STATE CENSORSHIP[1]

Mehdi Shadmehr[2] Dan Bernhardt[3]

July 26, 2013

[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, University of Miami finance department, University of Kansas economics department, and Berkeley Haas School of Business 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 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 strong media to a very weak one, but a good ruler prefers the opposite.

1 Introduction

We analyze the strategic choices by a ruler (an authoritarian state) of when to censor the information available to citizens to avoid revolution. In particular, we investigate how a ruler's expected payoffs are influenced by (1) his ability to commit to a censorship law (censorship strategy); (2) communication technologies that raise censorship costs; and (3) the media's ability to uncover news about the regime.

The ruler tries to manage information transmission to citizens to mitigate the likelihood of revolt. 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. 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 strong media that often uncovers news, a ruler censors only very bad news; while in a country with a weak 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 cannot commit to a censorship cutoff, and hence only weighs the direct consequences of censorship decisions. In particular, a ruler does not internalize that even though a 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 because citizens update less negatively following no news. This observation leads us to consider a ruler who can credibly set up institutions that allow him 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 are likely infeasible—for example, a bureaucrat has strong incentives to censor very bad news that would inevitably result in revolt and hence severe punishment of members of the regime, 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.

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*. 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, citizens would draw *more favorable* inferences about the regime following no news, and hence would be less likely to revolt. 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.

This analysis yields insights into how changes in censorship costs or in the media's ability to uncover news affect a ruler. One may think 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 initial 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 stronger media cause a ruler to censor less in equilibrium.

In fact, the consequences of higher censorship costs or of a stronger media for a ruler's expected payoffs 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. We show that, as a result, increases in censorship costs benefit a ruler (increase his ex-ante expected payoff) if and only if censorship costs are already high enough.

Finally, we prove that when the revolution payoff is low, a ruler prefers a weak media that uncovers almost no news—in effect censoring both bad news and good—to a strong media that uncovers almost all news. In contrast, when the revolution payoff is high, the ruler prefers a strong media that serves to commit the ruler to censoring almost nothing to a weak media. Intuitively, when revolution payoffs are high, a ruler *values* a strong media that might uncover *good* news about the status quo that then forestalls a revolution; but when revolution payoffs are low, a ruler *fears* a strong media that might uncover *bad* news about the status quo that then precipitates a revolution.

1.1 Literature Review

The closest paper is Besley and Prat (2006), which analyzes media capture in a democracy. The media can send a binary (good v. bad) signal to voters about the incumbent who can bribe the media to conceal the bad signal. They find that 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. Binary signals simplify characterizations, but are not rich enough to permit an analysis of the effects of commitment on citizen beliefs, or the effects of changes in censorship costs and media strength on a ruler's/incumbent's expected payoff, issues that are at the heart of our paper.

Based on Besley and Prat's model, 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. In their two state model, citizens can obtain a costly binary signal about the true state. They show that media bias is greater if the government owns the media and hence does not have to bribe it, or if the government cares less about advertising revenue.

In Egorov, Guriev, and Sonin (2009), a ruler can manipulate a binary (good v. 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. They 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.

Edmond (2011) studies information manipulation in dictatorships when 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 affects a ruler's payoffs.

Our finding that a ruler censors sufficiently bad news (Proposition 2) has analogues in the literature on the disclosure of accounting information (Dye 1985; Verrecchia 1983). In these papers, like ours, complete unraveling as in Grossman and Hart (1980), Grossman (1981), Milgrom (1981) or Milgrom and Roberts(1986) does not occur because there is a chance that nothing happened or that no information was uncovered. That literature studies pre-commitment to disclose information in the context of information sharing in oligopoly markets, a very different setting from ours (see Beyer et al. (2010), Dye (2001), and Verrecchia (2001) for comprehensive reviews).[1]

There are related analyses of information disclosure 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 the decision maker's. This leads the advisor to censor a bounded interval of signals. In particular, our result that, even when censorship is costless, the ruler does not censor all the bad news (Corollary 1) resonates with Dziuda (2011) who shows that a biased expert may provide information unfavorable to his preferences to convince the decision maker that the expert is not biased.

Our result that the ruler always gains if he can commit to censoring slightly less that

[1] Typically, two firms play a two-stage game. First, firms receive private information about demand or marginal costs of production. Then, they engage in Cournot or Bertrand competition. Pre-committment amounts to determining whether to disclose their information prior to receiving it (Vives 2001, p. 237-62).

the equilibrium level (Proposition 3) resonates with the result of Khalil (1997) in the context of optimal contracts with auditing/monitoring, who shows that the principal audits less often if he can pre-commit to an auditing strategy. However, the underlying reason in that paper is that "when threats are credible, they need not be used as often" (Khalil 1997, p. 630), while in our paper, pre-commitment allows the ruler to manipulate the citizen's beliefs following no news.

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 Balan, DeGraba, and Wickelgren (2004), Duggan and Martinelli (2011), or Anderson and McLaren (2012), the incentives of media outlets to selectively report news devolves from their partisan views, which lead them to have a preferred candidate or policy outcome; in Bernhardt, Krasa, and Polborn (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.[2] 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, DeGraba, and Wickelgren), and the role of endogenous media bias on electoral outcomes (Bernhardt, Krasa, and Polborn, Duggan and Martinelli) or policy outcomes (Anderson and McLaren).

2 Model

A representative citizen must decide whether or not to revolt against a ruler. If the citizen does not revolt, the status quo prevails; but if he revolts, the revolution succeeds with probability $r \in (0, 1]$. We normalize the ruler's payoff from preserving the status quo to 1 and his payoff after a successful revolution to 0. The citizen's net payoff from successful

[2]For a model where voters do not update about media bias (sequential rationality does not hold) see Mullainathan and Shleifer (2005).

revolution is $R - \gamma - \epsilon$, where R is known, but γ and ϵ are uncertain. We normalize the means of ϵ and γ to zero, so that R captures the ex-ante expected net payoff from successful revolution. Thus, a larger R can reflect (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. If the citizen revolts and the revolution fails, the citizen incurs an expected punishment $\mu > 0$, i.e., μ is the probability a failed revolter is caught times the punishment.

The representative citizen privately observes the realization of ϵ, which we assume is distributed according to a strictly positive, continuously differentiable density g on $\mathbb{R}$, with associated cdf G. The citizen does not directly observe γ, but has a prior that $\gamma \sim f$, where f is a strictly positive, continuously differentiable density on $\mathbb{R}$. We assume that ϵ and γ are independently distributed to ease exposition.[3]

The citizen learns the realization of γ if and only if the media observe γ and report it. With probability q, the media and ruler observe the realization of γ; but with probability $1 - q$ the realization is unobserved by the media.[4] Thus, q captures the probability that (i) an event occurs that is informative about γ *and* (ii) the media learn about it.

To focus on a ruler's censorship decisions, we abstract from strategic dissemination decisions by the media (see e.g., Duggan and Martinelli (2011) or the other papers in our literature review). Thus, we assume that when the media observes γ and the ruler does not censor, the media publicly reveals γ.[5] However, the ruler can censor the media at a cost $c \geq 0$, preventing the media from conveying γ to the representative citizen, in which case the citizen does not receive any further information about γ. For example, if a ruler censors a prison massacre of dissidents, the citizen never learns about it. Given his information, which includes γ only if the media learned about it and the ruler did not censor, the citizen decides whether to revolt.

[3]If we add parametric structure, allowing for correlation is straightforward (e.g., when ϵ and γ are jointly normally distributed).

[4]What is relevant is that the ruler observes γ any time the media does (else the ruler has no censorship decision). Our robustness analysis in Section 5 shows that our qualitative censorship findings are reinforced if the ruler sometimes directly observes γ when the media does not, and the ruler can sometimes choose whether to convey the realization of γ to the media, which then reports it.

[5]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.

The timing is as follows. First, γ and ϵ are realized and the representative citizen observes ϵ. Then with probability q the media learn γ, which they publicly report unless the ruler (who also observes the media's information) censors it. After the ruler makes his censorship decision, the citizen decides whether to revolt. Finally, payoffs are realized.

Analysis. The representative citizen's strategy is a function σ_c mapping his private information ϵ and any public information into a decision about whether to revolt, where $\sigma_c = 1$ indicates that he revolts, and $\sigma_c = 0$ indicates that he 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. We solve for Perfect Bayesian Equilibria.

We begin with the representative citizen's choice of whether to revolt. The citizen has private information ϵ and public information $\Omega \in \{\gamma, \emptyset\}$ about γ, where $\emptyset$ indicates there has not been a media report about γ. The citizen revolts if and only if the net expected payoffs from doing so are positive, i.e., $0 < r\ E[R - \gamma - \epsilon|\Omega] - (1-r)\ \mu$. Rearranging, the citizen revolts if and only if

$$\epsilon < R - E[\gamma|\Omega] - \frac{1-r}{r}\ \mu. \tag{1}$$

Denoting the right-hand side of (1) by $\rho(R, \Omega) = R - E[\gamma|\Omega] - \frac{1-r}{r}\ \mu$, it follows that the probability of successful revolution $P(\rho)$ is

$$P(\rho) = rG(\rho). \tag{2}$$

Good news for the ruler is any news that decreases the citizen's expectation of the net payoff from successful revolution, while bad news is the opposite. Thus, bad news raises ρ. Proposition 1, which directly follows from equation (2), characterizes the properties of the citizen's equilibrium behavior.

Proposition 1 *Bad news raises the likelihood of successful revolt:* $\frac{P(\rho)}{\partial \rho} > 0$. *The citizen almost always revolts when the news is very bad,* $\lim_{\rho \to \infty} P(\rho) = r$*; but almost never revolts when it is very good,* $\lim_{\rho \to -\infty} P(\rho) = 0$.

Proposition 1 contains all of the key structure that we exploit in our analysis of censorship. Specifically, the equilibrium likelihood of revolution is a differentiable function of the media reports about the net payoff from preserving the status quo versus successful revolution with the properties that (i) revolt almost always occurs following extremely bad news about the status quo, but (ii) almost never occurs after extremely good news. Our substantive findings only rely on this structure together with modest distributional assumptions such as log concavity. Appendix A presents a multi-citizen model, where citizens must coordinate for a revolution to succeed. The natural equilibrium to this multi-citizen model delivers exactly the *same* key properties.

More generally, Section 5 establishes the robustness of our findings to alternative assumptions. In our core model, a ruler who chooses to censor news always succeeds in concealing it. Section 5 shows that our qualitative conclusions extend when some news (e.g., very bad or very good news) is easier for the media to uncover than other news; or when the likelihood that censorship succeeds depends continuously on the resources that the ruler devotes to concealing the news; or when sometimes the ruler, but not the media, directly observes γ, and the ruler can sometimes selectively inform the media about γ.

Censorship Equilibrium. Let P_γ be the probability of revolution following a report γ, and $P_\emptyset$ be the probability of revolution following no report about γ: $P_\gamma = r\, G(R-\gamma-\frac{1-r}{r}\mu)$ and $P_\emptyset = r\, G(R-E[\gamma|\emptyset]-\frac{1-r}{r}\mu)$. The ruler censors a report γ if and only if his expected payoff from censoring, $1-P_\emptyset-c$, exceeds that of not censoring, $1-P_\gamma$. That is,

$$\sigma_r(\gamma) = 1 \text{ if and only if } P_\gamma - P_\emptyset > c. \tag{3}$$

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, $\bar{\gamma}$.

Lemma 1 *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.*

The ruler censors reports of γ whenever the likelihood of revolution following that report, P_γ, exceeds the probability of revolution following no report $P_\emptyset$ by a margin that exceeds the censorship cost c. Because in equilibrium, the citizen's beliefs must be consistent with the ruler's equilibrium strategy, $P_\emptyset$ implicitly depends on the equilibrium censorship cutoff $\bar{\gamma}_e$. Thus, $\bar{\gamma}_e$ is an equilibrium strategy of the ruler if and only if

$$P_\gamma - P_\emptyset(\bar{\gamma}_e) > c \text{ for all } \gamma < \bar{\gamma}_e, \text{ and } P_\gamma - P_\emptyset(\bar{\gamma}_e) \le c \text{ for all } \gamma \ge \bar{\gamma}_e. \tag{4}$$

The equilibrium threshold depends on how the representative citizen updates when there is no media report concerning γ. That is, $P_\emptyset(\bar{\gamma}_e)$ depends on $E[\gamma|\emptyset;\bar{\gamma}_e]$, where the citizen's beliefs are consistent with the ruler's censorship cutoff $\bar{\gamma}_e$. Given a censorship cutoff $\bar{\gamma}$, the representative citizen's beliefs about γ following no media report depend on $\bar{\gamma}$ via Bayes' rule:

$$\begin{aligned} E[\gamma|\emptyset,\bar{\gamma}] &= \frac{1-q}{1-q+q\ F(\bar{\gamma})}\ E[\gamma] + \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma<\bar{\gamma}] \\ &= \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma<\bar{\gamma}] \ = \ \frac{q}{1-q+q\ F(\bar{\gamma})} \int_{-\infty}^{\bar{\gamma}} \gamma f(\gamma) d\gamma, \end{aligned} \tag{5}$$

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 and when he *never* censors, i.e., when $\bar{\gamma} = \pm\infty$, the absence of a media report conveys no information about γ, so that $E[\gamma|\emptyset;\bar{\gamma} = \pm\infty] = E[\gamma] = 0$. If, instead, the ruler censors with positive probability less than one, the citizen updates negatively following no media report, because there may have been bad news that the ruler censored: $E[\gamma|\emptyset;\bar{\gamma}] < E[\gamma] = 0$.

Figure 1 illustrates two key features describing how the citizen updates following no media report: (1) how $E[\gamma|\emptyset,\bar{\gamma}]$ varies with the censorship threshold $\bar{\gamma}$, and (2) how the citizen would update were he to see the threshold news $\bar{\gamma}$. As a ruler censors more, the citizen updates more negatively following no 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, if the ruler only censors extremely bad news, the citizen does not update very negatively following no media report, because no news is more likely due to the media not observing γ rather than censorship. 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 no media report, censorship occurred, but it also raises the expectation of γ conditional on censorship occurred. The first effect reduces the expectation of γ following no media report, 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. If a ruler censors only bad news, the marginal news censored

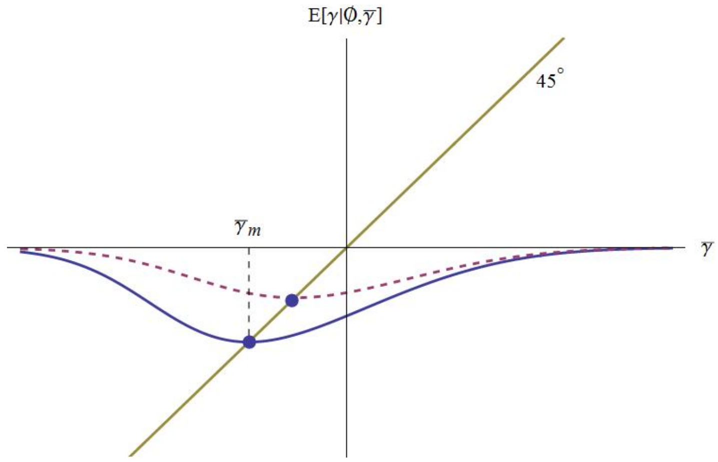

Figure 1: $E[\gamma|\emptyset;\bar{\gamma}]$ as a function of $\bar{\gamma}$. The solid curve corresponds to $q = 0.9$ and the dashed one corresponds to $q = 0.7$.

$\bar{\gamma}$ is worse than the average $E[\gamma|\emptyset;\bar{\gamma}]$, so that censoring slightly more than the threshold news *lowers* the citizen's expectations following no media report. As a ruler raises $\bar{\gamma}$, this expectation falls as long as $\bar{\gamma} < E[\gamma|\emptyset;\bar{\gamma}]$. Once the censorship cutoff exceeds $E[\gamma|\emptyset;\bar{\gamma}]$, censoring more begins to raise the citizen's belief following no media report. Since no reported news is bad news, it follows that $\bar{\gamma}_m = E[\gamma|\emptyset;\bar{\gamma}_m] < E[\gamma] = 0$. Formally,

Lemma 2 *The citizen's estimate $E[\gamma|\emptyset;\bar{\gamma}]$ of γ following no media report and given the ruler's censorship cutoff $\bar{\gamma}$ has 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 (a minimum) at* $\bar{\gamma}_m < 0$, *where* $E[\gamma|\emptyset;\bar{\gamma}_m] = \bar{\gamma}_m$. $E[\gamma|\emptyset;\bar{\gamma}] > \bar{\gamma}$ *if* $\bar{\gamma} < \bar{\gamma}_m$, *and* $E[\gamma|\emptyset;\bar{\gamma}] < \bar{\gamma}$ *if* $\bar{\gamma} > \bar{\gamma}_m$.

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

These features of the representative citizen's belief imply that censoring extremely bad news benefits a ruler because the citizen updates less negatively following no media report 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 the citizen rather than censoring it. It follows that there is a unique, finite, equilibrium censorship cutoff, $\bar{\gamma}_e$, as long as the censorship cost is low enough relative to the payoff from preventing successful revolution:

Proposition 2 *There is a level of censorship costs $\bar{c} \in (0, r)$ above which the ruler does not censor. If $c < \bar{c}$, the ruler censors all media reports worse than a threshold news $\bar{\gamma}_e \in \mathbb{R}$, where $\bar{\gamma}_e$ is the unique solution to $P_{\bar{\gamma}_e} - P_{\emptyset}(\bar{\gamma}_e) = c$.*

The next corollary highlights that it is the structure of the citizen's belief following no media report, and not the cost of censorship, that discourages a ruler from censoring all bad news: revealing modestly bad news that slightly lowers the citizen's expectation 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, $\bar{\gamma}_e(c = 0) = \bar{\gamma}_m < E[\gamma] = 0$.*

In the analysis that follows, we assume that $c < \bar{c}$. We next show how the primitives of the economy affect which media reports are censored.

Corollary 2 *Increases in (a) the costs of censorship or (b) the likelihood q that the media observes γ, both reduce a ruler's equilibrium censorship cutoff, $\bar{\gamma}_e$. Moreover, suppose that for $\bar{\gamma} < E[\gamma] = 0$, both $F(\bar{\gamma})$ and $E[\gamma|\gamma < \bar{\gamma}]$ decrease with the variance σ_0^2 of γ. 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 when facing a stronger media that is more likely to uncover γ. This is because for a given censorship cutoff, with a stronger media (higher q), when there is no report, it is more likely due to censorship, so the representative citizen updates more negatively. Similarly, a ruler censors less when there is more potential news: a higher σ_0^2 (e.g., when news is normally distributed) causes the citizen to update more negatively when there is no report because there is more probability mass on bad tail news. Therefore, the ruler ceases to censor more marginally bad news to prevent the citizen from drawing inferences that the news might have been far worse.

We next describe how the extreme limits of a very strong and very weak media affect a ruler's censorship decisions.

Corollary 3 *A ruler censors almost nothing if the media almost always uncover* γ*. Conversely, a ruler censors almost all bad news if censorship is almost costless and media almost never uncover* γ*. That is,* $\lim_{q\to 1} \bar{\gamma}_e(q,c) = -\infty$ *and* $\lim_{c,q\to 0} \bar{\gamma}_e(q,c) = 0$.

3 Commitment in Censorship

A ruler's equilibrium censorship decisions weigh the gains from reducing the likelihood of revolution when a (bad) media report is censored against the costs of censorship, ignoring the consequences for the level of the citizen's 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, because the citizen would then update more positively when there is no media report.

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

$$\begin{aligned} W(\bar{\gamma}, c, q) \;&\equiv\; q\Big(F(\bar{\gamma})\;(1 - P_\emptyset(\bar{\gamma}) - c) + (1 - F(\bar{\gamma}))\;\frac{\int_{\bar{\gamma}}^{\infty}(1 - P_\gamma)dF(\gamma)}{1 - F(\bar{\gamma})}\Big) \\ &\quad + (1 - q)\big(1 - P_\emptyset(\bar{\gamma})\big) \\ &= \; 1 - \Big([qF(\bar{\gamma}) + (1 - q)]P_\emptyset(\bar{\gamma}) + q\int_{\bar{\gamma}}^{\infty} P_\gamma dF(\gamma) + qF(\bar{\gamma})c\Big). \qquad (6) \end{aligned}$$

With probability $1 - q$ the media do not observe γ, so the ruler's expected payoff is $1 - P_\emptyset(\bar{\gamma})$. With probability q the media observe γ. But the ruler censors any report of $\gamma < \bar{\gamma}$. Thus, the ruler censors with probability $qF(\bar{\gamma})$, in which case his payoff is $1 - P_\emptyset(\bar{\gamma}) - c$. If γ exceeds $\bar{\gamma}$, the ruler does not censor, so the citizen learns γ. Thus, the probability that the citizen sees γ is $q(1 - F(\bar{\gamma}))$, and integrating over $\gamma > \bar{\gamma}$, the ruler's expected payoff is $\frac{\int_{\bar{\gamma}}^{\infty}(1-P_\gamma)dF(\gamma)}{1-F(\bar{\gamma})}$.

The citizen's belief about γ in the absence of a media report depends on the ruler's censorship policy. However, the ruler ignores this in his decision-making: his censorship decisions maximize his expected payoff *given* the citizen's (equilibrium) belief. If, instead, a ruler can *ex-ante* commit to a censorship level, he can "internalize" how his censorship rule affects updating by the citizen. 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 commit is limited. For example, a bureaucrat who fails to censor very bad news is likely to be punished (either by the ruler or by citizens following the likely resulting successful revolution); and a ruler with especially 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 becomes: Under what circumstances would a ruler want to increase or decrease his censorship cutoff marginally from its equilibrium level?

Remarkably, 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 3 *If censorship is costly, a ruler's ex-ante expected payoff 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 a report, the citizen now sees marginal realizations of γ that are slightly worse than 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 citizen's belief, 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 payoff 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 media report. To see this, recall that (a) as a ruler censors more news, the citizen updates more negatively following no media report, up to the point where the ruler censors all news worse than $\bar{\gamma}_m$ (see Figure 1); 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 $\bar{\gamma}_e$. Hence, the net effect of marginally reducing the censorship cutoff is to raise the ruler's ex-ante expected payoff.

The sole restriction in the proposition is that censorship be costly. The reason is that when censorship is costless, the citizen's belief becomes locally insensitive to changes in the 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 the citizen's belief about $E[\gamma|\emptyset;\bar{\gamma}]$ (see Lemma 2 and Figure 1), i.e., $\frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}} = 0$. This directly implies that $\frac{dW(\bar{\gamma}_e)}{d\bar{\gamma}} = 0$, which might lead one to conjecture that this censorship cutoff minimizes or maximizes the ruler's ex-ante expected payoff. However, this is not so:

Remark 1. When censorship is costless, i.e., when $c = 0$, a ruler's equilibrium censorship cutoff is an inflection point of his ex-ante expected payoff. That is, when $c = 0$, not only is $\frac{dW(\bar{\gamma}_e)}{d\bar{\gamma}} = 0$, but we also have $\frac{d^2W(\bar{\gamma}_e)}{d\bar{\gamma}^2} = 0$.

4 Censorship Costs and Media Strength

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 (Wheeler 2010) has led many to conclude that technologies that make censorship difficult, i.e., increase censorship costs, must work against dictators. We show that such assessments are only partially correct. To the contrary, higher costs of censorship can sometimes increase a ruler's chance of survival. In fact, these gains can be high enough to offset the greater direct costs of censorship and make the ruler better off. Similarly, one might suspect that dictators always prefer a weak media that rarely uncovers γ to a strong media. This conjecture is also wrong. In fact, when the ruler faces a serious risk of revolution, he prefers a strong media that almost always uncovers γ to a weak media.

Higher Censorship Costs. Higher censorship costs have direct and indirect effects. The direct effect reduces a ruler's ex-ante expected payoff 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 the citizen to update less negatively following no media report and hence benefit a ruler. It follows that the indirect effect of increases in censorship costs raise a ruler's ex-ante expected payoff. Proposition 4 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 payoff. However, when censorship costs are higher,

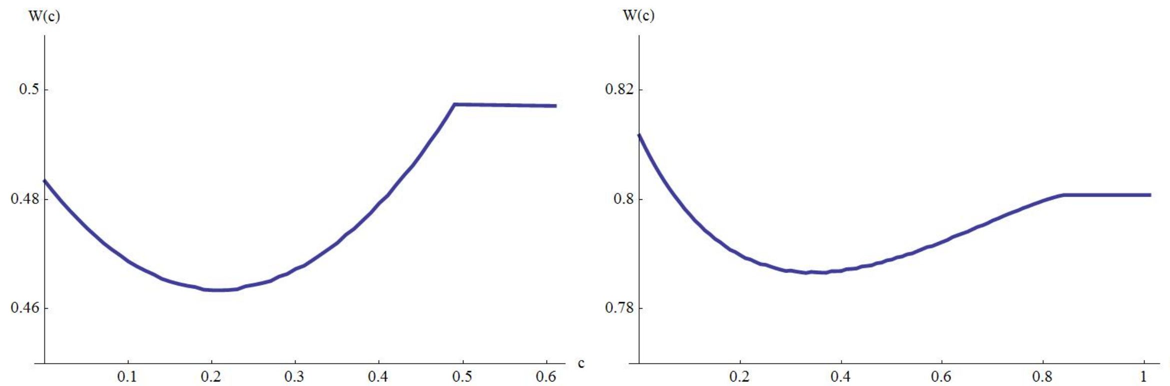

Figure 2: Parameters: $\gamma, \epsilon \sim N(0,1)$, $\mu = 0$, $q = 0.5$, $r = 1$, $R = 0$ (left panel), and $R = -1$ (right panel).

and the expected payoffs from successful revolution exceed the expected punishment from failed revolt, then the indirect effect eventually dominates, so that further increases in censorship costs benefit the ruler.

Proposition 4 *Suppose $F(\gamma)$ is log-concave, $G(\epsilon)$ is strictly unimodal with a median that does not exceed its mean, and $rR \geq (1-r)\,\mu$. Then there exists a critical censorship cost $\hat{c} > 0$ such that increases in a ruler's cost of censorship raise his ex-ante expected payoff if and only if $c > \hat{c}$.*

When censorship is costless, i.e., when $c = 0$, 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 the citizen's belief about γ following no media report. 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 the citizen's belief is modest.

However, the proposition identifies sufficient conditions such that once censorship costs are high enough, the indirect effect dominates: for censorship costs exceeding $\hat{c}$, further increases in censorship costs raise the ruler's expected payoff due to the less pessimistic updating by the citizen, 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 too likely to censor (and thus incur the higher direct censorship costs). See Figure 2. It follows that there exists a cost threshold $\tilde{c} < \hat{c}$ such that once censorship costs reach $\tilde{c}$, the ruler would be strictly better off with censorship costs that are so high that they discourage *all* censorship. Log-concavity of F alone is enough to guarantee that the ruler's ex-ante expected payoff rises once c is sufficiently high (close enough to $\bar{c} = r - P_{\bar{\gamma}=0}$). The remaining structure implies that $\frac{g(R-\mu_r-\bar{\gamma}_e)}{g(R-\mu_r-E[\gamma|\emptyset;\bar{\gamma}_e])}$ decreases in c, which is a sufficient condition that ensures $\frac{dW(\bar{\gamma}_e(c),c)}{dc}$ has a single-crossing property.

Media Strength. Increases in censorship costs work as a commitment device for the ruler, making him censor less in equilibrium. This logic suggests a similar role for the likelihood q that γ is uncovered by the media: (i) marginal increases in q benefit the ruler by making him censor marginally less in equilibrium, but (ii) increases in q raise the expected censorship costs $qF(\bar{\gamma}_e)c$. However, there is also a third effect: given any censorship cutoff $\bar{\gamma}$, increases in q cause the representative citizen to update more negatively following no media report. This third effect makes a comprehensive analysis intractable absent strong distributional assumptions. This leads us, instead, to analyze the ruler's ex-ante expected payoff in the two polar cases where the media almost always uncovers γ and when the media almost never uncovers it.

A ruler facing a weak media that almost never uncover γ is in the same situation as a ruler who commits to censoring almost everything—when censorship is costless. In both cases, the citizen updates only marginally negatively following no media report. Conversely, a ruler facing a strong media that almost always uncover γ is in the same situation as a ruler who commits to censor almost nothing. Proposition 5 compares a ruler's ex-ante expected payoff if he almost censored everything with his expected payoff if he almost censored nothing.

Proposition 5 *Suppose $G(\epsilon)$ is strictly unimodal and $f(\gamma)$ is symmetric. When R or c is sufficiently large, a ruler's ex-ante expected payoff 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 and c are sufficiently small, the opposite result obtains:* $\lim_{\bar{\gamma}\to\infty} W(\bar{\gamma}; R) > \lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma}; R)$.

When $g(\epsilon)$ is symmetric, one can strengthen the proposition: *when $c = 0$, a ruler is better off censoring nothing than censoring everything if and only if $rR > (1-r)\mu$.* The intuition is that when R is high enough that the citizen is likely to revolt based on his prior, a ruler is better off rolling the dice on media reports, 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, R is low, it takes unusually bad news to make the citizen revolt—the effect of good news on the probability of revolt is marginal, while the adverse effect of bad news is sharp, so the ruler prefers to censor everything rather than nothing.[6] Corollary 4 states the implications for a ruler's ex-ante preference for the strength of media.

Corollary 4 *Suppose $G(\epsilon)$ is strictly unimodal and $f(\gamma)$ is symmetric. Then there exists a critical threshold R^* such that if $R > R^*$, a ruler prefers a strong media that almost always uncover γ to a weak media that almost never does: if $R > R^*$, $\lim_{q\to 0} W(\bar{\gamma}_e(q); R) < \lim_{q\to 1} W(\bar{\gamma}_e(q); R)$. If, instead, $R < R^*$, a ruler prefers a weak media to a strong one: if $R < R^*$, then $\lim_{q\to 0} W(\bar{\gamma}_e(q); R) > \lim_{q\to 1} W(\bar{\gamma}_e(q); R)$.*

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 the citizen receives good news and updates positively. The chances that media uncovers and reports such good news is highest when the media is very strong ($q \approx 1$). Of course, a strong media may uncover bad news that further raises the likelihood of revolt, but because this likelihood is already very high, the possible gains from sharply positive news out weigh any loss from more bad news that can only raise the likelihood of revolt marginally. 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 1978 during the unfolding of the 1979 Iranian Revolution (see also Milani 1994, 117; Arjomand 1984, 115).

[6] 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).

5 Extensions

We study three extensions of the model. First, we allow the likelihood that the media discovers news to hinge on the "level" of the news. Second, we make the state's choice of censorship continuous, so that an attempt to censor news is more likely to succeed when the state devotes more resources to censoring. Third, we generalize the citizen's payoff (from successful revolution) from $R - \gamma - \epsilon$ to $R - \lambda\gamma - \epsilon$, to capture how the degree of relative importance of news in the citizen's payoff affects censorship decisions.

News Type and the Likelihood of Discovery. We first show that our findings extend if the likelihood that the media discover information about γ depends on γ, i.e., if q varies with γ. For example, the media may be more likely to uncover very bad or very good news. Crucially, the key properties of how citizens update following no news are unaffected:

$$
\begin{aligned}
E[\gamma|\emptyset;\bar{\gamma}] &= \frac{\int_{-\infty}^{\bar{\gamma}} \gamma \; dF(\gamma) + \int_{\bar{\gamma}}^{\infty} \gamma \; (1 - q(\gamma)) \; dF(\gamma)}{1 - \int_{\bar{\gamma}}^{\infty} q(\gamma) \; dF(\gamma)} \\
&= \frac{E[\gamma] - \int_{\bar{\gamma}}^{\infty} \gamma \; q(\gamma) \; dF(\gamma)}{1 - \int_{\bar{\gamma}}^{\infty} q(\gamma) dF(\gamma)} = -\frac{\int_{\bar{\gamma}}^{\infty} \gamma \; q(\gamma) \; dF(\gamma)}{1 - \int_{\bar{\gamma}}^{\infty} q(\gamma) \; dF(\gamma)} ,
\end{aligned}
$$

where the last equality follows from our normalization $E[\gamma] = 0$. Further,

$$
\frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}} = \frac{q(\bar{\gamma}) f(\bar{\gamma})}{1 - \int_{\bar{\gamma}}^{\infty} q(\gamma) dF(\gamma)} \; (\bar{\gamma} - E[\gamma|\emptyset;\bar{\gamma}]).
$$

Thus, $E[\gamma|\emptyset;\bar{\gamma}]$ again attains a minimum at $\bar{\gamma} = E[\gamma|\emptyset;\bar{\gamma}]$. Our propositions thus extend.

Indeed, it may be that the environment places enough structure on how $q(\gamma)$ varies with γ that we can characterize how the equilibrium level of censorship is affected. For example, as we observed earlier, the ruler may sometimes directly observe γ when the media do not; and when the ruler's information about γ is verifiable, the ruler may be able to convey it to the media at a cost. The ruler never has an incentive to convey bad news to the media, but he may have an incentive to convey good news. This has an effect analogous to increasing $q(\gamma)$ for γ sufficiently greater than $\bar{\gamma}$. It follows directly for any fixed censorship cutoff $\bar{\gamma}$ that $E[\gamma|\emptyset, \bar{\gamma}]$ is reduced—the citizen updates more negatively when no news is reported. In turn, this causes the ruler to censor less—when the ruler can convey good news to the media, bad news γ must be worse for the ruler to censor it.

Continuous Choice of Censorship. In practice, a ruler's attempts to censor a media report may fail. Advances in communication technologies have not only increased the cost of censorship, but have also raised the possibility that censorship may fail. For example, a ruler may not be able to conceal beatings of protesters if videos can be taken with cell phones and uploaded to Youtube. One can model this by assuming that a ruler who tries to censor succeeds in concealing the news with probability ζ. We focus on $\zeta = 1$, but the effects of $\zeta < 1$ are straightforward to derive. In particular, when ζ is smaller, the ruler reduces his cutoff, as censorship is less effective. Consequently, the citizen updates less negatively following no media report because it is less likely due to censorship. Qualitatively, the impact of a higher ζ is similar to that of a stronger media that uncovers γ with a higher probability q.

To illustrate more nuanced implications, we now suppose that the likelihood the ruler succeeds in censoring a media report depends on the level of resources he devotes to censoring that news. In particular, if the ruler expends c, he succeeds in censoring with probability $\zeta(c)$, where $\zeta(0) = 0$, $\zeta'(\cdot) > 0$, $\zeta''(\cdot) < 0$, and $\zeta(1) < 1$ means that it is never optimal to spend enough to guarantee that censorship succeeds. Thus, with probability $1 - \zeta(c) > 0$, censorship fails and the citizen sees the media report of γ. The ruler's expected payoff if he devotes c to censoring γ is $\zeta(c)(1 - P_\emptyset) + (1 - \zeta(c))(1 - P_\gamma) - c$. Hence, the ruler's optimization problem simplifies to

$$\max_{c \geq 0} \ \zeta(c)(P_\gamma - P_\emptyset) - c.$$

Let $C(\gamma)$ be the *function* indicating the resources devoted to censorship when the media would reveal γ, and let $P_\emptyset(C)$ be the equilibrium probability of revolution following no report. Focusing on the case where $\zeta'(0)$ is large enough that some efforts at censorship are optimal when the news is extremely bad, there exists a news threshold $\bar{\gamma}$ at which $\zeta'(0)\ (P_{\bar{\gamma}} - P_\emptyset) = 1$. Then, the resources $C(\gamma)$ that the ruler devotes to censoring γ solve

$$\zeta'(C(\gamma))\ (P_\gamma - P_\emptyset) = 1 \text{ for } \gamma < \bar{\gamma}, \text{ and } C(\gamma) = 0 \text{ for } \gamma \geq \bar{\gamma}. \tag{7}$$

Because P_γ decreases in γ, it follows that $C(\gamma)$ decreases in γ: the ruler devotes more resources to censoring worse news. Given the ruler's strategy $C(\gamma)$, the citizen's belief about the expected value of γ following no news becomes:

$$E[\gamma|\emptyset; C(\gamma)] = \frac{q\int_{-\infty}^{\infty} \gamma\zeta(C(\gamma))\ dF(\gamma)}{1 - q + q\int_{-\infty}^{\infty} \zeta(C(\gamma))\ dF(\gamma)}, \tag{8}$$

which determines the equilibrium $P_\emptyset(C)$. The ruler's ex-ante expected payoff is

$$\begin{aligned} W(C) &= (1-q)(1-P_\emptyset(C)) \\ &\quad +q\int_{-\infty}^{\bar{\gamma}} [\zeta(C(\gamma))(1-P_\emptyset(C)) + (1-\zeta(C(\gamma)))(1-P_\gamma)]dF(\gamma) \\ &\quad -q\int_{-\infty}^{\bar{\gamma}} C(\gamma)dF(\gamma) + q\int_{\bar{\gamma}}^{\infty} (1-P_\gamma)dF(\gamma). \end{aligned}$$

Algebra reveals that $1 - W(C)$ simplifies to

$$(1-q)P_\emptyset(C) - q\int_{-\infty}^{\bar{\gamma}} \{\zeta(C(\gamma))[P_\gamma - P_\emptyset(C)] - C(\gamma)\}dF(\gamma) + q\int_{-\infty}^{\infty} P_\gamma dF(\gamma).$$

Taking the functional derivative with respect to C yields:[7]

$$\begin{aligned} \frac{\delta W}{\delta C} &= q\int_{-\infty}^{\bar{\gamma}} \{\zeta'(C(\gamma))[P_\gamma - P_\emptyset(C)] - 1\}dF(\gamma) \\ &\quad - \left(q\int_{-\infty}^{\bar{\gamma}} \zeta(C(\gamma))\frac{dP_\emptyset(C)}{dC}dF(\gamma) + (1-q)\frac{dP_\emptyset(C)}{dC}\right) \\ &= -\left((1-q)\frac{dP_\emptyset(C)}{dC} + q\int_{-\infty}^{\bar{\gamma}} \zeta(C(\gamma))\frac{dP_\emptyset(C)}{dC}\ dF(\gamma)\right) \quad \text{at equilibrium,} \end{aligned}$$

where the second equality follows from the equilibrium condition (7). Moreover, from equation (8), $\frac{E[\gamma|\emptyset;C])}{dC} < 0$ for $\gamma < \bar{\gamma}$. Thus, $\frac{dP_\emptyset(C)}{dC} > 0$, and hence $\frac{\delta W}{\delta C} < 0$ at equilibrium. For each $\gamma < \bar{\gamma}$, the ruler equates the marginal benefit from greater censorship to its marginal cost, ignoring the adverse consequences for how the citizen updates following no report. It follows that, facing a bad report that would reveal $\gamma < \bar{\gamma}$, the ruler expends excessive resources trying to reduce the probability that the citizen receives that report: the ruler would be better off if he devoted slightly less resources to censoring bad news, i.e., if he slightly reduced $C(\gamma)$ pointwise for each $\gamma < \bar{\gamma}$.

Relative Importance of News. We assumed that the citizen values γ and ϵ equally. Now, we introduce a parameter $\lambda > 0$ that captures the relative importance of γ and

[7]Rewrite W as $W[C] = \int_a^b I(\gamma, C(\gamma), C'(\gamma))d\gamma$, and recall that $\frac{\delta W}{\delta C} = \frac{\partial I}{\partial C} - \frac{d}{d\gamma}\frac{\partial I}{\partial C'}$; because $\frac{\partial I}{\partial C'} = 0$, we have $\frac{\delta W}{\delta C} = \frac{\partial I}{\partial C}$.

ϵ for citizen payoffs: we modify the citizen's net payoff from successful revolution to $R - \lambda\gamma - \epsilon$. Then, the citizen's decision of whether to revolt hinges on whether $\epsilon < R - \frac{1-r}{r}\,\mu - \lambda\, E[\gamma|\Omega]$. Because the decision of whether to censor depends on the difference in the citizen's updating following news γ versus following no report (see Figure 1), we only need to derive the impact of λ on $\lambda\,(E[\gamma|\emptyset] - \gamma)$. But λ enters as a scaling factor: increases in λ widen the gap between γ and $E[\gamma|\emptyset]$, which, in turn, widen the gap between P_γ and $P_\emptyset$. It follows directly that raising the payoff weight on γ results in greater censorship, i.e., $\frac{d\bar{\gamma}_e}{d\lambda} > 0$.

6 Conclusion

Our paper investigates 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 through censorship, preventing the media from disseminating information about the regime that would raise the likelihood of revolt. Unfortunately for the ruler, citizens account for his 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, when bad news arrives, the ruler can increase his chances of survival by censoring it.

Remarkably, we show that a ruler's ex-ante expected payoffs are always increased if he can commit 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". This does not mean that less censorship always maximizes the ruler's ex-ante expected payoffs, only that slightly less censorship is always better than the equilibrium level of censorship. Ironically, when ex-ante expected revolution payoffs are high so that revolution is likely, a ruler is better off with a strong media that uncovers almost all news, causing the ruler to censor almost nothing, than with a weak one that uncovers almost nothing. However, when ex-ante expected revolution payoffs are low, so that revolution is unlikely, a ruler prefers a weak media, in effect costlessly censoring almost all news, to a strong media. So, too, 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 ex-ante expected payoffs first fall and then rise with the costs of censorship.

We focus on the censorship of media. However, our model can be reposed to study elections in dictatorships, where a dictator has incentives 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 choice of which news to censor: A ruler has private information about the extent of his support, and holding an election give potential protesters, etc. a signal about this support.

7 Appendix A: Multi-Citizen Game

Our previous analysis abstracts from coordination among citizens. 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. We show this framework also delivers the minimal structure (Proposition 1) that our results rely on, namely that (1) better net news about the status quo reduces the likelihood of revolution, and (2) 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 $\gamma + \epsilon$, 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 $\mu > 0$.

	revolt	*no revolt*
revolt	R, R	$\gamma + \epsilon - \mu, \gamma + \epsilon$
no revolt	$\gamma + \epsilon, \gamma + \epsilon - \mu$	$\gamma + \epsilon, \gamma + \epsilon$

Figure 3: Citizen Payoffs.

R is known, but γ and ϵ are uncertain: $\gamma \sim f$ and $\epsilon \sim g$, where f and g are strictly positive, continuously differentiable densities on $\mathbb{R}$. Without loss of generality, we normalize the means of ϵ 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 ϵ that is independent of γ.[8] Given their private signals about ϵ and the public information about γ, citizens simultaneously decide whether to revolt. A strategy for citizen i is a function σ_i mapping i's private signal s^i and 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.

[8] Outcomes would be identical if citizens received signals about R because optimal actions hinge on the *expected difference*, $E[\gamma + \epsilon - R|\Omega]$, not the *source* of the difference.

Let $\Omega \in \{\gamma, \emptyset\}$ be the public information about γ, where $\emptyset$ indicates that citizens do not observe any news. Then $\rho(R, \Omega) \equiv E[R - \gamma|\Omega]$ is the citizens' public knowledge about the expected value of $R - \gamma$. 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.

We characterize the equilibrium behavior of citizens (Lemma 5) using the following minimal structure on citizen signals:

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

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

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

Lemma 3 *The best response to a cutoff strategy is a cutoff strategy.*

Proof: Citizen i revolts if and only if

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

γ and ϵ are independently distributed, so the expected net payoff from revolt simplifies to

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

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

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

Substituting $\rho(R, \Omega) = E[R - \gamma|\Omega]$ and using the cutoff strategy structure yield

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

which can be rewritten as

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

We have (a) $0 < Pr(s^{-i} < k^{-i}|s^i)$, (b) $\rho(R, E[\gamma|\Omega])$ is finite, and (c) $E[\epsilon|s^i, s^{-i} < k^{-i}]$ 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,

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

so $\lim_{s^i \to +\infty} \Delta(s^i; k^{-i}) < 0$, and $\lim_{s^i \to -\infty} \Delta(s^i; k^{-i}) > 0$. These properties together with the continuity of $\Delta(s^i; k^{-i})$ imply 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 $s^i > k^i$. □

Lemma 4 *When the noise in private signals is vanishingly small, all cutoff equilibria are symmetric.*

Proof: Suppose to the contrary that an equilibrium exists in which citizens i and $-i$ adopt cutoff strategies with 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, i is almost sure that $s^{-i} > k^{-i}$ and hence almost sure that $-i$ will not revolt. But then i is almost surely punished if he revolts at s^i, so it is not a best response, a contradiction. □

Lemma 5 shows that the citizens' (sub)game has a unique equilibrium in finite-cutoff strategies when the noise in private signals is vanishingly small.[9]

Lemma 5 *Suppose the noise in citizens' private signals is vanishingly small: $\sigma_\nu^2 \to 0$. Then there is a unique, symmetric equilibrium in finite-cutoff strategies, characterized by the equilibrium cutoff $k = \rho(R, \Omega) - \mu$, where $\rho(R, \Omega) = E[R - \gamma|\Omega]$ is the citizens' forecast of $R - \gamma$ given public information.*

[9]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 following no report can drop discontinuously to zero as a function of the censorship cutoff. Lemma 5 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.

Proof: 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 $k^{-i} = k^i = k$:

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

From Assumption 1, $\lim_{\sigma_\nu \to 0} Pr(s^{-i} < k|s^i = k) = \frac{1}{2}$ and $\lim_{\sigma_\nu \to 0} E[\epsilon|s^i = k, s^{-i} < k] = k$, 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 $k = \rho(R,\Omega) - \mu$. $\square$

In addition to this finite-cutoff equilibrium, there is a cutoff strategy equilibrium 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, the never-revolt equilibrium does not describe the real world. This leads us to assume that citizens coordinate on the finite-cutoff equilibrium. Otherwise, citizens could coordinate on different equilibria after different public news reports so that, for example, if $\gamma = -1$, citizens coordinate on the equilibrium that features revolution (and high welfare), but if $\gamma = -10$, citizens coordinate on the no-revolution, low-welfare equilibrium. This would create perverse censorship incentives, possibly leading a ruler to censor $\gamma = -1$, but not $\gamma = -10$.

A revolution succeeds if and only if both citizens revolt, i.e., if and only if both $s^i < k(\rho)$ and $s^{-i} < k(\rho)$. Moreover, as the noise in signals goes to zero, i.e., as $\sigma_\nu^2 \to 0$, the probability of a successful revolution, $P(\rho)$, goes to the probability that $\epsilon < k(\rho)$. That is,

$$P(\rho) = G(k(\rho)), \text{ where } G(\epsilon) \text{ is the cdf of } \epsilon. \tag{9}$$

Good news is any news that raises citizens' expectations of the status quo payoff, while bad news is the opposite. Recall that $\rho(R,\Omega) = E[R - \gamma|\Omega]$ is the citizens' public knowledge about the expected payoff difference between revolution and the status quo. Thus, bad news increases ρ. Proposition 6, which directly follows from Lemma 5 and equation (9), states the properties of the equilibrium to the citizens' game that we exploit.

Proposition 6 *Bad news about the status quo raises the likelihood of revolution,* $\frac{P(\rho)}{\partial\rho} > 0$*. 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 obtain in settings where citizens know payoffs, but are uncertain about the probability a revolution would succeed if a fraction r of citizens revolt, or equivalently, about a 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 a private value structure (Bueno de Mesquita 2010).

The multi-citizen revolution framework endogenizes the probability that an attempted revolt succeeds because citizens must coordinate for successful revolution. As a result, this probability varies with the parameters of the model. Nonetheless, our results carry over from the representative citizen framework, where the probability r that an attempted revolt succeeded is exogenous.

8 Appendix B: Proofs

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

$$\begin{aligned}
\frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}} &= \frac{q\bar{\gamma}f(\bar{\gamma})\ (1-q+q\ F(\bar{\gamma})) - qf(\bar{\gamma})\ q\int_{-\infty}^{\bar{\gamma}}\gamma f(\gamma)d\gamma}{(1-q+q\ F(\bar{\gamma}))^2} \\
&= \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\left(\bar{\gamma} - \frac{q}{1-q+qF(\bar{\gamma})}\int_{-\infty}^{\bar{\gamma}}\gamma f(\gamma)d\gamma\right) \\
&= \frac{q\ f(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\left(\bar{\gamma} - E[\gamma|\emptyset;\bar{\gamma}]\right) \qquad (10) \\
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\ F(\bar{\gamma})(\bar{\gamma} - E[\gamma|\gamma<\bar{\gamma}])\ \right)
\end{aligned}$$

$$\begin{aligned}
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\ F(\bar{\gamma})\left(\frac{1}{F(\bar{\gamma})}\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma\right)\ \right) \\
&= \frac{q\ f(\bar{\gamma})}{(1-q+q\ F(\bar{\gamma}))^2}\left(\ (1-q)\bar{\gamma} + q\int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma\ \right), \qquad (11)
\end{aligned}$$

where the fifth equality follows from integration by parts. From equation (11),

$$\frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}} = 0 \text{ if and only if } -\frac{1-q}{q}\ \bar{\gamma} = \int_{-\infty}^{\bar{\gamma}}F(\gamma)d\gamma. \qquad (12)$$

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 $\bar{\gamma} = \bar{\gamma}_m$, and from equation (10), $\bar{\gamma}_m = E[\gamma|\emptyset;\bar{\gamma}_m]$. 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 $\bar{\gamma} \in (-\infty, \infty]$. Part 3 follows from equation (5). □

Proof of Proposition 2: From parts 1 and 2 of Lemma 2 and Proposition 1, $\bar{c} = \lim_{\bar{\gamma}\to-\infty} P_{\bar{\gamma}} - P_{\emptyset}(\bar{\gamma}) = r - P_{\bar{\gamma}=0} \in (0, r)$.

Proof of Corollary 2: Note that c does not affect P_γ or $P_\emptyset$, 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 recall that $E[\gamma] = 0$ and write

$$E[\gamma|\emptyset;\bar{\gamma}] = \frac{1-q}{1-q+q\ F(\bar{\gamma})}\ E[\gamma] + \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}\ E[\gamma|\gamma<\bar{\gamma}] \equiv A(\bar{\gamma})E[\gamma|\gamma<\bar{\gamma}],$$

where $A(\bar{\gamma}) = \frac{q\ F(\bar{\gamma})}{1-q+q\ F(\bar{\gamma})}$. Differentiating, we have $A'E[\gamma|\gamma<\bar{\gamma}] + AE'[\gamma|\gamma<\bar{\gamma}] < 0$, since both terms are negative (A increases in $F(\bar{\gamma})$ and $E[\gamma|\gamma<\bar{\gamma}] < 0$; and A is positive and $E[\gamma|\gamma<\bar{\gamma}]$ decreases in σ_0^2 for $\bar{\gamma} < 0$ since greater σ_0^2 places more probability weight on lower values of γ). Thus, to retrieve equality of $c = P_{\bar{\gamma}_e} - P_\emptyset(\bar{\gamma}_e)$, we must reduce $\bar{\gamma}_e$. □

Proof of Corollary 3: From equation (12), $\lim_{q\to1}\bar{\gamma}_m(q) = -\infty$. This together with $\bar{\gamma}_e \le \bar{\gamma}_m$ imply $\lim_{q\to1}\bar{\gamma}_e(q) = -\infty$. By continuity and Corollary 1, $\lim_{c\to0}\bar{\gamma}_e(q,c) = \bar{\gamma}_m(q)$, and from equation (12), $\lim_{q\to0}\bar{\gamma}_m(q) = 0$. □

Proof of Proposition 3: From equation (6),

$$\begin{aligned}\frac{dW}{d\bar{\gamma}} &= -qf(\bar{\gamma})\ P_\emptyset(\bar{\gamma}) - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} + qP_{\bar{\gamma}}f(\bar{\gamma}) - qf(\bar{\gamma})c \\ &= -qf(\bar{\gamma})\ [P_\emptyset(\bar{\gamma}) - P_{\bar{\gamma}} + c] - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} \\ &= qf(\bar{\gamma})\ [(P_{\bar{\gamma}} - P_\emptyset(\bar{\gamma})) - c] - [qF(\bar{\gamma}) + (1-q)]\ \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}}. \qquad (13)\end{aligned}$$

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

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

Proof of Remark 1: From equation (13),

$$\begin{aligned}\frac{d^2W}{d\bar{\gamma}^2} &= q\frac{df(\bar{\gamma})}{d\bar{\gamma}}\,[P_{\bar{\gamma}} - P_\emptyset(\bar{\gamma}) - c] + qf(\bar{\gamma})\,\left[\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} - \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}}\right] \\ &- \; qf(\bar{\gamma})\,\frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} - [qF(\bar{\gamma}) + 1 - q]\,\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2}. \end{aligned} \tag{14}$$

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 (14) simplifies to

$$\frac{d^2W}{d\bar{\gamma}^2} = qf(\bar{\gamma})\,\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} - [qF(\bar{\gamma}) + 1 - q]\,\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2}. \tag{15}$$

Moreover, letting $\mu_r = \frac{1-r}{r}\mu$,

$$\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}} = -rg(R - \mu_r - \bar{\gamma}); \quad \frac{dP_\emptyset(\bar{\gamma})}{d\bar{\gamma}} = -\frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}}\,rg(R - \mu_r - E[\gamma|\emptyset;\bar{\gamma}]). \tag{16}$$

Thus,

$$\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2} = -\,\frac{d^2E[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}^2}\,rg(R - \mu_r - E[\gamma|\emptyset;\bar{\gamma}]), \tag{17}$$

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 (16) and (17),

$$\frac{d^2P_\emptyset(\bar{\gamma})}{d\bar{\gamma}^2} = \frac{d^2E[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}^2}\,\frac{dP_{\bar{\gamma}}}{d\bar{\gamma}}. \tag{18}$$

Substituting equation (18) into equation (15) yields

$$\frac{d^2W}{d\bar{\gamma}^2} = \frac{dP_{\bar{\gamma}}}{d\bar{\gamma}}\left(qf(\bar{\gamma}) - [qF(\bar{\gamma}) + 1 - q]\,\frac{d^2E[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}^2}\right)\ (\text{at } \bar{\gamma} = \bar{\gamma}_e = \bar{\gamma}_m). \tag{19}$$

From equation (10),

$$\begin{aligned}\frac{d^2E[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}^2} &= \frac{d}{d\bar{\gamma}}\left(\frac{q\ f(\bar{\gamma})}{1 - q + q\ F(\bar{\gamma})}\right)\,(\bar{\gamma} - E[\gamma|\emptyset;\bar{\gamma}]) \\ &+ \frac{q\ f(\bar{\gamma})}{1 - q + q\ F(\bar{\gamma})}\,\left(1 - \frac{dE[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}}\right) \\ &= \frac{q\ f(\bar{\gamma})}{1 - q + q\ F(\bar{\gamma})}\ (\text{at } \bar{\gamma} = \bar{\gamma}_e = \bar{\gamma}_m,\ \text{when } c = 0). \end{aligned} \tag{20}$$

Substituting equation (20) for $\frac{d^2E[\gamma|\emptyset;\bar{\gamma}]}{d\bar{\gamma}^2}$ into equation (19) yields

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

Proof of Proposition 4:

$$\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} = \left.\frac{\partial W(\bar{\gamma}(c),c)}{\partial \bar{\gamma}}\right|_{\bar{\gamma}=\bar{\gamma}_e} \frac{d\bar{\gamma}_e(c)}{dc} + \frac{\partial W(\bar{\gamma}_e(c),c)}{\partial c}. \tag{21}$$

From Proposition 2,

$$\frac{d\bar{\gamma}_e(c)}{dc} = \left(\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma} - \frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}}\right)^{-1} = -\left(\frac{dP_{\emptyset}(\bar{\gamma}=\bar{\gamma}_e)}{d\bar{\gamma}} + \left|\frac{dP_{\gamma=\bar{\gamma}_e}}{d\gamma}\right|\right)^{-1}. \tag{22}$$

From equation (6),

$$\frac{\partial W(\bar{\gamma}_e(c),c)}{\partial c} = -qF(\bar{\gamma}_e). \tag{23}$$

Substituting from equations (13), (22), and (23) into equation (21) yields

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

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

Substituting from equations (16) and (10), $\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} > 0$ if and only if

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

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 (24) is zero, so $\frac{dW(\bar{\gamma}_e(c),c)}{dc} < 0$. If $c>0$, rearranging yields $\left.\frac{dW(\bar{\gamma}(c),c)}{dc}\right|_{\bar{\gamma}=\bar{\gamma}_e} > 0$ if and only if

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

As $c \to \bar{c}^- = r - rG(R-\mu_r)$, we have $\bar{\gamma}_e \to -\infty$, so the left-hand side goes to zero, but the right-hand side is positive and bounded away from zero (by logconcavity of F). Thus, $\frac{dW(\bar{\gamma}_e(c),c)}{dc} > 0$ for $c \in [\bar{c}-\delta, \bar{c})$, $\forall \delta > 0$ sufficiently small.

Next, we strengthen the result to prove that $\frac{dW(\bar{\gamma}_e(c),c)}{dc}$ has a single-crossing property. From Proposition 2, $\bar{\gamma}_e$ 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 1). If $R-\mu_r \geq 0$, then $R-\mu_r-E[\gamma|\emptyset;\bar{\gamma}_e]$ is positive and decreasing in c, while $R-\mu_r-\bar{\gamma}_e$ is positive and increasing in c. Then strict unimodality of $G(\epsilon)$ together with the median of $\epsilon \leq E[\epsilon]=0$ imply that $\frac{g(R-\mu_r-\bar{\gamma}_e)}{g(R-\mu_r-E[\gamma|\emptyset;\bar{\gamma}_e])}$ strictly decreases in c. $\square$

Proof of Proposition 5: From equation (6),

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

Define $\mu_r = \frac{1-r}{r}\,\mu$. Thus,

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

Strict unimodality of $G(\epsilon)$ 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 $c=0$ and R is sufficiently low, $\lim_{\bar{\gamma}\to+\infty} W(\bar{\gamma};R)-\lim_{\bar{\gamma}\to-\infty} W(\bar{\gamma};R)>0$. By continuity, this holds for sufficiently small c. Moreover, when g is symmetric, $\int_{-\infty}^{+\infty}[G(R-\mu_r-\gamma)-G(R-\mu_r)]dF(\gamma)$ is negative if and only if $R-\mu_r>0$. $\square$

Proof of Corollary 4: From Proposition 3, $\lim_{q\to0} P_{\emptyset}(\bar{\gamma};q)=P_{\gamma=0}$, $\lim_{q\to1}\bar{\gamma}_e(q)=-\infty$, and $\lim_{q\to1} P_{\emptyset}(\bar{\gamma}_e(q);q)=1$. Thus, from equation (6), $\lim_{q\to0} W(\bar{\gamma},c,q)=1-\lim_{q\to0} P_{\emptyset}(\bar{\gamma};q)=1-P_{\gamma=0}$ and

$$\begin{aligned}
\lim_{q\to1} W(\bar{\gamma}_e(q),c,q) &= 1-\lim_{q\to1} F(\bar{\gamma}_e(q))\ \lim_{q\to1} P_{\emptyset}(\bar{\gamma}_e(q);q)-\ \lim_{q\to1}\int_{\bar{\gamma}_e(q)}^{\infty} P_{\gamma}dF(\gamma) \\
&- \lim_{q\to1} F(\bar{\gamma}_e(q))c \\
&= 1-\int_{-\infty}^{\infty} P_{\gamma}dF(\gamma).
\end{aligned}$$

Thus, $\lim_{q\to0} W(\bar{\gamma}_e(q),c,q)-\lim_{q\to1} W(\bar{\gamma}_e(q),c,q)=\int_{-\infty}^{\infty}(P_{\gamma}-P_{\gamma=0})dF(\gamma)$. The result then follows from the proof of Proposition 5. $\square$

9 References

Anderson, Simon P., and John McLaren, "Media Mergers and Media Bias with Rational Consumers," *Journal of the European Economics Association*, 10 (2012), 831-859.

Arjomand, Said Amir, *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, 1984).

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

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

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

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

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

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

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

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

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

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

Dye, Ron A., "Disclosure of Non-proprietary Information," *Journal of Accounting Re-*

search, 23 (1985), 123-45.

Dye, Ron A., "An Evaluation of 'Essays on Disclosure' and the Disclosure Literature in Accounting," *Journal of Accounting and Economics*, 32 (2001), 181-235.

Dziuda, Wioletta, Strategic Argumentation, *Journal of Economic Theory*, 146 (2011), 1362-97.

Edmond, Chris, "Information Manipulation, Coordination and Regime Change," NBER Working Paper No. w17395, 2011.

Egorov, Georgy, Sergei Guriev, and Konstantin Sonin, "Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data," *American Political Science Review*, 103 (2009), 645-68.

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

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

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

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

Grossman, S. J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24 (1981), 461-84.

Grossman, S. J., and O. D. Hart, "Disclosure Laws and Takeover Bids," *Journal of Finance*, 35 (1980), 323-34.

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

Khalil, Fahad, "Auditing Without Commitment," *RAND Journal of Economics*, 28 (1997), 629-40.

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

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

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 17 (1981), 380-91.

Milgrom, Paul R., and J. Roberts, "Relying on the Information of Interested Parties," *Rand Journal of Economics* 17 (1986), 18-32.

Morris, Stephen, and Hyun S. Shin, "Global Games: Theory and Application," in *Advances in Economics and Econometrics, Eighth World Congress, Volume I,* M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, eds. (New York: Cambridge University Press, 2003).

Mullainathan, Sendhil, and Andrei Shleifer, "The Market for News," *American Economic Review*, 95 (2005), 1031-1053.

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

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

Shadmehr, Mehdi, and Dan Bernhardt, "Collective Action with Uncertain Payoffs: Coordination, Public Signals, and Punishment Dilemmas," *American Political Science Review*, 105 (2011), 829-851.

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

Verrecchia, Robert E., "Essays on Disclosure," *Journal of Accounting and Economics*, 32 (2001), 97-180.

Vives, Xavier, *Oligopoly Pricing: Old Ideas and New Tools*, (Cambridge, MA: MIT Press, 2001).

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