A Model of Censorship, Propaganda, and Repression

Scott Gehlbach
University of Chicago

Zhaotian Luo
University of Chicago

Anton Shirikov
University of Wisconsin–Madison

Dmitriy Vorobyev
Ural Federal University

This version: October 30, 2021

Abstract

We build on recent work on information design to explore the role of censorship, propaganda, and repression in autocratic rule. A government chooses a propaganda technology and a level of censorship to induce desired behavior by a representative citizen. Following receipt of a public signal, which censorship may render less informative, the government further decides whether to invest in repression. In equilibrium, censorship and repression coexist in the same regime; the government economizes on the former so that it is informed about when to invest in the latter. Propaganda, in turn, is uninformative—counter to the suggestion of standard models of persuasion. We show that the desired level of censorship can be implemented by a punishment strategy that induces self-censorship. Other forms of censorship may provide the government with more information than self-censorship, but the government cannot effectively employ such information to improve its welfare.

How do autocrats govern? Roughly speaking, the literature has identified two broad sets of strategies that autocrats use to hang onto power and pursue policy goals: they can manipulate information, through censorship and propaganda, and they can repress. These strategies interact with each other in important ways.

Consider, for example, propaganda. A fundamental constraint of information manipulation is that propaganda must be believed to be effective (Kamenica and Gentzkow, 2011). If propaganda is sufficiently biased, then it won't play a role in citizens' behavior; citizens may choose to ignore it entirely (Gehlbach and Sonin, 2014).

The presence of outside information tightens this constraint. Social media, independent media, printed books—at one time or another, these have all threatened autocrats' monopoly on information. Authoritarian leaders past and present have therefore invested not only in propaganda but in censorship. They remove social-media posts. They block the reporting of independent media. They banned printing presses.

Yet censorship itself is a blunt instrument. As Wintrobe (1998) famously suggested, fear of punishment for saying the wrong thing can discourage citizens from saying anything at all. Other forms of censorship may similarly leave autocrats in the dark. And so autocrats who censor have little information about, say, the level of discontent in society, which they need to know to efficiently invest in repression—a costly last resort that is to be avoided, if possible.

So how do autocrats govern? As Guriev and Treisman (2019, 2020, 2021) suggest in important recent work, autocrats may emphasize either information manipulation or repression; modern autocrats tend to focus more on the former than the latter. Yet many regimes are both informational and repressive. Think of Nazi Germany, for example, with Joseph Goebbels and Leni Riefenstahl on one side, and Heinrich Himmler on the other (Adena et al., 2015; Voigtländer and Voth, 2015).[1] Joseph Stalin's Soviet Union was similarly preoccupied

[1] Even the Night of Long Knives, often viewed as a quintessential act of autocratic repression, included important elements of censorship and propaganda: see Ullrich (2017, pp.

with propaganda, whether through posters, radio, or the new medium of film.[2] A complete theory of autocratic rule must explain such cases.

An additional empirical observation is that autocrats sometimes behave in ways inconsistent with theories that focus on one strategy to the exclusion of others. Models of Bayesian persuasion, for example, suggest that propaganda should mix enough fact with fiction to keep people guessing. This is quite clearly not always the case. State media in Russia, for example, have presented a very one-sided account of the conflict in Ukraine that seems to provide little real information.

Our view is that to understand phenomena such as these, we must examine censorship, propaganda, and repression jointly. This is what we do in this paper, with a model of information manipulation that incorporates some novel elements. In our setting, a receiver (here, the *citizen*) consumes a propaganda message and a public signal, following which she decides whether to take an action that is always in the sender's self interest but not necessarily the citizen's. The sender (the *government*) commits ex ante to a distribution over propaganda messages—it chooses a "propaganda technology"—that can be conditioned not only on the the state of the world, as in the canonical model of Bayesian persuasion, but also on the public signal. The interpretation is that the government designs propaganda institutions to respond to outside information—think of a cable news anchor, chosen for his ideology, who responds to news reported in independent media.

So far, all of this is what Bergemann and Morris (2016, 2019) call *information design with omniscient persuasion*—omniscient, as the propaganda technology is "aware" of the public signal. We provide two innovations. First, we assume that the government can filter negative signals, which is what we call censorship. A key assumption is that a filtered signal reaches neither the citizen nor the government. We can motivate this assumption in various ways, including by thinking of "censorship" as self-censorship by some outside

465–73).

[2] Though preoccupation did not always imply success: see Belodubrovskaya (2017).

actor in anticipation of government sanction—a perspective that we take up in an extension to the baseline model. Second, following receipt of the public signal, we assume that the government decides whether to invest in "repression," that is, the ability to force the citizen to take the desired action. We model this in a very reduced-form way; the key point is that outside information helps the government to invest efficiently.

This model generates some novel results. In equilibrium, the government censors *and* represses. Propaganda itself is uninformative, as when the media always report that the autocrat is competent or that the country is winning the war. Repression, in turn, is employed as a last resort, to be used when censorship fails.

Our model thus suggests that information manipulation and repression may coexist in the same regime. The extent to which the government relies on one versus the other depends not on the cost of repression but on the "difficulty of persuasion," that is, the extent to which the citizen is ex ante disinclined to take the action preferred by the government. As this difference grows, what appear to be "informational autocracies" may become increasingly oppressive. Vladimir Putin's Russia is perhaps an example.

Another interesting implication of our analysis is that what might seem to be a suboptimal propaganda strategy of always parroting the party line can in fact make sense alongside moderate censorship, with repression as a backstop. Autocrats who program the evening news with fawning but uninformative accounts of the leader's are not necessarily or only relying on the credulity of their citizens to persuade (Little, 2017), as would seem to be implied by standard models of propaganda in the Bayesian-persuasion tradition (e.g,. Gehlbach and Sonin, 2014). Rather, they are counting on censorship and repression to do the heavy lifting.

We extend this analysis in two directions. First, we examine a key assumption of our theoretical framework—that censorship deprives not only the citizen but also the government of information about the state of the world. We rationalize this assumption in a model with an informed and uninformed citizen, where the informed citizen self-censors—that is, sometimes chooses not to convey to the uninformed citizen, and thus the government, what it knows about the state of the world—in a manner that is functionally equivalent to censorship in the baseline model. Second, we show that relaxing the assumption that the government never observes the state of the world when it censors information to that effect cannot increase the government's expected payoff or increase the informativeness of propaganda—though for some parameter values it does imply the existence of a second equilibrium in which the government censors more than in the baseline model.

Our approach builds on a dynamic recent theoretical literature on censorship (Besley and Prat, 2006; Egorov, Guriev and Sonin, 2009; Duggan and Martinelli, 2011; Lorentzen, 2014; Shadmehr and Bernhardt, 2015; Guriev and Treisman, 2020), propaganda (Gehlbach and Sonin, 2014; Little, 2017; Guriev and Treisman, 2020), and repression (Svolik 2012, 2013; Tyson, 2018; Montagnes and Wolton, 2019, Paine, 2020, Rozenas, 2020). As a general rule, work in this tradition tends to treat these strategies in isolation. We are aware of two exceptions.[3]

First, as in our paper, Guriev and Treisman (2020) endogenize censorship, propaganda, and repression, broadly defined. From a modeling perspective, the main difference between their approach and ours is that our environment is a model of Bayesian persuasion, whereas theirs is a model of political agency. This implies a difference in the cost of censorship—in Guriev and Treisman, reduced consumption for the citizen; in our framework, less information for the autocrat—and therefore some key differences in results. The uninformativeness of propaganda in equilibrium is unique to our model, for example. Moreover, in Guriev and Treisman, information manipulation and repression are substitutes by assumption; they cannot occur together on the equilibrium path. In contrast, we show that information manipulation and repression can be employed by the same regime in equilibrium—and we demonstrate that even information manipulation, via self-censorship, may ultimately rest on

[3]In related work, Luo and Rozenas (2018) consider the interaction among ex ante and ex post electoral manipulation, which can be understood as alternative forms of information manipulation.

a repressive foundation.

Second, Horz (2021) analyzes a model in which, in contrast to our framework, citizens must exert effort to be Bayesian. His emphasis is on how various environmental factors, including repression, affect this decision. A key result is that repression reduces the importance of propaganda, but also increases its "extremeness," that is, the difference between the propaganda message and the Bayesian posterior. Thus, as in our model, propaganda may be maximal when repression is also available to the autocrat.

1 Model

There are two players: the government and a representative citizen. The government and the citizen share a prior belief over an unknown state of the world $\Theta \in \{B, G\}$, with $\Pr(\Theta = G) = p \in (0, 1)$. The citizen receives a state-dependent payoff from an action $a \in \{a_B, a_G\}$ (we will be explicit about timing and information), where the payoff from a is x if $\Theta = B$ and $a = a_B$, $1 - x$ if $\Theta = G$ and $a = a_G$, and 0 otherwise. The parameter $x \in (0, 1)$ thus measures the minimum probability that the citizen must assign to the event $\Theta = G$ to justify choosing $a = a_G$ in the absence of repression, of which more below. We assume $p < x$, so that the citizen ex ante prefers to take the action a_B.

The government prefers that the citizen always take the action a_G. In pursuit of this goal, the government can employ three instruments: censorship, propaganda, and repression. This works as follows. At the beginning of the game, the government publicly chooses a level of *censorship* $q \in [0, 1]$ and a *propaganda technology* $\tau = (\tau_{Bb}, \tau_{B\varnothing}, \tau_{Gb}, \tau_{G\varnothing})$, where the latter is a distribution over messages that can be conditioned both on the state and the public signal. Following this, the state of the world is realized and an public signal $s \in \{\varnothing, b\}$ is generated, where

$$\Pr(s = b \mid G) = 0,$$

$$\Pr(s = b \mid B) = 1 - q.$$

The variable q thus measures the probability that "bad news" is blocked, with the signal "no news" observed instead. We assume that the government and citizen each observe s but not (yet) the state Θ.

Having observed the public signal, the government publicly decides whether to invest in repression, $\rho \in \{0, 1\}$, at exogenous cost $\kappa \in (0, 1)$. If the government represses, the citizen bears a cost $c > x - p$ if and only if she chooses $a = a_B$, implying that for any posterior belief the citizen strictly prefers to choose the action $a = a_G$.

Following choice of ρ, the citizen observes a propaganda message $m \in \{b, g\}$, which is generated probabilistically according to τ:

$$\tau_{\Theta s} = \Pr(m = g \mid \Theta, s).$$

The propaganda message is thus conditioned both on the state of the world (now implicitly observed by the government) and the public signal—what Bergemann and Morris (2016, 2019) call "omniscient persuasion." Finally, the citizen chooses $a \in \{a_B, a_G\}$.

Summarizing, the timing of events is:

1. The government publicly chooses a level of censorship q and a propaganda technology $\tau = (\tau_{Bb}, \tau_{B\varnothing}, \tau_{Gb}, \tau_{G\varnothing})$.
2. The state of the world $\Theta \in \{B, G\}$ is realized and public signal $s \in \{\varnothing, b\}$ generated, given Θ and q.
3. The government decides whether to invest in repression, $\rho \in \{0, 1\}$, having observed s but not Θ.
4. A propaganda message $m \in \{b, g\}$ is generated according to τ, given the state Θ and public signal s.
5. The citizen chooses $a \in \{a_B, a_G\}$, having observed q, τ, s, ρ, and m.

The citizen's payoffs are defined above. The government maximizes the probability that the citizen takes the action $a = a_G$, net of the cost of repression κ. We assume that the government does not invest in repression upon observing the public signal if indifferent between doing and not doing so, which has no consequence for equilibrium behavior.

Our formalization captures a number of important features of autocratic governance. The assumption that the propaganda message can be conditioned on the public signal as well as the state reflects the idea that the government designs propaganda institutions to respond to outside information—think of a cable news anchor, chosen for his ideology, who responds to news reported online or in independent media. In Russia, for example, the talk-show host Dmitry Kiselyov responded to the Euromaidan Revolution in Ukraine—an event about which Russians were aware from various sources, including family ties—by speciously tying Ukrainian protesters' demands for European integration to the generally unpopular cause of gay rights.[4]

It is worth noting that there would be no change to our substantive results if we assumed that the government committed to a propaganda technology after rather than before observing the public signal, so long as we retained the assumption common to models of Bayesian persuasion that the message can be conditioned on the state of the world. We have more naturally assumed that the propaganda technology is chosen when the level of censorship is, which requires commitment only at a single point in time.

The assumption that the government decides whether to invest in repression when it has only imperfect information about the state of the world captures the need to mobilize security forces, position troops, and otherwise take costly actions that take time to implement. Governments that fail to take such measures can find themselves ill-equipped to deal with popular unrest—much as happened during the August Coup of 1991, when Soviet hardliners failed to arrest Boris Yeltsin and otherwise prepare for potential backlash to the putsch.

Implicitly, repression constitutes investment in the government's capacity to impose a cost on the citizen for choosing the government's less preferred action. On the equilibrium path, this cost is not actually borne, as when $\rho = 1$ the citizen strictly prefers to choose a_G. A similar logic may apply to censorship. As discussed above, fear of punishment may

[4]See, for example, Timothy Snyder, "Fascism, Russia, and Ukraine," *New York Review of Books*, March 20, 2014.

lead citizens to self-censor, thus depriving the regime of useful information (Wintrobe, 1998). "Censorship" may therefore be a complex web of institutions—not only the GlavLit censor with a blue pencil,[5] but also that portion of the security services responsible for monitoring communication and imposing sanctions for undesirable speech. We return to this interpretation in an extension to the baseline model further below.

2 Analysis

To establish a baseline, consider what would happen if neither the government nor the citizen had access to outside information. Absent repression, the citizen would choose a_B with certainty, as the prior $\Pr(\Theta = G) = p < x$ by assumption. Anticipating this, the government would instead invest in repression, which guarantees that the citizen takes the action a_G, for an (expected) payoff to the government of $1 - \kappa$. From the government's perspective, this is an inefficient outcome, as it represses not only when the $\Theta = B$, when the citizen would choose a_B if she knew the state of the world and the government did not repress, but also when $\Theta = G$, when the citizen would choose a_G if she knew the state of the world. The purpose of information manipulation—of the optimal use of censorship and propaganda—is to reduce the inefficient use of repression.

Central to our analysis is the assumption that repression ρ and the propaganda message m are both conditioned on the public signal $s \in \{\varnothing, b\}$. This implies an algorithm:

1. For a given level of censorship q, solve for the optimal propaganda technology $(\hat{\tau}_{Gs}, \hat{\tau}_{Bs})$, given the assumption of no repression at signal s.

2. For that same q, ask: Would the government repress at signal s, given the behavior implied by $(\hat{\tau}_{Gs}, \hat{\tau}_{Bs})$? If so, any (τ_{Gs}, τ_{Bs}) is optimal. If not, the optimal propaganda technology is $(\hat{\tau}_{Gs}, \hat{\tau}_{Bs})$.

[5] The Soviet institution GlavLit (the Main Directory for Literary and Publishing Affairs) was established in 1922 to centralize censorship of printed material within a single agency.

3. Solve for the optimal q, given the propaganda technology and repression strategy implied by (1) and (2).

We proceed accordingly, considering the first two steps jointly and then the third.

Propaganda versus repression

The public signal is either "bad news" ($s = b$) or "no news" ($s = \varnothing$). If the news is bad, then both government and citizen infer that the state $\Theta = B$. No propaganda message can change the citizen's belief, and the only way for the government to induce the citizen to take the action a_G is to repress. Doing so provides a payoff of $1 - \kappa$, versus a payoff of zero from not repressing.

In contrast, when there is no news, then the government and citizen share the posterior belief

$$\Pr(\Theta = G \mid s = \varnothing) = \frac{p}{p + (1-p)q}, \tag{1}$$

which is decreasing in the level of censorship q. If

$$q \leq \hat{q} \equiv \frac{p}{1-p} \cdot \frac{1-x}{x},$$

then the citizen would take the action a_G after observing the signal $s = \varnothing$ in the absence of any additional information. In this case, any propaganda technology $(\tau_{G\varnothing}, \tau_{B\varnothing})$ that has $\tau_{G\varnothing} = \tau_{B\varnothing}$ is optimal. An important example is $(\tau_{G\varnothing}, \tau_{B\varnothing}) = (1, 1)$, in which case propaganda is "maximal." The resulting message is uninformative, and the citizen takes the action a_G with certainty, alleviating the need for repression.

Thus far there is no tradeoff between propaganda and repression. A tradeoff arises if $q > \hat{q}$, such that the citizen's "interim posterior" belief does not justify taking the action a_G. To solve for the optimal propaganda technology in this case, we follow Bergemann and Morris (2016, 2019) in constructing an equilibrium in which the citizen "obeys" the propaganda message, choosing $a = a_G$ when $m = g$ and $a = a_B$ when $m = b$, given the public signal $s = \varnothing$. When the government has not repressed, the citizen obeys the message

$m = g$ if

$$\Pr(\Theta = G \mid s = \varnothing, m = g) \cdot (1 - x) \geq [1 - \Pr(\Theta = G \mid s = \varnothing, m = g)] \cdot x, \tag{2}$$

where the left-hand side of the "obedience constraint" is the expected payoff from choosing $a = a_G$ and the right-hand side is the expected payoff from choosing $a = a_B$. Bayes' rule implies

$$\Pr(\Theta = G \mid s = \varnothing, m = g) = \frac{p\tau_{G\varnothing}}{p\tau_{G\varnothing} + (1-p)q\tau_{B\varnothing}}.$$

Substituting into Condition 2 and simplifying gives

$$\tau_{G\varnothing} \geq q \cdot \frac{1-p}{p} \cdot \frac{x}{1-x} \cdot \tau_{B\varnothing}. \tag{3}$$

Similarly, the citizen obeys the message $m = b$ when $s = \varnothing$ if

$$[1 - \Pr(\Theta = G \mid s = \varnothing, m = b)] \cdot x \geq \Pr(\Theta = G \mid s = \varnothing, m = b) \cdot (1 - x),$$

which using Bayes' rule can be rewritten as

$$\tau_{G\varnothing} \geq q \cdot \frac{1-p}{p} \cdot \frac{x}{1-x} \cdot \tau_{B\varnothing} + \left(1 - q \cdot \frac{1-p}{p} \cdot \frac{x}{1-x} \cdot\right). \tag{4}$$

As $q > \hat{q}$, Condition 3 rather than Condition 4 binds. The optimal propaganda technology therefore follows from choosing $\tau_{G\varnothing} = 1$ (else the government could increase the probability that the citizen takes the action $a = a_G$ when $m = g$ without violating the obedience constraints above), which in turn pins down $\tau_{B\varnothing}$:

$$\hat{\tau}_{B\varnothing} = \frac{1}{q} \cdot \frac{p}{1-p} \cdot \frac{1-x}{x} = \frac{\hat{q}}{q}. \tag{5}$$

Per our algorithm above, for $q > \hat{q}$, we ask whether the government would prefer to repress at $s = \varnothing$, given the behavior implied by $(\hat{\tau}_{Gs}, \hat{\tau}_{Bs}) = \left(1, \frac{\hat{q}}{q}\right)$. For this case, it is optimal for the government to repress if

$$1 - \kappa > \Pr(\Theta = G \mid s = \varnothing) \cdot \hat{\tau}_{G\varnothing} + [1 - \Pr(\Theta = G \mid s = \varnothing)] \cdot \hat{\tau}_{B\varnothing},$$

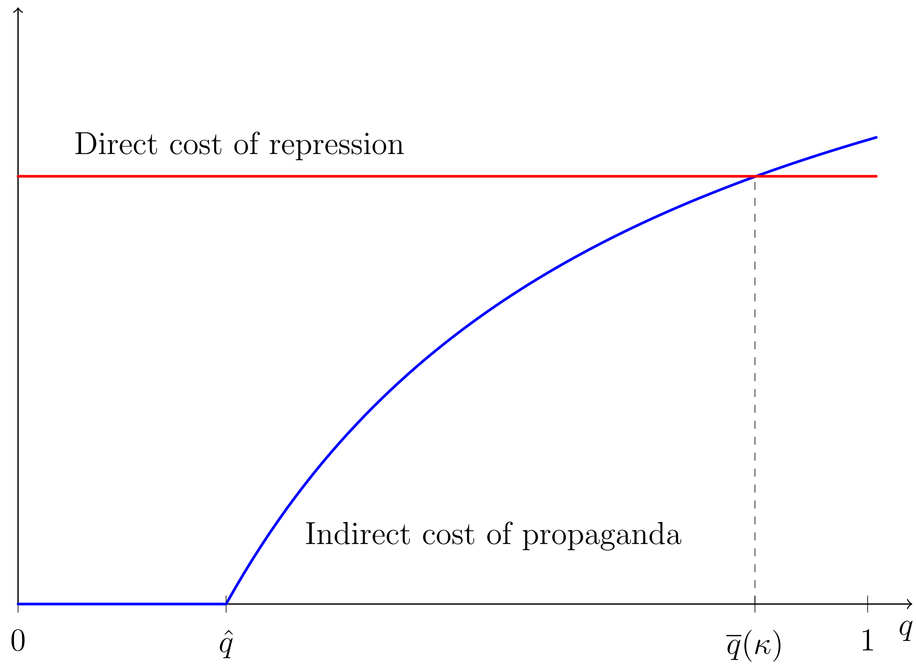

Figure 1: The tradeoff between propaganda and repression when $s = \varnothing$. As depicted, $p < (1 - \kappa)\, x$.

where $\Pr(\Theta = G \mid s = \varnothing)$ is given by Equation 1. Substituting for $(\hat{\tau}_{Gs}, \hat{\tau}_{Bs})$ and simplifying gives

$$q > \tilde{q} \equiv \frac{p}{1-p} \cdot \frac{[1 - (1-\kappa)x]}{(1-\kappa)x}.$$

Observe that $\tilde{q} < 1$ if and only if $p < (1 - \kappa)\, x$.

Figure 1 illustrates the tradeoff between propaganda and repression when the public signal is "no news." Repression entails a direct cost κ that does not depend on the level of censorship. In contrast, there is an indirect cost of propaganda in that, if the level of censorship is sufficiently high (i.e., $q > \hat{q}$), propaganda must be at least somewhat informative to compensate. This informativeness reduces the probability the citizen takes the desired action when the public signal $s = \varnothing$, thus reducing the attractiveness of propaganda as a means to that end. The indirect cost of propaganda surpasses the direct cost of repression as $q = \tilde{q}$, where the threshold $\tilde{q}$ is an increasing function of κ.

We summarize these observations in the following lemma.

Lemma 1. *Fix the level of censorship q.*

- *If $q \leq \hat{q}$, the government represses if and only if $s = b$. Any $(\tau_{G\varnothing}, \tau_{B\varnothing})$ such that $\tau_{G\varnothing} = \tau_{B\varnothing}$ (i.e., any uninformative propaganda technology) can be supported in equilibrium.*
- *If $\hat{q} < q \leq \tilde{q}$, the government represses if and only if $s = b$. The optimal propaganda technology has $(\tau_{G\varnothing}, \tau_{B\varnothing}) = \left(1, \frac{\hat{q}}{q}\right)$.*
- *If $q > \tilde{q}$, the government represses for any signal $s \in \{b, \varnothing\}$. Any $(\tau_{G\varnothing}, \tau_{B\varnothing})$ can be supported in equilibrium.*

Moreover, for any q, any (τ_{Gb}, τ_{Bb}) can be supported in equilibrium. On the equilibrium path (that is, for a given q), the citizen chooses the action a_G with certainty if $s = b$. If $s = \varnothing$, the citizen chooses the action a_G with certainty if $q \leq \hat{q}$ or $q > \tilde{q}$; otherwise, the citizen chooses the action $a = a_G$ with probability $\frac{p}{x}$.

Proof. See above. The probability the citizen chooses the action a_G if $\hat{q} < q \leq \tilde{q}$ and $s = \varnothing$ follows from substituting Equation 5 into

$$p + (1 - p)\, q\tau_{B\varnothing}.$$

□

Lemma 1 says that the citizen is least likely to take the action desired by the government when censorship is moderate. Of course, the government cares not only about the probability that the citizen takes the action a_G but also the cost of inducing that behavior. This takes us to the final step of our analysis.

Censorship

The third step of our algorithm solves for the optimal level of censorship q, given the propaganda technology and repression strategy given by Lemma 1. This involves a straightforward comparison of the government's expected utility for different values of q.

For $q \leq \hat{q}$, the government's expected payoff is increasing in q:

$$[p + (1-p)q] + (1-p)(1-q)(1-\kappa) = 1 - (1-p)(1-q)\kappa.$$

The first term is the probability that $s = \varnothing$, in which case the citizen takes the action a_G with certainty, given that the propaganda technology is uninformative. The second term is payoff to the government from repression, weighted by the probability that $s = b$.

In contrast, for $q \in (\hat{q}, \tilde{q}]$, the government's expected payoff is decreasing in the level of censorship:

$$\frac{p}{x} + (1-p)(1-q)(1-\kappa),$$

where we use the probability $\frac{p}{x}$ that the citizen takes the action a_G when $s = \varnothing$ from Lemma 1. In this region, higher censorship implies a lower probability that the citizen takes the desired action when the public signal is "no news," as discussed above: $\hat{\tau}_{B\varnothing}$ is decreasing in q. At the same time, censorship increases the probability $(1-p)q$ that the public signal is "no news." These two effects precisely cancel out, leaving only the probability-weighted payoff from "bad news" affected by q.

Finally, for $q > \tilde{q}$, the government's expected payoff is $1 - \kappa$, as in this region the government represses for any public signal.

Figure 2, to which we will return, illustrates this analysis. The optimal level of censorship is thus just shy of what would imply informative propaganda: $q^* = \hat{q}$. In equilibrium, the government represses if and only if the public signal is "bad news."

Proposition 1. *In equilibrium, the government chooses censorship* $q = \hat{q} = \frac{p}{1-p} \cdot \frac{1-x}{x}$ *and propaganda technology* $\tau = (\tau_{G\varnothing}, \tau_{Gb}, \tau_{B\varnothing}, \tau_{Bb})$*, where* $\tau_{G\varnothing} = \tau_{B\varnothing}$*. Any* (τ_{Gb}, τ_{Bb}) *can be supported in equilibrium. The government represses if and only if* $s = b$*.*

Proof. See above. □

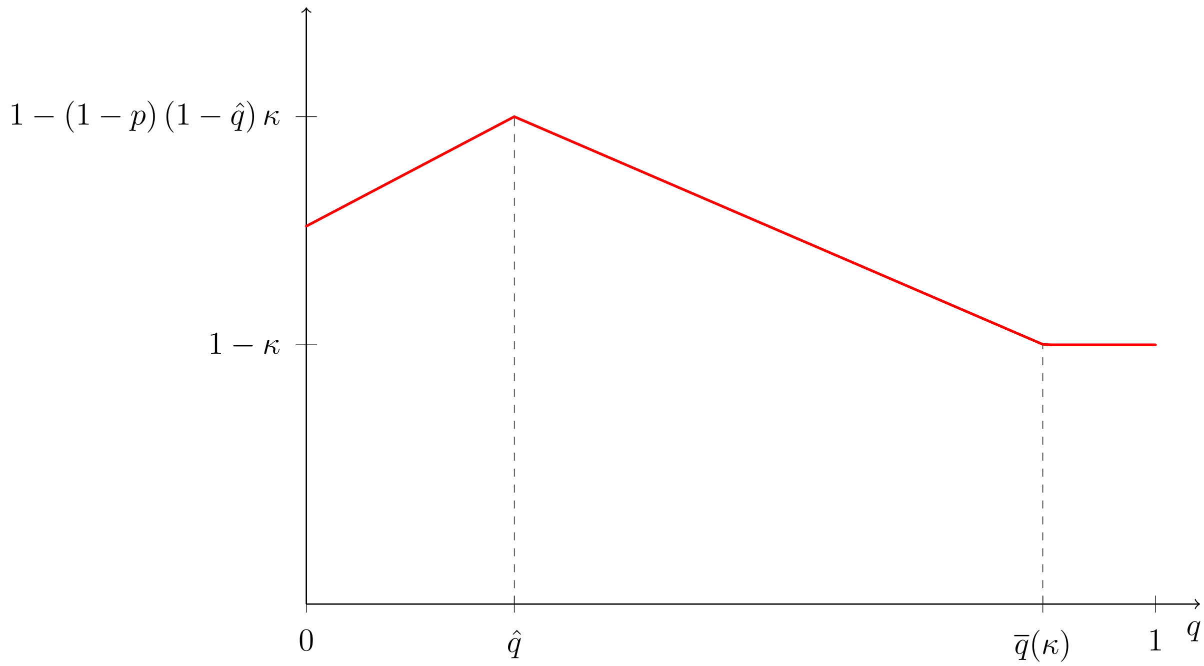

Figure 2: The government's expected payoff as a function of the level of censorship q. As depicted, $p < (1-\kappa)\,x$.

3 Implications

Proposition 1 has a number of implications for our understanding of authoritarian governance. Here we highlight a few that seem especially important.

First, authoritarian governments may employ both repression and information manipulation in pursuit of policy goals and political survival. This prediction most directly concerns the work of Guriev and Treisman (2019, 2020), who suggest that autocrats will typically rely on one or the other set of strategies. In our model, autocrats do both in equilibrium, but repression is costly and best avoided if possible: the point of information manipulation is to more efficiently invest in repression. The extent to which the government employs one versus the other depends on the "difficulty of persuasion," as measured by the relative difference between the prior belief p and the misalignment of preferences x:

$$\Pr(\text{Repression}) = \Pr(s=b) = (1-p)(1-\hat{q}) = \frac{x-p}{p},$$

where the equilibrium level of censorship $\hat{q}$ is given by Proposition 1. As this difference grows, what appear to be "informational autocracies" may become increasingly repressive. Vladimir Putin's Russia may be an example, with a wide crackdown on opposition figures in 2021 coinciding with a decline in Putin's popularity and nervousness about parliamentary elections.[6] Strikingly, the probability of repression does not depend on the cost of repression itself. One should not necessarily infer, for example, that the reason some autocrats repress and others don't is because the former can "get away with it."

Second, our theoretical framework also helps us to understand periodic waves of repression in autocratic regimes. New information technologies may temporarily place limits on the government's ability to censor outside information. Formally, we can think of the government as choosing a censorship level $q \in [0, \bar{q}]$. If $\hat{q} < \bar{q}$, as assumed above (the baseline model has $\bar{q} = 1$), then the government chooses censorship as in Proposition 1 and represses with probability $\frac{x-p}{x}$, as shown just above. If, however, $\bar{q} < \hat{q}$, then the government represses with probability

$$(1-p)(1-\bar{q}) > \frac{x-p}{p},$$

until such time as the government itself acquires the technological and administrative capacity to counter new modes of information transmission.

Third, what may seem to be a suboptimal propaganda strategy of always parroting the party line in fact can make sense alongside moderate censorship, with repression as a backstop. Recall that, in equilibrium, the government chooses a propaganda strategy such that $\tau_{G\varnothing} = \tau_{B\varnothing}$, which provides no additional information when the public signal is "no news." Although there are (infinitely) many ways to implement this strategy, a natural and easy-to-convey instruction is to always report that the state is G, whatever the actual state of the world. Viewed in isolation, this makes little sense: a basic lesson of models of Bayesian persuasion is that the sender must mix enough fact with fiction to keep the receiver guessing.

[6] The respected Levada Center regularly surveys Russians about Putin's popularity: see `www.levada.ru/en`.

The presence of censorship and repression alongside propaganda rationalizes the behavior. In equilibrium, the government censors just enough outside information that the citizen is indifferent between a_G and a_B when $s = \varnothing$. In the event that $s = b$, the government falls back on repression.

As a final implication, recall the observation above that censorship can also act as a sort of repression, discouraging citizens from voicing what they know about the state of the world. To the extent that is the case, our model suggests that today's autocratic regimes may be substantially more repressive than is conventionally understood. Although actual repression is largely off the equilibrium path, the threat of punishment hangs over individuals' speech and—sometimes—other behavior. Propaganda, in this telling, is a sort of sideshow, accomplishing little more than not rocking the boat. We formalize this idea in the following section.

4 Self-censorship, propaganda, and repression

A key assumption of our model is that censorship deprives not only the citizen but the also government of outside information. One can motivate this assumption on various grounds. Perhaps censorship is decentralized, and local agents do not report what they have actively censored to higher-level authorities. Censorship of online speech may prevent other netizens from reacting to (e.g., "liking") censored posts, thus depriving the government of information about their resonance. Online censorship may also push citizens onto parts of the Internet that are less traceable—the so-called "dark web." Finally, and perhaps most importantly, citizens may self-censor in anticipation of active censorship; clearly, the government does not observe self-censored signals.

In this section we explore the last of these possibilities. To do so, we modify our baseline model as follows. As before, there is a sender (government) and a receiver, though now we term the latter the "uninformed citizen." There is also an "informed citizen," who alone observes the state Θ upon its realization. If $\Theta = B$, the informed citizen subsequently decides whether to report "bad news" ($s = b$) or "no news" ($s = \varnothing$). In contrast, if $\Theta = G$, the informed citizen may only report $s = \varnothing$. Let q denote the (endogenous) probability that the informed citizen chooses $s = \varnothing$ when $\Theta = B$.

As in the baseline model, the uninformed citizen chooses an action $a \in \{a_B, a_G\}$. The informed citizen internalizes the payoff from this decision, receiving the payoff x if $\Theta = B$ and $a = a_B$, $1 - x$ if $\Theta = G$ and $a = a_G$, and 0 otherwise. In addition, the informed citizen receives a payoff $\mu > 0$ from reporting $s = b$ when $\Theta = B$, where the parameter μ represents an expressive benefit from "speaking truth to power," whatever the consequences. Our key results require only that this benefit be positive.

In addition to these payoffs, the informed citizen bears an endogenous cost that the government can condition on subsequent play of the game. Formally, prior to other play of the game, the government publicly chooses a *punishment scheme*, $r : H \to \Re^+$, where H is the set of terminal histories of the game. We say that a punishment scheme r *implements* self-censorship q if a) there is a corresponding propaganda technology and repression strategy such that the informed citizen is indifferent between sending the signal $s = b$ and sending the signal $s = \varnothing$ when $\Theta = B$, as is required for the mixed strategy q to be optimal, and b) this propaganda technology and repression strategy are optimal when $\Pr(s = \varnothing \mid \Theta = B) = q$.

A natural but not exclusive interpretation is that the informed citizen is a small elite, whereas the uninformed citizen represents the mass public. This interpretation supports various features of the formalization. Repression, for example, is a crude mechanism that imposes a fixed cost c on the uninformed citizen (mass public) from choosing $a = a_B$, whereas punishment can be fine-tuned to induce particular behavior by the informed citizen (small elite). Repression of the uninformed citizen, moreover, is more costly to government than is punishment of the informed citizen, as would be the case if the uninformed citizen represents a much larger group of actors.

We begin by demonstrating that a simple punishment scheme implements $\hat{q} = \frac{p}{1-p} \cdot \frac{1-x}{x}$, which as shown in the previous section is the government's preferred level of (self-)censorship.

By Lemma 1, when $q = \hat{q}$, the government represses if $s = b$ and adopts a propaganda technology such that the (uninformed) citizen chooses a_G with certainty if $s = \varnothing$. This, in turn, implies that a punishment scheme that imposes a cost μ if and only if the government represses leaves the informed citizen indifferent between these two signals when $\Theta = B$:

$$\overbrace{\mu - \mu}^{\text{Payoff from } s=b} = \underbrace{x \cdot \Pr\left(a = a_B \mid \Theta = B, s = \varnothing\right)}_{\text{Payoff from } s=\varnothing} = 0, \tag{6}$$

where the first expression is the payoff to the informed citizen from choosing $s = b$, which provides the expressive benefit μ at (endogenous) cost μ.

This punishment scheme does not, however, *uniquely* implement $\hat{q}$. By Lemma 1, the government behaves identically, and induces the same behavior by the uninformed citizen, for all $q \in [0, \hat{q}]$. For any q in this interval, the same punishment scheme therefore leaves the informed citizen indifferent between $s = b$ and $s = \varnothing$ when $\Theta = B$, as in Equation 6.

Proposition 2. *A punishment scheme r that imposes a cost μ if and only if the government invests in repression implements any self-censorship $q \in [0, \hat{q}]$. It implements no self-censorship $q > \hat{q}$.*

Proof. See above for $q \in [0, \hat{q}]$. Consider next $q \in (\hat{q}, \tilde{q}]$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression at $s = b$ and induces a_G with probability $\frac{\hat{q}}{q}$ at $\Theta = B$ and $s = \varnothing$. But then the informed citizen has an incentive to deviate to $q = 1$ at $\Theta = B$, as

$$\mu - \mu < x \cdot \left(1 - \frac{\hat{q}}{q}\right).$$

Finally, consider $q > \tilde{q}$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression for any $s \in \{b, \varnothing\}$. But then the informed citizen has an incentive to deviate to $q = 0$ at $\Theta = B$, as $\mu - \mu > -\mu$. □

The proof of Proposition 2 illustrates the importance of the expressive payoff μ to reporting "bad news" ($s = b$) when in fact the state $\Theta = B$. Absent this payoff—which can be arbitrarily small—the punishment scheme above would also implement very high levels of self-censorship ($q > \tilde{q}$).

Proposition 2 also draws a useful distinction between the government's behavior on and off the equilibrium path. On the equilibrium path, the government selectively punishes the informed citizen, imposing the cost μ if and only if the government invests in repression. Intuitively, repression is the informed citizen's "fault," as it follows the signal $s = b$. Off the equilibrium path, in contrast, punishment is uniform: the government imposes the cost μ for any signal $s \in \{b, \varnothing\}$. The latter case may capture situations in which the government has "lost control," repressing both the general population and informed elites, whatever the latter might choose to say.

Proposition 2 might seem to imply that the government would find it difficult to implement its preferred level of self-censorship. In fact, it can get arbitrarily close, as the following proposition demonstrates. The logic of the argument exploits the fact that there is a one-to-one correspondence between (self-)censorship q and the probability the citizen takes the action a_G for $q \in (\hat{q}, \tilde{q}]$.

Proposition 3. *A punishment scheme r that imposes a cost $\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q}+\epsilon}\right)$ if and only if the government invests in repression, with $\epsilon \in (0, \tilde{q} - \hat{q}]$, uniquely implements self-censorship $\hat{q} + \epsilon$.*

Proof. We show here that the punishment scheme implements $\hat{q}+\epsilon$. The proof for uniqueness is in the appendix. By Lemma 1, the government invests in repression at $s = b$ and induces a_G with probability $\frac{\hat{q}}{\hat{q}+\epsilon}$ at $\Theta = B$ and $s = \varnothing$. This, in turn, implies that the citizen is indifferent between sending the signal $s = b$ and $s = \varnothing$ at $\Theta = B$:

$$\mu - \left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q}+\epsilon}\right)\right] = x \cdot \left(1 - \frac{\hat{q}}{\hat{q}+\epsilon}\right).$$

□

Together, Propositions 2 and 3 illustrate that the government can implement its desired level of censorship—or arbitrarily close to it—with a simple punishment scheme that imposes a cost on informed actors at the same time that the government invests in repression against uninformed citizens. The capacity required to carry out this scheme seems comparatively limited, as all instruments of the state are pointed the same direction. The cost to the scheme is that explored in the baseline model: the government deprives itself as well as uninformed citizens of valuable outside information. In the following section we consider the consequences of alternative censorship mechanisms that may leave some information available to the government.

5 Other modes of censorship

In this section we analyze an extension of the model in which censorship blocks more information to the citizen than to the government itself. The environments captured by this extension are varied: direct censorship by top government officials, for example, who observe most of the information denied to the mass public, but also indirect censorship through social-media firms, who might periodically provide to the government a rough accounting of the content they have censored. We capture the diversity of such environments by parameterizing the probability that the government observes the state of the world when the citizen does not.

The setup is identical to the baseline model but for the following. In addition to the public signal s, which is observed by both the government and the citizen, the government receives a private signal $z \in \{\varnothing, b\}$, where

$$\Pr(z = b \mid \Theta = G) = 0$$

$$\Pr(z = b \mid \Theta = B, s = b) = 1$$

$$\Pr(z = b \mid \Theta = B, s = \varnothing) = \phi.$$

When $\Theta = B$ but the public signal s is "no news" ($s = \varnothing$) due to censorship, the government's private signal z is "bad news" with probability ϕ. The parameter $\phi \in (0, 1]$ thus measures the extent to which the government's ability to acquire information is immune to its own censorship effort. Our baseline model assumes $\phi = 0$, in which case the government has no information beyond that possessed by the citizen. We assume that the government can condition the propaganda message on z as well as on Θ and s.

The timing of the extension is as follows:

1. The government publicly chooses a level of censorship q and a propaganda technology $\tau = (\tau_{\Theta sz})$.

2. The state of the world $\Theta = \{B, G\}$ is realized; the public signal s is generated, given Θ and q; and the government's private signal z is generated, given Θ and s.

3. The decides whether to invest in repression, $\rho \in \{0, 1\}$, having observed s and z but not Θ.

4. A propaganda message $m \in \{b, g\}$ is generated according to τ, given the state Θ, public signal s, and private signal z.

5. The citizen chooses $a \in \{a_B, a_G\}$, having observed q, τ, s, ρ, and m.

To ease notation, let $\rho_{sz} \equiv \Pr(\rho = 1 \mid s, z)$ denote the probability that the government invests in repression after observing s and z. We generally suppress mention of ρ_{bb}, as the government trivially represses when $s = z = b$. Further, for simplicity we restrict attention to equilibria in which the obedience constraint holds when there exist multiple equilibria that generate the same behavior by the citizen; in practice this means that the government chooses maximal propaganda when propaganda is uninformative.

In the baseline model, the government may incur an informational cost of censorship. Censorship filters "bad news," such that the government sometimes fails to invest in repression when it should. Intuition suggests that this cost may be smaller, and the government's expected payoff correspondingly higher, when the government's information is less affected by censorship. To investigate whether this is the case, we examine in turn two cases: that in which the government does not exploit the additional information in z, instead repressing if and only if $s = \varnothing$ as in the baseline model, and that in which it does. The following proposition addresses the former case.

Lemma 2. *Fix (q, τ). For the subgame that follows this history, if $\rho_{\varnothing\varnothing} = \rho_{\varnothing b} = 0$ are equilibrium choices, the government's expected utility is bounded above by $1 - (1-p)(1-\hat{q})\kappa$.*

Proof. See online appendix. □

Lemma 2 says trivially that the government cannot do better than in the baseline model if it ignores its private signal z when deciding whether to invest in repression. When this is the case, the government's choice of censorship and propaganda technology is essentially equivalent to that in the baseline model: among all choices of q and τ that lead the government to ignore z, censorship $\hat{q}$ together with uninformative propaganda uniquely maximize the government's expected payoff.

What if the government does use its private information, conditioning its repression decision on z? The following lemma establishes that even in this case the government cannot do better than in the baseline model—though it may do as well by adopting a censorship level higher than in the baseline model.

Lemma 3. *Fix (q, τ). For the subgame that follows this history, if $\rho_{\varnothing\varnothing} > 0$ and/or $\rho_{\varnothing b} > 0$ are equilibrium choices, the government's expected utility is bounded above by $1 - (1-p)(1-\hat{q})\kappa$.*

Proof. See online appendix. □

The upper bounds established by Lemmas 2 and 3, which are identical to the government's expected payoff in the baseline model, are attained for particular choices of q and τ—and, for Lemma 3, only in part of the parameter space. From this, we can immediately characterize two classes of equilibria: one in which the government does not condition its repression decision on its private signal z, and one in which it does.

Proposition 4. *There exists an equilibrium in which* $q = \hat{q}$, $\tau_{\Theta\varnothing z} = 1$ *for all* Θ, z, *and* $\rho_{\varnothing\varnothing} = \rho_{\varnothing b} = 0$. *If* $\phi \leq 1 - \hat{q}$, *there exists an additional class of equivalent equilibria in which* $q = \frac{\hat{q}}{1-\phi}$, $\tau_{G\varnothing\varnothing} = \tau_{B\varnothing\varnothing} = 1$, $\tau_{B\varnothing b} \leq 1 - \kappa$, $\rho_{\varnothing\varnothing} = 0$, *and* $\rho_{\varnothing b} = 1$.

Proof. See online appendix. □

Proposition 4 implies that there may exist multiple equilibria when censorship blocks more information to the citizen than to the government itself. In the first equilibrium, the government behaves as in the baseline model, with $q = \hat{q}$ and investment in repression if and only if the public signal $s = b$. In the second, censorship is more stringent, with $q = \frac{\hat{q}}{1-\phi}$, and the government invests in repression if and only if its private signal $z = b$. Existence of the second equilibrium requires the government's information to be sufficiently limited by its censorship effort ($\phi \leq 1 - \hat{q}$), which says that the information asymmetry between the government and the citizen is not too large.

Although behaviorally distinct, the two equilibria here provide the government the same expected payoff as in the (essentially) unique equilibrium of the baseline model.[7] To understand this result, recall that the informational cost of censorship lies in the possibility that the signal realization $s = \varnothing$ misleads the government into not investing in repression when $\Theta = B$. In the baseline model, this is true if the government's expected payoff from *not* investing in repression when $\Theta = B$ is strictly greater when it knows only that $s = \varnothing$ than would be the case if it observed the state directly:

$$\begin{aligned} \Pr(a = a_G \mid s = \varnothing, \rho = 0) &= \Pr(\Theta = G \mid s = \varnothing)\, \tau_{G\varnothing} + \Pr(\Theta = B \mid s = \varnothing)\, \tau_{B\varnothing} \\ &> \tau_{B\varnothing} \qquad (7) \\ &= \Pr(a = a_G \mid \Theta = B, s = \varnothing, \rho = 0), \end{aligned}$$

[7]Proposition 1 admits a continuum of equilibria, each of which is characterized by an uninformative propaganda technology. Our restriction here to equilibria in which the obedience constraint holds pins down the propaganda technology in Lemma 2 to $\tau_{\Theta\varnothing z} = 1$ for all Θ, z.

that is, if $\tau_{B\varnothing} < \tau_{G\varnothing}$, which implies an informative propaganda technology. Yet as Lemma 1 establishes, propaganda is uninformative in equilibrium in the baseline model. The same is trivially true the first equilibrium in Proposition 4, in which the government ignores the private signal z. In the second class of equilibria, in which the government conditions the repression decision on z, propaganda is uninformative at $s = z = \varnothing$ (i.e., $\tau_{G\varnothing\varnothing} = \tau_{B\varnothing\varnothing}$), which is the only combination of signals for which the government does not repress.

6 Conclusion

We extend the canonical two-state, two-action model of Bayesian persuasion to study an autocratic government's choice among censorship, propaganda and repression. At issue is the action of a representative citizen, whom the government would like to take an action that may not be in the citizen's personal self-interest, were the citizen fully aware of the unknown state of the world. In our framework, the government commits to a propaganda technology—a probability distribution over propaganda messages that is conditioned on both an unobserved state and a public signal—and a level of censorship, which determines the probability that a negative public signal is blocked. Censorship is costly in that it also deprives the government of access to the public signal, which provides information useful in deciding whether to repress rather than persuade.

In equilibrium, the government employs both information manipulation and repression. The former takes the form of censorship; propaganda itself is uninformative. Repression serves as a backstop, to be employed when outside information suggests that the citizen is unlikely to take the government's desired action. The coincidence of information manipulation and repression in the same regime stands in contrast to other theoretical work that suggests that autocrats will tend to choose one or the other.

The presence in equilibrium of uninformative propaganda is an even sharper departure from prior work, which tends to emphasize the importance of mixing fact with fiction in pursuit of persuasion. In our model, censorship does enough work that the citizen is always persuaded to take the government's desired action when she receives a positive signal; optimal propaganda avoids stepping on that message. In an extension we show that such censorship can be indirect, as the threat of punishment encourages informed actors to self-censor. Viewed from this angle, the foundation of autocratic rule is the threat of violence directed either at elites (censorship) or mass publics (repression).

Our model thus helps to explain two important features of many autocratic regimes: the coexistence of information manipulation and repression, and the steady drumbeat of over-the-top propaganda. A seemingly important assumption of this model is that censorship filters signals to the government as effectively as it does to the citizen. In a second extension we show that this assumption can be relaxed without any change in the government's expected payoff or the informativeness of propaganda, though for some parameter values a second equilibrium emerges in which the government censors more than in the baseline model.

Stepping back from the particular results in this paper, we see our model as a template for the comparative analysis of autocratic institutions—as urged, for example, by Gehlbach, Sonin, and Svolik (2016). A particular contribution of our paper is to build on other models of Bayesian persuasion, allowing comparison with a benchmark in which only one institution is at play. Future work can further extend the model that we examine.

Appendix: Proof of Proposition 3

The proof in the main text establishes existence. Here we demonstrate uniqueness.

Consider first $q \in [0, \hat{q}]$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression at $s = b$ and induces a_G with probability 1 at $\Theta = B$ and $s = \varnothing$. But then the informed citizen has an incentive to deviate to $q = 0$ at $\Theta = B$, as

$$\mu - \left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q} + \epsilon}\right)\right] > x \cdot 0.$$

Now consider $q \in (\hat{q}, \hat{q} + \epsilon)$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression at $s = b$ and induces a_G

with probability $\frac{\hat{q}}{q}$ at $\Theta = B$ and $s = \varnothing$. But then the informed citizen has an incentive to deviate to $q = 0$ at $\Theta = B$, as

$$\mu - \left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q} + \epsilon}\right)\right] > x \cdot \left(1 - \frac{\hat{q}}{q}\right),$$

given $q < \hat{q} + \epsilon$.

Now consider $q \in (\hat{q} + \epsilon, \tilde{q}]$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression at $s = b$ and induces a_G with probability $\frac{\hat{q}}{q}$ at $\Theta = B$ and $s = \varnothing$. But then the informed citizen has an incentive to deviate to $q = 1$ at $\Theta = B$, as

$$\mu - \left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q} + \epsilon}\right)\right] < x \cdot \left(1 - \frac{\hat{q}}{q}\right),$$

given $q > \hat{q} + \epsilon$.

Finally, consider $q > \tilde{q}$, and assume to the contrary that the punishment scheme implements q. By Lemma 1, the government invests in repression for any $s \in \{b, \varnothing\}$. But then the informed citizen has an incentive to deviate to $q = 0$ at $\Theta = B$, as

$$\mu - \left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q} + \epsilon}\right)\right] > -\left[\mu - x \cdot \left(1 - \frac{\hat{q}}{\hat{q} + \epsilon}\right)\right].$$

□

References

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *Quarterly Journal of Economics*. 130(4):1885–1939.

Beazer, Quintin H., Charles Crabtree, Chris Fariss, and Holger Kern. 2021. "Censorship by Russian Private Media Firms: Evidence from a Correspondence Experiment." *British Journal of Political Science*. Forthcoming.

Belodubrovskaya, Maria. 2017. *Not According to Plan: Filmmaking under Stalin*. Ithaca, NY: Cornell University Press.

Bergemann, Dirk and Stephen Morris. 2016. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review*. 106(5):586–91.

Bergemann, Dirk and Stephen Morris. 2019. "Information Design: A Unified Perspective." *Journal of Economic Literature*. 57(1):44–95.

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

Chen, Yuyu and David Y. Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review*. 109(6):2294–2332.

Davenport, Christian. 2007. "State Repression and Political Order." *Annual Review of Political Science*. 10:1–23.

DellaVigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova and Ekaterina Zhuravskaya. 2014. "Cross-Border Media and Nationalism: Evidence from Serbian Radio in Croatia." *American Economic Journal: Applied Economics*. 6(3):103–32.

Duggan, John and Cesar Martinelli. 2011. "A Spatial Theory of Media Slant and Voter Choice." *Review of Economic Studies*. 78(2): 640–666.

Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. "Why Resource-Poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review*. 103(4):645–668.

Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies*. 80(4):1422–1458.

Enikolopov, Ruben, Maria Petrova and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review*. 101(7):3253–85.

Gehlbach, Scott. 2010. "Reflections on Putin and the Media." *Post-Soviet Affairs*. 26(1):77–87.

Gehlbach, Scott and Konstantin Sonin. 2014. "Government Control of the Media." *Journal of Public Economics*. 118:163–171.

Gehlbach, Scott, Konstantin Sonin, and Milan Svolik. 2016. "Formal Models of Nondemocratic Politics." *Annual Review of Political Science*. 19:565–584.

Gentzkow, Matthew and Jesse Shapiro. 2004. "Media, Education, and Anti-Americanism in the Muslim World." *Journal of Economic Perspectives*. 18(3):117–133.

Guriev, Sergei and Daniel Treisman. 2019. "Informational Autocrats." *Journal of Economic Perspectives*. 33(4):100–127.

Guriev, Sergei and Daniel Treisman. 2020. "A Theory of Informational Autocracy." *Journal of Public Economics*. Forthcoming.

Hill Jr., Daniel W. and Zachary M. Jones. 2014. "An Empirical Evaluation of Explanations for State Repression." *American Political Science Review*. 108(3):661–687.

Horz, Carlo. 2021. "Propaganda and Skepticism." *American Journal of Political Science*. 65(3):717-732.

Huang, Haifeng. 2015. "Propaganda as Signaling." *Comparative Politics*. 47(4):419–444.

Huber, John D. and Nolan McCarty. 2004. "Bureaucratic Capacity, Delegation, and Political Reform." *American Political Science Review*. 98(3):481–494.

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

Kern, Holger and Jens Hainmueller. 2009. "Opium for the Masses: How Foreign Media Can Stabilize Authoritarian Regimes." *Political Analysis*. 17(4):377–399.

King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review*. 107(2):326–343.

King, Gary, Jennifer Pan, and Margaret E. Roberts. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, not Engaged Argument." *American Political Science Review*. 111(3):484–501.

Little, Andrew. 2017. "Propaganda and Credulity." *Games and Economic Behavior*. 102:224–232.

Lorentzen, Peter. 2014. "China's Strategic Censorship." *American Journal of Political Sci-*

ence. 58(2):402–414.

Lukito, Josephine, Chris Wells, Yini Zhang, Larisa Doroshenko, Sang Jung Kim, Min-Hsin Su, Jiyoun Suk, Yiping Xia, and Deen Freelon. 2018. "The Twitter Exploit: How Russian Propaganda Infiltrated U.S. News." Mimeo, University of Wisconsin–Madison.

Luo, Zhaotian and Arturas Rozenas. 2018. "Strategies of Election Rigging: Tradeoffs, Determinants, and Consequences." *Quarterly Journal of Political Science*. 13(1):1–28

McMillan, John and Pablo Zoido. 2004. "How to Subvert Democracy: Montesinos in Peru." *Journal of Economic Perspectives*. 18(4):69–92.

Mechkova, Valeriya, Daniel Pemstein, Brigitte Seim, and Steven Wilson. 2019. "Measuring Internet Politics: Introducing the Digital Society Project." Digital Society Project Working Paper 2019:1.

Montagnes, B. Pablo and Stephane Wolton. 2019. "Mass Purges: Top-down Accountability in Autocracy. *American Political Science Review*. 113(4):1045–1059.

Paine, Jack. 2020. "Organizing Authoritarian Coercion: Trading Off Insider and Outsider Threats." Unpublished manuscript.

Pan, Jennifer. 2017. "How Market Dynamics of Domestic and Foreign Social Media Firms Shape Strategies of Internet Censorship." *Problems of Post-Communism*. 64(3-4):167–188.

Peisakhin, Leonid and Arturas Rozenas. 2018. "Electoral Effects of Biased Media: Russian Television in Ukraine." *American Journal of Political Science*. 62(3):535–550.

Pemstein, Daniel, Kyle L. Marquardt, Eitan Tzelgov, Yi-ting Wang, Juraj Medzihorsky, Joshua Krusell, Farhad Miri, and Johannes von Römer. 2021. "The V-Dem Measurement Model: Latent Variable Analysis for Cross-National and Cross-Temporal Expert-Coded Data." V- Dem Working Paper 2021:21.

Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda." *Journal of Economic Perspectives*. 31(1):117–140.

Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China's Great Fire-*

wall. Princeton, NJ: Princeton University Press.

Rozenas, Arturas. 2020. "A Theory of Demographically Targeted Repression." *Journal of Conflict Resolution.* 64(7–8): 1254–1278.

Rozenas, Arturas and Denis Stukal. 2019. "How Autocrats Manipulate Economic News: Evidence from Russia's State-Controlled Television." *Journal of Politics.* 81(3):982–996.

Sanovich, Sergey, Denis Stukal, and Joshua A. Tucker. 2018. "Turning the Virtual Tables: Government Strategies for Addressing Online Opposition with an Application to Russia." *Comparative Politics.* 50(3):435–482.

Shadmehr, Mehdi and Dan Bernhardt. 2015. "State Censorship." *American Economic Journal: Microeconomics.* 7(2):280–307.

Svolik, Milan. 2012. *The Politics of Authoritarian Rule.* New York: Cambridge University Press.

Svolik, Milan. 2013. "Contracting on Violence: The Moral Hazard in Authoritarian Repression and Military Intervention in Politics." *Journal of Conflict Resolution.* 57(5):765–794.

Tyson, Scott. 2018. "The Agency Problem Underlying Dictatorship." *Journal of Politics.* 80(4):1297–1310.

Ullrich, Volker. 2017. *Hitler: A Biography. Volume I: Ascent.* London: Vintage.

Voigtländer, Nico and Hans-Joachim Voth. 2015. "Nazi Indoctrination and Anti-Semitic Beliefs in Germany." *Proceedings of the National Academy of Sciences.* 112(26):7931–7936

Walker, Christopher and Robert W. Orttung. 2014. "Breaking the News: The Role of State-Run Media." *Journal of Democracy.* 25(1):71–85.

Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics.* 129(4):1947–1994.