A Theory of State Censorship[1]

Mehdi Shadmehr[2] Dan Bernhardt[3]

[1]We thank Raphael Boleslavsky, Ethan Bueno de Mesquita, Odilon Camara, Georgy Egorov, John Londregan, Paulina Marek, Aya Kachi, Patrick Kuhn, Arturas Rozenas, and participants at the 2010 and 2011 MPSA conference, and seminars at Caltech, UIUC departments of economics and political science, and University of Miami finance department for helpful suggestions and comments.

[2]Department of Economics, University of Miami, 5250 University Dr., Coral Gables, FL 33146. E-mail: shad@miami.edu

[3]Department of Economics, University of Illinois, 1407 W. Gregory Dr., Urbana, IL 61801. E-mail: danber@illinois.edu

Abstract

We characterize a ruler's decision of whether to censor media reports that convey public information to citizens who decide whether to revolt. We find: (1) a ruler gains (his ex ante expected payoff increases) by committing to censoring slightly less than he does in equilibrium: his equilibrium calculations ignore that censoring less causes citizens to update more positively following no news; (2) a ruler gains from higher censorship costs if and only if censorship costs exceed a critical threshold; (3) a bad ruler prefers a very active media to a very passive one, while a good ruler prefers the opposite.

1 Introduction

The recent Arab Spring has highlighted the complicated strategic calculus of revolution that citizens and unpopular rulers face. Citizens face great uncertainties about whether the possible rewards from an alternative regime are worth the risk of participating in a revolution, and use all the information at their disposal. In this paper, we analyze the strategic choices by a ruler (an authoritarian state) of when to manipulate citizens' information by censoring the media. In particular, we investigate how a ruler's welfare is influenced by (1) his ability to pre-commit to a censorship law (censorship strategy), (2) communication technologies that raise censorship costs, and (3) the strength/competence of the media that determines the likelihood that the media discovers news about the regime.

The ruler tries to manage information transmission to citizens to mitigate the likelihood of revolution. Censoring a news event can benefit a ruler whenever the likelihood of revolt following that news exceeds the likelihood of revolt following no news. Citizens understand a ruler's incentives to conceal bad news, so they update negatively about the regime when they see no news, inferring that there *might* have been bad news that was censored. However, when the media does not report politically-relevant news, citizens cannot distinguish whether there was news that the ruler censored, or whether there was simply no politically-relevant news. In equilibrium, provided that the direct censorship cost is not too high, there exists a unique threshold level of news such that the ruler censors a news event if and only if it is worse than that threshold. At the threshold, the gains from the reduced probability of revolution just equal the direct cost of censorship.

Moreover, when the media is more likely to uncover politically-relevant news, the news must be worse for a ruler to censor, as citizens update to conclude that an absence of news was more likely due to censorship. Thus, in a country with a vibrant media that often uncovers news, a ruler censors only very bad news; while in a country with a dormant, incompetent media, a ruler censors even modestly bad news. So, too, when there is more uncertainty about the possible news, a ruler ceases to censor marginally bad news in order to prevent citizens from drawing inferences that the news could have been far worse.

This initial analysis presumes that a ruler only weighs the immediate consequences of his censorship strategy. In particular, a ruler does not internalize that even though freer media is risky because it raises the probability of revolt following bad news, the gains from improving citizens' trust in media may offset those risks by causing citizens to update less negatively following no news. This observation leads us to consider a ruler who can credibly set up institutions that enforce a censorship law that allow the ruler to commit to censoring at other than the equilibrium level of censorship. By delegating censorship decisions to bureaucrats and threatening punishment if they deviate from censorship laws, a ruler might be able to commit to censoring slightly more or less than what he otherwise would without those institutions. Still, large deviations from the equilibrium level of censorship absent such commitment are likely infeasible—for example, a bureaucrat would have strong incentives to censor very bad news that would inevitably result in revolt and hence severe punishment of those associated with the ruler, or a ruler would find other ways to get good news out. We address whether and when a ruler would be better off if he could commit to censoring slightly more or less than he does in equilibrium.

Remarkably, we find that from an ex-ante perspective, a ruler would *always* be strictly better off if he censored slightly *less* than he does in the equilibrium where he cannot commit—a ruler would *always benefit* from a slightly *freer press*. This result reflects that a ruler's equilibrium tradeoff equates the marginal costs and benefits of censoring, *ignoring* the impact of his censorship cutoff on how citizens update when the media does not report politically-relevant news. Paradoxically, were a ruler to censor slightly less and censorship is even slightly costly, citizens would draw *more favorable* inferences about the regime following no news, and hence would be less likely to revolt. In particular, we show that (1) the likelihood citizens revolt following no news is a single-peaked function of a ruler's censorship cutoff, and (2) the equilibrium censorship cutoff is always to the left of the peak whenever censorship is costly.

Since the ruler censors more when censorship is less costly, the probability of revolution in the absence of news is *maximized* at the censorship cutoff chosen when censorship is *costless*. One might then conjecture that this cutoff maximizes or minimizes a ruler's ex-ante expected utility. In fact, when censorship is costless, marginal changes in the equilibrium cutoff have no effect on how citizens update, and hence have no effect on a ruler's welfare. However, we show that this cutoff is an inflection point of the ruler's ex-ante expected utility.

We then contrast a ruler's welfare from censoring everything versus censoring nothing. Ironically, when censorship is costless, and the revolution payoff is low (or citizens' priors are that the status quo is good), then censoring everything is better than censoring nothing. In contrast, when the revolution payoff is high (or citizens' priors are that the status quo is bad), then censoring nothing is better than censoring everything. Intuitively, if the potential revolution payoff is high, then revolution is likely unless a ruler can provide good public news that convinces citizens that the status quo is also pretty good. Conversely, if citizens believe that a "successful" revolution will result in anarchy, citizens must update very negatively about the status quo to revolt, and concealing all public news reduces the chances that this happens.

These welfare results have implications for how a ruler views changes in the cost of censorship or in the media's ability to uncover news. One's instinct might be that a ruler *must be harmed* by new technologies such as the Internet or cell phones that raise censorship costs and reduce the probability that censorship succeeds, or by the entry of a media organization such as Al

Jazeera that uncovers more news. However, our earlier analysis suggests that such changes could *benefit* a ruler: we showed that a ruler would be better off if he could commit to censoring less, and higher censorship costs or a more active media cause a ruler to censor less in equilibrium.

In fact, the consequences of higher censorship costs or of a more active media for a ruler's welfare are subtle. We prove that if censorship is almost costless, then slight increases in censorship costs *always harm* a ruler. Intuitively, when censorship is inexpensive, (1) a ruler censors a lot, and hence is likely to incur the incremental censorship cost; but (2) the probability of revolt given no news is insensitive to a marginal reduction in censorship. However, when censorship costs are higher, a ruler is less likely to censor, and the probability of revolt is more sensitive to the censorship threshold. In particular, there exists a threshold such that increases in censorship costs benefit a ruler (increase his ex ante expected payoff) if and only if censorship costs exceed that threshold.

Finally, we prove that when censorship is inexpensive, and (a) the revolution payoff is high, then a ruler prefers a very active media that uncovers almost all news to a very passive media that uncovers almost nothing; but if, instead, (b) the revolution payoff is lower, then a ruler's preferences are reversed. Intuitively, when revolution payoffs are high, a ruler *values* an active media that might uncover *good* news about the status quo that then forestalls a revolution; but when revolution payoffs are low, a ruler *fears* an active media that might uncover *bad* news about the status quo that then precipitates a revolution.

1.1 Literature Review

Our paper is related to the research on formal models of media freedom. Besley and Prat (2006) consider a setting in which an incumbent office holder is either good or bad; and if an incumbent is bad, media outlets may receive a signal conveying this information. If media outlets become informed, the incumbent chooses whether to bribe them not to reveal bad information; and given the news reported, voters decide whether to re-elect the incumbent. An incumbent bribes the media if and only if the costs of bribery (which reflect the number of outlets that must be bribed, and transactions costs) are not too large. The authors endogenize rent extraction by a bad incumbent, where greater rent extraction raises the probability that the media catch the incumbent. Their binary signal and action structure significantly simplifies equilibrium characterizations. However, this binary structure is not rich enough to permit the analysis of the effects of commitment (and hence of censorship costs and media strength) on a ruler's/incumbent's welfare that we provide.

In Egorov, Guriev, and Sonin (2009), a ruler can manipulate a binary (good or bad) signal sent by media about the regime's economic performance that directly affects the likelihood citizens revolt, and indirectly affects this likelihood via the effect on incentives of bureaucrats to exert effort that increases taxable economic output. The authors assume that the ruler does not see the actual economic performance, but, rather, only observes the possibly distorted media reports. As a result, a ruler trades off between inducing bureaucrats to perform well and directly discouraging citizens from revolting (see also Lorentzen 2012).

Edmond (2011) studies information manipulation in dictatorships in which a revolution succeeds if and only if the measure of citizens revolting exceeds the regime's "strength". Citizens decide whether to revolt after receiving private signals about the regime's strength, and the regime can take a costly hidden action that increases these private signals. In equilibrium, the probability of successful revolution falls discontinuously from one to zero at some regime strength threshold. Edmond's focus is on altering a signal, rather than concealing it, and he does not consider how commitment might affect ruler welfare.

Gehlbach and Sonin (2011) investigate media bias caused by a government that cares both about inducing citizens to take an action and generating advertising revenue. Their model features two states of the world, about which citizens can obtain a costly binary signal. Citizens choose (1) whether to watch the media, thereby observing a signal, and (2) whether to take action 0 or 1. Citizens want their action to match the state, but the government always wants citizens both to take action 1 and to watch the news in order to generate advertising revenue. They show that media bias is greater when (a) the government owns the media and hence does not have to bribe the media, (b) the (exogenous) costs of bribing private media is higher, and (c) the government cares less about advertising revenue. They also investigate whether the government should nationalize the media.

Our finding that a ruler censors sufficiently bad news (Proposition 2) has an analogue in the literature on the disclosure of accounting information (Dye 1985; Verrecchia 1983). However, the limited results in that literature on the effects of pre-commitment are derived in the context of information disclosure in oligopoly markets, with correspondingly different predictions (see Beyer et al. 2010 and Verrecchia 2001 for reviews). There are related analyses of the disclosure of information by an expert/advisor to a decision maker. For example, in Che and Kartik (2009), an advisor must decide whether to reveal his signal to a decision maker. The advisor and decision maker share the same preferences, but their priors over the state of the world have different means, so that the advisor's preferred action always differs by a constant from that of the decision maker. This leads the advisor to censor a bounded interval of signals.

A more distantly-related literature looks at the incentives of the media to selectively report news about candidates that affects how citizens update, and hence electoral outcomes. In Duggan and Martinelli (2011), Anderson and McLaren (2010), or Balan et al. (2004), the incentives of the media outlets to selectively report news devolves from their partisan views, which lead them to have a preferred candidate or policy outcome; in Bernhardt et al. (2008), media compete in their news mix for audiences that value hearing news that conforms with their views; in Chan and Suen (2008), media outlets commit to binary editorial recommendation cutoffs on the state of nature for recommending a left or right party that maximize their viewers' welfare, and these cutoffs feed back to influence party policy choice; and in Gentzkow and Shapiro (2006), the media care about reputation, which can lead to censoring of stories that do not conform with reader expectations.[1] This literature focuses on consumer choice of media outlet, the competition between media outlets for audiences, including the effects of mergers and merger policy (Anderson and McLaren, or Balan et al.), and the role of endogenous media bias on electoral outcomes (Bernhardt et al., Duggan and Martinelli) or policy outcomes (Anderson and McLaren).

2 Model

Authoritarian states censor news to avoid revolution. Thus, we need to specify how public information (news) influences citizens' decisions to revolt. The public nature of news suggests that its most salient effect is on citizens' coordination. To analyze strategic interactions among citizens, we modify the game analyzed in Shadmehr and Bernhardt (2011) by adding an uncertain element to citizen payoffs about which they may receive a public signal. However, we emphasize that our censorship results are robust to alternative ways of modeling citizens' interactions provided they yield the minimal requirements that better news reduces the likelihood of revolution; and that there is almost always a revolution following extremely bad news, but almost never a revolution following extremely good news.

Two representative citizens, A and B, can challenge a ruler by mounting a revolution. A revolution succeeds if and only if both citizens revolt; otherwise, the status quo prevails. The expected value to citizens of the status quo is

[1] For a model where voters do not update rationally about media bias see Mullainathan and Shleifer (2005).

θ, and their expected payoff from successful revolution is R. If a citizen revolts and revolution fails, she receives the status quo payoff minus an expected punishment cost μ.[2] See Figure 1. We normalize the net value to a ruler of preserving the status quo, i.e., of preventing successful revolution, to 1, and his payoff from a successful revolution to 0.

	revolt	*no revolt*
revolt	R, R	$\theta - \mu, \theta$
no revolt	$\theta, \theta - \mu$	θ, θ

Figure 1: Citizen Payoffs.

Citizens receive information that leads them to update about the status quo payoff θ.[3] More specifically, we suppose that the status quo payoff consists of two independent components, $\theta = \gamma + s$, where $\gamma \sim f$ and $s \sim g$, with f and g being strictly positive, continuously differentiable densities on $\mathbb{R}$. Without loss of generality, we normalize the means of s and γ to zero because only considerations of the net expected payoff from revolt (vs. the status quo) determine whether citizens revolt. Each citizen $i \in \{A, B\}$ receives a noisy private signal s^i about s that is independent of γ. We characterize the equilibrium behavior of citizens (Lemma 1) using the following minimal structure on citizen signals:

Assumption 1 *s^i, s^{-i}, and s are strictly affiliated with a strictly positively, continuously differentiable density on $\mathbb{R}^3$. Let σ_ν^2 be the variance of $s^i|s$ and $s^{-i}|s$. For every i and k,*

$$(a)\ \lim_{s^i \to \infty} E[s|s^{-i} < k, s^i] = \infty,\ \ \lim_{s^i \to -\infty} E[s|s^{-i} < k, s^i] = -\infty$$

$$(b)\ \lim_{s^i \to \infty} Pr(s^{-i} < k|s^i) = 0,\ \ \lim_{s^i \to -\infty} Pr(s^{-i} < k|s^i) = 1$$

$$(c)\ \lim_{\sigma_\nu \to 0} Pr(s^{-i} < k|s^i = k) = \frac{1}{2},\ \ \lim_{\sigma_\nu \to 0} E[s|s^i = k, s^{-i} < k] = k.$$

[2]Thus, μ is the probability a failed revolter is caught times the punishment.

[3]Outcomes would be identical if citizens received signals about R because optimal actions hinge on the citizens' *expected difference*, $E[\theta - R]$, not the *source* of the difference.

Assumption 1 is obviously satisfied by the additive, normal noise signal structure with $s^i = s + \nu^i$, where the error ν^i is normally distributed, $\nu^i \sim N(0,\ \sigma_\nu)$, and is independent from both s and the error ν^{-i} in the other citizen's signal.

Citizens do not directly observe γ. However, with probability q, the media and ruler observe the realization of γ; while with probability $1-q$ either γ is degenerate, equal to zero, or equivalently its realization is unobserved by the media. Thus, q has two interpretations: (1) the probability a news event occurs, and (2) the probability the media learn about an event, and hence can report its payoff consequences γ. This latter interpretation links q with the vibrancy of the media sector, in particular, its ability to uncover news. When the media observes γ and the ruler does not censor, the media publicly reveals γ to all citizens.[4] However, a ruler can censor the media at a cost $c \in [0,\ 1)$, preventing the media from conveying γ to any citizen. That is, if the ruler censors the media, then citizens cannot even discern whether a politically-relevant news event occurred. For example, if a ruler censors a prison massacre of political prisoners, citizens do not learn about it.

Given their information about the status quo, which includes γ only if the media learned about it and the ruler did not censor, citizens decide whether to revolt. Thus, a larger R can reflect both (a) a society in which the revolution payoff is higher; and/or (b) a society in which, from an ex-ante perspective, citizens have a lower assessment of the status quo.

The timing is as follows. First, s and γ are realized, and citizens observe their private signals. Then with probability q the media learn γ, which they publicly report to citizens unless the ruler, who also observes the media's information, censors it. After the ruler makes his censorship decision, citizens decide whether to revolt. Finally, payoffs are realized.

A strategy for citizen i is a function σ_i mapping i's private signal s^i and

[4] Our findings are qualitatively unaffected if the media and the ruler observe γ with noise; or if citizens observe the media's report with noise.

any public information into a decision about whether to revolt, where $\sigma_i = 1$ indicates that citizen i revolts, and $\sigma_i = 0$ indicates that i does not. A strategy for the ruler is a function σ_r mapping γ into a decision about whether to censor, where $\sigma_r(\gamma) = 1$ indicates that the ruler censors γ, and $\sigma_r(\gamma) = 0$ indicates that he does not. The equilibrium concept is Bayes Nash, i.e., an equilibrium is a strategy profile, $(\sigma_A^*, \sigma_B^*, \sigma_r^*)$, of mutual best responses, where citizens update according to the Bayes' rule.

Citizens' Game. We first analyze the choices by citizens of whether to revolt. Because optimal actions hinge on the public information *difference* between the status quo and revolution payoffs, the role of public information about θ in citizen decisions is indistinguishable from that of R. Let $\Omega \in \{\gamma, \emptyset\}$ be the public information about θ, where $\emptyset$ indicates that the citizens do not observe any news. Then $\rho(R, \Omega) \equiv E[R - \theta|\Omega]$ is the citizens' public knowledge about the expected value of $R - \theta$. In this setting, a citizen's natural strategy takes a cutoff form: citizen i revolts if and only if his private signal is less than some threshold $k^i(\rho)$. We focus on this class of strategies in our analysis.

Given his private signal about s and the public information $\Omega \in \{\gamma, \emptyset\}$ about γ, a citizen decides whether to revolt. Lemma 1 shows that the citizens' (sub)game features a unique equilibrium in finite-cutoff strategies when the noise in their private signals is vanishingly small.[5]

Lemma 1 *Suppose the noise in the citizens' private signals is vanishingly small:* $\sigma_\nu^2 \to 0$*. There is a unique, symmetric equilibrium in finite-cutoff strategies, characterized by the equilibrium cutoff* $k = \rho(R, \Omega) - \mu$*, where* $\rho(R, \Omega) = E[R - \theta|\Omega]$ *is the citizens' public knowledge about the expected value of* $R - \theta$*.*

[5] With noisier signals, there can be two equilibria in finite-cutoff strategies. The equilibrium featuring more revolution (the one consistent with our equilibrium) is the sole stable one (Shadmehr and Bernhardt 2011). Strategic behavior is qualitatively unaffected as long as the ex-ante revolution payoff is not so low that the probability of revolution given no news can drop discontinuously to zero as a function of the censorship cutoff. Lemma 1 looks similar to uniqueness results in Carlsson and van Damme (1993) and Morris and Shin (2003). However, our game does not feature two-sided limit dominance, which is central to their arguments.

In addition to this finite-cutoff equilibrium, there is a unique equilibrium in infinite-cutoff strategies in which citizens never revolt, i.e., $k = -\infty$. The finite-cutoff equilibrium yields higher citizen welfare than the no-revolution equilibrium. Moreover, as recent events highlight, revolutions occur in practice. Thus, the no-revolution equilibrium is an implausible candidate for describing the real world. This leads us to assume that citizens always coordinate on the finite-cutoff equilibrium. Otherwise, citizens could coordinate on different equilibria after different public news realizations so that, for example, if $\gamma = 0.01$, citizens coordinate on the equilibrium that features revolution (and high welfare), but if $\gamma = -1$ or $\gamma = -10$, citizens coordinate on the no-revolution, low-welfare equilibrium. Such equilibria create perverse censorship incentives, possibly leading a ruler to censor $\gamma = 0.01$, but not $\gamma = -10$.

A revolution succeeds if and only if both citizens revolt, i.e., if and only if both citizens receive a private signal below their revolution threshold: $s^i < k(\rho)$ and $s^{-i} < k(\rho)$. Moreover, as the noise goes to zero, i.e., as $\sigma_\nu^2 \to 0$, signals approaches s, i.e., $\lim_{\sigma_\nu^2 \to 0} s^i = s$. Thus, the probability of a successful revolution, $P(\rho)$, is the probability that $s < k(\rho)$. That is,

$$P(\rho) = G(k(\rho)), \text{ where } G(s) \text{ is the cdf of } s. \tag{1}$$

Good news about the status quo is any news that raises citizens' expectations of the status quo payoff θ, while bad news is the opposite. Recall that $\rho(R, \Omega) = E[R - \theta | \Omega]$ is the citizens' public knowledge about the expected payoff difference between revolution and the status quo. Thus, bad news increases ρ. Proposition 1, which directly follows from Lemma 1 and equation (1), states the properties of the equilibrium to the citizens' game that we use in this paper.

Proposition 1 *Bad news about the status quo raises the likelihood of revolution,* $\frac{P(\rho)}{\partial \rho} > 0$. *Moreover, citizens almost always revolt following extremely bad news,* $\lim_{\rho \to \infty} P(\rho) = 1$*; and almost never revolt following extremely good news about the status quo,* $\lim_{\rho \to -\infty} P(\rho) = 0$.

One can show that analogous results would arise in settings where citizens know their payoffs, but are uncertain about the probability a revolution would succeed if a fraction r of citizens revolt, or equivalently, about the regime's ability to suppress revolt (Boix and Svolik 2009; Edmond 2011); or in settings with standard global games structure that feature super-modularity (Persson and Tabellini 2009) or private value structure (Bueno de Mesquita 2010).

Censorship Equilibrium. In equilibrium, a ruler's censorship strategy takes a cutoff form in which he censors any news that is worse than a threshold $\bar{\gamma}$ (Lemma 4 in the Appendix). The ruler censors news event γ whenever the likelihood of revolution following that news P_γ exceeds the probability of revolution following no news $P_\emptyset$ by a margin that exceeds the censorship cost c. Because in equilibrium, citizens' beliefs must be consistent with the ruler's equilibrium strategy, $P_\emptyset$ also depends on the equilibrium censorship cutoff $\bar{\gamma}_e$. Thus, $\bar{\gamma}_e$ is an equilibrium strategy of the ruler if and only if

$$P_\gamma - P_\emptyset(\bar{\gamma}_e) > c \text{ for all } \gamma < \bar{\gamma}_e, \text{ and } P_\gamma - P_\emptyset(\bar{\gamma}_e) \leq c \text{ for all } \gamma \geq \bar{\gamma}_e. \qquad (2)$$

The equilibrium threshold depends on the likelihood of revolution following no news, $P_\emptyset(\bar{\gamma}_e)$, which, in turn, depends on how citizens update when they do not observe any news, i.e., it depends on $E[\theta|\emptyset, \bar{\gamma}_e]$, where citizen beliefs are consistent with the ruler's strategy $\bar{\gamma}_e$. When a ruler chooses a censorship cutoff $\bar{\gamma}$, citizen beliefs about the value of γ following no news depend on $\bar{\gamma}$ via Bayes' rule:

$$\begin{aligned} E[\gamma|\emptyset, \bar{\gamma}] &= \frac{1-q}{1-q+q\ F(\bar{\gamma})}\ E[\gamma] + \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma<\bar{\gamma}] \\ &= \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma<\bar{\gamma}] \ = \ \frac{q}{1-q+q\ F(\bar{\gamma})} \int_{-\infty}^{\bar{\gamma}} \gamma f(\gamma) d\gamma, \qquad (3) \end{aligned}$$

where f and F are the pdf and cdf of the prior distribution over γ, and the second equality reflects the normalization of $E[\gamma]$ to zero. Both when a ruler *always* censors, i.e., when $\bar{\gamma} = \infty$, and when he *never* censors, i.e., when $\bar{\gamma} =$

$-\infty$, not seeing news conveys no information about γ, and hence $E[\gamma|\emptyset, \bar{\gamma} = \pm\infty] = E[\gamma] = 0$. If, instead, the ruler censors with positive probability less than one, citizens update negatively when they do not receive news, because there may have been bad news that the ruler censored: $E[\gamma|\emptyset, \bar{\gamma}] < E[\gamma] = 0$.

Figure 2 illustrates two important features describing how citizens update following no news: (1) how $E[\theta|\emptyset, \bar{\gamma}]$ varies with the censorship threshold $\bar{\gamma}$, and (2) how citizens would update were they to see the threshold news $\bar{\gamma}$ instead of no news. As a ruler censors more, citizens update more negatively following no public news, up to the point where all news worse than $\bar{\gamma}_m$ is censored, where $\bar{\gamma}_m$ solves $E[\gamma|\emptyset, \bar{\gamma}_m] = \bar{\gamma}_m$. Intuitively, when the ruler only censors extremely bad news, citizens do not update very negatively following no public news, because with high probability the lack of publicly reported news was due to there being no politically-relevant news event, which is not informative about γ. Thus, $\lim_{\bar{\gamma}\to-\infty} E[\gamma|\emptyset, \bar{\gamma}] = E[\gamma] = 0$. Raising the censorship cutoff $\bar{\gamma}$ raises the probability that, conditional on not observing news, censorship occurred, but it also raises the expectation of γ when no news is reported. The first effect reduces the expectation of γ when no news is reported, while the second one raises it.

That the minimizer of $E[\gamma|\emptyset, \bar{\gamma}]$ is at $E[\gamma|\emptyset, \bar{\gamma}_m] = \bar{\gamma}_m$ reflects the relationship between the average, $E[\gamma|\gamma < \bar{\gamma}]$, and the marginal, $\bar{\gamma}$, of a variable. When a ruler censors only bad news, the marginal news to be censored $\bar{\gamma}$ is worse than the average $E[\gamma|\emptyset, \bar{\gamma}]$, and hence censoring slightly more than the threshold news *lowers* citizens' expectations following no news. As a ruler raises $\bar{\gamma}$, these expectations fall as long as $\bar{\gamma} < E[\gamma|\emptyset, \bar{\gamma}]$. Once the censorship cutoff is raised past $E[\gamma|\emptyset, \bar{\gamma}]$, censoring more begins to raise citizen beliefs following no news. Since no reported news is bad news, it follows that $\bar{\gamma}_m = E[\gamma|\emptyset, \bar{\gamma}_m] < E[\gamma] = 0$.

These features of citizen beliefs imply that censoring extremely bad news

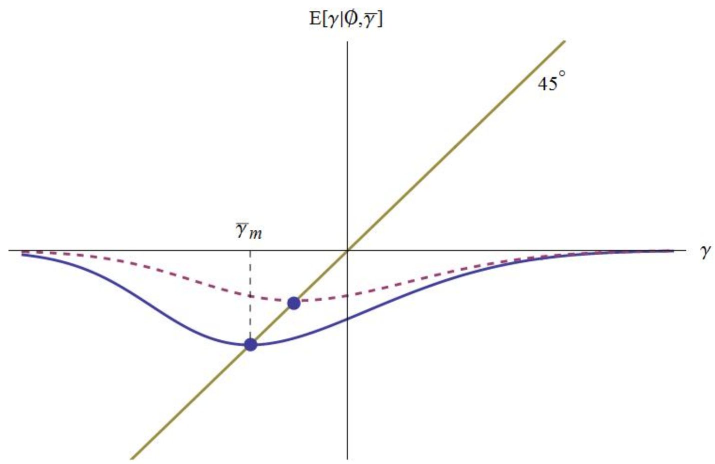

Figure 2: $E[\gamma|\emptyset, \bar{\gamma}]$ as a function of $\bar{\gamma}$. The solid curve corresponds to $q = 0.9$ and the dashed one corresponds to $q = 0.7$.

benefits a ruler because citizens update less negatively following no news than following extremely bad news. As a ruler censors more, this gain falls until it becomes zero at $\bar{\gamma} = \bar{\gamma}_m$. Beyond this point, a ruler is better off allowing the threshold news $\bar{\gamma}$ to reach citizens rather than censoring it. It follows that as long as the censorship cost is less than the payoff from preventing successful revolution, there is a unique, finite equilibrium censorship cutoff, $\bar{\gamma}_e \leq \bar{\gamma}_m < 0$:

Proposition 2 *If $c < 1$, then there is a unique, finite equilibrium censorship cutoff $\bar{\gamma}_e$ such that the ruler censors all news $\gamma < \bar{\gamma}_e$, where $\bar{\gamma}_e$ solves $P_{\bar{\gamma}_e} - P_{\emptyset}(\bar{\gamma}_e) = c$. Otherwise, the ruler does not censor, i.e., $\bar{\gamma}_e = -\infty$.*

The next corollary highlights that it is the structure of citizen beliefs following no news, and not the cost of censorship, that discourages a ruler from censoring all bad news. It reflects that revealing modestly bad news that slightly lowers citizen beliefs below the prior (i.e., $\bar{\gamma}_m < \gamma < E[\gamma] = 0$) is still better for a ruler than no news, because news that is concealed *could* be far worse.

Corollary 1 *Even when censorship is costless, a ruler never censors all bad news. That is, $\bar{\gamma}_e(c = 0) = \bar{\gamma}_m < E[\gamma] = 0$.*

To simplify presentation, we focus our exposition on the interesting case where the censorship cost does not exceed the ruler's net payoff from preventing revolution, i.e., $c < 1$. We next derive how the primitives of the economy affect which news events are censored.

Proposition 3 *Increases in (a) the costs of censorship or (b) the likelihood q that the media observes/uncovers a politically-relevant news event, both reduce a ruler's equilibrium censorship cutoff, $\bar{\gamma}_e$. Moreover, suppose for $\bar{\gamma} < 0$, $F(\bar{\gamma})$ and $E[\gamma|\gamma < \bar{\gamma}]$ are decreasing in the variance of γ, σ_0^2. Then (c) increases in σ_0^2 (which captures the amount of potential news about γ), reduce a ruler's equilibrium censorship cutoff, $\bar{\gamma}_e$.*

That higher censorship costs reduce censorship is obvious. A ruler also censors less in a country with a more vibrant and competent media that is more likely to uncover news. This is because for a given censorship cutoff, with a more active media (a higher q), when no news is reported, it is more likely due to censorship, causing citizens to update more negatively. So, too, a ruler censors less when there is more potential news: a higher σ_0^2 causes citizens to update more negatively when there is no news because there is more probability mass on bad tail news about the status quo. Therefore, the ruler ceases to censor more marginally bad news to prevent citizens from drawing inferences that the news could have been far worse.

We next describe how the extreme limits of almost completely competent and incompetent media affect a ruler's censorship decisions.

Proposition 4 *A ruler censors almost nothing if there is almost always politically-relevant news that a vibrant media almost always uncovers: $\lim_{q\to 1} \bar{\gamma}_e(q) = -\infty$. Conversely, a ruler censors almost all bad news if either there is almost never politically-relevant news or the media is almost completely passive, and censorship is almost costless: $\lim_{c,q\to 0} \bar{\gamma}_e(q, c) = 0$.*

3 Commitment and Ruler Welfare

A ruler's equilibrium censorship decisions compare the gains from reducing the likelihood of revolution when a news event is censored with the costs of censorship, ignoring the consequences for the level of citizens' trust in the media. We now show that a ruler would always be better off were he able to commit to censoring slightly less than he does in equilibrium, as citizens would then update more positively when no news is reported.

We first calculate a ruler's ex-ante expected utility from a given censorship cutoff $\bar{\gamma}$:

$$
\begin{aligned}
W(\bar{\gamma}, c, q) &\equiv q\left(F(\bar{\gamma})\ (1 - P_{\emptyset}(\bar{\gamma}) - c) + (1 - F(\bar{\gamma}))\ \frac{\int_{\bar{\gamma}}^{\infty}(1 - P_{\gamma})dF(\gamma)}{1 - F(\bar{\gamma})}\right) \\
&\quad +(1 - q)\,(1 - P_{\emptyset}(\bar{\gamma})) \\
&= 1 - \left([qF(\bar{\gamma}) + (1 - q)]P_{\emptyset}(\bar{\gamma}) + q\int_{\bar{\gamma}}^{\infty} P_{\gamma}dF(\gamma) + qF(\bar{\gamma})c\right) \quad (4)
\end{aligned}
$$

With probability $1 - q$ there is no news, and the ruler's expected payoff is $1 - P_{\emptyset}(\bar{\gamma})$. With probability q there is some news γ, which the ruler censors whenever $\gamma < \bar{\gamma}$. If he censors the news, which happens with probability $F(\bar{\gamma})$, then his payoff is $1 - P_{\emptyset}(\bar{\gamma}) - c$. If he does not censor, which happens with probability $1 - F(\bar{\gamma})$, then his expected payoff is $\frac{\int_{\bar{\gamma}}^{\infty}(1-P_{\gamma})dF(\gamma)}{1-F(\bar{\gamma})}$.

Citizen beliefs about the status quo payoff in the absence of news vary with a ruler's censorship policy. However, the ruler ignores this in his decision-making: his censorship decisions maximize his expected utility *given* citizens' (equilibrium) beliefs. If, instead, a ruler can *ex-ante* commit to a censorship level, he can "internalize" the effects of his chosen censorship rule on how citizens update. For example, a ruler may be able to do this by passing censorship laws and delegating enforcement to bureaucrats, threatening them with punishment if they adopt a different cutoff. Still, a ruler's ability to do this is limited. For example, a bureaucrat who fails to censor very bad public news is likely to be punished (either by the ruler or by citizens following the likely resulting successful revolution); and a ruler with particularly good news may feel compelled to reveal it. Thus, such commitment mechanisms only allow a ruler to commit to censorship levels that differ slightly from what he would do without those mechanisms. The question then becomes: Under what circumstances would a ruler want to increase or decrease his censorship cutoff marginally from its equilibrium level?

Remarkably, we now show that if censorship has *any* costs, i.e., if $c > 0$, then a ruler's equilibrium censorship choice is *never* optimal from an ex-ante perspective.

Proposition 5 *If censorship is costly, a ruler's ex-ante expected utility would be raised if he could commit to censoring slightly less than he does in equilibrium:* $\frac{dW(\bar{\gamma})}{d\bar{\gamma}}\Big|_{\bar{\gamma}=\bar{\gamma}_e} < 0$.

Small reductions in the censorship cutoff below the equilibrium threshold have three effects: (1) When there is news, citizens now see marginal news realizations that are slightly below the equilibrium cutoff $\bar{\gamma}_e$. For these marginal news events, the probability of revolution rises from $P_{\emptyset}(\bar{\gamma})$ to $P_{\bar{\gamma}}$. However, (2) the ruler does not incur censorship costs. At equilibrium, fixing the beliefs of citizens, these two effects just offset each other, as a ruler's equilibrium cutoff $\bar{\gamma}_e$ equates these marginal costs and benefits of censorship. Thus, the net impact on a ruler's ex-ante expected utility equals the third effect: (3) A slight reduction in the censorship cutoff reduces the probability of revolution $P_{\emptyset}(\bar{\gamma})$ when there is no news. To see this, recall that (a) as a ruler censors more news, citizens update more negatively following no news, up to the point where the ruler censors all news worse than $\bar{\gamma}_m$ (see Figure 2); and (b) $\bar{\gamma}_e < \bar{\gamma}_m$ when censorship is costly. Therefore, the probability of revolution $P_{\emptyset}(\bar{\gamma})$ is rising in $\bar{\gamma}$ at the equilibrium cutoff $\bar{\gamma}_e$, and hence the net effect of marginally reducing the censorship cutoff is to raise a ruler's ex-ante expected utility.

The sole restriction in the proposition is that censorship be costly. The reason is that when censorship is costless, citizens' beliefs become locally insensitive to changes in the ruler's equilibrium censorship cutoff: When $c = 0$, we have $\bar{\gamma}_e = \bar{\gamma}_m$, i.e., the ruler's equilibrium level of censorship minimizes $E[\gamma|\emptyset, \bar{\gamma}]$. Hence, marginal changes in $\bar{\gamma}$ do not affect citizen beliefs about $E[\gamma|\emptyset, \bar{\gamma}]$ (see Lemma 5 and Figure 2), i.e., $\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}} = 0$. This directly implies that $\frac{dW(\bar{\gamma}_e)}{d\bar{\gamma}} = 0$. This might lead one to conjecture that this censorship cutoff minimizes or maximizes the ruler's ex-ante expected utility. Proposition 6 shows that this is not so:

Proposition 6 *When censorship is costless, i.e., $c = 0$, a ruler's equilibrium censorship cutoff is an inflection point of his ex-ante expected welfare. That is, $\frac{dW(\bar{\gamma}_e)}{d\bar{\gamma}} = 0$ and $\frac{d^2W(\bar{\gamma}_e)}{d\bar{\gamma}^2} = 0$.*

When censorship is costless, the equilibrium level of censorship *maximizes* the probability of revolution when politically-relevant news is not publicly reported. Proposition 6 says that deviating marginally from $\bar{\gamma}_e = \bar{\gamma}_m$ does not affect a ruler's welfare, but it says nothing about the welfare impacts of censorship cutoffs $\bar{\gamma}$ that are further from $\bar{\gamma}_m$. To understand those impacts, recall that a higher level of R corresponds to a higher revolution payoff and/or a lower ex-ante assessment of status quo payoffs. Given any censorship cutoff $\bar{\gamma} \neq \bar{\gamma}_e$, Proposition 7 identifies a critical $\bar{\gamma}$-contingent cutoff $R^*(\bar{\gamma})$ on R that determines whether reducing the censorship cutoff from $\bar{\gamma}$ raises or lowers a ruler's ex-ante expected utility. In particular, when R is higher (so revolution is more likely absent favorable public news), a ruler would gain if he could commit to censoring less; and when $R < R^*(\bar{\gamma})$ (so revolution is less likely absent unfavorable public news), a ruler would benefit if he could commit to censoring more.

Proposition 7 *Suppose $G(s)$ is strictly unimodal.*[6] *If censorship is costless*

[6] $G(s)$ is strictly unimodal if and only if there exists a unique s^* such that $G(s)$ is strictly convex for $s < s^*$ and strictly concave for $s > s^*$.

and $\bar{\gamma} \neq \bar{\gamma}_e$, then there exists a critical cutoff $R^*(\bar{\gamma})$ that determines whether increased censorship raises or lowers a ruler's welfare: $\frac{dW(\bar{\gamma};R)}{d\bar{\gamma}} > 0$ if $R < R^*(\bar{\gamma})$ and $\frac{dW(\bar{\gamma};R)}{d\bar{\gamma}} < 0$ if $R > R^*(\bar{\gamma})$.

A similar logic lets us compare a ruler's ex-ante expected utility if censored everything, i.e., $\bar{\gamma} = +\infty$, with his utility if he censored nothing, i.e., $\bar{\gamma} = -\infty$.

Proposition 8 *Suppose $G(s)$ is strictly unimodal and $f(\gamma)$ is symmetric. When R <u>or</u> c are sufficiently large, a ruler's welfare is higher when he censors nothing than when he censors everything: $\lim_{\bar{\gamma}\to\infty} W(\bar{\gamma};R) < \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma};R)$. Conversely, if R <u>and</u> c are sufficiently small, the opposite welfare result obtains: $\lim_{\bar{\gamma}\to\infty} W(\bar{\gamma};R) > \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma};R)$.*

Moreover, if $G(s)$ is symmetric, one can strengthen the proposition: *when $c = 0$, a ruler is better off censoring nothing than censoring everything if and only if $R > \mu$.* When the revolution payoff is high so that citizen and ruler interests are far from being "aligned", then from an ex-ante perspective, the ruler is better off censoring nothing rather than everything. The intuition is that when the revolution payoff is high enough, the ruler would be better off rolling the dice on public news, hoping for good news that would forestall revolution: unimodality of G and symmetry of f imply that when R is high, bad news only marginally increases the likelihood of revolt, but good news can sharply reduce that probability. If, instead, the revolution payoff is low so that citizens are unlikely to revolt and citizen and ruler interests are more aligned, the ruler prefers to censor everything rather than nothing.[7]

[7] Kamenica and Gentzkow (2011) analyze a sender's design of a signal technology that provides a decision maker (receiver) a signal about the state of the world. Our result contrasts with their finding that it can be optimal for the sender to design a signal technology that reveals everything to the receiver if interests are closely aligned, and to reveal nothing if interests are less aligned (p. 2604-6).

4 Censorship Costs, Active Media, and Ruler Welfare

New technologies such as the Internet and cell phones have significantly increased the costs of censorship. The extensive use of these technologies in the Arab Spring, the Green Movement in Iran, and the Orange Movement in Kuwait (e.g., Wheeler 2010) has led many to conclude that technologies that make censorship difficult, i.e., increase censorship costs, always work against dictators. We show that these assessments are not completely correct. To the contrary, higher costs of censorship can sometimes increase a ruler's chance of survival. In fact, these gains can be so high to offset the direct costs of censorship and make the ruler better off. Similarly, one might suspect that dictators always prefer a dormant, passive media that uncovers no politically-relevant news to a vibrant, active media that uncovers almost all news. This conjecture is also wrong. Surprisingly, when the revolution payoff is high so that the ruler faces a serious danger of revolution, he prefers a very active media that uncovers almost all news to a very passive media that reports nothing important.

Higher Censorship Costs. Higher censorship costs have direct and indirect welfare effects. The direct effect reduces a ruler's ex-ante expected utility because a ruler must pay more each time he censors. The indirect effect is that higher censorship costs cause a ruler to censor less in equilibrium. We have established that decreases in the equilibrium level of censorship cause citizens to update less negatively following no news and hence benefit a ruler. It follows that the indirect effect of increases in censorship costs raise a ruler's ex-ante expected utility. Proposition 9 shows that when the costs of censorship are small, the direct effect always dominates, so that marginal increases in censorship costs reduce the ruler's ex-ante expected utility. However, when censorship costs are higher, the indirect effect dominates, so that further increases in censorship costs benefit the ruler.

Proposition 9 *Suppose $G(s)$ is strictly unimodal with a median that does not exceed its mean, $f(\gamma)$ is log-concave, and $R \geq \mu$. Then there exists a critical censorship cost $\hat{c}$ such that increases in a ruler's cost of censorship raise his ex-ante expected utility if and only if $c > \hat{c}$.*

When censorship is costless, i.e., $c = 0$, the equilibrium level of censorship is at the minimum of $E[\gamma|\emptyset, \bar{\gamma}]$, i.e., $\bar{\gamma}_e = \bar{\gamma}_m$. Hence, marginal changes in the extent of censorship do not affect citizen beliefs about the status quo payoff when they do not see news. Therefore, a marginal increase in censorship cost from zero only has the direct effect of imposing positive censorship costs on the ruler. This reasoning extends whenever censorship costs are sufficiently low, in which case the ruler is likely to censor (thereby incurring the higher censorship cost), but the improvement in citizen beliefs about the status quo when they do not see news is modest.

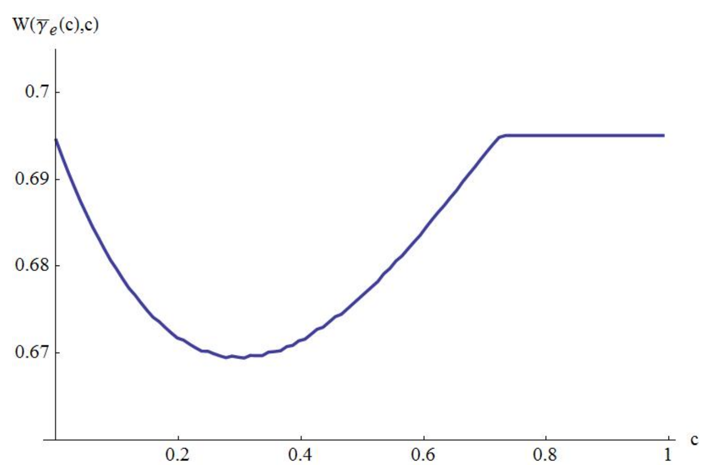

Figure 3: The ruler's ex-ante expected utility, $W(\bar{\gamma}_e(c), c)$, as a function of c.

However, the proposition shows that there is a critical censorship cost threshold such that the indirect effect dominates under the sufficient distributional assumptions placed on G and f: for censorship costs exceeding the threshold, further increases in censorship costs raise the ruler's expected utility due to the less pessimistic updating by citizens, and to the fact that the ruler is less likely to incur the censorship cost. Intuitively, the indirect effect can dominate once censorship costs are intermediate so that $P_{\emptyset}(\bar{\gamma} = \bar{\gamma}_e(c))$ is sensitive to extent of censorship, $\bar{\gamma}$, and the censorship cutoff is low enough that the ruler is not *that* likely to censor (and thus incur the higher direct censorship costs). See Figure 3. Moreover, it follows that there exists a cost threshold $\bar{c} < \hat{c}$ such that once censorship costs reach $\bar{c}$, the ruler would be strictly better off with censorship costs of one, which discourages *all* censorship.

More Active Media. Related reasoning suggests that a ruler may prefer a very active media that uncovers almost all politically-relevant news, to a passive media that uncovers almost nothing. Proposition 10 shows that when revolution payoffs are high so that revolution is likely, a ruler prefers to have a media that uncovers all news to a passive media that uncovers nothing.

Proposition 10 *Suppose $G(s)$ is strictly unimodal and $f(\gamma)$ is symmetric. There exists a critical cutoff R^* such that if $R > R^*$ then a ruler's ex-ante expected utility is higher if he faces an aggressive media that uncovers almost all news, than if he faces a passive media that uncovers almost nothing: if $R > R^*$, then $\lim_{q\to 0} W(\bar{\gamma}_e(q); R) < \lim_{q\to 1} W(\bar{\gamma}_e(q); R)$. Conversely, if $R < R^*$, a ruler's ex-ante expected utility is higher when the situation is reversed so that he faces a very passive media: if $R < R^*$, then $\lim_{q\to 0} W(\bar{\gamma}_e(q); R) > \lim_{q\to 1} W(\bar{\gamma}_e(q); R)$.*

A ruler facing a passive media that uncovers almost nothing is in the same situation as a ruler who commits to censoring everything (when censorship is costless). In both cases, citizens update only marginally negatively when no news is reported. Conversely, a ruler facing a passive media that uncovers almost nothing is in the same situation as a ruler who commits to censor nothing. Then Proposition 8 reveals that a ruler prefers a very active media to a very passive media if and only if the revolution payoff exceeds a critical threshold.

When revolution is likely, i.e., when the revolution payoff appears high relative to the ex-ante status quo payoff, a ruler's best hope for survival is that citizens receive good news about the status quo and update positively. The chances that media uncovers and reports such good news is highest when the media is very active ($q \approx 1$). Of course, an active media may also uncover bad news that further raises the likelihood of revolt, but because this likelihood is already very high, the possible gains from sharply positive public news may exceed any loss from more bad news that can only raise the likelihood of revolt marginally from, say, 98% to 99%. This result offers a rationale for why dictators sometimes relax censorship amidst unfolding revolutions. For example, Milani (2011, 388) notes that the Shah relaxed censorship in the second half of 1978 during the unfolding of the 1979 Iranian Revolution (see also Milani 1994, 117; Arjomand 1984, 115).

Related results would obtain if we allowed for the possibility that a ruler may fail in his efforts to censor a news event. For example, recent technological innovations in communication between citizens via the internet or cell phones have not only raised the cost of censorship, but introduced the possibility that some types of information cannot be censored. For example, a ruler may not be able to conceal beatings of protesters because videos can be taken with cell phones and uploaded to Youtube. One can model this by assuming that if a ruler attempts to censor at a cost c, then he still fails to conceal the news with probability ζ. Our model focuses on $\zeta = 0$, but the effects of $\zeta > 0$ are straightforward to derive. In particular, a larger ζ causes a ruler to reduce his cutoff, as censorship is less effective, and citizens update less negatively following no news because it is more likely that no news occurred (both because the ruler censors less, and because censorship is less likely conditional on no news because it would have required the censorship of bad news to succeed.) Qualitatively, the impact of $\zeta > 0$ is similar to that of a more active media that uncovers a greater fraction q of news events.

5 Conclusion

We investigate the dilemma of an unpopular ruler facing citizens who are deciding whether or not to revolt based on the public and private information that they gather. A ruler can manipulate their information-gathering process by censoring the media, thereby preventing the media from disseminating information about the regime that would raise the likelihood of revolt. Unfortunately for the ruler, citizens take into account a ruler's incentives to conceal bad news. Not knowing whether an absence of news was due to censorship or because the media failed to uncover politically-relevant news, citizens update negatively when no news is reported—"no news" becomes "bad news". And yet, because citizens are not sure whether an absence of news is due to censorship, the ruler can increase his chances of survival by censorship. Indeed, in equilibrium, provided censorship costs do not exceed the ruler's gain from preventing revolution, a unique cutoff determines which news a ruler censors.

Remarkably, we show that a ruler would always like to commit himself to censoring slightly *less* than what he does in equilibrium in order to induce citizens to update less negatively in the absence of news, i.e., so that "no news" becomes "not quite so bad news". In contrast to conventional wisdom, we show that: (1) Higher censorship costs (induced by new communication technologies) can make a ruler better off, increasing his ex-ante chances of survival. Indeed, a ruler's welfare first falls and then rises with the costs of censorship. (2) Ironically, when revolution payoffs are high so that revolution is likely, a ruler is better off with an active media that uncovers almost all news than with a passive one that uncovers almost nothing. In contrast, a ruler prefers a passive media when revolution is unlikely.

Our paper focuses on the censorship of media. However, our model can be reposed to study elections in dictatorships, where the dictator has an incentive to conceal public information. Researchers have pointed out that one role of elections in dictatorships is to send a public signal to potential protesters (Egorov and Sonin 2011), opposition (Rozenas 2010), or government officials (Gehlbach and Simpser 2011) about a regime's support. Electoral fraud aside, a decision of whether to hold an election is related to the decision of which news to censor: A ruler has (possibly noisy) private knowledge of the extent of his support, and a decision to hold an election is equivalent to giving potential protesters, opposition, or government officials a (possibly noisy) signal about this support.

6 Appendix

Proof of Lemma 1: Citizen i revolts if and only if

$$Pr(\sigma_{-i} = 1|s^i, \Omega)\, R + Pr(\sigma_{-i} = 0|s^i, \Omega)\, (E[\theta|s^i, \Omega, \sigma_{-i} = 0] - \mu) - E[\theta|s^i, \Omega] > 0.$$

Since γ and s are independently distributed, the expected net payoff from revolt simplifies to

$$Pr(\sigma_{-i} = 1|s^i, \Omega)\, R + Pr(\sigma_{-i} = 0|s^i, \Omega)\, (E[\gamma|\Omega] + E[s|s^i, \sigma_{-i} = 0] - \mu) - E[\gamma|\Omega] - E[s|s^i].$$

Collecting terms that include $E[\gamma|\Omega]$ with R, write this expected net payoff as

$$Pr(\sigma_{-i} = 1|s^i, \Omega)\, (R - E[\gamma|\Omega]) + Pr(\sigma_{-i} = 0|s^i, \Omega)\, (E[s|s^i, \sigma_{-i} = 0] - \mu) - E[s|s^i].$$

Substituting $\rho(R, \Omega) = E[R - \theta|\Omega]$, we write the expected net payoffs from revolt as

$$Pr(s^{-i} < k^{-i}|s^i)\, \rho(R, E[\gamma|\Omega]) + Pr(s^{-i} \geq k^{-i}|s^i)\, (E[s|s^i, s^{-i} \geq k^{-i}] - \mu) - E[s|s^i].$$

Next, we prove two lemmas.

Lemma 2 *The best response to a cutoff strategy is a cutoff strategy.*

Proof of Lemma 2: Rewrite the expected net payoffs from revolt as

$$\Delta(s^i; k^{-i}) \equiv Pr(s^{-i} < k^{-i}|s^i)\ (\rho(R, E[\gamma|\Omega]) - E[s|s^i, s^{-i} < k^{-i}] + \mu) - \mu.$$

$0 < Pr(s^{-i} < k^{-i}|s^i)$, $\rho(R, E[\gamma|\Omega])$ is finite, and $E[s|s^i, s^{-i} < k^{-i}]$ is strictly increasing in s^i. Thus, if $\Delta(x; k^{-i}) \geq 0$, then $\Delta(s^i; k^{-i}) > 0$ for all $s^i < x$. Thus, $\Delta(s^i; k^{-i}) = 0$ has the single-crossing property. Moreover,

$$\lim_{s^i \to +\infty} E[s|s^i, s^{-i} < k^{-i}] = +\infty, \quad \lim_{s^i \to -\infty} E[s|s^i, s^{-i} < k^{-i}] = -\infty,$$

hence $\lim_{s^i \to +\infty} \Delta(s^i; k^{-i}) < 0$, $\lim_{s^i \to -\infty} \Delta(s^i; k^{-i}) > 0$. These properties together with the continuity of $\Delta(s^i; k^{-i})$ imply that there is a unique $k^i(k^{-i})$ such that $\Delta(s^i; k^{-i}) > 0$ for all $s^i < k^i$, $\Delta(s^i = k^i; k^{-i}) = 0$, and $\Delta(s^i; k^{-i}) < 0$ for all $s^i > k^i$. $\square$

Lemma 3 *When the noise in private signals is vanishingly small, all cutoff equilibria are symmetric.*

Proof of Lemma 3: Suppose to the contrary that an equilibrium exists in which citizens i and $-i$ adopt cutoff strategies with respective cutoffs k^i and k^{-i}, where without loss of generality $k^i > k^{-i}$. Suppose $s^i \in [(k^i + k^{-i})/2, k^i]$. Then, when the noise in private signals is vanishingly small, citizen i is almost sure that $s^{-i} > k^{-i}$ and hence almost sure that citizen $-i$ will not revolt. But then citizen i is almost surely punished if he revolts at s^i, so it is not a best response, a contradiction. $\square$

Since all cutoff strategy equilibria are symmetric, equilibria in which citizens revolt with positive probability are characterized by the zeros (roots) of the symmetric expected net payoff function $\Delta_1(k)$, where citizens adopt the common cutoff $k^{-i} = k^i = k$:

$$\begin{aligned}\Delta_1(k) &\equiv Pr(s^{-i} < k|k)\ \rho(R, \Omega) + Pr(s^{-i} \geq k|k)\ (E[s|k, s^{-i} \geq k] - \mu) - E[s|k] \\ &= Pr(s^{-i} < k|s^i = k)\ (\rho(R, \Omega) - E[s|s^i = k, s^{-i} < k] + \mu) - \mu.\end{aligned}$$

From Assumption 1, $\lim_{\sigma_\nu \to 0} Pr(s^{-i} < k | s^i = k) = \frac{1}{2}$ and $\lim_{\sigma_\nu \to 0} E[s|s^i = k, s^{-i} < k] = k$, so

$$\lim_{\sigma_\nu^i \to 0} \Delta_1(k) = Pr(s^{-i} < k | s^i = k)\ (\rho(R,\Omega) - k + \mu) - \mu = \frac{1}{2}\ (\rho(R,\Omega) - \mu - k),$$

and hence $\Delta_1(k)$ has a unique root at $k = \rho(R,\Omega) - \mu$. □

Lemma 4 *The ruler's equilibrium censorship strategy takes a cutoff form: there exists a* $\bar{\gamma} \in \mathbb{R} \cup \{\pm\infty\}$ *such that* $\sigma_r(\gamma) = 1$ *if and only if* $\gamma < \bar{\gamma}$ *for some* $\bar{\gamma} \in \mathbb{R} \cup \{\pm\infty\}$*, where* $\bar{\gamma} = -\infty$ *corresponds to never censor, and* $\bar{\gamma} = \infty$ *corresponds to censor everything.*

Proof of Lemma 4: The ruler censors γ if and only if his expected utility from censoring, $1 - P_\emptyset - c$, exceeds that of not censoring, $1 - P_\gamma$. That is,

$$\sigma_r(\gamma) = 1 \text{ if and only if } P_\gamma - P_\emptyset > c. \tag{5}$$

P_γ decreases monotonically in γ and $P_\emptyset$ does not depend on γ. Thus, the ruler adopts a cutoff strategy in equilibrium, censoring γ if and only if it is below some critical cutoff, $\bar{\gamma}$. □

Proof of Proposition 2: Note that $c^* \equiv \lim_{\bar{\gamma} \to -\infty} P_{\bar{\gamma}} - P_\emptyset(\bar{\gamma}) = 1^- - 0^+ = 1$. The result follows directly from parts 1 and 2 of the following Lemma together with Proposition 1.

Lemma 5 *Citizen estimates* $E[\gamma|\emptyset, \bar{\gamma}]$ *of* γ *when they see no news and the ruler sets censorship cutoff* $\bar{\gamma}$ *have the following properties:*

1. $E[\gamma|\emptyset, \bar{\gamma}] \leq E[\gamma] = 0$*, with equality only for* $\bar{\gamma} = \pm\infty$*. Moreover,* $E[\gamma|\emptyset, \bar{\gamma}]$ *is continuous in* $\bar{\gamma}$*, with* $\lim_{\bar{\gamma} \to \pm\infty} E[\gamma|\emptyset, \bar{\gamma}] = E[\gamma] = 0$*.*

2. $E[\gamma|\emptyset, \bar{\gamma}]$ *has a unique extremum at* $\bar{\gamma}_m < 0$*, which is a minimum. Further,* $E[\gamma|\emptyset, \bar{\gamma}_m] = \bar{\gamma}_m$*, with* $E[\gamma|\emptyset, \bar{\gamma}] > \bar{\gamma}_m$ *if* $\bar{\gamma} < \bar{\gamma}_m$*, and* $E[\gamma|\emptyset, \bar{\gamma}] < \bar{\gamma}_m$ *if* $\bar{\gamma} > \bar{\gamma}_m$*.*

3. *$E[\gamma|\emptyset,\bar{\gamma};q]$ falls in q for a fixed $\bar{\gamma}$. Thus, its minimizer, $\bar{\gamma}_m$ decreases in q.*

Proof of Lemma 5: Part 1 is immediate from equation (3). To prove part 2, differentiate equation (3) with respect to $\bar{\gamma}$:

$$\begin{aligned}
\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}} &= \frac{q\bar{\gamma}f(\bar{\gamma})\ (1-q+q\ F(\bar{\gamma})) - qf(\bar{\gamma})\ q\int_{-\infty}^{\bar{\gamma}}\gamma f(\gamma)d\gamma}{(1-q+q\ F(\bar{\gamma}))^2} \\
&= \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\left(\bar{\gamma} - \frac{q}{1-q+qF(\bar{\gamma})}\ \int_{-\infty}^{\bar{\gamma}}\gamma f(\gamma)d\gamma\right) \\
&= \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\left(\bar{\gamma} - E[\gamma|\emptyset,\bar{\gamma}]\right) \qquad (6) \\
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\ F(\bar{\gamma})(\bar{\gamma} - E[\gamma|\gamma<\bar{\gamma}])\ \right) \\
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\ F(\bar{\gamma})\left(\frac{1}{F(\bar{\gamma})}\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma\right)\ \right) \\
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\ \int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma\ \right), \qquad (7)
\end{aligned}$$

where the fifth equality follows from integration by parts. From equation (7),

$$\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}} = 0 \text{ if and only if } \ -\frac{1-q}{q}\ \bar{\gamma} = \int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma. \qquad (8)$$

The left-hand side is strictly decreasing and onto, and the right-hand side is strictly increasing. Thus, there exists a unique $\bar{\gamma}_m$ such that $\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}} = 0$ at $\bar{\gamma} = \bar{\gamma}_m$. Moreover, $\lim_{\bar{\gamma}\to\pm\infty} -\frac{1-q}{q}\ \bar{\gamma} = \mp\infty$, $\lim_{\bar{\gamma}\to+\infty}\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma > 0$, and $\lim_{\bar{\gamma}\to-\infty}\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma = 0$. Therefore, $\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}}\Big|_{\bar{\gamma}=\bar{\gamma}_m} < 0$ if $\bar{\gamma} < \bar{\gamma}_m$, and $\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}}\Big|_{\bar{\gamma}=\bar{\gamma}_m} > 0$ if $\bar{\gamma} > \bar{\gamma}_m$. That $\bar{\gamma}_m < 0$ is immediate from the facts that $-\frac{1-q}{q}\ \bar{\gamma} > 0$ if and only if $\bar{\gamma} < 0$, and $\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma > 0$ for all $\gamma \in (-\infty,\infty]$. Finally, from equation (8), $\lim_{q\to 0^+}\bar{\gamma}_m = 0$. This, together with part 1, yields part 2. Part 3 follows from equation (6). Part 4 follows directly from (3). □

Proof of Proposition 3: Note that c does not affect $\Delta P(\bar{\gamma})$, and q and σ_0^2 only affect $P_\emptyset$. In particular, increases in q and σ_0^2 decrease $E[\gamma|\emptyset,\bar{\gamma}]$, and hence raise $P_\emptyset(\bar{\gamma})$, for a fixed $\bar{\gamma}$. To see the claim for σ_0^2 write $E[\gamma|\emptyset,\bar{\gamma}] =$

$\frac{1-q}{1-q+q\ F(\bar{\gamma})}\ E[\gamma] + \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma < \bar{\gamma}] \equiv A(\bar{\gamma})E[\gamma|\gamma < \bar{\gamma}]$. Differentiating, we have $A'E[\gamma|\gamma < \bar{\gamma}] + AE'[\gamma|\gamma < \bar{\gamma}] < 0$, since both terms are negative (A is increasing in $F(\bar{\gamma})$, and $E[\gamma|\gamma < \bar{\gamma}] < 0$, while A is positive, and $E[\gamma|\gamma < \bar{\gamma}]$ is decreasing in σ_0^2 for $\bar{\gamma} < 0$ since increases in σ_0^2 place more probability weight on lower values of γ). Hence, to retrieve equality of $c = \Delta P(\bar{\gamma}_e)$, it must be that $\bar{\gamma}_e$ is reduced. $\square$

Proof of Proposition 4: From equation (3),

$$\lim_{q\to 1} E[\gamma|\emptyset, \bar{\gamma}_e(q)] = \lim_{q\to 1} E[\gamma|\gamma < \bar{\gamma}_e(q)], \text{ and } \lim_{q\to 0} E[\gamma|\emptyset, \bar{\gamma}] = 0.$$

Moreover, from equation (8), $\lim_{q\to 1} \bar{\gamma}_m(q) = -\infty$, and from Proposition 2, $\bar{\gamma}_e \leq \bar{\gamma}_m$, and hence $\lim_{q\to 1} \bar{\gamma}_e(q) = -\infty$. Thus,

$$\lim_{q\to 1} P_\emptyset(\bar{\gamma}_e(q); q) = \lim_{\rho\to\infty} P(\rho) = 1, \text{ and } \lim_{q\to 0} P_\emptyset(\bar{\gamma}_e(q); q) = \lim_{\rho\to R} P(\rho) = P_{\gamma=0}.$$

From the proof of Proposition 2 and Corollary 1, $\lim_{c\to 0} \bar{\gamma}_e(q, c) = \bar{\gamma}_m(q)$, and from equation (8), $\lim_{q\to 0} \bar{\gamma}_m(q) = 0$. $\square$

Proof of Proposition 5: From equation (4),

$$\begin{aligned} \frac{dW}{d\bar{\gamma}} &= -qf(\bar{\gamma})\ P_\emptyset(\bar{\gamma}) - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} + qP_{\bar{\gamma}}f(\bar{\gamma}) - qf(\bar{\gamma})c \\ &= -qf(\bar{\gamma})\ [P_\emptyset(\bar{\gamma}) - P_{\bar{\gamma}} + c] - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} \\ &= qf(\bar{\gamma})\ [(P_{\bar{\gamma}} - P_\emptyset(\bar{\gamma})) - c] - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}}. \end{aligned} \tag{9}$$

From Proposition 2, $P_{\bar{\gamma}_e} - P_\emptyset(\bar{\gamma}_e) - c = 0$. Since $c > 0$, we have $\bar{\gamma}_e < \bar{\gamma}_m$, so $\frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} > 0$. Thus,

$$\frac{dW}{d\bar{\gamma}} = -[qF(\bar{\gamma}_e) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}}\bigg|_{\bar{\gamma}=\bar{\gamma}_e} < 0. \quad \square$$

Proof of Proposition 6: From equation (9),

$$\begin{aligned} \frac{d^2W}{d\bar{\gamma}^2} &= q\frac{df(\bar{\gamma})}{d\bar{\gamma}}\ [P_{\bar{\gamma}} - P_\emptyset(\bar{\gamma}) - c] + qf(\bar{\gamma})\ \left[\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} - \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}}\right] \\ &- qf(\bar{\gamma})\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} - [qF(\bar{\gamma}) + 1 - q]\ \frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2}. \end{aligned} \tag{10}$$

At equilibrium, $P_{\bar{\gamma}_e} - P_\emptyset(\bar{\gamma}_e) = c$. Further, when $c = 0$, $\bar{\gamma}_e = \bar{\gamma}_m$, so $\frac{dP_\emptyset(\bar{\gamma}_e)}{d\bar{\gamma}} = 0$. Thus, equation (10) simplifies to

$$\frac{d^2W}{d\bar{\gamma}^2} = qf(\bar{\gamma})\ \frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} - [qF(\bar{\gamma}) + 1 - q]\ \frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2}. \tag{11}$$

Moreover,

$$\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} = -g(R-\mu-\bar{\gamma}),\ \text{and}\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} = -\frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}}\ g(R-\mu-E[\gamma|\emptyset,\bar{\gamma}]). \tag{12}$$

Thus,

$$\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2} = -\ \frac{d^2E[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}^2}\ g(R-\mu-E[\gamma|\emptyset,\bar{\gamma}]), \tag{13}$$

where we used the fact that when $c = 0$, $\bar{\gamma}_e = \bar{\gamma}_m$, and hence $\frac{dE[\gamma|\emptyset,\bar{\gamma}_e]}{d\bar{\gamma}} = 0$. Moreover, when $c = 0$, $\bar{\gamma}_e = E[\gamma|\emptyset,\bar{\gamma}_e]$. Thus, from equations (12) and (13),

$$\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2} =\ \frac{d^2E[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}^2}\ \frac{dP_{\bar{\gamma}}}{d\bar{\gamma}}. \tag{14}$$

Substituting equation (14) into equation (11) yields

$$\frac{d^2W}{d\bar{\gamma}^2} = \frac{dP_{\bar{\gamma}}}{d\bar{\gamma}}\left(qf(\bar{\gamma}) - [qF(\bar{\gamma}) + 1 - q]\ \frac{d^2E[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}^2}\right)\ (\text{at } \bar{\gamma} = \bar{\gamma}_e = \bar{\gamma}_m). \tag{15}$$

From equation (6),

$$\begin{aligned}\frac{d^2E[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}^2} &= \frac{d}{d\bar{\gamma}}\left(\frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\right)(\bar{\gamma} - E[\gamma|\emptyset,\bar{\gamma}]) \\ &+ \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ \left(1 - \frac{dE[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}}\right) \\ &= \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ (\text{at } \bar{\gamma} = \bar{\gamma}_e = \bar{\gamma}_m,\ \text{when } c = 0).\end{aligned} \tag{16}$$

Substituting equation (16) for $\frac{d^2E[\gamma|\emptyset,\bar{\gamma}]}{d\bar{\gamma}^2}$ into equation (15) yields

$$\left.\frac{d^2W}{d\bar{\gamma}^2}\right|_{\bar{\gamma}=\bar{\gamma}_e} = \frac{dP_{\bar{\gamma}_e}}{d\bar{\gamma}}\left(qf(\bar{\gamma}_e) - qf(\bar{\gamma}_e)\right) = 0. \quad \square$$

Proof of Proposition 7: From equations (9), (12), and (6)

$$\begin{aligned}\frac{dW}{d\bar{\gamma}} &= qf(\bar{\gamma})\ [\ G(R-\mu-\bar{\gamma}) - G(R-\mu-E[\gamma|\emptyset,\bar{\gamma}])\ ] \\ &+ qf(\bar{\gamma})\,(\bar{\gamma} - E[\gamma|\emptyset,\bar{\gamma}])\ g(R-\mu-E[\gamma|\emptyset,\bar{\gamma}])\end{aligned}$$

We have $\bar{\gamma} < \bar{\gamma}_m$, which implies $\bar{\gamma} - E[\gamma|\emptyset, \bar{\gamma}] < 0$. Define $\delta \equiv -(\bar{\gamma} - E[\gamma|\emptyset, \bar{\gamma}]) > 0$, and $z \equiv R - \mu - E[\gamma|\emptyset, \bar{\gamma}]$. Then, $R - \mu - \bar{\gamma} = z + \delta$. Rewrite the expression for $\frac{dW}{d\bar{\gamma}}$ as

$$\frac{dW}{d\bar{\gamma}} = qf(\bar{\gamma}) \left[G(z+\delta) - G(z) - \delta \frac{dG(z)}{dz}\right].$$

The second and third terms in the bracket are the first order (linear) Taylor approximation to the first term around z. Because $G(s)$ is strictly unimodal, it has a unique inflection point s^* at which it switches from being strictly convex to being strictly concave. The tangent line to a convex (concave) curve is below (above) the curve. When $z \in [s^*, \infty)$, the bracket is always negative because at $z+\delta$ the tangential line is always above the curve. When $z \in (-\infty, s^*)$, if $z+\delta$ is also on the convex segment of the curve, i.e., $z+\delta < s^*$, then the tangent line is under the curve and the bracket is positive. However, if $z+\delta$ is on the concave segment of the curve, i.e., $z+\delta > s^*$, then the tangent line can be above the curve at $z+\delta$, in which case the bracket is negative. Nevertheless, if this happens at z^*, then the tangent line is above the curve at $z+\delta$ for all $z > z^*$. That is, there exists a $z^* \leq s^*$ such that the bracket is positive for $z < z^*$ and negative for $z > z^*$. The result follows from $z = R - \mu - E[\gamma|\emptyset, \bar{\gamma}]$. □

Proof of Proposition 8: From equation (4),

$$\begin{aligned}
\lim_{\bar{\gamma}\to+\infty} W(\bar{\gamma}; R) &= 1 - P_\emptyset(+\infty) - qc = 1 - P_{\gamma=0} - qc \\
\lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma}; R) &= 1 - (1-q)P_\emptyset(-\infty) - q\int_{-\infty}^{+\infty} P_\gamma dF(\gamma) \\
&= 1 - P_{\gamma=0} - q\int_{-\infty}^{+\infty} (P_\gamma - P_{\gamma=0}) dF(\gamma).
\end{aligned}$$

Thus,

$$\begin{aligned}
\lim_{\bar{\gamma}\to+\infty} W(\bar{\gamma}; R) - \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma}; R) &= q\int_{-\infty}^{+\infty} (P_\gamma - P_{\gamma=0}) dF(\gamma) - qc \\
&= q\int_{-\infty}^{+\infty} [G(R-\mu-\gamma) - G(R-\mu)] dF(\gamma) - qc.
\end{aligned}$$

Therefore, the strict unimodality of $G(s)$ and symmetry of $f(\gamma)$ imply that (1) when R or c are sufficiently large, $\lim_{\bar{\gamma}\to+\infty} W(\bar{\gamma};R) - \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma};R) < 0$; and (2) when R is sufficiently low and $c = 0$, then $\lim_{\bar{\gamma}\to+\infty} W(\bar{\gamma};R) - \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma};R) > 0$. By continuity, the same relationship holds for sufficiently small c. Moreover, when G is symmetric, $\int_{-\infty}^{+\infty}[G(R-\mu-\gamma) - G(R-\mu)]dF(\gamma)$ is negative if and only if $R - \mu > 0$. □

Proof of Proposition 9:

$$\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} = \left.\frac{\partial W(\bar{\gamma}(c),c)}{\partial\bar{\gamma}}\right|_{\bar{\gamma}=\bar{\gamma}_e} \frac{d\bar{\gamma}_e(c)}{dc} + \frac{\partial W(\bar{\gamma}_e(c),c)}{\partial c}. \tag{17}$$

From Proposition 2,

$$\frac{d\bar{\gamma}_e(c)}{dc} = \left(\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma} - \frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}}\right)^{-1} = -\left(\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}} + \left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right|\right)^{-1}. \tag{18}$$

From equation (4),

$$\frac{\partial W(\bar{\gamma}_e(c),c)}{\partial c} = -qF(\bar{\gamma}_e). \tag{19}$$

Substituting from equations (9), (18), and (19) into equation (17) yields

$$\begin{aligned}\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} &= \frac{(qF(\bar{\gamma}_e)+(1-q))\,\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}}}{\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}} + \left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right|} - qF(\bar{\gamma}_e) \\ &= \frac{-qF(\bar{\gamma}_e)\left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right| + (1-q)\,\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}}}{\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}} + \left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right|},\end{aligned}$$

which implies $\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} > 0$ if and only if $F(\bar{\gamma}_e)\,\frac{q}{1-q} < \frac{\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}}}{\left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right|}$.

Substituting from equations (12) and (6), $\frac{dW(\bar{\gamma}(c),c)}{dc}\Big|_{\bar{\gamma}=\bar{\gamma}_e} > 0$ if and only if

$$F(\bar{\gamma}_e)\,\frac{q}{1-q} < \frac{qf(\bar{\gamma}_e)}{1-q+qF(\bar{\gamma}_e)}\,(E[\gamma|\emptyset,\bar{\gamma}_e] - \bar{\gamma}_e)\,\frac{g(R-\mu-E[\gamma|\emptyset,\bar{\gamma}_e])}{g(R-\mu-\bar{\gamma}_e)}.$$

When $c = 0$, $\bar{\gamma}_e = \bar{\gamma}_m$, so $\bar{\gamma}_e - E[\gamma|\emptyset, \bar{\gamma}_e] = 0$, and hence the right-hand side of equation (20) is zero, implying that $\frac{dW(\bar{\gamma}_e(c),c)}{dc} < 0$. If $c > 0$, rearranging yields $\frac{dW(\bar{\gamma}(c),c)}{dc}\Big|_{\bar{\gamma}=\bar{\gamma}_e} > 0$ if and only if

$$\frac{1}{E[\gamma|\emptyset, \bar{\gamma}_e] - \bar{\gamma}_e} \frac{g(R - \mu - \bar{\gamma}_e)}{g(R - \mu - E[\gamma|\emptyset, \bar{\gamma}_e])} < \frac{f(\bar{\gamma}_e)}{F(\bar{\gamma}_e)} \frac{1 - q}{1 - q + qF(\bar{\gamma}_e)}.$$

$\lim_{c\to 1^-} \bar{\gamma}_e = -\infty$, and hence for $c \to 1^-$, the right-hand side is positive and bounded away from zero (by logconcavity of $f(\gamma)$), while the left-hand side goes to 0. Thus, $\frac{dW(\bar{\gamma}_e(c),c)}{dc} > 0$ for sufficiently large c. Next, we prove $\frac{dW(\bar{\gamma}_e(c),c)}{dc}$ has a single-crossing property. From Proposition 3, $\bar{\gamma}_e$ decreases in c. Thus, logconcavity of $f(\gamma)$ implies that the right-hand side strictly increases in c; moreover, $E[\gamma|\emptyset, \bar{\gamma}_e] - \bar{\gamma}_e$ also strictly increases in c (see Figure 2). If $R - \mu \geq 0$, then $R - \mu - E[\gamma|\emptyset, \bar{\gamma}_e]$ is positive and decreasing in c, while $R - \mu - \bar{\gamma}_e$ is positive and increasing in c. Then strict unimodality of $G(s)$ together with the median of $s \leq E[s] = 0$ imply that $\frac{g(R-\mu-\bar{\gamma}_e)}{g(R-\mu-E[\gamma|\emptyset,\bar{\gamma}_e])}$ strictly decreases in c. □

Proof of Proposition 10: From Proposition 4, $\lim_{q\to 0} P_\emptyset(\bar{\gamma}; q) = P_{\gamma=0}$, $\lim_{q\to 1} \bar{\gamma}_e(q) = -\infty$, and $\lim_{q\to 1} P_\emptyset(\bar{\gamma}_e(q); q) = 1$. Therefore, from equation (4), $\lim_{q\to 0} W(\bar{\gamma}, c, q) = 1 - \lim_{q\to 0} P_\emptyset(\bar{\gamma}; q) = 1 - P_{\gamma=0}$ and

$$\begin{aligned} \lim_{q\to 1} W(\bar{\gamma}_e(q), c, q) &= 1 - \lim_{q\to 1} F(\bar{\gamma}_e(q)) \lim_{q\to 1} P_\emptyset(\bar{\gamma}_e(q); q) - \lim_{q\to 1} \int_{\bar{\gamma}_e(q)}^{\infty} P_\gamma dF(\gamma) \\ &- \lim_{q\to 1} F(\bar{\gamma}_e(q))c \\ &= 1 - \int_{-\infty}^{\infty} P_\gamma dF(\gamma). \end{aligned}$$

Thus, $\lim_{q\to 0} W(\bar{\gamma}_e(q), c, q) - \lim_{q\to 1} W(\bar{\gamma}_e(q), c, q) = \int_{-\infty}^{\infty} (P_\gamma - P_{\gamma=0}) dF(\gamma)$. The result then follows from the proof of Proposition 8. □

7 References

Anderson, Simon P., and John McLaren. 2010. "Media Mergers and Media Bias with Rational Consumers." University of Virginia. Mimeo.

Arjomand, Said Amir. 1984. *The Shadow of God and the Hidden Imam: Religion, Political Order, and Societal Change in Shi'ite Iran from the Beginning to 1890.* Chicago: University of Chicago Press.

Balan, David J., Patrick DeGraba, and Abraham L. Wickelgren. 2004. "Ideological Persuasion in the Media." University of Texas. Mimeo.

Bernhardt, Dan, Stefan Krasa and Mattias Polborn. 2008. "Political Polarization and the Electoral Effects of Media Bias." *Journal of Public Economics* 92: 1092-1104.

Besley, Timothy, and Andrea Prat. 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *American Economic Review* 96: 720-736.

Beyer, Anne, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther. 2010. "The Financial Reporting Environment: Review of the Recent Literature." *Journal of Accounting and Economics* 50: 296-343.

Boix, Carles, and Milan Svolik. 2009. "The Foundations of Limited Authoritarian Government: Institutions and Power-sharing in Dictatorships." Princeton University. Mimeo.

Bueno de Mesquita, Ethan. 2010. "Regime Change and Revolutionary Entrepreneurs." *American Political Science Review* 104: 446-66.

Carlsson, Hans and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica*, 61: 989-1018.

Chan, Jimmy, and Wing Suen. 2008. "A Spatial Theory of News Consumption and Electoral Competition." *Review of Economic Studies* 75: 699-728.

Che, Yeon-Koo, and Navin Kartik. 2009. "Opinions as Incentives." *Journal of Political Economy* 117: 815-60.

Duggan, John, and Cesar Martinelli. 2011. "A Spatial Theory of Media Slant and Voter Choice." *Review of Economic Studies* 78: 640-666.

Dye, Ron A. 1985. "Disclosure of Non-proprietary Information." *Journal of Accounting Research* 23: 123-45.

Edmond, Chris. 2011. "Information Manipulation, Coordination and Regime Change." NBER Working Paper 17395.

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: 645-68.

Egorov, Georgy, and Konstantin Sonin. 2011. "Incumbency Advantage in Non-democracies." Northwestern University. Working Paper.

Gehlbach, Scott, and Konstantin Sonin. 2011. "Government Control of the Media." University of Wisconsin-Madison. Mimeo.

Gehlbach, Scott, and Arberto Simpser. 2011. "Electoral Manipulation as Bureaucratic Control." University of Wisconsin-Madison. Mimeo.

Gentzkow, Matthew, and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy* 114: 280-316.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuation." *American Economic Review* 101: 2590-2615.

Lorentzen, Peter L. 2012. "Strategic Censorship." Mimeo.

Milani, Abbas. 2011. *The Shah.* New York: Palgrave MacMillan.

Milani, Mohsen M. 1994. *The Making of Iran's Islamic Revolution.* Boulder: Westview Press.

Morris, Stephen, and Hyun S. Shin. 2003. "Global Games: Theory and Application." In *Advances in Economics and Econometrics, Eighth World Congress, Volume I,* ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky. NY: Cambridge University Press.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95: 1031-53.

Persson, Torsten, and Guido Tabellini. 2009. "Democratic Capital: The Nexus of Political and Economic Change." *American Economic Journal: Macroeconomics* 1: 88-126.

Rozenas, Arturas. 2010. "Forcing Consent: Information and Power in Non-Democratic Elections." Duke University. Mimeo.

Shadmehr, Mehdi, and Dan Bernhardt. 2011. "Collective Action with Uncertain Payoffs: Coordination, Public Signals, and Punishment Dilemmas." *American Political Science Review* 105: 829-51.

Verrecchia, Robert E. 1983. "Discretionary Disclosure." *Journal of Accounting and Economics* 5: 365-80.

Wheeler, Deborah L. 2010. "Information (without) Revolution? Ethnography and the Study of New Media Enabled Change in the Middle East." *APSA 2010 Annual Meeting Paper.*