exists if and (up to an indifference condition) only if at least one of the conditions does not hold.

Proposition 1. *In the simple model with a nonstrategic citizen, there exists an equilibrium in which the bureaucrat chooses $e = 1$ if and only if $\hat{v} = 1$. If*

$$\frac{q}{q + (1 - q)\,h} > \frac{\eta - \beta w}{w - \beta w}$$

and

$$\kappa < (1 - q)\,h\beta\zeta,$$

the ruler manipulates in this equilibrium. If at least one of these conditions does not hold (weakly), the ruler does not manipulate.

Proposition 1 generates a number of comparative statics that suggest conditions under which electoral manipulation might be used as a means of bureaucratic control. First, given the assumed restriction to a particular region of the parameter space, electoral manipulation is more likely when the bureaucrat's "wage" w is large, as then the bureaucrat is willing to tolerate more uncertainty about the ruler's popularity. In particular, electoral manipulation will be ineffective in inducing bureaucratic effort when the bureaucrat's current employment or future career concerns are unaffected by the ruler's survival. To see this, assume that the bureaucrat receives a payoff z with probability l if the ruler is replaced, whereas he receives z with certainty if the ruler survives. Then the marginal benefit from the ruler's survival is $w \equiv (1 - l)\,z$, and the analysis goes through as before. Clearly, the first condition in Proposition 1 is less likely to hold if l is

large. Thus, we might expect civil service reform to reduce the incentives for electoral manipulation—and also bureaucratic effort, unless some other mechanism for incentivizing effort could be found to replace that analyzed here. This analysis highlights a key condition for our central finding to hold: the dependence of the bureaucrat on his political superior.

Second, manipulation is more likely when the ruler's expected popularity (represented by the parameter q, which measures the probability that the citizen is a proponent) is neither too high nor too low:

$$\frac{h(\eta - \beta w)}{w(1 - h\beta) - \eta(1 - h)} \leq q \leq 1 - \frac{\kappa}{h\beta\zeta}.$$

Intuitively, if ex ante the citizen is likely to be an opponent, then the bureaucrat is unlikely to believe that the ruler actually "won" the election (i.e., that $v = 1$) when the election outcome $\hat{v} = 1$ and so will not exert effort, regardless of the election outcome. In contrast, when the probability is high that the citizen is a proponent, then there is little incentive for the ruler to manipulate, as the bureaucrat is likely to exert effort in any event. The latter logic helps to explain the decision not to manipulate the 2013 Moscow mayoral elections. Assured by pre-election polls that showed him as the strong favorite, the incumbent mayor Sergei Sobyanin evidently felt that he had more to lose than gain from electoral manipulation. Alas, the polls failed to capture the enthusiasm among supporters of opposition activist Alexei Navalny, and Sobyanin won only narrowly, thus returning him to the mayoralty a diminished figure less likely to command the allegiance of the Kremlin and Moscow bureaucracy.[10]

Third, the ruler is more likely to manipulate when manipulation is moderately effective:

$$\frac{\kappa}{(1 - q)\beta\zeta} \leq h \leq \frac{q}{1 - q}\left(\frac{w - \eta}{\eta - \beta w}\right). \qquad (9)$$

On the one hand, manipulation increases the probability that the bureaucrat observes an election outcome supportive of the ruler. On the other hand, manipulation decreases the importance the bureaucrat attaches to this outcome, as by design there is some probability that the voter's signal has been jammed. We consider the implications of making the effectiveness of manipulation h a choice variable later in the article.

Finally, electoral manipulation is more likely when β, which measures the probability that the ruler survives when $e = 1$ and the citizen is an opponent, is large. To interpret this result, it is useful to have a better understanding of the role of the citizen in determining the ruler's survival. The model of the following section provides a framework for analyzing this role.

[10]See Timothy Frye, "Moscow Mayoral Election: The Risks of Using 'Relatively' Free Elections to Gain Legitimacy," *Monkey Cage* (blog, http://bit.ly/1jBou9p), September 9, 2013, and Scott Gehlbach, "Why Do Leaders Manipulate (or Not) Elections? Revisiting the 2013 Moscow Mayoral Elections," *Monkey Cage* (blog, http://wapo.st/1l2jmiR) [*Washington Post*], October 2, 2013.

Extended Model: Strategic Voting and Retention

Environment

We adapt the environment of the previous section to model the citizen as a strategic actor. To capture the idea that the citizen's incentives depend on her type, we now refer to type p as a "pragmatist"; as before, type o is an "opponent." The key assumption is that a pragmatist is more likely than an opponent to retain the ruler if the bureaucrat exerts effort on the ruler's behalf.

Formally, the information structure and timing of events are now as follows:

1. Nature determines the citizen's type $t \in \{p, o\}$, which is observed by neither the ruler nor the bureaucrat.
2. Simultaneously and independently, the ruler chooses whether to manipulate the election, $m \in \{0, 1\}$, and the citizen votes against or for the ruler, $v(t) \in \{0, 1\}$. The election outcome $\hat{v}(m, v) \in \{0, 1\}$ is jointly determined by the ruler's and citizen's actions, according to the process previously described. The bureaucrat observes the manipulation choice m and the election outcome $\hat{v}$ but not the citizen's action v.
3. The bureaucrat chooses an effort level $e \in \{0, 1\}$, which is observed by the citizen.
4. The citizen chooses whether to retain the ruler (e.g., whether to refrain from a "revolution"), $\rho \in \{0, 1\}$, where $\rho = 1$ indicates retention.
5. Output $y \in \{0, 1\}$ is realized, with $\Pr(y = 1) = \rho e$.

Implicitly, we assume that the ruler is unable to contract on output, other than to guarantee the bureaucrat's continued employment or other career concerns in the event of the ruler's survival. Output, in turn, depends not only on the bureaucrat's effort (deterministically, for simplicity) but also on the ruler's survival. Intuitively, one can think of an investment project that is abandoned if the ruler is replaced, either because the ruler's successor has other priorities or because the ruler's specific human capital is essential for project completion (Robinson and Torvik 2005).

We assume that a citizen of type $t \in \{p, o\}$ has preferences represented by

$$u_t = y\psi + (1 - v)\,\gamma_t + (1 - \rho)\,\delta_t. \tag{10}$$

The citizen receives a payoff $\psi > 0$ if output $y = 1$, which as discussed above occurs if and only if the bureaucrat exerts effort and the ruler survives. We normalize $\gamma_p = 0$ and assume that $\gamma_o = \gamma > 0$, where γ is sufficiently large that an opponent strictly prefers to oppose the ruler in the election (i.e., regardless of the consequences for subsequent play of the game). In addition, we assume $\delta_p = \delta$, where $0 < \delta < \psi$, and $\delta_o = \delta + \mu$, where $\mu \in \{0, \bar{\mu}\}$ is a random variable realized after the bureaucrat's effort choice but before the citizen's decision to retain or replace the ruler. We let $\Pr(\mu = 0) = \beta$, and we assume $\delta + \bar{\mu} > \psi$. It immediately follows that a pragmatist prefers to retain the ruler if and only if $e = 1$, whereas an opponent prefers to retain the ruler if and only if $e = 1$ and $\mu = 0$, where the latter event occurs with probability β.

As in the previous section, the assumption that the election does not directly determine the ruler's survival can be interpreted in a variety of ways, including that the ruler can only be removed by force. In this interpretation, the direct payoff to the citizen of voting can be thought of as expressive rather than material. By assumption, voting in support of the ruler is prohibitively costly, in expressive terms, for an opponent but not a pragmatist, as with Kuran's (1991) suggestion that "preference falsification" may be more costly for certain types. Physically replacing the ruler, in contrast, entails material costs and benefits. In particular, the cost of replacing the ruler is prohibitively large for a pragmatist, so long as the bureaucrat has supported the project, but not necessarily for an opponent. In other words, the pragmatist is "pragmatic"—she always prefers to retain the ruler if doing so ensures successful completion of the ruler's project—whereas with probability $1 - \beta$, the opponent discovers himself to be "dogmatic," preferring to replace the ruler at any cost.[11]

Clearly, there are many ways to capture the assumption that the pragmatist's willingness to support the ruler depends on whether the bureaucrat has exerted effort on the ruler's behalf. Our results are robust to any formulation that assumes a complementarity between bureaucratic effort and the citizen's type, so long as effort increases the payoff from retaining the ruler.

All other elements of the game are the same as in the simple model presented previously.

[11] A related interpretation is that $1 - \beta$ is the probability that opponents overcome their collective action problems.

Analysis

Given the assumptions on citizen preferences and the dependence of output on both the bureaucrat's effort and the ruler's survival, it immediately follows that the ruler's survival probability is precisely that given by Equation (1). Further, the bureaucrat's incentives are unchanged from before: It is a best response for the bureaucrat to exert effort if and only if the posterior probability $\hat{q}$ that the citizen is a pragmatist is at least $\bar{q} = \frac{\eta - \beta w}{w - \beta w}$.

We now consider the conditions for existence of pooling and separating equilibria, where by "pooling" and "separating" we refer to the voting decision for the pragmatist and opponent in the election. Recalling that $q < \bar{q}$ by assumption, this immediately implies that in any pooling equilibrium, in which the posterior $\hat{q}$ equals the prior q, the bureaucrat exerts no effort and the ruler is removed with certainty. (This is true even if the ruler manipulates and $v = 0$ for both types. In that case, both $\hat{v} = 0$ and $\hat{v} = 1$ are observed with positive probability, but regardless of the election result, there is no updating of the bureaucrat's beliefs.) Does such an equilibrium exist? Observe that in equilibrium the pragmatist and opponent cannot pool on $v = 1$, as by assumption the opponent strictly prefers to choose $v = 0$. If there is a pooling equilibrium, with or without manipulation, it must therefore be the case that both types pool on $v = 0$.

Consider first the case of an equilibrium without manipulation in which the types pool on $v = 0$. For this to be a best response for the pragmatist, who receives a payoff of δ in equilibrium, off the equilibrium path (i.e., having observed $\hat{v} = 1$) the bureaucrat must believe that the citizen is an opponent with sufficiently high probability; otherwise she would exert effort, providing an incentive for the pragmatist to deviate and receive a payoff of $\psi > \delta$. Any such equilibrium, however, does not satisfy the Intuitive Criterion of Cho and Kreps (1987), as deviation to $v = 1$ is equilibrium dominated for the opponent but not the pragmatist.[12]

Now consider the case of an equilibrium with manipulation in which the types pool on $v = 0$. In this case, both $\hat{v} = 0$ and $\hat{v} = 1$ are observed with positive probability, but in either case there is no updating of beliefs, implying that the bureaucrat exerts no effort and the ruler is removed with certainty. As this is the worst possible outcome for the ruler, the ruler can profitably deviate

[12] We emphasize the Intuitive Criterion due to familiarity with this equilibrium refinement, but these beliefs are also ruled out by the weaker requirement that probability zero be assigned to types for which a message is strictly (rather than equilibrium) dominated (Cho and Kreps 1987, 199–201).

to $m = 0$, thus saving the cost of manipulation κ. This is therefore not an equilibrium.

Thus, the only pooling equilibrium is ruled out by natural restrictions on beliefs off the equilibrium path, including the Intuitive Criterion of Cho and Kreps (1987). We summarize this result with the following proposition.

Proposition 2. *In the extended model with a strategic citizen, there does not exist a pooling equilibrium that satisfies the Intuitive Criterion.*

We now turn to separating equilibria. Given that the opponent strictly prefers to vote against the ruler, any such equilibrium must have $v = 1$ for the pragmatist and $v = 0$ for the opponent. In principle, two such equilibria are possible: one where the ruler does not manipulate ($m = 0$) and one where he does ($m = 1$). In either case, the ruler's and bureaucrat's incentives are identical to those in the simple model with a nonstrategic citizen, as equilibrium separation implies beliefs identical to those with nonstrategic sincere voting. We therefore need only check that the pragmatist behaves optimally, given the ruler's manipulation decision and the bureaucrat's anticipated effort choice.

Consider first the case where the ruler does not manipulate, $m = 0$. In equilibrium, the pragmatist receives ψ, given that the bureaucrat exerts effort following $\hat{v} = 1$ and thus guarantees the ruler's survival. If instead the pragmatist deviated to $v = 0$, the bureaucrat would choose $e = 0$, producing a payoff to the pragmatist of δ. By assumption, $\psi > \delta$, so the pragmatist has no incentive to deviate.

Now consider the case where the ruler manipulates, $m = 1$. As before, if $\tilde{q} = \frac{q}{q+(1-q)h} \geq \bar{q} = \frac{\eta-\beta w}{w-\beta w}$, then it is a best response for the bureaucrat to exert effort if and only if he observes $\hat{v} = 1$. If this condition holds, then the pragmatist is playing a best response by choosing $v = 1$, as doing so induces bureaucratic effort. In contrast, it cannot be an equilibrium for the ruler to manipulate if $\tilde{q} < \bar{q}$, as manipulation is costly and the bureaucrat exerts no effort.

The model with a strategic citizen therefore generates equilibrium behavior identical to that assumed in the simple model with a nonstrategic citizen. In particular, the conditions for a separating equilibrium with manipulation are identical to those identified previously.

Proposition 3. *In the extended model with a strategic citizen, there exists a separating equilibrium in which a pragmatist chooses $v = 1$, an opponent chooses $v = 0$, and the bureaucrat chooses $e = 1$ if and only if $\hat{v} = 1$. The conditions for manipulation in this equilibrium are identical to those in Proposition 1.*

The predictions of the model with a strategic citizen are thus equivalent to those in the simple model with a nonstrategic citizen. The setting here helps to illustrate one result in particular: that electoral manipulation is more likely when the parameter β is large. As modeled, the parameter β measures the probability that the opponent behaves like a pragmatist when deciding whether to retain the ruler. The larger this probability—the more likely the opponent is to reward effort by retaining the ruler—the greater the incentive for the ruler to manipulate to induce bureaucratic effort.

The explicit modeling of citizen preferences also facilitates discussion of the welfare effects of electoral manipulation. Clearly, a pragmatist is no better off in an equilibrium with manipulation (i.e., when the conditions of Proposition 1 are satisfied), as in any event she supports the incumbent in the election, thus inducing bureaucratic effort. In expectation, however, an opponent is strictly better off, as absent manipulation the bureaucrat would not exert effort when the opponent votes $v = 0$. As shown by Equation (10), an opponent rewards such effort when the random variable $\mu = \bar{\mu}$, thus producing the payoff ψ, which is greater than the payoff δ when the ruler does not manipulate.

Manipulation also improves the ruler's utility; otherwise, he would not opt to manipulate. But the bureaucrat is strictly worse off with manipulation than without. To see this, observe that the bureaucrat's expected payoff in an equilibrium without manipulation is

$$q\,(w - \eta) + (1 - q) \cdot 0.$$

With probability q, the citizen is a pragmatist who votes in support of the incumbent, inducing the bureaucrat to exert effort and guarantee the ruler's survival, whereas with probability $1 - q$, the citizen is an opponent, in which case the bureaucrat exerts no effort and the ruler is replaced. In contrast, the bureaucrat's expected payoff when the ruler manipulates is

$$q\,(w - \eta) + (1 - q)\,(h\beta w - \eta)\,.$$

With manipulation, the bureaucrat exerts effort not only when the citizen is a pragmatist, but also (with probability h) when the citizen is an opponent; with probability β, this results in the ruler's retention. The difference between the second expression and the first is $(1 - q)\,(h\beta w - \eta)$, which is less than zero, given the assumption $\beta w < \eta$. Intuitively, when the citizen is an opponent, manipulation encourages the bureaucrat to take an action that is not in his individual best interest.

We summarize these observations with the following proposition.

Proposition 4. *In the extended model with a strategic citizen, electoral manipulation results in a transfer of utility in expectation from the bureaucrat to the ruler and the citizen.*

Notwithstanding this proposition, we would not claim based on our model that manipulation is normatively desirable, even were we to discount its impact on the bureaucracy's welfare. A more complete analysis would take account of other factors. For example, bureaucratic effort could take the form of diversion of public money to provide particularistic goods. In such instances, even if electoral manipulation—via bureaucratic effort—increased the welfare of the representative citizen (as modeled, a kind of swing voter), this could come at the expense of other, unmodeled citizens. Further, it is important to stress that manipulation in our model is costly to the ruler, implying a possible efficiency loss relative to the case where the ruler is able to induce bureaucratic compliance through other mechanisms.

Alternative Technologies of Electoral Manipulation

In what follows, we show that our central result—that electoral manipulation can function as an instrument of bureaucratic control—holds under various alternative technologies of manipulation, though differences in technology imply some differences in specific predictions.

Partial Observability of Manipulation

In some settings, there may be ambiguity about whether manipulation was attempted. It may be difficult to tell, for example, whether campaign expenditure limits were exceeded, the media were bribed, or voter rolls were padded with fictitious voters. The extent to which this is the case may depend on the presence of election monitors or other independent observers. Here we explore the implications of assuming that manipulation is only partially observable.

Consider the following modification of the model in the previous section. Following the election, the bureaucrat observes the election outcome $\hat{v}$, determined as before, and a signal $\hat{m} \in \{0, 1\}$ of whether manipulation has taken place. The signal-generating process is as follows: If $m = 0$ (i.e., the ruler has not manipulated), then $\hat{m} = 0$ with certainty. In contrast, if $m = 1$, then $\hat{m} = 1$ with probability s and $\hat{m} = 0$ with probability $1 - s$. Our baseline model is a special case, with $s = 1$.

As before, we look for pooling and separating equilibria of the model, beginning with the former. Consider first a pooling equilibrium without manipulation. In principle, there are four possible information sets, $(\hat{m}, \hat{v}) \in \{0, 1\} \times \{0, 1\}$, at which the bureaucrat could find himself. In practice, with pooling and no manipulation, only $(\hat{m}, \hat{v}) = (0, 0)$ is on the equilibrium path, so the bureaucrat's posterior belief is equal to his prior, and there is no effort. But then, as before, natural restrictions on beliefs off the equilibrium path imply that the pragmatist would prefer to deviate to $v = 1$.

Consider next a pooling equilibrium with manipulation. All four information sets $(\hat{m}, \hat{v})$ are on the equilibrium path—the bureaucrat may or may not receive a signal that the ruler has manipulated, and the election outcome may or may not reflect support for the ruler—but at each information set, the bureaucrat does not update his prior belief about the citizen's type and therefore does not exert effort. Given this, the ruler has an incentive to deviate and save the cost of manipulation.

We now turn to separating equilibria, focusing first on the case of no manipulation. In the baseline model, a sufficient condition for the ruler not to deviate in this equilibrium is $\tilde{q} < \bar{q}$, which says that the bureaucrat would not exert effort upon observing that the ruler has manipulated and the election outcome $\hat{v} = 1$. With partial observability, this condition is no longer sufficient. To see this, observe that if the ruler deviates to $m = 1$, then with probability $1 - s$, the bureaucrat finds himself at either $(\hat{m}, \hat{v}) = (0, 0)$ or $(\hat{m}, \hat{v}) = (0, 1)$. As these information sets are both on the equilibrium path, the bureaucrat infers that the ruler has not manipulated and exerts effort if and only if $\hat{v} = 1$. By manipulating, the ruler therefore receives an expected payoff of

$$(1 - s)\,[q + (1 - q)h\beta]\,\zeta - \kappa,$$

where the expression in brackets is the probability that the ruler survives, conditional on manipulation having been unobserved. This is less than the equilibrium payoff $q\zeta$ if

$$\kappa \geq [(1 - s)(1 - q)h\beta - sq]\,\zeta, \tag{11}$$

which may or may not hold, given assumed restrictions on parameters. Thus, a separating equilibrium with no manipulation exists for a smaller region of the parameter space with partial observability than with perfect observability.[13]

[13] Observe that $\kappa \geq (1 - q)h\beta\zeta$ continues to be a sufficient condition for the ruler not to deviate in this equilibrium, as this rules out deviation not only when $\tilde{q} < \bar{q}$ (see Condition 11), but also when $\tilde{q} \geq \bar{q}$. To see the latter point, observe that if $\tilde{q} \geq \bar{q}$, then the bureaucrat exerts effort when the ruler deviates to $m = 1$, regardless of the signal $\hat{m}$. The condition for this not to be profitable is therefore identical to that in the baseline model.

Finally, consider a separating equilibrium with manipulation. As before, a necessary condition for this to be an equilibrium is $\tilde{q} \geq \bar{q}$. Assume this to be true. With partial observability, the bureaucrat infers that the ruler has manipulated, regardless of which signal $\hat{m}$ he receives. By deviating to $m = 0$, the ruler therefore reduces the probability that $\hat{v} = 1$ but does not change the bureaucrat's belief about the citizen's type, conditional on observing that outcome. In particular, given $\tilde{q} \geq \bar{q}$, the bureaucrat chooses $e = 1$ following $\hat{v} = 1$, which now occurs with probability q. The condition for the ruler's deviation not to be profitable is therefore identical to that in the baseline model: $\kappa \leq (1 - q)h\beta\zeta$.

We summarize these observations in the following proposition.

Proposition 5. *In the model with partial observability, there does not exist a pooling equilibrium that satisfies the Intuitive Criterion. A separating equilibrium without manipulation exists for a smaller region of the parameter space than in the baseline model, whereas the conditions for existence of a separating equilibrium with manipulation are identical to those in the baseline model.*

Importantly, an increase in the probability that manipulation is observed (e.g., as might result, from the actions of election monitors) has no impact on the existence of an equilibrium with manipulation.

Choice of Manipulation Effectiveness

Some technologies of electoral manipulation are more reliable than others. Removing a challenger from the race, for example, may do more to ensure a victory than relying on local poll workers to stuff ballot boxes. When the ruler has a menu of manipulation options available, he may be able to choose the effectiveness of manipulation.

To explore the implications of this argument, we may modify the baseline model to assume that the ruler chooses the effectiveness of manipulation h. In this case, the ruler would choose h to be as large as possible and still satisfy Condition (9), so long as there is some h that does so; that is, he would choose $h = \frac{q}{1-q}(\frac{w-\eta}{\eta-\beta w})$. (Observe that the ruler could not equivalently mix over $m = 0$ and $m = 1$, as by assumption the bureaucrat observes whether the ruler has manipulated before choosing effort.) It is then a best response for the ruler to manipulate if Condition (8), evaluated at this h, holds, that is, if

$$\kappa \leq q\left(\frac{w-\eta}{\eta-\beta w}\right)\beta\zeta.$$

Relative to our baseline model, there is one key difference in comparative statics: Electoral manipulation is now unambiguously more likely when the ruler is more popular. Intuitively, when the ruler has a choice of manipulation technologies, he can compensate for low expected popularity by reducing the effectiveness of manipulation, thus ensuring that the bureaucrat's posterior belief upon observing $\hat{v}$ is high enough to justify effort.

"Election-Night" Fraud

Our baseline model implicitly assumes that electoral manipulation takes place primarily before votes are counted, as the ruler must choose whether to manipulate before he observes the vote v. The literature, however, documents cases of "election-night" fraud, where incumbents use privileged information about the count to change the official vote total before it is released to the general public (e.g., Christensen and Colvin 2005). How would our results change if the ruler had the option to change the election outcome $\hat{v}$ after observing v?

To answer this question, we change the information structure and timing of events such that the ruler observes v, following which he chooses an election outcome $\hat{v} \in \{0, 1\}$, where the ruler bears a cost κ from reporting $\hat{v} = 1$ when $v = 0$. We assume that the bureaucrat observes $\hat{v}$ but not v. Further, for simplicity, we revert to the assumption in the simple model initially presented that the citizen (again, proponent or opponent) votes sincerely, with the ruler's survival probability given by Equation (1). The model thus takes the form of a simple signaling game, in which a sender (ruler) of unknown "type" (i.e., information—here, whether the citizen is a proponent or an opponent) sends one of two messages $\hat{v} \in \{0, 1\}$, following which the receiver (bureaucrat) chooses an action $e \in \{0, 1\}$. To further simplify the analysis, we assume that $\kappa < \beta\zeta$, which implies that the ruler would always choose to misreport the election result when $v = 0$ if he were certain that the bureaucrat would subsequently exert effort.

We look for an informative equilibrium, in which the ruler's action conveys information about the citizen's vote and therefore type. It is natural to assume that such an equilibrium takes the following (semi-separating) form: The ruler reports $\hat{v} = 1$ with certainty when $v = 1$, whereas he reports $\hat{v} = 1$ with probability $\theta \in (0, 1)$ when $v = 0$. In this equilibrium, the bureaucrat infers that the citizen is an opponent (and therefore exerts no effort) when $\hat{v} = 0$, whereas he believes the citizen to be a proponent with probability $\frac{q}{q+(1-q)\theta}$ upon observing $\hat{v} = 1$.

For this strategy to be a best response for the ruler, the ruler's expected payoff from choosing $\hat{v} = 1$ and from choosing $\hat{v} = 0$ must be identical when $v = 0$:

$$\alpha\beta\zeta - \kappa = 0,$$

where α is the (endogenous) probability that the bureaucrat chooses $e = 1$ when he observes $\hat{v} = 1$. The ruler survives and receives the rent ζ if and only if the bureaucrat exerts effort and the citizen—an opponent—subsequently rewards it, where the latter event occurs with probability β. Solving for the equilibrium α gives

$$\alpha^* = \frac{\kappa}{\beta\zeta},$$

which is bounded by 0 and 1, given the assumption $0 < \kappa < \beta\zeta$.

Similarly, for the bureaucrat to be willing to mix between $e = 0$ and $e = 1$ when $\hat{v} = 1$, he must be indifferent between the two actions, which is equivalent to saying that the posterior probability that the citizen is a proponent must equal $\bar{q}$, which we previously defined as the posterior belief that justifies effort:

$$\frac{q}{q + (1-q)\theta} = \frac{\eta - \beta w}{w - \beta w}.$$

Solving for the equilibrium θ gives

$$\theta^* = \frac{q(w - \eta)}{(1-q)(\eta - \beta w)},$$

which is bounded by 0 and 1, given the assumptions $\beta w < \eta < w$ and $q < \bar{q}$.

The only remaining step is to show that the ruler is optimizing by choosing $\hat{v} = 1$ with certainty when $v = 1$. In this case, the ruler's expected payoff is $\alpha^*\zeta > 0$, where we recall that there is no cost to reporting $\hat{v} = 1$ when in fact $v = 1$.

Proposition 6. *In the model with "election-night" fraud, there exists an equilibrium in which the ruler reports $\hat{v} = 1$ with certainty when $v = 1$ and reports $\hat{v} = 1$ with probability*

$$\frac{q(w - \eta)}{(1-q)(\eta - \beta w)}$$

when $v = 0$, and the bureaucrat chooses $e = 1$ with probability $\frac{\kappa}{\beta\zeta}$ when $\hat{v} = 1$ and chooses $e = 0$ with certainty when $\hat{v} = 0$.

Qualitatively, the comparative statics here are similar to those in the baseline model, though there are some interesting differences. First, as before, electoral manipulation is greater when the bureaucrat's "wage" w is large. Second, electoral manipulation is strictly increasing in the ruler's expected popularity q, in contrast to the baseline model, where the relationship is nonmonotonic. The difference between the two predictions relates to the fact that here the ruler knows whether he is popular (i.e., he knows the citizen's type) before deciding whether to manipulate. Third, there is no analogous prediction with respect to the effectiveness of manipulation (h in the baseline model), as the nature of fraud in this model is that the ruler simply reports a number. Finally, as in the baseline model, electoral manipulation is increasing in β, which measures the probability that the ruler survives when $e = 1$ and the citizen is an opponent.

Conclusion

We have shown that electoral manipulation can be used to elicit bureaucratic compliance, and we have demonstrated that this result is robust to alternative assumptions about the technology of manipulation that mirror common practices.

Our analysis suggests that electoral manipulation can be viewed as a novel solution to the problem of "coalitional drift" (Horn and Shepsle 1989): the possibility that a political principal may fall from power, thus jeopardizing the implementation of a task that has been delegated to an agent. In our model, a bureaucrat has an incentive to exert effort on a ruler's behalf, to the extent that a representative citizen is likely to respond to such effort by returning the ruler to power. Manipulation in this setting provides the ruler with an opportunity to persuade the bureaucrat that the citizen is of a type most likely to respond to bureaucratic effort.

We note that, while this formalization emphasizes the impact of electoral manipulation on beliefs about the ruler's popularity, in many settings, a ruler's hold on power depends also on factors such as discretion to utilize state resources for partisan ends, capacity to break the law with impunity, and willingness to abuse power. Electoral manipulation can potentially signal the ruler's hold on power through these alternative mechanisms (Simpser 2013).

Looking beyond the electoral environment on which we focus, the idea of costly actions taken to convince agents that the principal's hold on power is secure may generalize to other settings. Employees may be more likely to comply with specific requests when they believe their manager's job to be safe, and executives may respond to directives with greater alacrity when they believe that major shareholders will be around for the long run. Future work should seek to identify the analogues to electoral manipulation in these and related settings.

References

Alesina, Alberto, and Guido Tabellini. 2007. "Bureaucrats or Politicians? Part I: A Single Policy Task." *American Economic Review* 97(1): 169–79.

Alesina, Alberto, and Guido Tabellini. 2008. "Bureaucrats or Politicians? Part II: Multiple Policy Tasks." *Journal of Public Economics* 92(3–4): 426–47.

Boix, Carles, and Milan W. Svolik. 2013. "The Foundations of Limited Authoritarian Government: Institutions and Power-Sharing in Dictatorships." *Journal of Politics* 75(2): 300–16.

Cho, In-Koo, and David Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102(2): 179–221.

Christensen, Ray, and Kyle Colvin. 2005. "Stealing Elections on Election Night: A Comparison of Statistical Evidence from Japan, Canada and the United States." Paper presented at the annual meeting of the American Political Science Association.

de Figueiredo, Rui J. P., Jr. 2002. "Electoral Competition, Political Uncertainty, and Policy Insulation." *American Political Science Review* 96(2): 321–33.

Egorov, Georgy, and Konstantin Sonin. 2009. "Dictators and Their Viziers: Endogenizing the Loyalty-Competence Tradeoff." *Journal of the European Economic Association* 9(5): 903–30.

Egorov, Georgy, and Konstantin Sonin. 2012. "Incumbency Advantage in Non-Democracies." Northwestern University and New Economic School.

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

Gandhi, Jennifer. 2008. *Political Institutions under Dictatorship*. New York: Cambridge University Press.

Gandhi, Jennifer, and Adam Przeworski. 2007. "Authoritarian Institutions and the Survival of Autocrats." *Comparative Political Studies* 40(11): 1279–1301.

Gehlbach, Scott, and Philip Keefer. 2012. "Private Investment and the Institutionalization of Collective Action in Autocracies: Ruling Parties and Legislatures." *Journal of Politics* 74(2): 621–35.

Horn, Murray. 1995. *The Political Economy of Public Administration: Institutional Choice in the Public Sector*. New York: Cambridge University Press.

Horn, Murray J., and Kenneth A. Shepsle. 1989. "Commentary on 'Administrative Arrangements and the Political Control of Agencies': Administrative Process and Organizational Form as Legislative Responses to Agency Costs." *Virginia Law Review* 75(2): 499–508.

Howard, Marc Morje, and Philip G. Roessler. 2006. "Liberalizing Electoral Outcomes in Competitive Authoritarian Regimes." *American Journal of Political Science* 50(2): 365–81.

Hyde, Susan D., and Nikolay Marinov. 2012. "Which Elections Can Be Lost?" *Political Analysis* 20(2): 191–210.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101(6):2590–2615.

Kuran, Timur. 1991. "Now out of Never: The Element of Surprise in the East European Revolution of 1989." *World Politics* 44(1): 7–48.

Levitsky, Steven, and Lucan A. Way. 2002. "The Rise of Competitive Authoritarianism." *Journal of Democracy* 13(2): 51–65.

Levitsky, Steven, and Lucan A. Way. 2010. *Competitive Authoritarianism: Hybrid Regimes after the Cold War*. New York: Cambridge University Press.

Little, Andrew T. 2012. "Elections, Fraud, and Election Monitoring in the Shadow of Revolution." *Quarterly Journal of Political Science* 7(4): 249–83.

Little, Andrew T. 2013. "Fraud and Monitoring in Noncompetitive Elections." Cornell University.

Lust-Okar, Ellen. 2006. "Elections under Authoritarianism: Preliminary Lessons from Jordan." *Democratization* 13(3): 456–71.

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

McCubbins, Matthew D., Roger G. Noll, and Barry R. Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, and Organization* 3(2): 243–77.

McCubbins, Matthew D., Roger G. Noll, and Barry R. Weingast. 1989. "Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies." *Virginia Law Review* 75(2): 431–82.

Meirowitz, Adam, and Joshua A. Tucker. 2007. "Run Boris Run: Strategic Voting in Sequential Elections." *Journal of Politics* 69(1): 88–99.

Moe, Terry M. 1990. "The Politics of Structural Choice: Toward a Theory of Public Bureaucracy." In *Organization Theory: From Chester Barnard to the Present and Beyond*, ed. Oliver E. Williamson . Oxford: Oxford University Press, 116–153.

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

Reuter, Ora John, and Graeme Robertson. 2012. "Sub-National Appointments in Authoritarian Regimes: Evidence from Russian Gubernatorial Appointments." *Journal of Politics* 74(4): 1023–37.

Robinson, James A., and Ragnar Torvik. 2005. "White Elephants." *Journal of Public Economics* 89(2–3): 197–210.

Rothenberg, Lawrence S. 1994. *Regulation, Organization, and Politics: Motor Freight Policy at the Interstate Commerce Commission*. Ann Arbor: University of Michigan Press.

Rozenas, Arturas. 2013. "Forcing Consent: Information and Power in Non-Democratic Elections." New York University.

Schedler, Andreas Paul, ed. 2006. *Electoral Authoritarianism*. Boulder, CO: Lynne Rienner.

Simpser, Alberto. 2005. *Making Votes Not Count: Strategic Incentives for Electoral Corruption*. Ph.D. dissertation, Stanford University.

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

Svolik, Milan. 2012. "Third-Party Actors and the Success of Democracy: How Electoral Commissions, Courts, and Observers Shape Incentives for Election Manipulation and

Post-Election Protest." University of Illinois at Urbana-Champaign.

Vlaicu, Razvan, and Alex Whalley. 2013. "Hierarchical Accountability in Government: Theory and Evidence." University of Maryland and University of California at Merced.

Way, Lucan. 2006. "Authoritarian Failure: How Does State Weakness Strengthen Electoral Competition?" In *Electoral Authoritarianism: The Dynamics of Unfree Competition*, ed. Andreas Schedler. Boulder, CO: Lynne Rienner, 167–180.