Informational Autocrats, Diverse Societies*

A. Arda Gitmez† Pooya Molavi‡

August 25, 2023

This paper presents a theoretical model of an autocrat who controls the media in an attempt to persuade society of his competence. We base our analysis on a Bayesian persuasion framework in which citizens have heterogeneous preferences and beliefs about the autocrat. We characterize the autocrat's information manipulation strategy when society is monolithic and when it is divided. When the preferences and beliefs in society are more diverse, the autocrat engages in less information manipulation. Our findings thus suggest that the diversity of attitudes and opinions can act as a bulwark against information manipulation by hostile actors.

*We would like to thank Daron Acemoglu, Raphael Boleslavsky, Georgy Egorov, Emin Karagözoğlu, Nicola Persico, Mehdi Shadmehr, Emilie Sartre, Alex Wolitzky and Leifu Zhang for helpful discussions, and various seminar participants for their comments. Earlier versions of this paper were circulated under the title "Media Capture: A Bayesian Persuasion Approach" and "Polarization and Media Bias."

†Department of Economics, Bilkent University. E-mail: arda.gitmez@bilkent.edu.tr.

‡Northwestern University. E-mail: pmolavi@kellogg.northwestern.edu.

1 Introduction

Over the past two decades, many democracies have devolved into hybrid regimes and outright autocracies (Bermeo, 2016; Levitsky and Ziblatt, 2018; Haggard and Kaufman, 2021). From Venezuela's Hugo Chávez to Hungary's Victor Orbán and Russia's Vladimir Putin, politicians who came to power through democratic means have consolidated their control and undermined democratic institutions. Unlike the dictators of the 20th century, this new breed of autocrat does not resort to overt violence. Instead, they maintain power by building support among the masses and winning elections that appear to be democratic. To cultivate their image as competent leaders, they manipulate information by controlling state media (Rozenas and Stukal, 2019), co-opting or pressuring independent media outlets (McMillan and Zoido, 2004; Szeidl and Szucs, 2021), and covertly censoring unfavorable news (Lorentzen, 2014).[1] They are, as Guriev and Treisman (2019, 2020) put it, *informational autocrats*.

But not all autocracies are alike. Even setting aside those that adhere to the 20th-century playbook and completely control the media (such as North Korea), there is still a wide variation in media freedom across informational autocracies. As Egorov and Sonin (2022) note, "media freedom varies a lot across nondemocratic regimes, from levels comparable to mature democracies, to that of totalitarian regimes."[2] The extent of information manipulation within a country also varies substantially over time. For instance, Tai (2014) shows that from 2007 to 2013, the Chinese "propaganda apparatus has banned fewer reports and guided more of them." As Guriev and Treisman (2019) observe, "in Taiwan, an overt dictatorship under Chiang Kai-Shek evolved into an informational autocracy under his son, Chiang Ching-Kuo, in his later years, before transitioning to full democracy in the 1990s."[3] This raises the question of why some societies are capable of preserving a degree of media freedom under autocratic rule while others are totally dominated by information manipulation.

This paper establishes a theoretical link between the distribution of attitudes and opinions in a society and its vulnerability to information manipulation. It argues that autocrats engage in less information manipulation in more diverse societies. Informational autocrats need to fine-tune their manipulation strategies to citizens' attitudes and opinions; greater diversity complicates this task. As attitudes and opinions become more dispersed, it becomes harder for autocrats to con-

[1] Attempts to control media and manipulate information extend beyond autocracies. Examples in democracies abound, from Argentina (Di Tella and Franceschelli, 2011) to Italy (Durante and Knight, 2012), and from Mexico (Stanig, 2015) to the United States (Qian and Yanagizawa-Drott, 2017; Gentzkow and Shapiro, 2008).

[2] A close look at the Global Media Freedom Dataset (GMFD) of Whitten-Woodring and Van Belle (2017) reveals the extent of variation. Among the 196 countries from 1948 to 2010 included in GMFD, 5275 country-year pairs are labeled as "non-democracies" based on their Polity Score. Interestingly, 4456 of such observations are classified as having "Not Free Media" by the authors, whereas the remaining observations have "Imperfectly Free/Free Media" (Whitten-Woodring and Van Belle, 2017, Fig.1, p.184).

[3] Yet another example is the immense pressure on media outlets during Rafael Correa's rule in Ecuador from 2007 to 2017, and the backtracking of these policies by his vice-president and successor, Lenín Moreno: https://www.cjr.org/analysis/ecuador-moreno-correa-supercom-press-freedom.php.

vince their opponents without alienating their supporters. They respond optimally by manipulating information less and allowing for a more free media landscape.

We present this insight using a Bayesian persuasion model, à la Kamenica and Gentzkow (2011), with a population of heterogeneous citizens. The autocrat commits to a public communication policy. Citizens observe the message drawn according to the policy and decide whether to support the autocrat or not. The extent to which the autocrat manipulates information depends on the distribution of citizens' attitudes and opinions. We analyze the autocrat's decision and characterize the optimal information manipulation policy.

At the model's core is the trade-off faced by an informational autocrat. To gain support, he must convince citizens of his competence through information manipulation. He would do so by sending the message that things are "good" as frequently as possible. However, citizens understand that information is manipulated and only act based on the autocrat's communication when they find it informative. As the autocrat manipulates information more, fewer people act based on the autocrat's messages but those who do lend him a higher support. The optimal policy balances these two effects.

The paper's main result establishes that when citizens' attitudes and opinions are more dispersed, the autocrat finds it optimal to engage in less information manipulation. The intuition is best illustrated by introducing some heterogeneity to a Bayesian persuasion model where citizens share identical attitudes and opinions. In the homogeneous model, the optimal strategy involves sending the "bad" message just frequently enough to make citizens indifferent between supporting the autocrat and not upon receiving the "good" message. However, in a society with dispersed attitudes and opinions centered around those in the homogeneous society, this strategy would only secure the support of half the citizens when the "good" message is sent. To gain broader support, the autocrat must appeal to more skeptical citizens, which requires a reduction in information manipulation.

Our main contribution is to illustrate how this reasoning generalizes under a novel partial order on distributions, which captures the idea of dispersion in attitudes and opinions. We start with a model in which citizens are heterogeneous along two dimensions: preferences and initial beliefs about the autocrat. We collapse these two dimensions of heterogeneity into a one-dimensional distribution, which we call the *virtual density*. The virtual density is a sufficient statistic for the distribution of opinions and attitudes when it comes to citizens' support for the autocrat. When the virtual density is single-peaked, the opinions and attitudes in the society are similar to each other, i.e., the society is *monolithic*. Conversely, when the virtual density is single-dipped, there are two large groups with opposing attitudes toward the autocrat, i.e., the society is *divided*. We characterize the optimal persuasion policy in the cases where the virtual density is single-peaked and single-dipped. This allows us to consistently define what it means for the

autocrat to engage in *more/less information manipulation*. It also allows us to introduce a partial order on one-dimensional distributions that captures the idea of *dispersion*. Our comparative statics results establish that when the virtual density is more dispersed, the autocrat engages in less information manipulation under the optimal policy. This result holds in both monolithic and divided societies.

Related Literature. First and foremost, our model contributes to the growing literature on informational autocrats (Guriev and Treisman, 2019, 2020, 2022; Egorov and Sonin, 2022; Gehlbach, Luo, Shirikov and Vorobyev, 2022).[4] A closely related literature studies *media capture*, the idea that politicians exert control over media by co-opting private media (Besley and Prat, 2006), controlling state media (Gehlbach and Sonin, 2014), censoring news (Shadmehr and Bernhardt, 2015; Boleslavsky, Shadmehr and Sonin, 2021), or controlling media's access to information (Ozerturk, 2022); see Prat (2015) and Enikolopov and Petrova (2015) for two comprehensive reviews.[5] We contribute to this literature by establishing that the vulnerability of a society to media capture depends not only on the opinion of the median citizen but also on the dispersion of opinions in society.

Another literature focuses on understanding the variation in information manipulation. As a source of variation, Egorov, Guriev and Sonin (2009) study the natural resource endowment, Von-Doepp and Young (2013) study the threats that governments face, while McGreevy-Stafford (2020) study protests. A closely related literature identifies factors that limit information manipulation. Di Tella, Galiani and Schargrodsky (2012) emphasize first-hand experiences, Durante and Knight (2012), Gläßel and Paula (2020), Knight and Tribin (2022), and Enikolopov, Rochlitz, Schoors and Zakharov (2023) emphasize the existence of alternative media outlets, Qin, Strömberg and Wu (2018) emphasize market competition, and Knight and Tribin (2019) emphasize citizens' ability to "tune out." Our findings contribute to this literature by highlighting the role of diversity of citizens' attitudes and opinions in limiting the extent of information manipulation.

We also contribute to the literature on political consequences of diversity in a society. Unlike the literature on the impact of heterogeneity on conflict (Desmet, Ortuño-Ortín and Wacziarg, 2017; Arbatlı, Ashraf, Galor and Klemp, 2020) and political institutions and governance (Alesina, Baqir and Easterly, 1999; La Porta, Lopez-de Silanes, Shleifer and Vishny, 1999; Lindqvist and Östling, 2010; Galor and Klemp, 2018), our focus is on the impact of population heterogeneity on information manipulation.

Our model of information manipulation follows the Bayesian persuasion approach (Kamenica and Gentzkow, 2011). We contribute to this literature by allowing for receivers with both heteroge-

[4] Also related is the literature on *democratic authoritarianism* (Brancati, 2014) and *competitive authoritarianism* (Levitsky and Way, 2002) in political science. However, those works are less focused on information manipulation and more on the dismantling of democratic institutions.

[5] Corneo (2006); Petrova (2008, 2012), and Alonso and Padró i Miquel (2022) present models of media capture by special interest groups.

neous preferences (Wang, 2015; Alonso and Câmara, 2016a; Kolotilin, Mylovanov, Zapechelnyuk and Li, 2017; Bardhi and Guo, 2018; Chan, Gupta, Li and Wang, 2019; Arieli and Babichenko, 2019; Kerman, Herings and Karos, 2022; Sun, Schram and Sloof, 2022) and heterogeneous priors (Alonso and Câmara, 2016b; Laclau and Renou, 2017; Kosterina, 2022; Innocenti, 2022; Gitmez and Sonin, 2023).[6] To collapse our two-dimensional primitives onto a single dimension, we introduce an object that we call virtual density. Our comparative statics results rely on the partial order we introduce on the virtual density. Kolotilin (2015), Kolotilin, Mylovanov and Zapechelnyuk (2022), Sun, Schram and Sloof (2022), and Curello and Sinander (2022) also conduct comparative statics exercises in Bayesian persuasion settings. Whereas Kolotilin (2015) focuses on changes in welfare, we analyze how the optimal policy changes with parameters of the model. Sun, Schram and Sloof (2022) derive comparative statics results with respect to the sender's preferences. Our comparative statics result complement theirs by focusing on changes in the receivers' characteristics. Kolotilin, Mylovanov and Zapechelnyuk (2022)'s comparative statics are with respect to the receivers' inclination to be persuaded, whereas ours are with respect to the receivers' heterogeneity. Finally, in parallel work, Curello and Sinander (2022) examine the comparative statics of Bayesian persuasion. However, our approach differs from theirs in two key ways. Firstly, while they focus on changes in the sender's value function, we focus on variations in the distribution of receiver types. Secondly, our approach employs a dispersion partial order on distributions, enabling us to view changes in receiver types as changes in the extent of their heterogeneity.

A crucial assumption in the Bayesian persuasion literature is commitment by the sender. In our model, the sender can commit to a public communication policy, which is observed by all receivers. This assumption can be defended on several grounds. First of all, in our setup, persuasion satisfies the credibility assumption of Lin and Liu (2022).[7] Second, the autocrat's policy can viewed as an "editorial policy," which describes the general attitude of media sources, with the details of the coverage to be decided by reporters and editors (Gehlbach and Sonin, 2014). Finally, the outcome under commitment can be seen as a benchmark that describes the best-case scenario for the sender. Under this interpretation, our results characterize an "ideal media landscape" for a politician in a diverse society.

[6] Also related is the literature on information design, which studies the optimal information structure in a game to be played among multiple players (Bergemann and Morris, 2019; Taneva, 2019; Mathevet, Perego and Taneva, 2020; Inostroza and Pavan, 2022).

[7] In particular, we can allow for undetectable deviations by the sender. Since the sender's payoff in our model is additively separable, there is no profitable deviation that gives the same message distribution as the optimal policy. It should be noted that, with heterogeneous priors, the set of undetectable deviations is different for each agent. In such a case, one has to require deviations to be undetectable given the sender's prior.

2 Setup

2.1 The Model

There are two types of agents: an autocrat and a unit measure of citizens, indexed by $r \in [0,1]$. Each citizen takes an action $a_r \in \{0,1\}$, the action being whether she supports the autocrat.

There is an underlying state of the world: $\theta \in \{0,1\}$. We call the $\theta = 1$ state the "good" state. When the state is good, supporting the autocrat is in the citizens' best interest. The $\theta = 0$ state is the "bad" state where it is optimal for citizens not to support the autocrat. Citizen r's payoff when she chooses action a_r and the state is θ is given by

$$u_r(a_r, \theta) = a_r(\theta - c_r), \tag{1}$$

where $c_r \in [0,1]$ is citizen r's cost of supporting the autocrat. If citizen r knew the state, she would support the autocrat (i.e., $a_r = 1$) in the good state and not support her (i.e., $a_r = 0$) in the bad state. Hence, the good state may be interpreted as the state where the autocrat is competent and the bad state as the one where the autocrat is incompetent.

The autocrat wants to persuade citizens to support him regardless of the state of the world. His payoff when citizen r provides support a_r and the state is θ is given by

$$u_s(\{a_r\}_r) = \int_0^1 a_r dr. \tag{2}$$

When the state is good, the autocrat and citizens have common interests, whereas when the state is bad, their interests are opposed. We denote the autocrat's prior that the state is good by $p_s \in (0, 1)$.

Citizens do not learn the state of the world until after they have decided on whether to support the autocrat. Since citizens do not observe the state, they can only act based on their beliefs. The autocrat can influence those beliefs (and the resulting actions) by sending informative messages. To simplify the analysis, we assume that the autocrat can commit to a public communication strategy $\sigma : \{0,1\} \to \Delta(M)$, where $\sigma(\theta)[m]$ is the probability that public message $m \in M$ is generated when the state is θ. The communication strategy represents the policies followed by media controlled by the autocrat and used by him to influence the views of citizens.

Citizens are heterogeneous in both their preferences and their prior beliefs. The heterogeneity of priors captures the idea that even people with identical payoffs may have different perspectives about the likelihood that the autocrat is competent. We let p_r denote citizen r's prior that the state is good and, and let $f(c,p)$ denote the joint density of costs and priors in the population. We take f as a primitive of the model and study how changing this distribution affects the autocrat's optimal policy. We assume that f is common knowledge and continuously differentiable and bounded over its support.

The heterogeneity of perspectives poses a challenge for an autocrat who wants to garner broad support. Convincing different citizens with different preferences and beliefs requires different communication strategies. Yet, communication is public, so the autocrat cannot tailor his messaging strategy to citizens' diverse perspectives. The main characterization result of the paper concerns the optimal way of resolving the inherent tension in convincing different segments of the population.

Timing. The timing of the communication game is as follows:

1. The prior and cost of each citizen are drawn, and citizen r observes (p_r, c_r).
2. The autocrat commits to a strategy σ, which is observed by all citizens.
3. The state is realized, and the autocrat sends the message drawn according to σ.
4. Each citizen r updates her prior and chooses an action a_r.
5. Payoffs are realized.

The solution concept we adopt is the Perfect Bayesian Equilibrium.

2.2 An Equivalent Representative-Citizen Problem

The fact that the autocrat is communicating with a population of heterogeneous citizens complicates her problem. However, the autocrat's optimal strategy can be found by solving a related persuasion problem with a *representative citizen* whose prior coincides with the autocrat's prior.

The key simplification comes from Proposition 1 of Alonso and Câmara (2016b). Consider citizens r and r' with priors p_r and $p_{r'} = p_s$. Since the two citizens observe the same (public) message, their posteriors are related through the following expression:

$$\mu_r = \frac{\mu_{r'} \frac{p_r}{p_{r'}}}{\mu_{r'} \frac{p_r}{p_{r'}} + (1 - \mu_{r'}) \frac{1-p_r}{1-p_{r'}}}, \tag{3}$$

where μ_r and $\mu_{r'}$ denote the posteriors of r and r', respectively.[8] This coupling of posteriors holds regardless of the communication strategy employed by the autocrat. It uniquely pins down the posterior μ_r of *every* citizen r as a function of the posterior of citizen r'—who will be our representative citizen.

Citizen r supports the autocrat if and only if her posterior that the state is good is at least as large as his cost of action; that is, $a_r = 1$ if and only if

$$c_r \leq c(\mu_s, p_r) \equiv \frac{\mu_s \frac{p_r}{p_s}}{\mu_s \frac{p_r}{p_s} + (1 - \mu_s) \frac{1-p_r}{1-p_s}}, \tag{4}$$

[8] Throughout the paper, we use posterior to mean subjective posterior probability of state $\theta = 1$ given an agent's information.

where μ_s denotes the posterior of the representative citizen (who has the same prior as the autocrat). The payoff to the autocrat is the share of the population who supports him:

$$v(\mu_s) = \int_0^1 \int_0^{c(\mu_s,p)} f(p,c)dcdp. \tag{5}$$

The autocrat's problem is thus equivalent to a standard Bayesian persuasion problem with a representative citizen. The autocrat and representative citizen share the common prior p_s that the state is good. The payoff to the autocrat, when he induces a posterior of μ_s for the representative citizen, is given by $v(\mu_s)$, defined in equation (5). Following Kamenica and Gentzkow (2011), we refer to $v(\mu_s)$ as the autocrat's *value function*. Whenever there is no risk of confusion, we drop the s subscript and simply write $v(\mu)$ for the value to the autocrat of inducing posterior μ for the representative citizen.

The value function has several useful properties. First, $v(\mu)$ is increasing in μ. Inducing a higher posterior for the representative citizen results in a higher posterior for every citizen, thus increasing the share of citizens who support the autocrat. Second, $v(0) = 0$ and $v(1) = 1$. When the representative citizen is certain that the state is bad, so is every other citizen. Therefore, no citizen supports the autocrat. Likewise, when the representative citizen is certain that the state is good, every other citizen is also certain that the state is good and supports the autocrat. Finally, $v(\mu)$ is differentiable in μ due to the differentiability of f.

The value function can thus be seen as a differentiable cumulative distribution function. We let $h(\mu) \equiv v'(\mu)$ denote the corresponding density and refer to it as the *virtual density* of the persuasion problem with heterogeneous citizens. The virtual density has an intuitive interpretation: $h(\mu)$ is the density of citizens who are indifferent between taking the two actions whenever the representative citizen's posterior is equal to μ. Importantly, the construction of h allows us to collapse a two-dimensional object into a one-dimensional one. It defines a *belief threshold* for every citizen such that citizen r supports the autocrat if and only if the representative citizen's posterior is above citizen r's threshold. Note that the virtual density is a primitive of the problem—the shape of f determines the shape of h. For example, when citizens have a common prior that coincides with the autocrat's prior, the virtual density reduces to the probability density $f(c)$ of costs in the population.

3 Information Manipulation in Monolithic Societies

3.1 Single-Peaked Distributions

The solution to the autocrat's persuasion problem takes a particularly simple form when the distribution of citizen types satisfies the following condition:

Definition 1. The virtual density $h(\mu)$ is *single-peaked* if there exists some $\tilde{\mu} \in [0,1]$ such that $h'(\mu) > 0$ for all $\mu < \tilde{\mu}$ and $h'(\mu) < 0$ for all $\mu > \tilde{\mu}$.

Single-peakedness is an assumption on the joint distribution of citizens' costs and prior beliefs. It requires a large share of citizens to have moderate preferences and beliefs, with fewer and fewer people having extreme preferences or beliefs. We thus consider single-peaked virtual densities to be representative of *monolithic* societies.

The significance of Definition 1 rests on the following observation: when the virtual density is single-peaked, the autocrat's value function is first convex and then concave. Therefore, Corollary 2 of Kamenica and Gentzkow (2011) implies the following characterization of the optimal strategy:

Proposition 1. *If the virtual density is single-peaked, the optimal strategy uses only two messages, and one of the messages fully reveals the bad state.*

We maintain the assumption of single-peakedness throughout this section. We do so in part for tractability. However, single-peaked distributions also constitute a natural and widely used class of distribution functions. In Section 4, we show that the optimal strategy in the case where the virtual density is instead single-dipped is the mirror image of the optimal strategy in the single-peaked case.

Whether the virtual density is single-peaked only depends on the distribution of types, f, and the autocrat's prior, p_s. In the remainder of this subsection, we find a set of easy-to-check sufficient conditions for the virtual density to the single-peaked. If citizens have a common prior that coincides with the autocrat's prior, then the single-peakedness of the virtual density is equivalent to the single-peakedness of the density of costs:

Proposition 2. *Suppose $p_r = p_s$ for all r. The virtual density $h(\mu)$ is single-peaked in μ if and only if the density of costs $f(c)$ is single-peaked in c.*

If citizens instead have a common cost, then the single-peakedness of the virtual density is implied by a condition that is weaker than the log-concavity of the density of priors:

Proposition 3. *Suppose $c_r = c \in (0,1)$ for all r. The virtual density $h(\mu)$ is single-peaked if the density of priors $f(p)$ is strictly positive for all $p \in (0,1)$ and satisfies*

$$\frac{d^2}{dp^2}\log f(p) < 2(\gamma-1)^2 \min\left\{1, \frac{1}{\gamma^2}\right\} \quad \text{for all } p \in (0,1), \tag{6}$$

where $\gamma \equiv \frac{1-c}{c}\frac{1-p_s}{p_s} \geq 0$.

The following corollary of Proposition 3 is a straightforward consequence of the facts that the left-hand side of equation (6) is negative if $f(p)$ is log-concave, while its right-hand side is always non-negative:[9]

Corollary 1. *Suppose $c_r = c$ for all r. The virtual density $h(\mu)$ is single-peaked in μ if the density of priors $f(p)$ is strictly log-concave in p.*

[9] See Bagnoli and Bergstrom (2005) for a list of well-known distributions satisfying log-concavity.

3.2 A Measure of Information Manipulation

In light of Proposition 1, we can assume without loss that the autocrat uses only two messages. We label the messages $m \in M = \{0,1\}$, with $m = 1$ the "good" message, which is suggestive of $\theta = 1$, and $m = 0$ the "bad" message, which is suggestive of $\theta = 0$. The autocrat's strategy can be represented by a pair of numbers:

$$\sigma = (\sigma^0, \sigma^1) \in [0,1]^2,$$

where $\sigma^\theta \equiv \sigma(\theta)[m = 1]$ is the probability of sending the good message in state $\theta \in \{0,1\}$. Throughout, we assume without loss of generality that $\sigma^1 \geq \sigma^0$.

The autocrat *manipulates information* if he sends the good message when the state is bad or sends the bad message when the state is good. By Proposition 1, when the virtual density is single-peaked, the bad message fully reveals the bad state; this entails sending the good message whenever the state is good, i.e., $\sigma^1 = 1$. Therefore, in the single-peaked case, the extent of information manipulation is conveniently summarized by the probability σ^0 of sending the good message when the state is bad. We use the following notion of information manipulation in this case:

Definition 2. Consider single-peaked virtual densities h_1 and h_2 with the corresponding optimal strategies $\sigma_1 = (\sigma_1^0, \sigma_1^1)$ and $\sigma_2 = (\sigma_2^0, \sigma_2^1)$ for the autocrat. The autocrat *manipulates information less* given h_1 than given h_2 if $\sigma_1^0 \leq \sigma_2^0$.

3.3 A Measure of Dispersion

To study how diversity affects information manipulation, we need to introduce a measure of dispersion. Our measure is a novel partial order on probability distributions:

Definition 3. Consider two single-peaked distributions with densities f_1 and f_2 supported on a common compact set. f_1 is *more dispersed* than f_2 if

$$f_2(x) = \alpha\left(f_1(x)\right) \qquad \text{for all } x, \tag{7}$$

for some strictly increasing and convex function $\alpha : \mathbb{R}^+ \to \mathbb{R}^+$.

This partial order has an intuitive interpretation. Since α is increasing, f_2 is single-peaked whenever f_1 is single-peaked. Since α is convex and f_1 and f_2 both have to integrate to one, transforming f_1 by α magnifies the parts of f_1 with larger values and shrinks the parts with smaller values. Moving from f_1 to f_2 thus moves mass from parts of the distribution that initially have smaller mass to parts with larger initial mass. In other words, f_2 looks like f_1, but with higher peaks and deeper troughs. But since f_1 is single-peaked, most of its mass is concentrated around its peak. Therefore, f_2 has even more mass in the center and even less mass in the periphery

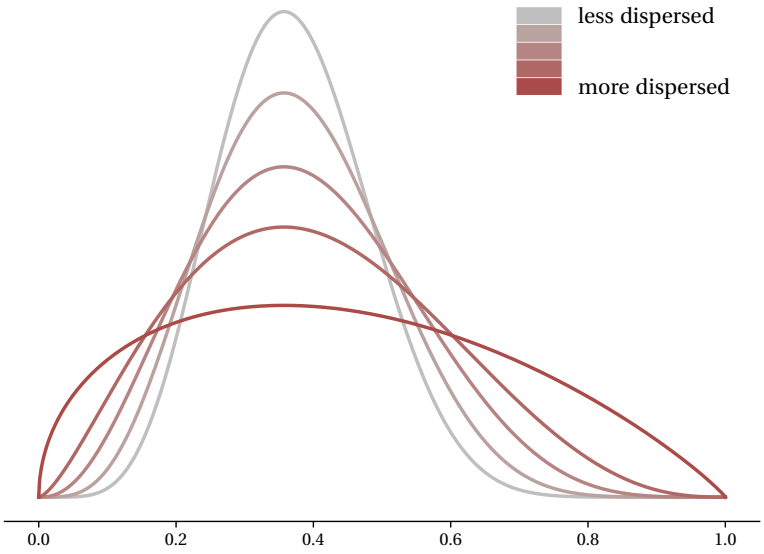

Figure 1. The dispersion order on single-peaked densities.

relative to f_1; that is, f_1 is more dispersed than f_2. Figure 1 illustrates the probability density functions for a set of single-peaked Beta distributions that are ranked in the dispersion order.

Members of many parametric families of distributions can be ordered in the dispersion order. Two examples follow:

Example 1. Consider two single-peaked Beta distributions:

$$f_1 = \text{Beta}(\alpha_1, \beta_1),$$
$$f_2 = \text{Beta}(\alpha_2, \beta_2),$$

where $\frac{\alpha_1-1}{\alpha_1+\beta_1-2} = \frac{\alpha_2-1}{\alpha_2+\beta_2-2}$. If $\alpha_1 \geq \alpha_2$, then f_2 is more dispersed than f_1, while if $\alpha_1 \leq \alpha_2$, then f_1 is more dispersed than f_2. In particular, any two single-peaked Beta distributions with the same mode are ranked according to the dispersion partial order.

Example 2. Consider the following truncated normal distributions on $[0, 1]$:

$$f_1 = \text{TruncatedNormal}(\mu, \sigma_1^2),$$
$$f_2 = \text{TruncatedNormal}(\mu, \sigma_2^2).$$

If $\sigma_2^2 \geq \sigma_1^2$, then f_2 is more dispersed than f_1.

Johnson and Myatt (2006)'s *rotation order* is a related partial order, which also ranks distributions in terms of their dispersion or heterogeneity. The main difference between the two is that Johnson and Myatt (2006) consider rotations of a cumulative distribution function around a given point, whereas in our partial order the rotation point itself depends on the distribution function. The endogeneity of the rotation point to the distribution function is crucial for our comparative

statics results. It ensures that the rotation point is always in the appropriate range for an increase in dispersion to have an unambiguous effect on the extent of information manipulation.

3.4 Dispersion and Information Manipulation in Monolithic Societies

We are now ready to examine how dispersion affects information manipulation. Our main result establishes that information manipulation is less severe in more diverse societies:

Theorem 1. *Let h_1 and h_2 be two single-peaked virtual densities. If h_1 is more dispersed than h_2, then the autocrat manipulates information less given h_1 than h_2.*

To gain some intuition, recall that each citizen is assigned a belief threshold, which ranks them by their inclination to support the autocrat—those with lower thresholds are more supportive. The autocrat's optimal strategy targets a marginal citizen who, upon receiving the positive message, is just indifferent between supporting the autocrat and not. Citizens with lower thresholds act in line with the message, while those with higher thresholds never support the autocrat. The optimal strategy balances the autocrat's goal of maximizing the mass of supporters with that of maximizing the support frequency. In a less dispersed society, the belief thresholds are tightly concentrated around the modal citizen's threshold. Therefore, targeting a citizen whose threshold is slightly above the mode ensures the support of almost every citizen. But in a diverse society with more dispersed thresholds, this same approach yields too few supporters. To counter this, the autocrat needs to increase the informativeness of the media, appealing to those with belief thresholds further from the mode.

Theorem 1 describes the impact of dispersion on information manipulation while maintaining the assumption that the society is monolithic, and so, the virtual density is single-peaked. In the next section, we study persuasion in highly divided societies, in which there are more people in the extremes of preference and belief distribution than at its center.

4 Divided Societies

Throughout this section, we study the properties of the optimal persuasion strategy when the virtual density is the polar opposite of single-peaked:

Definition 4. The virtual density $h(\mu)$ is *single-dipped* if there exists some $\tilde{\mu} \in [0,1]$ such that $h'(\mu) < 0$ for all $\mu < \tilde{\mu}$ and $h'(\mu) > 0$ for all $\mu > \tilde{\mu}$.

In a society with a single-dipped virtual density, there are fewer moderates than those with extreme preferences or beliefs. Therefore, we consider single-dipped virtual densities to be representative of *divided* societies.[10]

[10] Following Fiorina and Abrams (2008, Figure 1), one may also call such a society *polarized*. We refrain from adopting this terminology, because polarization is typically visualized as the society having a small number of groups, with high homogeneity within groups and high heterogeneity across groups (Esteban and Ray, 1994).

As in the single-peaked case, the single-dippedness of the virtual density only depends on the distribution of types and the autocrat's prior. For example, if citizens have a common prior that coincides with the autocrat's prior, the single-dippedness of the virtual density is equivalent to the single-dippedness of the density of costs:

Proposition 4. *Suppose $p_r = p_s$ for all r. The virtual density $h(\mu)$ is single-dipped in μ if and only if the density of costs $f(c)$ is single-dipped in c.*

If citizens have a common cost, the single-dippedness of the virtual density is implied by a condition that is stronger than the log-convexity of the density of priors:

Proposition 5. *Suppose $c_r = c$ for all r. The virtual density $h(\mu)$ is single-dipped if*

$$\frac{\partial^2}{\partial p^2} \log f(p) > 2(\gamma - 1)^2 \max\left\{1, \frac{1}{\gamma^2}\right\} \quad \text{for all } p \in [0,1], \tag{8}$$

where $\gamma \equiv \frac{1-c}{c}\frac{1-p_s}{p_s} \geq 0$.

Note that the right-hand side of equation (8) is positive, and the left-hand side is positive if $f(p)$ is log-convex. Therefore, condition (8) can be interpreted as "$f(p)$ being sufficiently log-convex." If $p_s + c = 1$, then $\gamma = 1$, and condition (8) reduces to the log-convexity of the distribution of priors.

When the virtual density is single-dipped, the autocrat's value function is first concave and then convex. Our next result characterizes the optimal persuasion strategy in this case.

Proposition 6. *If the virtual density is single-dipped, the optimal strategy uses only two messages, and one of the messages fully reveals the good state.*

A comparison of Propositions 1 and 6 reveals that the optimal persuasion strategies are qualitatively different in monolithic and divided societies. In a divided society, there are many strong supporters (i.e., those with belief thresholds close to zero) and many strong skeptics (i.e., those with belief thresholds close to one). The autocrat's challenge is to convince the skeptics without alienating his supporters. The optimal strategy is to use a media source that frequently sends the bad message, so that the rare but credible occurrence of the good message is sufficient to convince even the most skeptical citizens.[11][12] The strong supporters then have no reason to follow the media, because they support the autocrat even when the bad message is realized.

[11] One may interpret the optimal persuasion strategy as the existence of a limited number of independent media that are often critical of the autocrat. The generally critical coverage by such media lends them credibility, allowing the autocrat to benefit from their positive coverage in times of crisis. Such strategies are indeed employed by informational autocrats from time to time. For instance, following the anti-government protests and riots in Zhanaozen in December 2011, Kazakhstan's President Nursultan Nazarbayev suffered from a lack of credibility of the state broadcasting outlets. When all else failed to calm the public, the government invited six well-known bloggers, most labeling themselves as "independent," to make a two-day visit to Zhanaozen. The bloggers carried a sense of credibility that the government sources lacked, and they were "quite effective at reassuring readers that the city was outwardly calm, that rumors of morgues or hospitals full of corpses were unfounded and that shops were well-stocked and inhabitants able to buy food and drink" (Lewis, 2016, p.267, also see Guriev and Treisman, 2022, p.79). In a similar episode, Vladimir Putin utilized the liberal Russian radio station Echo of Moscow to cover a credible account of a large pro-government demonstration in the capital in early 2012, thereby discouraging participation in opposition rallies elsewhere (Sobolev, 2023).

[12] Baum and Groeling (2009), Ladd and Lenz (2009), and Chiang and Knight (2011) document evidence of the persuasive power of communication when messages are sent by actors least expected to send them.

When the virtual density is single-dipped, the autocrat's optimal strategy entails sending the bad message whenever the state is bad. Then, the extent of information manipulation is summarized by the probability σ^1 of sending the good message when the state is good:

Definition 5. Consider single-dipped virtual densities h_1 and h_2 with the corresponding optimal strategies $\sigma_1 = (\sigma_1^0, \sigma_1^1)$ and $\sigma_2 = (\sigma_2^0, \sigma_2^1)$ for the autocrat. The autocrat *manipulates information less* given h_1 than given h_2 if $\sigma_1^1 \geq \sigma_2^1$.

We now examine the impact of increased dispersion on information manipulation. The following partial order is the appropriate adaptation of the partial order defined in Section 3.3 for single-peaked densities to the set of single-dipped densities:

Definition 6. Consider two single-dipped distributions with densities f_1 and f_2 supported on a common compact set. f_2 is *more dispersed* than f_1 if

$$f_2(x) = \alpha\left(f_1(x)\right) \qquad \text{for all } x, \tag{9}$$

for some strictly increasing and convex function $\alpha : \mathbb{R}^+ \to \mathbb{R}^+$.

Figure 2 illustrates the dispersion partial order on a set of a single-dipped Beta distributions. As the distribution becomes more dispersed, mass is moved from the center of the distribution to its tails.

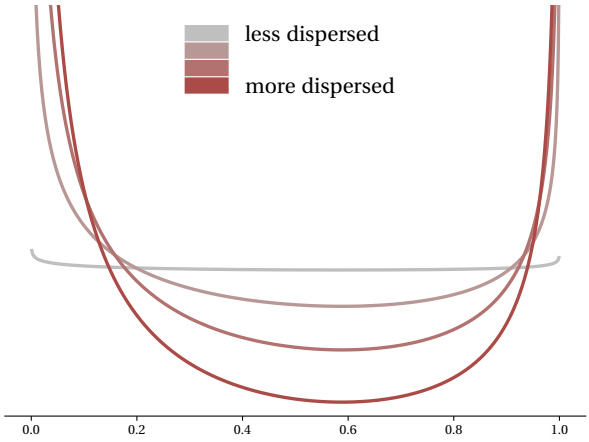

Figure 2. The dispersion order on single-dipped densities.

Our next result establishes that, here, as in the single-peaked case, dispersion reduces information manipulation:

Theorem 2. *Let h_1 and h_2 be two single-dipped virtual densities. If h_1 is more dispersed than h_2, then the autocrat manipulates information less given h_1 than h_2.*

Theorem 2 shows that the main message of Theorem 1 continues to hold in divided societies: Dispersion of attitudes and opinions reduces the extent of information manipulation.

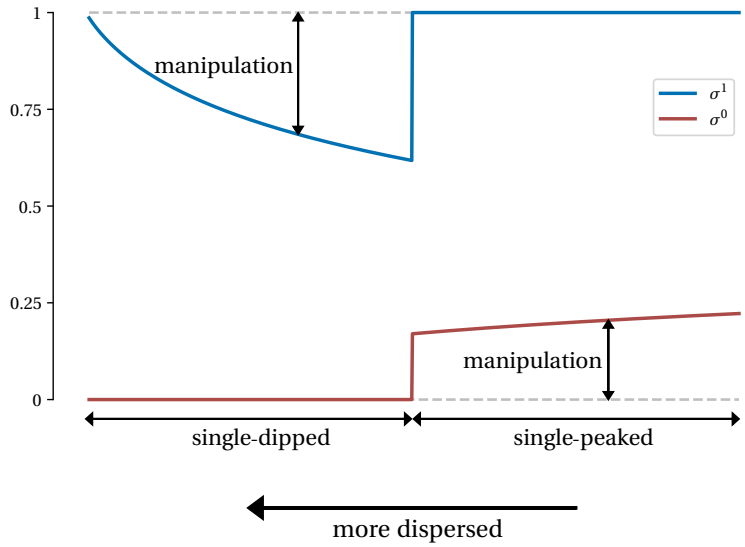

Figure 3. Autocrat's information manipulation as a function of dispersion in society.

The effect of dispersion on information manipulation can be succinctly summarized in a single figure by considering a parametric family of distributions that span both single-peaked and single-dipped cases. Figure 3 illustrates the effect of dispersion on information manipulation for the case where the virtual density is a (symmetric) Beta$(1+\alpha, 1+\alpha)$ distribution, and the autocrat's prior is given by $p_s = 0.4$. The figure plots how the autocrat's optimal strategy changes as α ranges from -1 to $+1$. In the right half of the figure, $\alpha > 0$, the distribution is single-peaked, and so, by Proposition 1, the optimal policy has the form $(\sigma_\alpha^0, \sigma_\alpha^1) = (\sigma_\alpha^0, 1)$. As the society becomes more dispersed, by Theorem 1, σ_α^0 decreases and the autocrat manipulates information less. On the left half of the figure, $\alpha < 0$, the distribution is single-dipped, the optimal policy has the form $(\sigma_\alpha^0, \sigma_\alpha^1) = (0, \sigma_\alpha^1)$ (by Proposition 6), and σ_α^1 increases and information manipulation decreases with dispersion (by Theorem 2).[13]

5 Conclusion

The growing literature on the rise of informational autocrats (Guriev and Treisman, 2022) discusses the modern autocrats' tendency to manipulate information. A natural question that follows from this research is about the conditions that make a society more susceptible to information manipulation. In this paper, we show that the dispersion of opinions puts a limit on an informational autocrat's ability to manipulate information.

To provide empirical support for this prediction, one needs to find variables that capture the dispersion of attitudes and opinions in a society. One readily available, albeit imperfect measure

[13] Transitioning from a single-peaked to a single-dipped virtual density changes the nature of the autocrat's optimal policy. This makes it hard to compare the extent of information manipulation between the single-peaked and single-dipped cases.

of heterogeneity is the Gini coefficient.[14] Petrova (2008) provides evidence that suggests a link between income inequality and media freedom in autocracies. Figure 2 of Petrova (2008) shows that, within autocracies (classified as countries with Democracy score ≤ 1 in Polity IV dataset), there is a positive association between the Gini coefficient and Freedom House's media freedom index in 1994–2003. Reassuringly, the corresponding association is negative for countries classified as democracies in that period (Figure 1 of Petrova, 2008), suggesting that the lack of functioning democratic institutions is an essential part of this story.

Throughout our analysis, we considered the distribution of opinions and attitudes to be exogenous, and we remained agnostic about the forces that may increase its dispersion. Two channels that may lead to increased dispersion are independent media and online media. In a recent working paper, Enikolopov, Rochlitz, Schoors and Zakharov (2023) demonstrate that access to independent online TV in Russia before the 2016 elections had asymmetric effects on individuals who rely on news from social media. Specifically, it bolstered the support among supporters of the regime while leading to a decline in support among those who opposed the regime. Motivated by their findings, and in light of the discussion here, one can argue that online media not only affect the attitudes of citizens but also have an impact on the effectiveness of traditional state-controlled media. In particular, online media do not have to convince every citizen—as long as they influence the opinions of *some* citizens, they could make it harder for the informational autocrat to engage in information manipulation.

In this paper, we focused on information manipulation as the only tool available to an autocrat. In reality, many autocrats have other tools at their disposal, such as repression and indoctrination (Gitmez and Sonin, 2023; Gehlbach, Luo, Shirikov and Vorobyev, 2022), even if they do not always use them. The question of how the mix of tools used by autocrats is affected by the distribution of opinions is a fruitful avenue for future research.

[14] Although only a measure of income heterogeneity, the Gini coefficient has been shown to be related to social conflict (Rodrik, 1999) and lack of social cohesion (Easterly, Ritzen and Woolcock, 2006).

Appendices

A Proofs for Section 3

Because f is continuously differentiable over its support and bounded, $h'(\mu)$ exists and is continuous. We begin by noting that the virtual density is single-peaked if and only if $h'(\mu)$ satisfies the *strict single-crossing-from-above property*. The strict single-crossing property is adapted from (Milgrom and Shannon, 1994, p.160) and is as follows:

$$\text{If } h'(\mu) \geq 0 \text{ for some } \mu \in [0,1], \text{ then } h'(\tilde{\mu}) > 0 \text{ for all } \tilde{\mu} < \mu.$$

In our proofs, we rely on the equivalence of this condition with single-peakedness of h.

Proof of Proposition 1. If $h'(\mu)$ satisfies the strict single-crossing-from-above condition, by definition, so does $v''(\mu)$. Therefore, whenever $v(\mu)$ is convex at μ, it is strictly convex at any $\tilde{\mu} < \mu$. This means that $v(\mu)$ is first strictly convex and then strictly concave. Therefore, the set where the concave closure of $v(\mu)$—call it $V(\mu)$—coincides with $v(\mu)$ has the following form:

$$\{\mu \in [0,1] : V(\mu) = v(\mu)\} = \{0\} \cup [\hat{\mu}, 1],$$

for some $\hat{\mu} \in [0,1]$.

When $p_s < \hat{\mu}$, by Corollary 2 of Kamenica and Gentzkow (2011), the optimal policy generates two posteriors: $\mu \in \{0, \hat{\mu}\}$. This is achieved by two messages, with one perfectly revealing the bad state.

When $p_s \geq \hat{\mu}$, the optimal policy is not revealing any information. This can also be achieved by two messages, $m \in \{0,1\}$, and an information structure where $\Pr(m = 1|\theta = 0) = \Pr(m = 1|\theta = 1)$. Message $m = 0$ will occur with probability zero, and the posterior beliefs following $m = 0$ will be free in a Perfect Bayesian Equilibrium. One can assign the posterior $\Pr_r(\theta = 0|m = 0) = 1$ assign $m = 0$ as the message that perfectly reveals the bad state. □

Proof of Proposition 2. Since $p_r = p_s$ for all r, equation (5) simplifies to

$$v(\mu) = \int_0^{c(\mu,p_s)} f(c)dc.$$

On the other hand, by definition, $c(\mu, p_s) = \mu$ for all μ. Therefore, $h(\mu) = v'(\mu) = f(c(\mu, p_s)) = f(\mu)$, and so, h is single-peaked if and only if f is single-peaked. □

Proof of Proposition 3. When $c_r = c$ for all r, equation (5) simplifies to

$$v(\mu) = \int_{p(\mu,c)}^1 f(p)dp, \tag{10}$$

where

$$p(\mu, c) \equiv \frac{1-\mu}{(1-\mu)+\mu\frac{1-c}{c}\frac{1-p_s}{p_s}}. \tag{11}$$

The virtual density $h(\mu)$ is then given by

$$h(\mu) = v'(\mu) = -f(p(\mu, c)) \cdot \frac{\partial}{\partial \mu} p(\mu, c),$$

and so,

$$h'(\mu) = -f'(p(\mu, c)) \left(\frac{\partial}{\partial \mu} p(\mu, c)\right)^2 - f(p(\mu, c)) \cdot \frac{\partial^2}{\partial \mu^2} p(\mu, c).$$

Therefore, the sign of $h'(\mu)$ is the same as the sign of

$$-\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c)\right)^2},$$

where $f(p(\mu, c)) > 0$ by assumption and $\partial p(\mu, c)/\partial \mu > 0$ follows $\gamma > 0$. Using (11) and substituting $\gamma = \frac{1-c}{c}\frac{1-p_s}{p_s}$, we get

$$-\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c)\right)^2} = -\frac{\partial}{\partial p} \log f(p(\mu, c)) - 2\frac{\gamma - 1}{\gamma}(1 + (\gamma - 1)\mu). \tag{12}$$

Substituting the value of γ into equation (11) gives: $p(\mu, c) = \frac{1-\mu}{1+(\gamma-1)\mu}$. Solving for μ,

$$\mu = \frac{1 - p(\mu, c)}{1 + (\gamma - 1)p(\mu, c)}. \tag{13}$$

Substituting for μ in equation (12), we get

$$\begin{aligned} -\frac{f'(p(\mu, c))}{f(p(\mu, c))} - \frac{\frac{\partial^2}{\partial \mu^2} p(\mu, c)}{\left(\frac{\partial}{\partial \mu} p(\mu, c)\right)^2} &= -\frac{\partial}{\partial p} \log f(p(\mu, c)) - 2\frac{\gamma - 1}{1 + (\gamma - 1)p(\mu, c)} \\ &= -\frac{\partial}{\partial p} \log f(p(\mu, c)) + g\left(p(\mu, c)\right), \end{aligned}$$

where $g(p) \equiv -2\frac{\gamma-1}{1+(\gamma-1)p}$. Note that $g(p)$ is increasing in p, is convex in p if $\gamma \leq 1$, and is concave in p if $\gamma \geq 1$. Therefore,

$$\min_{p\in[0,1]} g'(p) = \begin{cases} g'(0) & \text{if } \gamma \leq 1, \\ g'(1) & \text{if } \gamma \geq 1 \end{cases} = \begin{cases} 2(\gamma-1)^2 & \text{if } \gamma \leq 1, \\ 2\frac{(\gamma-1)^2}{\gamma^2} & \text{if } \gamma \geq 1 \end{cases} = 2(\gamma-1)^2 \min\left\{1, \frac{1}{\gamma^2}\right\}. \tag{14}$$

If condition (6) holds, then

$$\frac{\partial^2}{\partial p^2} \log f(p) < \min_{p\in[0,1]} g'(p),$$

which implies

$$\frac{\partial^2}{\partial p^2}\log f(p) < g'(p) \quad \forall p \in [0,1]. \tag{15}$$

Our claim is that, under condition (6), $h'(\mu)$ satisfies the strict single-crossing-from-above condition. To see this, take any two $\mu, \hat{\mu}$ with $\hat{\mu} < \mu$ and $h'(\mu) \geq 0$. Because $p(\mu, c)$ is strictly decreasing in μ, $p(\mu, c) < p(\hat{\mu}, c)$. Since $h'(\mu) \geq 0$, $\frac{\partial}{\partial p}\log f(p(\mu, c)) - g\left(p(\mu, c)\right) \leq 0$. Then,

$$\begin{aligned}
&\frac{\partial}{\partial p}\log f(p(\hat{\mu}, c)) - g\left(p(\hat{\mu}, c)\right) \\
&\quad = \frac{\partial}{\partial p}\log f(p(\mu, c)) - g\left(p(\mu, c)\right) + \int_{p(\mu,c)}^{p(\hat{\mu},c)} \underbrace{\left(\frac{\partial^2}{\partial p^2}\log f(p) - g'(p)\right)}_{<0 \text{ by } (15)} dp \\
&\quad < \frac{\partial}{\partial p}\log f(p(\mu, c)) - g\left(p(\mu, c)\right) \leq 0.
\end{aligned}$$

Therefore, $h'(\hat{\mu}) > 0$. The result follows. □

We continue with some notation and preliminary results for the proof of Theorem 1.

Lemma 1. *The value function $v(\mu)$ satisfies*

$$\lim_{\mu\to 0} \mu v'(\mu) = \lim_{\mu\to 1}(1-\mu)v'(\mu) = 0.$$

Proof. First, note that

$$\begin{aligned}
v'(\mu) &= \int_0^1 f(p, c(\mu, p)).\frac{\partial}{\partial \mu}c(\mu, p)dp \\
&= \int_0^1 f(p, c(\mu, p))\frac{p(1-p)p_s(1-p_s)}{(p_s(1-p) + \mu(p-p_s))^2}dp,
\end{aligned}$$

But since f is bounded, there exist some $C > 0$ such that

$$\begin{aligned}
|v'(\mu)| &\leq C\int_0^1 \frac{p(1-p)p_s(1-p_s)}{(p_s(1-p) + \mu(p-p_s))^2}dp \\
&= C\frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s) - (\mu(1-p_s) + p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right].
\end{aligned}$$

On the other hand,

$$\begin{aligned}
&\lim_{\mu\to 0}\mu \cdot \frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s) - (\mu(1-p_s) + p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right] \\
&\quad = \lim_{\mu\to 0}\frac{-(1-p_s)}{p_s}\mu\log(\mu) = 0,
\end{aligned}$$

and

$$\lim_{\mu\to 1}(1-\mu) \cdot \frac{p_s(1-p_s)}{(p_s-\mu)^3}\left[2(\mu-p_s) - (\mu(1-p_s) + p_s(1-\mu))\log\left(\frac{\mu(1-p_s)}{p_s(1-\mu)}\right)\right]$$

$$= \lim_{\mu \to 1} \frac{-p_s}{1-p_s}(1-\mu)\log(1-\mu) = 0.$$

Therefore,

$$\lim_{\mu \to 0} \mu v'(\mu) = \lim_{\mu \to 1}(1-\mu)v'(\mu) = 0.$$

This completes the proof of the Lemma. □

Consider a single-peaked virtual density $h(\mu)$. As discussed in the proof of Proposition 1, $\{\mu \in [0,1] : V(\mu) = v(\mu)\} = \{0\} \cup [\hat{\mu}, 1]$ for some $\hat{\mu} \in [0,1]$. Note that:

- $v'(\mu)\mu < v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat{\mu} = 0$.
- $v'(\mu)\mu > v(\mu)$ for all $\mu \in (0,1)$ if and only if $\hat{\mu} = 1$.
- When $\hat{\mu} \in (0,1)$, it satisfies:

$$v'(\hat{\mu})\hat{\mu} = v(\hat{\mu}). \tag{16}$$

Let

$$y(\mu) \equiv v'(\mu)\mu - v(\mu) = h(\mu)\mu - \int_0^{\mu} h(\tilde{\mu})\tilde{\mu}, \qquad \forall \mu \in [0,1]. \tag{17}$$

Then, $\hat{\mu} \in (0,1)$ is characterized by the equation: $y(\hat{\mu}) = 0$. We start with some remarks that will be used in the proof of Theorem 1.

Remark 1. $\lim_{\mu \to 0} y(\mu) = 0$. This follows from Lemma 1 and the fact that $v(0) = 0$.

Remark 2. *$y(\mu)$ is continuous in μ over $(0,1)$.* This is because f is continuous over its support.

Remark 3. *$y(\mu)$ is first strictly increasing and then strictly decreasing.* This is because $y'(\mu) = v''(\mu)\mu + v'(\mu) - v'(\mu) = v''(\mu)\mu = h'(\mu)\mu$. Since $h'(\mu)$ satisfies strict single crossing from above, so does $y'(\mu)$, and the remark follows.

Proof of Theorem 1. Take two single-peaked virtual densities $h_1(\mu)$ and $h_2(\mu)$ that satisfy equation (7). The cdf of h_2 is:

$$H_2(\mu) \equiv \int_0^{\mu} h_2(x)dx = \int_0^{\mu} \alpha\left(h_1(x)\right)dx.$$

For $k \in \{1,2\}$, let

$$y_k(\mu) \equiv h_k(\mu)\mu - H_k(\mu).$$

Since both h_1 and h_2 are single-peaked, Remarks 1, 2 and 3 imply that the set $\mathcal{U}_{y_k} \equiv \{\mu \in [0,1] : y_k(\mu) \geq 0\}$, for $k \in \{1,2\}$, has the following form:

$$\mathcal{U}_{y_k} = [0, \hat{\mu}_k].$$

The proof goes through showing that $\hat{\mu}_2 \leq \hat{\mu}_1$. If $\hat{\mu}_1 = 1$, this inequality is satisfied. If $\hat{\mu}_1 < 1$, $y_1(\hat{\mu}_1) = 0$, which implies: $\frac{H_1(\hat{\mu}_1)}{\hat{\mu}_1} = h_1(\hat{\mu}_1)$. Then,

$$\frac{H_2(\hat{\mu}_1)}{\hat{\mu}_1} = \frac{\int_0^{\hat{\mu}_1} \alpha\left(h_1(x)\right) dx}{\hat{\mu}_1} \geq \alpha\left(\frac{\int_0^{\hat{\mu}_1} h_1(x) dx}{\hat{\mu}_1}\right) = \alpha\left(\frac{H_1(\hat{\mu}_1)}{\hat{\mu}_1}\right) = \alpha\left(h_1(\hat{\mu}_1)\right) = h_2(\hat{\mu}_1),$$

where the inequality follows from the integral form of Jensen's inequality (e.g., Dragomir et al. 2016). Therefore, $H_2(\hat{\mu}_1) \geq h_1(\hat{\mu}_1)\hat{\mu}_1$ and $y_2(\hat{\mu}_1) \leq 0$.

Since $\mathcal{U}_{y_2} = [0, \hat{\mu}_2]$, $y_2(\mu)$ crosses 0 from above at $\hat{\mu}_2$, and $y_2'(\hat{\mu}_2) \leq 0$. Since $y_2'(\mu)$ satisfies strict single-crossing from above, $y_2(\tilde{\mu}) < 0$ for any $\tilde{\mu} > \hat{\mu}_1$. We conclude that $\hat{\mu}_2 \leq \hat{\mu}_1$.

To conclude the proof, consider three cases:

1. If $p_s \geq \hat{\mu}_1$, the optimal policy does not reveal any information in either case, and we pick $\sigma_1^0 = \sigma_2^0 = 1$.

2. If $\hat{\mu}_1 > p_s \geq \hat{\mu}_2$, the optimal policy under $h_2(\mu)$ does not reveal any information. In this case, we pick $\sigma_2^0 = 1$ and $\sigma_1^0 < 1$.

3. If $p_s < \hat{\mu}_2$, the optimal policies σ_1^0 and σ_2^0 satisfy:

$$\frac{p_s}{p_s + (1 - p_s)\sigma_1^0} = \hat{\mu}_1, \qquad \frac{p_s}{p_s + (1 - p_s)\sigma_2^0} = \hat{\mu}_2.$$

   Then, $\hat{\mu}_1 \geq \hat{\mu}_2$ implies $\sigma_1^0 \leq \sigma_2^0$.

In any case, $\sigma_1^0 \leq \sigma_2^0$, and the result follows. □

B Proofs for Section 4

Note that single-dippedness of the virtual density is equivalent to the following *strict single-crossing-from-below property* for $h'(\mu)$:

$$\text{If } h'(\mu) \geq 0 \text{ for some } \mu \in [0, 1], \text{ then } h'(\tilde{\mu}) > 0 \text{ for all } \tilde{\mu} > \mu.$$

Proof of Proposition 6. If $h'(\mu)$ satisfies the strict single-crossing-from-below condition, by definition, so does $v''(\mu)$. Therefore, whenever $v(\mu)$ is convex at μ, it is strictly convex at any $\hat{\mu} \geq \mu$. This means that $v(\mu)$ is first strictly concave and then strictly convex. Therefore, the set where the concave closure of $v(\mu)$ coincides with $v(\mu)$ has the following form:

$$\{\mu \in [0, 1] : V(\mu) = v(\mu)\} = [0, \hat{\mu}] \cup \{1\}.$$

When $p_s < \hat{\mu}$, the optimal policy is not revealing any information. This can be achieved by two messages, $m \in \{0, 1\}$, and an information structure where $\Pr(m = 1|\theta = 0) = \Pr(m = 1|\theta = 1) = 0$. Message $m = 1$ will occur with probability zero, and the posterior beliefs following $m = 1$ will be

free in a Perfect Bayesian Equilibrium. One can assign posteriors $\Pr_r(\theta = 1|m = 1) = 1$ to make $m = 1$ as the message that perfectly reveals the good state.

When $p_s \geq \hat{\mu}$, by Corollary 2 of Kamenica and Gentzkow (2011), the optimal policy generates two posteriors: $\mu \in \{\hat{\mu}, 1\}$. This is achieved by two messages, $m \in \{0, 1\}$, where message $m = 1$ perfectly reveals the good state. □

Proof of Proposition 4. The proof of Proposition 4 is identical to the proof of Proposition 2. □

Proof of Proposition 5. The proof follows identical steps to that of Proposition 3 until equation (14). The rest of the argument is provided below.

Recall that $g(p) \equiv -2\frac{\gamma-1}{1+(\gamma-1)p}$ and $g(p)$ is increasing in p, convex in p if $\gamma \leq 1$, and concave in p if $\gamma \geq 1$. Therefore,

$$\max_{p\in[0,1]} g'(p) = \begin{cases} g'(1) & \text{if } \gamma \leq 1 \\ g'(0) & \text{if } \gamma \geq 1 \end{cases} = \begin{cases} 2\frac{(\gamma-1)^2}{\gamma^2} & \text{if } \gamma \leq 1 \\ 2(\gamma-1)^2 & \text{if } \gamma \geq 1 \end{cases} = 2(\gamma-1)^2 \max\left\{1, \frac{1}{\gamma^2}\right\}. \tag{18}$$

If condition (8) holds,

$$\frac{\partial^2}{\partial p^2} \log f(p) > \max_{p\in[0,1]} g'(p),$$

and so

$$\frac{\partial^2}{\partial p^2} \log f(p) > g'(p) \quad \forall p \in [0, 1]. \tag{19}$$

Our claim is that, under condition (8), $h'(\mu)$ satisfies the strict single-crossing-from-below condition. To see this, take any two $\mu, \hat{\mu}$ with $\hat{\mu} > \mu$ and $h'(\mu) \geq 0$. Because $p(\mu, c)$ is strictly decreasing in μ, $p(\hat{\mu}, c) < p(\mu, c)$. Since $h'(\mu) \geq 0$, $\frac{\partial}{\partial p} \log f(p(\mu, c)) - g\,(p(\mu, c)) \leq 0$. Then,

$$\begin{aligned} &\frac{\partial}{\partial p} \log f(p(\hat{\mu}, c)) - g\,(p(\hat{\mu}, c)) \\ &\quad = \frac{\partial}{\partial p} \log f(p(\mu, c)) - g\,(p(\mu, c)) - \int_{p(\hat{\mu},c)}^{p(\mu,c)} \underbrace{\left(\frac{\partial^2}{\partial p^2} \log f(p) - g'(p)\right)}_{>0 \text{ by } (19)} dp \\ &\quad < \frac{\partial}{\partial p} \log f(p(\mu, c^*)) - g\,(p(\mu, c)) \leq 0. \end{aligned}$$

Therefore, $h'(\hat{\mu}) > 0$. The result follows. □

We continue by introducing some notation and preliminary results for the remaining proofs.

Consider a single-dipped virtual density $h(\mu)$. As discussed in the proof of Proposition 6, $\{\mu \in [0, 1] : V(\mu) = v(\mu)\} = [0, \hat{\mu}] \cup \{1\}$ for some $\hat{\mu} \in [0, 1]$. Note that:

- $v'(\mu)(1 - \mu) > 1 - v(\mu)$ for all $\mu \in (0, 1)$ if and only if $\hat{\mu} = 1$.
- $v'(\mu)(1 - \mu) < 1 - v(\mu)$ for all $\mu \in (0, 1)$ if and only if $\hat{\mu} = 0$.

- When $\hat{\mu} \in (0,1)$, it satisfies:

$$v'(\hat{\mu})(1-\hat{\mu}) = 1 - v(\hat{\mu}). \tag{20}$$

Let

$$z(\mu) \equiv v'(\mu)(1-\mu) - (1-v(\mu)) = h(\mu)(1-\mu) - \int_{\mu}^{1} h(\tilde{\mu})\tilde{\mu}, \qquad \forall \mu \in [0,1]. \tag{21}$$

Then, $\hat{\mu} \in (0,1)$ is characterized by the equation: $z(\hat{\mu}) = 0$. We start with some remarks.

Remark 4. $\lim_{\mu \to 1} z(\mu) = 0$. This follows Lemma 1 and the fact that $1 - v(1) = 0$.

Remark 5. *$z(\mu)$ is continuous in μ over $(0,1)$.* This is because f is continuous over its support.

Remark 6. *$z(\mu)$ is first strictly decreasing and then increasing.* This is because $z'(\mu) = v''(\mu)(1-\mu) - v'(\mu) + v'(\mu) = v''(\mu)(1-\mu) = h'(\mu)(1-\mu)$. Since $h'(\mu)$ satisfies strict single crossing from below, so does $z'(\mu)$, and the remark follows.

Proof of Theorem 2. Take two single-peaked virtual densities $h_1(\mu)$ and $h_2(\mu)$ that satisfy equation (9). For $k \in \{1,2\}$, let

$$z_k(\mu) \equiv h_k(\mu)(1-\mu) - (1 - H_k(\mu)).$$

Based on Remarks 4, 5 and 6, the set $\mathcal{L}_{z_k} \equiv \{\mu \in [0,1] : z_k(\mu) \leq 0\}$ has the following form:

$$\mathcal{L}_{z_k} = [\hat{\mu}_k, 1].$$

The proof goes through showing that $\hat{\mu}_2 \leq \hat{\mu}_1$. If $\hat{\mu}_1 = 1$, this inequality is satisfied. If $\hat{\mu}_1 < 1$, $z_1(\hat{\mu}_1) \leq 0$, which implies: $\frac{1-H_1(\mu))}{1-\hat{\mu}_1} \geq h_1(\hat{\mu}_1)$. Then,

$$\frac{1-H_2(\hat{\mu}_1)}{1-\hat{\mu}_1} = \frac{\int_{\hat{\mu}_1}^{1} \alpha\left(h_1(x)\right) dx}{1-\hat{\mu}_1} \geq \alpha\left(\frac{\int_{\hat{\mu}_1}^{1} h_1(x) dx}{1-\hat{\mu}_1}\right) = \alpha\left(1 - \frac{H_1(\hat{\mu}_1)}{1-\hat{\mu}_1}\right) \geq \alpha\left(h_1(\hat{\mu}_1)\right) = h_2(\hat{\mu}_1).$$

Therefore, $z_2(\hat{\mu}_1) \leq 0$. This means $\hat{\mu}_1 \in \mathcal{L}_{z_2} = [\hat{\mu}_2, 1]$, and therefore $\hat{\mu}_2 \leq \hat{\mu}_1$. Repeating the same argument in the proof of Theorem 1, we conclude that $\sigma_2^1 \geq \sigma_1^1$. □

References

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics*, 114(4): 1243–1284.

Alonso, Ricardo, and Odilon Câmara. 2016a. "Persuading Voters." *American Economic Review*, 106(11): 3590–3605.

Alonso, Ricardo, and Odilon Câmara. 2016b. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory*, 165 672–706.

Alonso, Ricardo, and Gerard Padró i Miquel. 2022. "Competitive Capture of Public Opinion." *Working Paper*.

Arbatlı, Cemal Eren, Quamrul H. Ashraf, Oded Galor, and Marc Klemp. 2020. "Diversity and Conflict." *Econometrica*, 88(2): 727–797.

Arieli, Itai, and Yakov Babichenko. 2019. "Private Bayesian Persuasion." *Journal of Economic Theory*, 182 185–217.

Bagnoli, Mark, and Ted Bergstrom. 2005. "Log-Concave Probability and Its Applications." *Economic Theory*, 26(2): 445–469.

Bardhi, Arjada, and Yingni Guo. 2018. "Modes of Persuasion Toward Unanimous Consent." *Theoretical Economics*, 13(3): 1111–1149.

Baum, Matthew A., and Tim Groeling. 2009. "Shot by the Messenger: Partisan Cues and Public Opinion Regarding National Security and War." *Political Behavior*, 31 157–186.

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

Bermeo, Nancy. 2016. "On Democratic Backsliding." *Journal of Democracy*, 27(1): 5–19.

Besley, Timothy, and Andrea Prat. 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *American Economic Review*, 96(3): 720–736.

Boleslavsky, Raphael, Mehdi Shadmehr, and Konstantin Sonin. 2021. "Media Freedom in the Shadow of a Coup." *Journal of the European Economic Association*, 19(3): 1782–1815.

Brancati, Dawn. 2014. "Democratic Authoritarianism: Origins and Effects." *Annual Review of Political Science*, 17 313–326.

Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang. 2019. "Pivotal Persuasion." *Journal of Economic Theory*, 180 178–202.

Chiang, Chun-Fang, and Brian Knight. 2011. "Media Bias and Influence: Evidence from Newspaper Endorsements." *The Review of Economic Studies*, 78(3): 795–820.

Corneo, Giacomo. 2006. "Media Capture in a Democracy: The Role of Wealth Concentration." *Journal of Public Economics*, 90(1-2): 37–58.

Curello, Gregorio, and Ludvig Sinander. 2022. "The Comparative Statics of Persuasion." *Working Paper, arXiv preprint arXiv:2204.07474*.

Desmet, Klaus, Ignacio Ortuño-Ortín, and Romain Wacziarg. 2017. "Culture, Ethnicity, and Diversity." *American Economic Review*, 107(9): 2479–2513.

Di Tella, Rafael, and Ignacio Franceschelli. 2011. "Government Advertising and Media Coverage of Corruption Scandals." *American Economic Journal: Applied Economics*, 3(4): 119–51.

Di Tella, Rafael, Sebastian Galiani, and Ernesto Schargrodsky. 2012. "Reality versus Propaganda in the Formation of Beliefs about Privatization." *Journal of Public Economics*, 96(5-6): 553–567.

Dragomir, Silvestru Sever, Muhammad Adil Khan, and Addisalem Abathun. 2016. "Refinement of the Jensen Integral Inequality." *Open Mathematics*, 14(1): 221–228.

Durante, Ruben, and Brian Knight. 2012. "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy." *Journal of the European Economic Association*, 10(3): 451–481.

Easterly, William, Jozef Ritzen, and Michael Woolcock. 2006. "Social Cohesion, Institutions, and Growth." *Economics and Politics*, 18(2): 103–120.

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.

Egorov, Georgy, and Konstantin Sonin. 2022. "The Political Economics of Non-democracy." *Journal of Economic Literature*, Forthcoming.

Enikolopov, Ruben, and Maria Petrova. 2015. "Media Capture: Empirical Evidence." In *Handbook of Media Economics*. 1: Elsevier, 687–700.

Enikolopov, Ruben, Michael Rochlitz, Koen J.L. Schoors, and Nikita Zakharov. 2023. "The Effect of Independent Online Media in an Autocracy." *Working Paper, Available at SSRN 4346225*.

Esteban, Joan-María, and Debraj Ray. 1994. "On the Measurement of Polarization." *Econometrica*, 62(4): 819–851.

Fiorina, Morris P., and Samuel J. Abrams. 2008. "Political Polarization in the American Public." *Annual Review of Political Science*, 11 563–588.

Galor, Oded, and Marc Klemp. 2018. "Roots of Autocracy." *Working Paper*.

Gehlbach, Scott, Zhaotian Luo, Anton Shirikov, and Dmitriy Vorobyev. 2022. "A Model of Censorship, Propaganda, and Repression." *Working Paper*.

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

Gentzkow, Matthew, and Jesse M. Shapiro. 2008. "Competition and Truth in the Market for News." *Journal of Economic Perspectives*, 22(2): 133–154.

Gitmez, A. Arda, and Konstantin Sonin. 2023. "The Dictator's Dilemma: A Theory of Propaganda and Repression." *University of Chicago, Becker Friedman Institute for Economics Working Paper 2023-67*.

Gläßel, Christian, and Katrin Paula. 2020. "Sometimes Less Is More: Censorship, News Falsification, and Disapproval in 1989 East Germany." *American Journal of Political Science*, 64(3): 682–698.

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*, 186 104–158.

Guriev, Sergei, and Daniel Treisman. 2022. *Spin Dictators: The Changing Face of Tyranny in the 21st Century*. Princeton, NJ: Princeton University Press.

Haggard, Stephan, and Robert Kaufman. 2021. "The Anatomy of Democratic Backsliding." *Journal of Democracy*, 32(4): 27–41.

Innocenti, Federico. 2022. "Can Media Pluralism Be Harmful to News Quality?" *Working Paper, Available at SSRN 4257390*.

Inostroza, Nicolas, and Alessandro Pavan. 2022. "Adversarial Coordination and Public Information Design." *Working Paper*.

Johnson, Justin P., and David P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review*, 96(3): 756–784.

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

Kerman, Toygar T., P. Jean-Jacques Herings, and Dominik Karos. 2022. "Persuading Strategic Voters." *Working Paper*.

Knight, Brian, and Ana Tribin. 2019. "The Limits of Propaganda: Evidence from Chavez's Venezuela." *Journal of the European Economic Association*, 17(2): 567–605.

Knight, Brian, and Ana Tribin. 2022. "Opposition Media, State Censorship, and Political Accountability: Evidence from Chavez's Venezuela." *The World Bank Economic Review*, 36(2): 455–487.

Kolotilin, Anton. 2015. "Experimental Design to Persuade." *Games and Economic Behavior*, 90 215–226.

Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk. 2022. "Censorship as Optimal Persuasion." *Theoretical Economics*, 17(2): 561–585.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica*, 85(6): 1949–1964.

Kosterina, Svetlana. 2022. "Persuasion with Unknown Beliefs." *Theoretical Economics*, 17(3): 1075–1107.

La Porta, Rafael, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny. 1999. "The Quality of Government." *Journal of Law, Economics, and Organization*, 15(1): 222–279.

Laclau, Marie, and Ludovic Renou. 2017. "Public Persuasion." *Working Paper*.

Ladd, Jonathan McDonald, and Gabriel S. Lenz. 2009. "Exploiting a Rare Communication Shift to Document the Persuasive Power of the News Media." *American Journal of Political Science*, 53(2): 394–410.

Levitsky, Steven, and Lucan A. Way. 2002. "Elections without Democracy: The Rise of Competitive Authoritarianism." *Journal of Democracy*, 13(2): 51–65.

Levitsky, Steven, and Daniel Ziblatt. 2018. *How Democracies Die*. New York, NY: Crown Publishing Group.

Lewis, David. 2016. "Blogging Zhanaozen: Hegemonic Discourse and Authoritarian Resilience in Kazakhstan." *Central Asian Survey*, 35(3): 421–438.

Lin, Xiao, and Ce Liu. 2022. "Credible Persuasion." *Working Paper*.

Lindqvist, Erik, and Robert Östling. 2010. "Political Polarization and the Size of Government." *American Political Science Review*, 104(3): 543–565.

Lorentzen, Peter. 2014. "China's Strategic Censorship." *American Journal of Political Science*, 58(2): 402–414.

Mathevet, Laurent, Jacopo Perego, and Ina Taneva. 2020. "On Information Design in Games." *Journal of Political Economy*, 128(4): 1370–1404.

McGreevy-Stafford, Eoghan Wallace. 2020. "First, the Bad News: Opposition Media in Authoritarian Regimes." Ph.D. dissertation, University of California, Los Angeles.

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

Milgrom, Paul, and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica* 157–180.

Ozerturk, Saltuk. 2022. "Media Access, Bias and Public Opinion." *European Economic Review*, 147, p. 104161.

Petrova, Maria. 2008. "Inequality and Media Capture." *Journal of Public Economics*, 92(1-2): 183–212.

Petrova, Maria. 2012. "Mass Media and Special Interest Groups." *Journal of Economic Behavior & Organization*, 84(1): 17–38.

Prat, Andrea. 2015. "Media Capture and Media Power." In *Handbook of Media Economics*. eds. by Simon P. Anderson, Joel Waldfogel, and David Strömberg, 1 of Handbook of Media Economics: North-Holland, 669 – 686.

Qian, Nancy, and David Yanagizawa-Drott. 2017. "Government Distortion in Independently Owned Media: Evidence from US News Coverage of Human Rights." *Journal of the European Economic Association*, 15(2): 463–499.

Qin, Bei, David Strömberg, and Yanhui Wu. 2018. "Media Bias in China." *American Economic Review*, 108(9): 2442–76.

Rodrik, Dani. 1999. "Where Did All the Growth Go? External Shocks, Social Conflict, and Growth Collapses." *Journal of Economic Growth*, 4(4): 385–412.

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

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

Sobolev, Anton. 2023. "How Can Free Media Help an Autocrat to Deal with Mass Protests? A Curious Case of Russia." *Working Paper*.

Stanig, Piero. 2015. "Regulation of Speech and Media Coverage of Corruption: An Empirical Analysis of the Mexican Press." *American Journal of Political Science*, 59(1): 175–193.

Sun, Junze, Arthur Schram, and Randolph Sloof. 2022. "Public Persuasion in Elections: Single-Crossing Property and the Optimality of Censorship." *Working Paper, Available at SSRN 4028840*.

Szeidl, Adam, and Ferenc Szucs. 2021. "Media Capture through Favor Exchange." *Econometrica*, 89(1): 281–310.

Tai, Qiuqing. 2014. "China's Media Censorship: A Dynamic and Diversified Regime." *Journal of East Asian Studies*, 14(2): 185–210.

Taneva, Ina. 2019. "Information Design." *American Economic Journal: Microeconomics*, 11(4): 151–85.

VonDoepp, Peter, and Daniel J. Young. 2013. "Assaults on the Fourth Estate: Explaining Media Harassment in Africa." *The Journal of Politics*, 75(1): 36–51.

Wang, Yun. 2015. "Bayesian Persuasion with Multiple Receivers." *Working Paper*.

Whitten-Woodring, Jenifer, and Douglas A. Van Belle. 2017. "The Correlates of Media Freedom: An Introduction of the Global Media Freedom Dataset." *Political Science Research and Methods*, 5(1): 179–188.