

Media Freedom in the Shadow of a Coup

Raphael Boleslavsky* Mehdi Shadmehr[†] Konstantin Sonin[‡]

Abstract

Popular protests and palace coups are the two domestic threats to dictators. We show that free media, which informs citizens about their rulers, is a double-edged sword that alleviates one threat, but exacerbates the other. Informed citizens may protest against a ruler, but they may also protest to restore him after a palace coup. In choosing media freedom, the leader trades off these conflicting effects. We develop a model in which citizens engage in a regime change global game, and media freedom is a ruler's instrument for Bayesian persuasion, used to manage the competing risks of coups and protests. A coup switches the status quo from being in the ruler's favor to being against him. This introduces convexities in the ruler's Bayesian persuasion problem, causing him to benefit from an informed citizenry. Rulers tolerate freer press when citizens are pessimistic about them, or coups signal information about them to citizens.

Keywords: authoritarian politics, media freedom, protest, coup, global games, Bayesian persuasion, signaling

JEL Classification: H00, D82

*Department of Economics, University of Miami, 5250 University Dr., Coral Gables, FL 33146. E-mail: r.boleslavsky@miami.edu

[†]Harris School of Public Policy, University of Chicago, and Department of Economics, University of Calgary, 1155 East 60th Street, Chicago 60637. E-mail: mshadmeh@gmail.com

[‡]Harris School of Public Policy, University of Chicago, 1155 East 60th Street, Chicago 60637. E-mail: ksonin@uchicago.edu

1 Introduction

Modern autocratic regimes consider media freedom a major threat to their stability. Numerous studies have analyzed the use of censorship or sophisticated manipulation of the media (King, Pan, and Roberts 2013, 2014; Lorentzen 2014; Chen and Yang 2018). For regimes that routinely jail or murder their opponents, would it not be natural to prevent the media from disseminating all negative news about the regime? If not, is it true that media freedom is an anathema to dictators, who tolerate it only because of the efficiency costs of censorship (Egorov, Guriev, and Sonin 2009; Lorentzen 2013)?

In this paper, we argue that partial media freedom serves an important role for dictators by protecting them against palace coups—removal at the hands of a leader’s current supporters. In the past seven decades, elite coups have been the most frequent reason for a non-democratic leader’s departure (Svolik 2009, 2012). Dorsch and Maarek (2018) calculated that 75% of irregular power transitions in autocracies, barring foreign interventions, are results of coup d’états. With a partially free media, the dictator might lose power as a result of protests that are more likely than under total censorship. However, the increased likelihood of counter-protest in favor of the deposed leader offers a layer of protection from a coup by the elite. Thus, the leader faces a trade off with respect to the freedom of the media.

In 1964, Nikita Khrushchev, the powerful leader of the Soviet Union, was displaced in a palace coup led by Leonid Brezhnev, number two in the party hierarchy. The Soviet people learned about the coup after the power transition was completed, Khrushchev allies were sidelined, and the palace guards replaced (Taubman 2003). By contrast, the military coup that removed Mikhail Gorbachev, the last Soviet leader, in 1991 failed because partial media freedom allowed the supporters of Gorbachev to start mass protests the same evening (Taubman 2017). In Venezuela in 2002, President Chavez, forced out by the military, succeeded to reverse the course.¹ Other examples of such “coup reversals,” in which the mass protests helped the incumbents regain full control, include France in April 1961, Spain in April 1981, Russia in October 1993, and Turkey in July 2016.

Two observations underlie our argument for why media freedom provides a layer of protection for rulers against palace coups. First, the incumbent leader wants a popular

¹Ten years before that, Chavez was on a different side of the story: a military unit led by him removed President Peres from power. However, Peres was able to address the country via local television and reverse the coup.

protest to follow a coup, but does not want a popular protest if there is no coup. Thus, when choosing the level of media freedom, the leader must take account of its effects on the likelihood of both an initial protest against him and a counter-protest in the event of a coup by the elites. Second, citizens' acquiescence has inertia: protesting is costly. Thus, there is an inherent asymmetry between information that favors inaction and information that favors action. Once information favoring inaction is sufficiently strong, strengthening it further has little influence on citizens' behavior. By contrast, stronger information favoring action always has a substantial effect: it keeps compelling some citizens to act, thereby raising the likelihood of a successful protest. These two countervailing forces generate the contrast in the leader's choice of media freedom in the presence and absence of a coup threat.

If the leader's subordinates do not have any informational advantage over ordinary citizens (i.e., they have the same prior beliefs about the incumbent's type), then the officer's coup attempt has no effect on citizens' beliefs. Thus, citizens' decision to protest depends on their costs and prior beliefs about the ruler. If citizens are pessimistic about the incumbent, then a counter-protest is less likely to materialize and a coup is more likely to succeed. By contrast, if citizens are optimistic about the incumbent, then they are likely to come to his aid if a coup is attempted. Thus, if the initial belief about the ruler is pessimistic and he faces an imminent threat from the officer, then he may wish to allow some degree of media freedom: positive information from a free press has the potential to improve citizens' beliefs about the ruler, quashing the officer's ambitions.

By contrast, when the officer is informed about the incumbent's type, citizens make inferences about the quality of the ruler from the very fact that he is dismissed in a coup. In this case, coup attempts occur with a positive probability even when citizens are optimistic about the incumbent's type. The difference is a result of the fact that the elites have a stronger incentive to attempt a coup against a bad ruler than a good one. Thus, a coup attempt conveys bad information about the ruler and makes citizens less inclined to counter-protest. As a result, in equilibrium, the officer always mounts a coup against a bad ruler, but also "bluffs" by sometimes deposing a good ruler. To mitigate the likelihood of a coup, the ruler must allow a greater degree of media freedom, even when citizens are optimistic about his type. Indeed, the fact citizens infer the ruler's type from his subordinates' actions against him makes the incumbent more vulnerable.

Our paper combines insights from regime change global games (Morris and Shin 1998,

2003) and Bayesian persuasion (Kamenica and Gentzkow 2011). A growing literature takes a global game approach to analyze coordination problems in regime change settings (Angeletos, Hellwig, and Pavan 2007; Bueno de Mesquita 2010; Boix and Svolik 2013; Edmond 2013; Casper and Tyson 2014; Egorov and Sonin 2014; Chen, Lu, and Suen 2016; Rundlett and Svolik 2016; Barbera and Jackson 2017; Smith and Tyson 2018). Although coordination considerations tend to generate multiple equilibria, the global game approach introduces small correlated asymmetric information that selects the risk-dominant equilibrium of the complete information game (Carlsson and van Damme 1993; Morris and Shin 1998, 2003). Applications in revolution settings often place uncertainty on the regime’s strength—the minimum fraction of citizens required to overturn the status quo. We follow Persson and Tabellini (2009) by placing uncertainty on the costs of protest, which allows us to sufficiently disentangle the coordination and information design aspects of the model to provide a transparent characterization of media freedom. However, unlike Persson and Tabellini (2009), we require protest costs to be positive.

Bayesian persuasion literature studies how an information designer, who can ex-ante commit to an information disclosure policy, should optimally choose the policy to induce his desired behavior in a single Bayesian player (Kamenica and Gentzkow 2011). This framework has been applied to study elections and media freedom in autocracies (Gehlbach and Sonin 2014; Gehlbach and Simpson 2015; Gentzkow, Shapiro, and Stone 2016; Luo and Rozenas 2018).² Although we restrict our attention to a truth-or-noise signal structure, our key insight, which remains valid for any signal design (see Section 4), is that the threat of coup induces convexities in the leader’s expected payoff as a function of citizens’ prior beliefs about his type, thereby inducing him to design informative signals. A recent literature investigates information design in multi-player models under a variety of assumptions (Bergemann and Morris 2016, 2017). In particular, Goldstein and Huang (2016) and Inostroza and Pavan (2018) study information design in regime change global games settings. In these papers, the information designer chooses the information policy about the same state that players will receive private information in the subsequent global game, and focus on the worst equilibrium for the designer (adversarial equilibrium). In contrast, in our paper the information designer (ruler) chooses the information policy about a state (ruler’s type) that is uncorrelated with the state about which players subsequently receive private information (costs). As a result,

²See Prat (2015) for a survey of models of the media.

the global games approach suffices to select a unique equilibrium in our setting.³

Our paper contributes to the formal work on intra-elite power struggles in autocracies (Bueno de Mesquita, Morrow, Siverson, and Smith 2003; Acemoglu and Robinson 2006; Gandhi and Przeworski 2006, 2007; Besley and Kudamatsu 2008; Guriev and Treisman 2015). Recent contributions have focused on the civilian control of the military (Svolik 2008, 2012; Acemoglu, Ticchi, and Vindigni 2010) and information content of protests for conspiring officers (Egorov and Sonin 2011; Casper and Tyson 2014). By adapting Padro i Miquel’s (2007) “politics of fear” to elite infighting, Hollyer, Rosendorf, and Vreeland (2018) consider a complete information Stackelberg game played by a ruler and an elite challenger. By assumption, higher “transparency” increases the probability that the regime collapses, but with a larger marginal effect when the challenger overthrows the leader. The ruler allows transparency to destabilize the regime: because a destabilized regime is less likely to survive infighting, transparency may deter an overthrow attempt by the elite challenger. By contrast, we present a Bayesian persuasion model in which the leader’s choice of media freedom determines the probability with which citizens will learn his type, ahead of a coordination game in which they must decide whether to overthrow the ruler themselves or restore him to power following a coup. Thus, in our model, an unpopular leader allows media freedom in the hope of improving citizens’ beliefs about his type, which can simultaneously improve stability in the absence of a coup and mobilize citizens if a coup occurs.

Finally, there are important reasons why dictators might allow partial media freedom even in the absence of a coup threat. In Egorov, Guriev, and Sonin (2009), a resource-poor dictator allows media freedom as he is concerned with providing his bureaucrats with proper incentives. Lorentzen (2013, 2014) considers similar efficiency-protests trade off in models of strategic protest restrictions and censorship. In Shadmehr and Bernhardt (2015), the state does not censor modestly bad news to prevent citizens from making inferences from the absence of news that the news could have been far worse. These models do not allow for a double-edged threat of public protests, both for and against the incumbent.

The rest of the paper is organized as follows. Section 2 presents the model, Section 3

³That the citizens’ game always has an equilibrium in which no one revolts may give the impression that we also have to focus on the adversarial equilibrium. However, this is the result of assuming non-negative costs, which preclude lower dominance region in the global game. If we allow costs to be sufficiently negative, and interpret these negative “costs” as expressive payoff that citizens get from protesting, this no revolution equilibrium disappears. Then, we will have a unique equilibrium following any information design of the ruler. See footnote 5.

establishes the relationship between media freedom and popular protests in the absence of a coup threat, Section 4 introduces the threat of coups, and Section 5 discusses the setting in which citizens infer information from the fact that the leader was overthrown. Section 6 concludes.

2 Model

There is a ruler, an officer, and a continuum 1 of citizens. The ruler is one of two possible types, good (g) and bad (b). Under the common prior, the ruler is bad with probability $p \in (0, 1)$. Citizens want a good ruler to stay in power and want a bad ruler to be removed.

The game proceeds as follows. The ruler chooses the level of media freedom $m \in [0, 1]$ before knowing his type. The ruler's type is privately realized, and the (non-strategic) media generates a public report about the ruler's type consistent with the level of media freedom. In particular, with probability m , the media accurately reports the ruler's type and with a complementary probability, the media sends a null message that conveys no information to the citizens and officer. Next, the officer decides whether to attempt a coup (C) against the ruler or not (N). If the officer attempts a coup, then the citizens decide whether to protest against the coup. If the protest succeeds, then the ruler is restored; if it fails or there is no protest, then the coup succeeds. If the officer does not attempt a coup, then the citizens decide whether to protest against the regime. If the protest succeeds, then both the ruler and the officer are removed from power and the ruler is replaced by an opponent of the opposite type.⁴ We begin by focusing on the case in which the officer has no private information about the ruler's type. We then use the results as an intermediate step in the general case in which the officer has private information about the ruler and his decision to attempt a coup can signal this information to the citizens.

The ruler's payoff is 1 if he retains office and 0 if he is removed. The officer prefers to mount a coup if he is sure that it will succeed. In particular, the officer's payoff of remaining in power under the ruler is $\beta > 0$ and his payoff from mounting a successful coup is $\alpha > \beta$. If the officer mounts a failed coup or is removed from power by the citizens' protest against the regime, his payoff is 0.

The essential feature of the strategic interaction that we study is that citizens' incentives

⁴The assumption that the replacement's type is the opposite of the current ruler's streamlines the analysis considerably. It can be relaxed to allow an imperfect negative correlation.

		Outcome				Outcome	
		$n > \theta$	$n \leq \theta$			$n > \theta$	$n \leq \theta$
Protest		$(1 + \delta) u_r - c_i$	$u_o - c_i$			$(1 + \delta) (-u_r) - c_i$	$u_r - c_i$
No protest		u_r	u_o			$-u_r$	u_r

Figure 1: Left Panel: Citizens' game following the coup. Right Panel: Citizens' game following no coup. Parameters: u_o is the payoff under the officer and u_r , with $r \in \{b, g\}$, is the payoff under the original ruler. $c_i \geq 0$ is citizen i 's private costs from protesting against the status quo. $\delta > 0$ captures that a citizen who participates in the protest will receive the payoff from the outcome with a higher intensity, where we recall that $u_g = -u_b = u > 0$.

to protest depend on their beliefs about the ruler and on the status quo (i.e., whether the power is held by the ruler or the officer). Following a coup, citizens have an incentive to protest if they believe that the original ruler was good, whereas they do not have an incentive to protest if they believe that he was bad. By contrast, if there is no coup, then citizens have an incentive to protest if they believe that the ruler is bad, whereas they do not if they believe that he is good.

We model the strategic interactions between citizens as a coordination game with private costs. Citizens simultaneously decide whether to protest against the *status quo* (either the current ruler or the officer). In both cases, the strength of the *status quo* is reflected in a commonly known parameter $\theta \in (0, 1)$. If the measure of protesters is sufficiently large, $n > \theta$, then the status quo is overturned; otherwise, it is maintained. A coup changes the *status quo*, but does not fundamentally alter the structure of citizens' protest game.

The left panel of Figure 1 presents citizens' protest game following a coup, while the right panel presents the protest game if no coup occurs. A citizen's payoff depends on the type of the final ruler and on the citizen's protest decision. The payoff of a citizen who chooses not to protest is simply the type of the ultimate ruler. If the officer is left in charge, then a citizen's payoff is u_o . If the original ruler retains power, then such a citizen's payoff is u_r , where $r \in \{g, b\}$ represents the original ruler's type. By contrast, if citizens replace the ruler in the absence of a coup, then the replacement's type is opposite that of the ruler and a non-participating citizen's payoff is $-u_r$. We maintain the simplifying assumptions that $u_g = u > 0$ and $u_b = -u$.

The payoff of a protester is different from that of a non-protester in two ways. First, citizens have correlated private costs of protesting against the *status quo*. In particular,

$c_i = \bar{c} + \sigma \nu_i$, with \bar{c} and ν_i 's being independent, $\bar{c} \sim G$, and $\nu_i \sim \text{i.i.d. } F$ with full support on \mathbb{R}_+ . To simplify the analysis, we assume that $G = U[0, 1]$. Second, if the protest succeeds, then protesting citizens receive the payoff from the replacement of the *status quo* with a higher intensity. That is, a protest participant's payoff from the final ruler's type is multiplied by $1 + \delta$, where $\delta > 0$. Thus, a citizen who participates in a protest that leaves a good ruler in charge experiences a "warm glow" or "pleasure in agency" (Wood 2003; Morris and Shadmehr 2017), whereas the same citizen experiences regret if a bad ruler is left in charge.

We impose two restrictions on the model parameters to streamline our analysis. Our first assumption requires that the upper bound of the common value component of the protest costs ($\bar{c} \sim U[0, 1]$) is sufficiently large relative to the gains of participating in a successful protest.

Assumption 1 (*Upper limit dominance*) $\delta u < 1$.

When \bar{c} is at its maximum of 1, the costs of all citizens are above 1 for any positive noise, no matter how small. Thus, Assumption 1 implies that as the noise vanishes, when $\bar{c} = 1$, citizens have a dominant strategy not to protest. At the same time, Assumption 1 does not imply that citizens have a dominant strategy to protest when $\bar{c} = 0$. For example, even if the protest has no cost, after a coup attempt, citizens do not want to contribute to bringing back a bad ruler and they will choose not to protest if their beliefs about the ruler are sufficiently pessimistic.

Our second assumption requires that the officer's rent, β , from being part of the incumbent's inner circle is sufficiently high that he does not attempt a coup against a good ruler. As we see later (inequality (5)), Assumption 2 states this required restriction on the parameter space.

Assumption 2 *The officer does not attempt a coup if the ruler is known to be good:*
 $\beta/\alpha > 1 - G((1 - \theta)\delta u)$.

Relaxing these parametric restrictions complicates the analysis without adding substantive insights.

3 Media Freedom and Popular Protests

The two coordination games among citizens, presented in Figure 1, can be concisely rewritten as a single game. In particular, because each citizen is infinitesimal, she takes the probability of protest success as exogenous and her strategic decision depends only on the difference in row payoffs. Thus, the second row can be subtracted from the first. Defining $\hat{u}_r \equiv \delta u_r \cdot (\mathbf{1}\{\text{Coup}\} - \mathbf{1}\{\text{No Coup}\})$ allows us to represent citizens' coordination game as in Figure 2.

		outcome	
		$n > \theta$	$n \leq \theta$
Citizen i	Protest	$\hat{u}_r - c_i$	$-c_i$
	No Protest	0	0

Figure 2: Protest Game: $\hat{u}_r = \delta u_r \cdot (\mathbf{1}\{\text{Coup}\} - \mathbf{1}\{\text{No Coup}\})$.

Consider the case in which a coup attempt has been made, so that $\hat{u}_r = \delta u_r$. If citizen i is certain that the ruler is bad, then he has a dominant strategy not to protest. If citizen i is certain that the ruler is good, then his decision depends on his private cost of protest c_i as well as his belief about the likelihood of success P_i , so that he protests if and only if $P_i \delta u > c_i$. Of course, citizens may be uncertain about the ruler's type. They use the available information (the prior, the media's report, and the officer's action) to form a posterior about whether the ruler is good or bad. From that posterior, they calculate the expectation of $E[u_r]$:

$$E[u_r] = \Pr(\text{ruler is good}) u + \Pr(\text{ruler is bad}) (-u). \quad (1)$$

Thus, a citizen protests if and only if $P_i \delta E[u_r] > c_i$.

Regardless of \hat{u}_r , if a citizen expects that no one else will protest, then she also does not protest. Thus, there is always an equilibrium in which no citizen revolts.⁵ As in Bueno de Mesquita (2010) and Shadmehr and Bernhardt (2011), we focus on the equilibria in symmetric cutoff strategies in which protests can succeed with a positive probability, namely equilibria in which a citizen protests if and only if her cost is below some threshold $c^* > -\infty$.

Assumption 1 guarantees that following a coup, even if citizens are sure that the previous ruler is good, the realization of the cost fundamental might be sufficiently high so that the

⁵This is because the game has one-sided limit dominance, which can be ruled out if we assume that \bar{c} can be sufficiently negative. Persson and Tabellini (2009) assumed that \bar{c} has full support on \mathbb{R} , implying that c_i is sometimes not a cost, but an expressive benefit of protesting that a citizen obtains from protesting independent of the outcome.

coup succeeds. The same condition ensures that if the officer decides not to attempt a coup, then the ruler maintains power with a positive probability, even if citizens are sure that he is bad. Thus, there is no belief that citizens could hold about the ruler's type under which either the officer's coup fails for certain or a protest against the ruler succeeds for certain. In essence, the regime is sufficiently stable that even a ruler who is known to be bad may survive a popular protest; however, it is sufficiently unstable that a ruler who is known to be good may not survive a coup.

Proposition 1 *In the limit when the noise goes to zero, the status quo collapses if and only if*

$$\bar{c} < (1 - \theta) E[\hat{u}_r],$$

where the expectation is conditioned on all the information available to citizens.

Intuitively, the status quo collapses if citizens' cost fundamental is lower than a threshold value. This threshold decreases in the strength of the status quo, θ , and increases in a citizen's (updated) expectation about the net payoff of overturning the status quo, $E[\hat{u}_r]$. Thus, as citizens become more optimistic about the alternative to the status quo, the status quo is more likely to collapse. For example, if a coup occurred, then as citizens become more optimistic about the previous ruler ($E[u_r]$ increases), the coup is less likely to succeed.

To ease exposition, let $\lambda \equiv (1 - \theta)\delta u$. With this notation, if citizens believe that the ruler is bad with probability p in the protest stage, then the probability that the ruler is restored following an attempted coup is

$$R(p) \equiv G((1 - \theta)\delta E[u_r]) = G((1 - \theta)\delta(p(-u) + (1 - p)(u))) = G(\lambda(1 - 2p)), \quad (2)$$

where we recall that $G(\cdot)$ is the uniform prior of the cost fundamental \bar{c} . Now, the probability that the regime survives a popular protest in the absence of a coup is

$$\begin{aligned} S(p) &\equiv 1 - G((1 - \theta)\delta E[-u_r]) \\ &= 1 - G(\lambda(2p - 1)). \end{aligned} \quad (3)$$

Before analyzing the full game, we first study the media freedom choice assuming that the only threat to the regime comes from a popular protest (i.e., there is no possibility of a coup). This benchmark provides a counterfactual that allows us to isolate how the threat of a coup affects media freedom.

At the beginning of the protest stage, every citizen holds the same belief about the ruler. If no media report was generated, then the citizens maintain the prior belief that the ruler is bad with probability p . In this case, the regime survives with probability $S(p)$. If a media report was generated, then the citizens know the ruler's type. A good ruler always survives a popular protest, $S(0) = 1$, and a bad ruler survives with probability $S(1)$. Thus, the ruler chooses m to maximize the likelihood of the regime surviving:

$$\max_{m \in [0,1]} m(pS(1) + (1-p)) + (1-m)S(p). \quad (4)$$

Because the ruler's payoff is linear in m , we have $m^* = 0$ if and only if $S(p) \geq pS(1) + (1-p)$. This inequality is equivalent to

$$p \geq \frac{G(\lambda(2p-1))}{G(\lambda)} = \begin{cases} 0 & \text{if } p \leq 1/2 \\ 2p-1 & \text{if } p \geq 1/2, \end{cases}$$

where we use $G(x) = x$ because G is the cdf of $U[0,1]$ as well as Assumption 1, which implies $G(\lambda) \leq \delta u < 1$. The inequality always holds, and hence $m^* = 0$. Thus, we have the following proposition.

Proposition 2 *In a benchmark model in which there is no threat of a coup, the dictator allows no media freedom ($m^* = 0$).*

Proposition 2 has a natural intuition. Observe that the ruler's payoff in the absence of media freedom, given by

$$S(p) = 1 - G(\lambda(2p-1)) = \begin{cases} 1 & \text{if } p \leq \frac{1}{2} \\ 1 - \lambda(2p-1) & \text{if } p \geq \frac{1}{2}, \end{cases}$$

is concave in p . Thus, the ruler prefers $S(p)$ to the convex combination $pS(1) + (1-p)S(0)$ that he obtains with full media freedom, as illustrated in Figure 3. The concavity of $S(p)$ is linked both to the fundamental strategic forces and to the substantive meaning of Assumption 1. First, if pessimism about the original ruler is low, $p \leq 1/2$, then citizens have no reason to protest and the regime always survives (this holds independently of Assumption 1). In this case, the ruler does not benefit from further reductions in citizens' pessimism. By contrast, if citizens are more pessimistic about the ruler, $p > 1/2$, then the probability of survival is decreasing and linear in citizens' pessimism and Assumption 1 guarantees that the probability of survival does not reach 0, even at the maximum pessimism, $p = 1$. Thus, when




Figure 3: Illustration of Proposition 2. The blue curve represents $S(p)$, the ruler's payoff with no media freedom, and the red line represents the ruler's payoff from allowing complete media freedom. Parameters: $\lambda = 1/2$.

pessimism is low, $p \leq 1/2$, further reductions in pessimism are not beneficial; however, when pessimism is high, $p > 1/2$, increases in pessimism are costly, resulting in a concave $S(\cdot)$.

More broadly, the asymmetry between high and low levels of pessimism is connected to the idea that any *status quo* is difficult to overturn in a protest game. If citizens are optimistic about the ruler, then they have no reason to overturn that status quo; however, even if citizens are sure that the ruler is bad, there is a chance that he will remain in power (Assumption 1). Thus, the “easy” action for citizens—not protesting—is also the one that the ruler would like them to choose. This situation is reversed following a coup: the ruler wants citizens to overturn the new *status quo* and restore him to power. To that end, the ruler will want to provide information about himself, as we see in the next section.

4 Media Freedom Under the Threat of a Coup

We now consider the ruler's media freedom choice in the full game. To analyze this question, we first derive the officer's optimal strategy.

Officer's Decision. The officer has a direct benefit from carrying out a coup against the ruler, but his decision also depends on the likelihood that his coup attempt will succeed and on the likelihood that he (and the ruler) will be removed from power by the citizens in the absence of a coup. Thus, our model allows for the possibility of a preemptive coup, whose

purpose is to reduce the probability of a regime collapse.

The officer has no private information about the ruler's type and, Therefore, his decision to attempt a coup has no effect on citizens' belief about the ruler.⁶ Let p' be citizens' posterior belief that the ruler is bad after observing the media report. This belief can take three possible values: $p' = 0$ (media revealed the ruler to be good), $p' = 1$ (media revealed the ruler to be bad), and $p' = p$ (media revealed no new information). From Equations (2) and (3), the officer attempts a coup if and only if $\alpha(1 - R(p')) > \beta S(p')$, which is equivalent to

$$\frac{1 - R(p')}{S(p')} = \frac{1 - G(\lambda(1 - 2p'))}{1 - G(\lambda(2p' - 1))} > \gamma \equiv \frac{\beta}{\alpha}. \quad (5)$$

Assumption 2 implies that if citizens are sure that the ruler is good, $p' = 0$, then the likelihood of a subsequent protest in support of the ruler is sufficient to prevent a coup. Thus, being known as a good type is valuable for the ruler because it ensures that he survives.

By implication, the officer always attempts a coup if the ruler is known to be bad:

$$\frac{1 - R(1)}{S(1)} = \frac{1}{1 - G(\lambda)} > \frac{1}{\gamma} > \gamma.$$

Moreover, the ratio in Equation (5) is strictly increasing in p' . When the ruler is more likely to be bad, he is less likely to be restored by the citizens in the event of a coup, which implies that the numerator is increasing in p' . Furthermore, the regime is less likely to survive if the officer does not attempt a coup and hence the denominator is decreasing in p' . By implication, the officer's incentive to attempt a coup becomes stronger as the belief about the ruler becomes worse. Thus, there exists a unique threshold belief, $\frac{1 - G(\lambda(1 - 2p'))}{1 - G(\lambda(2p' - 1))} = \gamma$, at which the officer's best response switches from no coup to attempting a coup.

Proposition 3 *In equilibrium, the officer attempts a coup if and only if $p' > P(\lambda, \gamma)$, where $P(\lambda, \gamma) \equiv \frac{1}{2} - \frac{1 - \gamma}{2\lambda}$.*

In the absence of a coup attempt, the regime survives when citizens are optimistic about the ruler. In particular, Equation (3) shows that when citizens believe that the ruler is bad with a probability less than $1/2$, the regime is sure to survive in the absence of a coup. This result together with Proposition 3 allows us to identify the incentives that drive the officer's decision. When $p' < P(\lambda, \gamma)$, citizens believe, upon observing the media report, that the incumbent is likely to be good. Consequently, they would not protest against

⁶We relax this assumption in Section 5.

the regime if there is no coup, but they are likely to respond to a coup attempt with a counter-protest. Thus, the officer prefers to collect his payoff of being part of the regime without risking a coup, which is likely to fail. In the middle region, $P(\lambda, \gamma) < p' < \frac{1}{2}$, the officer's behavior is a "power grab" against a stable ruler. That is, citizens are sufficiently confident that the incumbent is good that they would never protest against him directly; however, if there is a coup, they are unlikely to protest either. This provides the officer with the opportunity to topple the incumbent and then, likely, not face a counter-protest. For $p' > \frac{1}{2}$, the officer anticipates both a lower probability of the current regime's survival and no counter-protest. Both these factors contribute to the officer's decision to take power: if he does not act, he is likely to lose his position with the unpopular leader; if he acts, there will be no counter-protest as people expect an improvement relative to the *status quo*.

Media Freedom. If the ruler reveals his type, then he survives if he is good (because there will be no coup and no protest) and loses power if he is bad (because the officer will carry out a coup and the citizens will not protest against the coup). If the ruler does not reveal his type, then his payoff depends on both the officer's decision and the citizens' response. Obviously, if $p < P(\lambda, \gamma)$, then the ruler prefers not to reveal information. Indeed, in this case, the officer is sufficiently worried about a counter-protest that he does not attempt a coup and the citizens are sufficiently optimistic about the ruler they do not protest against him. By contrast, if $p > P(\lambda, \gamma)$, then the ruler is truly "hanging by a thread": in the absence of additional positive information, the officer will attempt to overthrow him. In this case, the ruler's expected payoff from choosing media freedom $m \in [0, 1]$ is

$$m(1 - p) + (1 - m)R(p).$$

Because the ruler's payoff is linear in m , we have $m^* = 1$ if and only if $1 - p \geq R(p)$, which is, in turn, equivalent to

$$p \leq 1 - G(\lambda(1 - 2p)) = \begin{cases} 1 & \text{if } p \geq \frac{1}{2} \\ 1 - \lambda(2p - 1) & \text{if } p < \frac{1}{2}. \end{cases}$$

Obviously, this inequality holds for $p \geq 1/2$. For $p < 1/2$, the functions on both sides of the inequality are linear and, therefore, it suffices to check that the inequality holds at $p = 0$ and $p = 1/2$.




Figure 4: Illustration of Proposition 4. The blue curve represents the ruler's payoff with no media freedom and the red line represents the ruler's payoff from allowing complete media freedom. The dotted red line is the expected payoff that the ruler can achieve if he uses a more complex media structure. This is the upper bound of what the ruler can achieve with Bayesian citizens. Parameters: $\lambda = 1/2$, $P(\lambda, \gamma) = 3/10$.

Proposition 4 *When there is a threat of a coup and a popular protest, in equilibrium, the ruler allows complete media freedom, $m^* = 1$, if the citizens are sufficiently pessimistic about him, $p > P(\lambda, \gamma)$, whereas he allows no media freedom, $m^* = 0$, if the citizens are sufficiently optimistic about him, $p < P(\lambda, \gamma)$.*

Figure 4 illustrates this proposition graphically. The blue curve represents the ruler's expected payoff as a function of the citizens' prior belief. For $p < P(\lambda, \gamma)$, the officer does not attempt a coup and the ruler is concerned only with surviving the citizens' protest; therefore, in this domain, the ruler's payoff is equal to $S(p)$. For $p > P(\lambda, \gamma)$, the officer does attempt a coup and, therefore, the ruler is concerned with the probability that he will be restored following a coup; his payoff in this domain is thus $R(p)$. The red line is the ruler's expected payoff with complete media freedom, which is higher than the blue curve when the ruler is under threat and lower when he is not.

Together, Figures 3 and 4 illustrate why the potential for a status quo reversal that accompanies a coup affects the ruler's choice of media freedom. If no coup is possible, then the ruler is only worried about surviving a popular protest and his payoff is therefore the concave function $S(p)$. However, when a coup is possible, the officer will attempt it whenever $p > P(\lambda, \gamma)$, generating a reversal of the *status quo*. Thus, for large p , the ruler is worried about being restored following the coup, and his payoff, therefore, is the convex function $R(p)$. Hence, the possibility of a *status quo* reversal generates a payoff function for

the ruler that is not concave, opening the possibility that complete media freedom could be beneficial. Furthermore, given that the officer prefers to mount a coup when the citizens are unlikely to support the ruler, for $p > P(\lambda, \gamma)$, the ruler's payoff $R(p)$ is low (or zero for $p > 1/2$). Thus, there is little to lose from being revealed to be the bad type by the media; however, there is a lot to gain from being revealed to be a good type. Therefore, complete media freedom is optimal.

Finally, suppose that the ruler can design a more complex policy that the media will use when reporting the ruler's type. In particular, suppose that the ruler can design a set of messages M and a pair of reporting rules $\pi(\cdot|g)$ and $\pi(\cdot|b)$ that specify the probability of each possible message conditional on the ruler's underlying type; moreover, the media's reporting policy is observed by citizens along with the message. By using the arguments in the literature on Bayesian persuasion (Kamenica and Gentzkow 2011), the ruler's maximized payoff can be found by forming the concave envelope of his payoff function and evaluating it at the prior belief.

For $p < P(\lambda, \gamma)$, the ruler would still prefer to reveal no information. However, for $p > P(\lambda, \gamma)$, the concave envelope of the ruler's payoff function (the dotted red line) lies strictly above the solid red line (which represents his payoff from a fully revealing reporting rule); therefore, the ruler can do better than allowing the media to fully report his type. The optimal signal generates the posterior beliefs $p = 1$ and $p = P(\lambda, \gamma)$, where the concave closure meets the ruler's payoff function (Kamenica and Gentzkow 2011). Thus, the optimal reporting policy consists of two reports, bad and good. The bad report perfectly reveals that the ruler is bad ($p_B = 1$), whereas a good report improves beliefs sufficiently that the officer prefers not to attempt a coup ($p_G = P(\lambda, \gamma)$). Thus, the good report is always transmitted if the ruler is good, but both the good and the bad reports are sometimes transmitted if the ruler is bad. Such a reporting policy has a natural interpretation as the (stochastic) censorship of bad news: if the ruler is good, the news is always reported; on the contrary, if the ruler is bad, the media sometimes replaces the (truthful) bad report with a good one.

5 Signaling and an Informative Coup

In addition to changing the status quo, a coup may convey information to citizens about the ruler's type, a consideration that was absent from our previous analysis. In this section,

we allow the officer to have private information about the ruler's type and his coup decision to potentially signal his information to citizens. We show that, in a natural setting, these additional signaling considerations strengthen our main finding that the threat of coups creates incentives for the ruler to allow media freedom.

When the officer decides whether to attempt a coup, suppose he knows the ruler's type and that his payoff from the ruler remaining in power depends on that type. In particular, as before, let α be the officer's payoff from mounting a successful coup. However, now, let β_i be the officer's payoff if a type $i \in \{g, b\}$ ruler remains in power. We focus on the natural case of $\alpha > \beta_g > \beta_b > 0$, meaning that the officer's payoff under a good ruler is larger than that under a bad one.

Denote the probability that the officer attempts a coup against a type- i ruler by σ_i , citizens' belief that the ruler is bad before observing the officer's decision by p' , and citizens' updated beliefs that the ruler is bad following a coup and no coup by p'_C and p'_{NC} . Unlike in the preceding sections, which assumed $p'_C = p'_{NC}$ (because the officer was uninformed), here these probabilities could be different because the officer's decision may signal his private information to citizens.

The officer's incentives are similar to those described in Section 4. The officer compares the relative payoff of allowing the ruler to remain in power (risking a popular protest) and an attempted coup (risking counter-protest). The difference here is that the payoff of not attempting a coup also depends on the ruler's type. The officer attempts a coup against the type- i ruler (i.e., $\sigma_i = 1$) if and only if $\alpha(1 - R(p'_C)) > \beta_i S(p'_{NC})$, which is equivalent to

$$\frac{1 - R(p'_C)}{S(p'_{NC})} = \frac{1 - G(\lambda(1 - 2p'_C))}{1 - G(\lambda(2p'_{NC} - 1))} > \gamma_i \equiv \frac{\beta_i}{\alpha}. \quad (6)$$

Because $\gamma_g > \gamma_b$, the officer has a weaker incentive to mount a coup against a good ruler. This observation is important both for the possible equilibrium configurations and for the refinement of the off-path beliefs in our analysis of possible pooling equilibria.

Citizens' beliefs must be consistent with Bayes' rule applied to the officer's strategy:

$$p'_C = \frac{p'\sigma_b}{p'\sigma_b + (1 - p')\sigma_g} \quad (7)$$

$$p'_{NC} = \frac{p'(1 - \sigma_b)}{p'(1 - \sigma_b) + (1 - p')(1 - \sigma_g)}. \quad (8)$$

Increases in the probability of a coup attempt against a bad ruler (σ_b) make citizens more pessimistic about the ruler and, therefore, they are less inclined to engage in a

counter-protest. Conversely, increases in the probability of a coup attempt against a good ruler (σ_g) make citizens more inclined to engage in a counter-protest.

First, consider a separating equilibrium in which the officer's decision to attempt a coup completely reveals the ruler's type. In such an equilibrium, $\sigma_g = 0$ and $\sigma_b = 1$, and hence citizens can infer the ruler's type: $p'_C = 1$ and $p'_{NC} = 0$. In this case, the left-hand side of (6) is 1. Therefore, for such a strategy to be optimal for the officer, the parameter values should satisfy $\gamma_b < 1 < \gamma_g$, which is equivalent to $\beta_b < \alpha < \beta_g$. Thus, an equilibrium in which the officer's decision conveys the ruler's type exists if and only if the officer's preferences are aligned sufficiently well with those of the citizens. If he is sure that a coup attempt will succeed, then the officer attempts it if and only if the ruler is bad. However, given our assumption that $\alpha > \beta_g$, no such equilibrium exists.

Next, consider the possibility of a pooling equilibrium in which both types of officers play the same pure strategy. In such an equilibrium, the action *not* chosen by any type in equilibrium can, in general, have any beliefs associated with it. We use the *D1* (Fudenberg and Tirole 2000, p. 452) refinement to pin down the following off-path beliefs.⁷

Remark 5 The D1 refinement delivers the following off-path beliefs: (1) If both types of officers do not attempt a coup, $\sigma_g = \sigma_b = 0$, then upon observing an off-path coup, citizens believe that the ruler is bad, $p'_C = 1$; (2) If both types of officers do attempt a coup, $\sigma_g = \sigma_b = 1$, then upon observing no coup off the equilibrium path, citizens believe that the ruler is good, $p'_{NC} = 0$.

These beliefs immediately rule out the possibility of a pooling equilibrium in which both types do not attempt a coup ($\sigma_g = \sigma_b = 0$). In such an equilibrium, an off-path coup would be interpreted as confirmation that the ruler is bad and no citizen would support him in a counter-protest. Therefore, any coup attempt would succeed for certain and the officer would prefer to deviate by mounting it. An equilibrium in which the officer attempts a coup against both types of rulers ($\sigma_g = \sigma_b = 1$) exists if p is sufficiently large. In such an equilibrium,

⁷The D1 refinement assigns beliefs to an off-path action by considering the incentives of the underlying types to select the action. In particular, the off-path action is treated as if it was selected only by the player type for whom the action is a best response for a larger set of strategies. In our setting, the officer has a stronger incentive to mount a coup against a bad ruler. In particular, whenever the ruler is good and the officer mounts a coup, he would also mount a coup if the ruler were bad. Therefore, the D1 refinement implies that when the off-path action is a coup, it reveals that the ruler is bad and when the off-path action is no coup, it reveals that the ruler is good.

$p'_C = p'$ because a coup attempt reveals no information and the above remark yields that $p'_{NC} = 0$. For both officers to be willing to attempt a coup, the parameters should satisfy

$$\frac{1 - R(p'_C)}{S(p'_{NC})} \geq \gamma_g,$$

or, equivalently, $p' \geq P(\lambda, \gamma_g)$.

Finally, consider the possibility of a semi-separating equilibrium. Given that the officer has a stronger incentive to mount a coup against a bad ruler, the only possibility is that $\sigma_b = 1$ and $\sigma_g \in (0, 1)$. In such an equilibrium, the officer's decision not to mount a coup reveals that the ruler is good, $p'_{NC} = 0$. For the officer to mix strategies when the ruler is good, the following condition should be fulfilled: $1 - R(p'_C) = \gamma_g$. Equivalently, $p'_C = P(\lambda, \gamma_g)$.

By substituting the officer's strategy into (7) and solving the equation, we find

$$\sigma_g = \frac{p'}{1 - p'} \frac{1 - P(\lambda, \gamma_g)}{P(\lambda, \gamma_g)}.$$

Therefore, we have the following characterization of the officer's equilibrium strategy and citizens' posterior beliefs following the officer's decision.

Proposition 6 *Suppose that after observing the media's message, citizens believe that the ruler is bad with probability p' . In equilibrium, the officer's strategy and citizens' beliefs are uniquely determined as follows.*

- (i) *If $p' \geq P(\lambda, \gamma_g)$, then the officer attempts a coup against both types of rulers, $\sigma_g = \sigma_b = 1$, a coup attempt conveys no information, $p'_C = p$, and (consistent with the D1 criterion) no coup attempt reveals that the ruler is good, $p'_{NC} = 0$.*
- (ii) *If $p' < P(\lambda, \gamma_g)$, then the officer always attempts a coup against the bad ruler, $\sigma_b = 1$, and attempts a coup against a good ruler with probability*

$$\sigma_g = \frac{p'}{1 - p'} \frac{1 - P(\lambda, \gamma_g)}{P(\lambda, \gamma_g)}.$$

A coup attempt conveys bad information about the ruler, but does not reveal his type completely, $p'_C = P(\lambda, \gamma_g)$, and no coup reveals that the ruler is good, $p'_{NC} = 0$.

To understand the structure of the equilibrium, note first that both the officer and the citizens would like the officer to replace a bad ruler. Thus, an increase in the likelihood




Figure 5: Illustration of Proposition 7. The blue curve represents the ruler's payoff with no media freedom and the red line represents the ruler's payoff from allowing complete media freedom. Parameters: $\lambda = 1/2$, $P(\lambda, \gamma_g) = 4/10$.

of a coup against the bad ruler makes citizens update more negatively about the ruler, reducing their incentives to protest following a coup, which in turn reduces the officer's risk of attempting a coup. As a result, the officer always attempts to overthrow the bad ruler. By contrast, when the ruler is good, the officer still wants to replace him (although less intensely than he wants to replace a bad ruler), whereas the citizens want the good ruler to stay in power. In this case, an increase in the likelihood of a coup against the good ruler makes citizens update less negatively about the ruler, increasing their incentives to protest following a coup, which in turn increases the officer's risk of attempting a coup. When citizens are pessimistic about the ruler, this updating is too weak to cause them to protest following a coup and the officer always stages a coup against the good ruler. When citizens are more optimistic about the ruler, the updating is stronger and they sometimes protest following a coup and succeed in thwarting it. Thus, the officer attempts to overthrow the good ruler less often.

Media Freedom. To analyze media freedom when the coup is informative, we construct the ruler's payoff as a function of p' , namely the citizens' belief following the media's message. For every p' , this payoff function can be written as

$$\Pr(\text{coup})R(p'_C) + \Pr(\text{no coup})S(p'_{NC}).$$

From Proposition 6, we know that $p'_{NC} = 0$. That is, in the event of no coup, the ruler is revealed to be good and hence the regime always survives, $S(p'_{NC}) = 1$. Next, for $p' >$

$P(\lambda, \gamma_g)$, the probability of a coup is 1 and the probability of regime survival is simply $R(p')$. Finally, for $p' < P(\lambda, \gamma_g)$, the probability of a coup is $p' + (1 - p')\sigma_g = p'/P(\lambda, \gamma_g)$. Following a coup, the citizens believe that the ruler is bad with probability $p'_C = P(\lambda, \gamma_g)$, and he is therefore restored with probability $R(P(\lambda, \gamma_g))$. Thus, the ruler's payoff as a function of p' is

$$U(p') = \begin{cases} \frac{p'}{P(\lambda, \gamma_g)} R(P(\lambda, \gamma_g)) + (1 - \frac{p'}{P(\lambda, \gamma_g)}) & \text{if } p' < P(\lambda, \gamma_g) \\ R(p') & \text{if } p' \geq P(\lambda, \gamma_g). \end{cases}$$

It is straightforward to verify that $U(p')$ is continuous, piecewise linear, and convex, as illustrated in Figure 5. By using Corollary 1 and Proposition 3 in Kamenica and Gentzkow (2011), we find that the ruler always allows complete media freedom in equilibrium, even if he can design any reporting policy for the media. Proposition 7 states this result formally.

Proposition 7 *In the equilibrium of the game with an informed officer, the ruler allows complete media freedom, $m^* = 1$ for all prior beliefs.*

Unlike in the case without signaling, where the officer only attempts a coup when beliefs about the ruler are sufficiently bad, the probability of a coup attempt is positive for all $p' > 0$ in the current environment. Furthermore, citizens infer bad information about the ruler from the officer's attempted coup and they are less likely to stage a counter-protest in support of the ruler. In particular, when citizens' beliefs about the ruler are initially favorable, citizens are strictly less likely to support the ruler following a coup in the environment with a privately informed officer (and if beliefs are initially unfavorable, then the ruler is restored with the same probability with and without signaling). To mitigate the increased likelihood of a successful coup, the ruler allows a greater degree of media freedom when citizens are optimistic about his type.

6 Conclusion

Although dictators do not like protesters in the streets in normal circumstances, they might want to have them after being deposed in a coup. Even the threat of subsequent protests makes subordinates hesitate before attempting a coup, and after the coup, such protests can help restore deposed rulers. This paper studies the implications of this observation for media freedom in dictatorships in a simple setup that combines the global game technique with Bayesian persuasion. Whether there is a threat of a coup or a threat of a popular protest,

should not the leader raise media freedom when he believes that the media will mostly report positive information about him and reduce media freedom when he believes that the media will mostly report negative information about him? We show that the answer is no. When a ruler perceives a high likelihood of a popular protest and no coup, he opts for a less free media, which reveals little information about what kind of ruler he is. In sharp contrast, he sometimes does the opposite when there is a high likelihood of a coup and when citizens infer little about the incumbent's quality from the fact that he was removed in a coup.

References

- Acemoglu, Daron, and James Robinson. 2006. *Political Origins of Dictatorship and Democracy*. New York, NY: Cambridge University Press.
- Acemoglu, Daron, Davide Ticchi, and Andrea Vindigni. 2010. “A Theory of Military Dictatorships.” *American Economic Journal: Macroeconomics* 2 (1): 1-42.
- Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. “Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks.” *Econometrica* 75 (3): 711-56.
- Barbera, Salvador, and Matthew Jackson. 2017. “A Model of Protests, Revolution, and Information.” Mimeo.
- Bergemann, Dirk, and Stephen Morris. 2016. “Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium.” *American Economic Review: Papers & Proceedings* 106 (5): 586-91.
- Bergemann, Dirk, and Stephen Morris. 2017. “Information Design: A Unified Perspective.” Mimeo.
- Boix, Carles, and Milan Svolik. 2013. “The Foundations of Limited Authoritarian Government: Institutions and Power-sharing in Dictatorships.” *Journal of Politics* 75 (2): 300-16.
- Bueno de Mesquita, Bruce, Alastair Smith, Randolph Siverson, and James Morrow. 2003. *The Logic of Political Survival*. Cambridge, MA: MIT Press.
- Bueno de Mesquita, Ethan. 2010. “Regime Change and Revolutionary Entrepreneurs.” *American Political Science Review* 104 (3): 446-66.
- Casper, Brett, and Scott Tyson. 2014. “Popular Protest and Elite Coordination in a Coup Detat.” *Journal of Politics* 76 (2): 548-64.
- Carlsson, Hans, and Eric van Damme. 1993. “Global Games and Equilibrium Selection.” *Econometrica* 61 (5): 989-1018.
- Chen, Heng, Yang K. Lu, and Wing Suen. 2016. “The Power of Whispers: A Theory of Rumor, Communication and Revolution.” *International Economic Review* 57 (1): 89-116.
- Chen, Yuyu, and David Yang. 2018. “The Impact of Media Censorship: 1984 or Brave New World?” Mimeo.

- Dorsch, Michael, and Paul Maarek. 2018. "Rent Extraction, Revolutionary Threat, and Coups in Non-democracies." *Journal of Comparative Economics*, forthcoming.
- Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. "Why Resource-Poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review* 103 (4): 645-68.
- Egorov, Georgy, and Konstantin Sonin. 2011. "Dictators and Their Viziers: Endogenizing the Loyalty-Competence Trade-off." *Journal of European Economic Association* 9 (5): 903-30.
- Egorov, Georgy, and Konstantin Sonin. 2014. "Elections in Non-Democracies." Mimeo.
- Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80 (4): 1422-58.
- Fudenberg, Drew, and Jean Tirole. 2000. *Game Theory*. Cambridge, MA: MIT Press.
- Gandhi, Jennifer, and Adam Przeworski. 2006. "Cooperation, Cooptation, and Rebellion under Dictatorships." *Economics and Politics* 18 (1): 1-26.
- Gandhi, Jennifer, and Adam Przeworski. 2007. "Authoritarian Institutions and the Survival of Autocrats." *Comparative Political Studies* 40 (11): 1279-1301.
- Gehlbach, Scott, and Alberto Simpser. 2015. "Electoral Manipulation as Bureaucratic Control." *American Journal of Political Science* 59 (1): 212-24.
- Gehlbach, Scott, and Konstantin Sonin. 2014. "Government Control of the Media." *Journal of Public Economics* 118: 163-71.
- Gentzkow, Matthew, Jesse Shapiro, and Daniel Stone. 2016. "Media Bias in the Marketplace: Theory." In *Handbook of Media Economics*, vol. 2. Ed. Anderson, Waldfogel, and Stromberg. North Holland.
- Goldstein, Itay, and Chong Huang. 2016. "Bayesian Persuasion in Coordination Games." *American Economic Review: Papers & Proceedings* 106 (5): 592-96.
- Guriev, Sergei, and Daniel Treisman. 2015. "How Modern Dictators Survive: An Informational Theory of the New Authoritarianism." NBER Discussion Paper No. 21136.
- Hollyer, James, Peter Rosendorff, and James Vreeland. 2018. "Why do Autocrats Disclose? Economic Transparency and Inter-Elite Politics in the Shadow of Mass Unrest." *Journal of Conflict Resolution*, forthcoming.

- Inostroza, Nicolas, and Alessandro Pavan. 2018. "Persuasion in Global Games with Application to Stress Testing." Mimeo.
- Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590-2615.
- King, Gary, Jennifer Pan, and Margaret Roberts. 2013. "How Censorship in China Allows Government Criticism, but Silences Collective Expression." *American Political Science Review* 107 (2): 326-43.
- King, Gary, Jennifer Pan, and Margaret Roberts. 2014. "Reverse-engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345, 1251722.
- Lorentzen, Peter. 2013. "Regularizing Rioting: Permitting Public Protest in an Authoritarian Regime." *Quarterly Journal of Political Science* 8: 127-58.
- Lorentzen, Peter. 2014. "China's Strategic Censorship." *American Journal of Political Science* 58 (2): 402-14.
- Luo, Zhaotian, and Arturas Rozenas. 2018. "Strategies of Election Rigging: Trade-Offs, Determinants, and Consequences." *Quarterly Journal of Political Science* 13 (1): 1-28.
- Morris, Stephen and Mehdi Shadmehr. 2017. "Inspiring Regime Change." Mimeo.
- Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3): 587-97.
- Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Application." In *Advances in Economics and Econometrics, Theory and Applications, Eighth World Congress, Vol. I*. Ed. Dewatripont, Hansen, and Turnovsky. New York, NY: Cambridge University Press.
- Padró i Miquel, Padro. 2007. "The Control of Politicians in Divided Societies: The Politics of Fear." *Review of Economic Studies* 74: 1259-74.
- Persson, Torsten, and Guido Tabellini. 2009. "Democratic Capital: The Nexus of Political and Economic Change." *American Economic Journal: Macroeconomics* 1 (2): 88-126.
- Prat, Andrea. 2015. "Media Capture and Media Power." In *Handbook of Media Economics, Vol. 1b*. Ed. Anderson, Waldfogel, and Stromberg. North Holland.
- Rundlett, Ashlea, and Milan Svolik. 2016. "Deliver the Vote! Micromotives and Macrobehavior in Electoral Fraud." *American Political Science Review* 110 (1): 180-97.

- Shadmehr, Mehdi, and Dan Bernhardt. 2011. "Collective Action with Uncertain Payoffs: Coordination, Public Signals and Punishment Dilemmas." *American Political Science Review* 105 (4): 829-51.
- Shadmehr, Mehdi, and Dan Bernhardt. 2015. "State Censorship." *American Economic Journal: Microeconomics* 7 (2): 280-307.
- Smith, Alastair, and Scott Tyson. 2018. "Dual-Layered Coordination and Political Instability." *Journal of Politics* 80 (1): 44-58.
- Svolik, Milan. 2008. "Authoritarian Reversals and Democratic Consolidation." *American Political Science Review* 102 (2): 153-68.
- Svolik, Milan. 2009. "Power-sharing and Leadership Dynamics in Authoritarian Regimes." *American Journal of Political Science* 53 (2): 477-94.
- Svolik, Milan. 2012. *The Politics of Authoritarian Rule*. New York, NY: Cambridge University Press.
- Taubman, William. 2003. *Khrushchev: The Man and His Era*. New York, NY: W. W. Norton & Company.
- Taubman, William. 2017. *Gorbachev: His Life and Times*. New York, NY: W.W. Norton & Company.
- Wood, Elisabeth J. 2003. *Insurgent Collective Action and Civil War in El Salvador*. New York, NY: Cambridge University Press.

Appendix

Proof of Proposition 1: We look for equilibria in which a citizen with cost c_i protests if and only if $c_i < c^* \in (0, 1]$. As discussed in the text, there is always an equilibrium in which citizens never revolt. This equilibrium corresponds to any $c^* \leq 0$. We discuss the case of $c^* > 1$ later. Given the realization of \bar{c} and a cutoff strategy $c^* \in (0, 1]$, the measure of protesters is $\Pr(c_i < c^* | \bar{c}) = F\left(\frac{c^* - \bar{c}}{\sigma}\right)$. Thus, the protest succeeds if and only if $F\left(\frac{c^* - \bar{c}}{\sigma}\right) > \theta$. In the limit when $\sigma \rightarrow 0$,

$$\lim_{\sigma \rightarrow 0} \Pr(c_i < c^* | \bar{c} = 0) = \lim_{\sigma \rightarrow 0} F\left(\frac{c^*}{\sigma}\right) = 1 > \theta > 0 = \lim_{\sigma \rightarrow 0} F\left(\frac{c^* - 1}{\sigma}\right) = \lim_{\sigma \rightarrow 0} \Pr(c_i < c^* | \bar{c} = 1).$$

Moreover, $\Pr(c_i < c^* | \bar{c})$ is decreasing in \bar{c} . Thus, given $c^* \in (0, 1]$, there exists a threshold $\tilde{c} \in (0, 1)$ such that a protest succeeds if and only if $\bar{c} < \tilde{c}$, where

$$F\left(\frac{c^* - \tilde{c}}{\sigma}\right) = \theta. \quad (9)$$

This also implies

$$\lim_{\sigma \rightarrow 0} \tilde{c} = \lim_{\sigma \rightarrow 0} [c^* - \sigma F^{-1}(\theta)] = \lim_{\sigma \rightarrow 0} c^*. \quad (10)$$

Next, because costs are correlated, a citizen's belief about the likelihood of a protest succeeding depends on her type c_i . Thus, citizen i believes that the protest will succeed with probability $P_i = \Pr(\bar{c} < \tilde{c} | c_i)$ and she protests if and only if

$$\Pr(\bar{c} < \tilde{c} | c_i) \delta E[u_r] > c_i. \quad (11)$$

Because $c_i \geq 0$, no citizen protests and the status quo survives whenever $E[u_r] \leq 0$. Moreover, if $\tilde{c} = 0$, then the left-hand side is 0 and no citizen protests. Now, consider the case in which both $E[u_r] > 0$ and $\tilde{c} > 0$. Still, because the left-hand side of (11) is bounded, a citizen with a sufficiently high cost c_i does not revolt. However,

$$\Pr(\bar{c} < \tilde{c} | c_i = 0) \delta E[u_r] > c_i = 0.$$

Thus, for any $\tilde{c} > 0$, there is a unique $c^* > 0$ such that a citizen with cost c_i protests if and only if $c_i < c^*$ and a citizen with cost $c_i = c^*$ is indifferent between protesting and not protesting,

$$\Pr(\bar{c} < \tilde{c} | c_i = c^*) \delta E[u_r] = c^*, \quad \text{for } \tilde{c} > 0. \quad (12)$$

Any pair (c^*, \tilde{c}) , with $\tilde{c} > 0$ and $c^* \in (0, 1]$, that satisfies (9) and (12) is an equilibrium.

Moreover, when \tilde{c} , $c^* > 0$, in the limit when the noise in private signals vanishes,

$$\begin{aligned}
\Pr(\bar{c} < \tilde{c} | c_i = c^*) &= \int_{\bar{c}=0}^{\tilde{c}} pdf(\bar{c} | c_i = c^*) d\bar{c} \\
&= \int_{\bar{c}=0}^{\tilde{c}} \left[\frac{pdf(c_i = c^* | \bar{c}) pdf(\bar{c})}{\int_{\bar{c}=0}^1 pdf(c_i = c^* | \bar{c}) pdf(\bar{c}) d\bar{c}} \right] d\bar{c} \quad (\text{the term in brackets is Bayes' rule}) \\
&= \int_{\bar{c}=0}^{\tilde{c}} \frac{f\left(\frac{c^* - \bar{c}}{\sigma}\right) g(\bar{c})}{\int_{\bar{c}=0}^1 f\left(\frac{c^* - \bar{c}}{\sigma}\right) g(\bar{c}) d\bar{c}} d\bar{c} \quad (\text{we canceled } 1/\sigma, \text{ and used } g(\cdot) \text{ as the pdf of } \bar{c}) \\
&= \int_{z=\frac{c^* - \tilde{c}}{\sigma}}^{\frac{c^*}{\sigma}} \frac{f(z) g(c^* - \sigma z)}{\int_{z=\frac{c^* - 1}{\sigma}}^{\frac{c^*}{\sigma}} f(z) g(c^* - \sigma z) dz} dz, \quad (\text{a change of variables to } z = \frac{c^* - \bar{c}}{\sigma}) \\
&= \int_{z=\frac{c^* - \tilde{c}}{\sigma}}^{\frac{c^*}{\sigma}} \frac{f(z)}{\int_{z=\frac{c^* - 1}{\sigma}}^{\frac{c^*}{\sigma}} f(z) dz} dz, \quad (\text{because } \bar{c} \sim U[0, 1] \text{ implies } g(\cdot) = 1) \\
&= \frac{F\left(\frac{c^*}{\sigma}\right) - F\left(\frac{c^* - \tilde{c}}{\sigma}\right)}{F\left(\frac{c^*}{\sigma}\right) - F\left(\frac{c^* - 1}{\sigma}\right)} \\
&= \frac{F\left(\frac{c^*}{\sigma}\right) - \theta}{F\left(\frac{c^*}{\sigma}\right) - F\left(\frac{c^* - 1}{\sigma}\right)} \quad (\text{from equation (9)}). \tag{13}
\end{aligned}$$

Because $c^* \in (0, 1]$, Equation (13) implies that

$$\lim_{\sigma \rightarrow 0} \Pr(\bar{c} < \tilde{c} | c_i = c^*) = 1 - \theta. \tag{14}$$

By combining (10) with Equations (12) and (14),

$$\lim_{\sigma \rightarrow 0} \tilde{c} = \lim_{\sigma \rightarrow 0} c^* = (1 - \theta) \delta E[u_r] < 1. \tag{15}$$

It remains to investigate whether there can be an equilibrium with $c^* > 1$. However, Assumption 1 together with (11) rules this out.

Finally, Equation (15) also captures the case of $E[u_r] \leq 0$, in which case $\lim_{\sigma \rightarrow 0} \tilde{c} = \lim_{\sigma \rightarrow 0} c^* \leq 0$, implying that no one protests and the regime survives. \square