



Business School / School of Economics

UNSW Business School Working Paper

UNSW Business School Research Paper No. 2016 ECON 21

Persuasion of a Privately Informed Receiver

Anton Kolotilin
Tymofiy Mylovanov
Andriy Zapechelnyuk
Ming Li

This paper can be downloaded without charge from
The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=2913916>

PERSUASION OF A PRIVATELY INFORMED RECEIVER

ANTON KOLOTILIN, TYMOFIY MYLOVANOV, ANDRIY ZAPECHELNYUK
AND MING LI

ABSTRACT. We study persuasion mechanisms in linear environments. A privately informed receiver chooses between two actions. A sender designs a persuasion mechanism that can condition the information disclosed to the receiver on the receiver's report about his type. We establish the equivalence of implementation by persuasion mechanisms and by experiments. We also characterize the optimal persuasion mechanisms. In particular, if the density of the receiver's type is log-concave, then the optimal persuasion mechanism reveals the state if and only if the state is below a threshold. We apply our results to the design of media censorship policies.

JEL Classification: D82, D83, L82

Keywords: Bayesian persuasion, information disclosure, information design, mechanism design without transfers, experiments, persuasion mechanisms, media

Date: November 19, 2016.

Kolotilin: School of Economics, UNSW Australia, Sydney, NSW 2052, Australia. E-mail: akolotilin@gmail.com.

Mylovanov: University of Pittsburgh, Department of Economics, 4714 Posvar Hall, 230 South Bouquet Street, Pittsburgh, PA 15260, USA. E-mail: mylovanov@gmail.com.

Zapechelnyuk: Adam Smith Business School, University of Glasgow, University Avenue, Glasgow G12 8QQ, UK. E-mail: azapech@gmail.com.

Li: Concordia University and CIREQ, 1455 Boulevard de Maisonneuve Ouest, Department of Economics, Concordia University, Montréal, H3G 1M8, Canada. Email: ming.li@concordia.ca.

The authors would like to express their gratitude to anonymous referees for their very helpful comments and to Navik Kartik for a thoughtful discussion at the Boston ASSA meetings. They are also grateful to Ricardo Alonso, Dirk Bergemann, Patrick Bolton, Alessandro Bonatti, Rahul Deb, Péter Esö, Johannes Hörner, Florian Hoffman, Roman Inderst, Emir Kamenica, Daniel Kräehmer, Marco Ottaviani, Mallesh Pai, Ilya Segal, Larry Samuelson, and Joel Sobel for their excellent comments. Kolotilin acknowledges financial support from the Australian Research Council.

1. INTRODUCTION

In numerous situations, an uninformed sender (she) wishes to influence an action of a receiver (he) who privately knows his preference type but is uncertain about a payoff-relevant state. The sender has full control over information disclosure about the state but cannot use any other incentive tools such as monetary transfers. If the sender knew the receiver's type, she would tailor information disclosure to the type. However, the sender can only ask the receiver to report his type, which the receiver may misreport, to affect the information he receives. This begs the question of whether the sender can benefit by designing a complex *persuasion mechanism* that tailors information disclosure to the receiver's report, as compared to an *experiment* that provides the same information about the state, regardless of the receiver's report.

In our model, the receiver must choose one of two actions. The sender's and receiver's utilities depend linearly on the state and receiver's type.¹ The state and receiver's type are one-dimensional independent random variables. The sender and receiver share a common prior about the state; the receiver privately knows his type. The sender commits to a persuasion mechanism that asks the receiver to report his type and returns a stochastic message that depends on the state and report. After receiving the message, the receiver updates his beliefs about the state and takes an action that maximizes his expected utility.

We start by characterizing the set of the receiver's interim utilities implementable by persuasion mechanisms.² The upper and lower bounds are achieved by *full disclosure* and *no disclosure* of information about the state. For any mapping from messages to actions under a persuasion mechanism, the receiver's expected utility is linear in his type. Because each type chooses an optimal mapping, the receiver's interim utility is an upper envelope of linear functions and, hence, is convex. To sum up, an implementable interim utility is necessarily a convex function that lies between the upper and lower bounds generated by full and no disclosure.

Our main theorem shows that (1) this necessary condition is also sufficient and that (2) any implementable interim utility can be attained by an experiment. Moreover, in our model the receiver's interim utility uniquely determines the sender's expected utility. Therefore, there is no loss of generality in restricting attention to experiments. In particular, the sender need not consider more complex persuasion mechanisms that condition information disclosure on the receiver's report.

We now outline the argument behind this result. Because the utilities are linear in the state, the only relevant content of a message about the state is the posterior mean. Therefore, an experiment can be fully described by the distribution of the posterior mean state H that it generates. By Blackwell (1953), there exists an experiment that generates H if and only if the prior distribution of the state is a mean-preserving spread of H . By linearity, the receiver's interim utility U can be represented as an

¹For our analysis to apply, we only need to assume that the utilities are linear in some transformation of the state and are arbitrary functions of the receiver's type.

²The receiver's interim utility is the mapping from the receiver's type to his expected utility given optimal behavior under a persuasion mechanism.

appropriate integral of H . Using the integral form of the mean-preserving spread condition as in Blackwell (1951), we show that the prior is a mean-preserving spread of H if and only if U is convex and lies between the upper and lower bounds generated by full and no disclosure. Therefore, any U that satisfies the necessary conditions that we identified above is implementable by an experiment.

This characterization allows us to formulate the sender's problem as the maximization of a linear functional on a bounded set of convex functions. We derive a general structure of optimal persuasion mechanisms and use it to simplify the sender's problem into a finite-variable optimization problem.

We then consider a special case in which the sender's expected utility is a weighted sum of the expected action and the receiver's expected utility. In this case, censorship mechanisms are generally optimal. A persuasion mechanism is called an upper-censorship mechanism if it reveals the states below some threshold and pools the states above this threshold. We show that the optimal persuasion mechanism is an upper-censorship mechanism, regardless of the prior distribution of the state and regardless of the weight the sender puts on the receiver's utility, if and only if the probability density of the receiver's type is log-concave.

If an upper-censorship mechanism is optimal, the comparative statics analysis is particularly tractable. We show that the optimal censorship threshold increases, and, thus, the sender discloses more information, if the sender's and receiver's utilities become more aligned or if each type of the receiver experiences a preference shock detrimental to the sender.

We apply our results to the problem of media control by the government. We consider a stylized setting where the government wishes to influence people's actions by regulating the media content. We extend our model to allow for a continuum of heterogeneous receivers. This extension permits the average action of the receivers to affect their utility but not their optimal actions. For example, an election outcome impacts all people but does not change their preferences over candidates. We provide conditions under which the optimal media control policy prescribes to permit all sufficiently loyal media outlets to broadcast and to prohibit all other media outlets. Our main result of implementation equivalence between persuasion mechanisms and experiments implies that the optimal media control policy does not depend on whether the government can restrict people to follow only one media outlet of their choice.

We conclude the paper by exploring the limits of implementation equivalence between persuasion mechanisms and experiments. We show that the equivalence does not hold if the receiver has more than two actions (Section 6) or if the receiver's utility is nonlinear in the state (Section 7).

Related Literature. Our model is a variation of Kamenica and Gentzkow (2011), who show that an optimal experiment corresponds to a point on the concave closure of the sender's value function. This concavification approach does not rely on any structure of the problem such as the linearity of utility functions or the two actions of the receiver. However, as Gentzkow and Kamenica (2016d) point out, this approach has limited applicability when the state space is infinite because it requires to calculate

a concave closure on the infinite-dimensional space of distributions over the state space.

Rayo and Segal (2010) impose more structure on the receiver’s utility and the distribution of the receiver’s type than our paper. At the same time, they allow the sender’s utility to be nonlinear in the state. Rayo and Segal (2010) partially characterize optimal experiments and represent the sender’s problem as the maximization problem of a concave objective function subject to linear constraints. In our setting, their assumptions about the receiver would imply that either full or no disclosure is optimal.

Kolotilin (2016) allows for nonlinear sender’s and receiver’s utility functions and an arbitrary joint distribution of the state and the receiver’s type. The linear-programming approach in Kolotilin (2016) permits to verify whether a candidate experiment is optimal. But this approach has limited applicability because it does not allow to directly characterize the optimal experiments.

The three papers above study experiments.³ In contrast, we consider persuasion mechanisms, which can tailor information disclosure to the receiver’s report.⁴ Linear utility functions with two possible actions of the receiver enable us to use the envelope representation of incentive compatibility as in Myerson (1981). However, the characterization of implementable mechanisms in our setting differs from Myerson (1981) because there are no transfers between the sender and receiver, and there are obedience constraints instead of participation constraints.⁵ The equivalence between persuasion mechanisms and experiments relies on the majorization theory of Hardy et al. (1929) and thus relates to the equivalence between Bayesian and dominant-strategy implementation in Manelli and Vincent (2010) and Gershkov et al. (2013). The optimality of upper-censorship mechanisms relies on the single-crossing property of the integrand in the optimization problem and, thus, relates to the optimality of reserve-price auctions in Myerson (1981).

³There is a rapidly growing Bayesian persuasion literature that studies optimal experiments. Bayesian persuasion with a privately informed sender is considered in Gill and SgROI (2008, 2012), Perez-Richet (2014), and Alonso and Camara (2016b). Bayesian persuasion with multiple senders is analyzed in Ostrovsky and Schwarz (2010), Board and Lu (2016), Gentzkow and Kamenica (2016a, 2016b), and Li and Norman (2016). Dynamic Bayesian persuasion is examined in Au (2015), Ely et al. (2015), and Ely (2016). Information acquisition and the value of information in Bayesian persuasion are explored in Gentzkow and Kamenica (2014, 2016c), Kolotilin (2015), Alonso and Camara (2016a), and Bergemann and Morris (2016).

⁴Focusing on experiments is not without loss of generality in settings with multiple receivers in Bergemann and Morris (2013), Alonso and Camara (2016c), Mathevet et al. (2016), and Taneva (2016).

⁵Bayesian persuasion with monetary transfers is investigated in Bergemann and Pesendorfer (2007), Eso and Szentes (2007), Bergemann et al. (2015, 2016a, 2016b), Li and Shi (2016), and Horner and Skrzypacz (2016).

2. MODEL

2.1. Setup. There are two players: sender and receiver. The receiver makes a choice $a \in A = \{0, 1\}$ between action ($a = 1$) and inaction ($a = 0$). There are two payoff-relevant random variables: the state of the world $\omega \in \Omega$ and the receiver's private type $r \in R$, where Ω and R are intervals in the real line. Random variables ω and r are independent and have distributions F and G .

Let the receiver's and sender's utilities be

$$\begin{aligned} u(\omega, r, a) &= a \cdot (\omega - r), \\ v(\omega, r, a) &= a + \rho(r) u(\omega, r, a), \end{aligned} \tag{1}$$

where ρ is a bounded measurable function. In Appendix A, we extend the analysis to the case in which the receiver's and sender's utility functions are linear in the state ω . The intercepts and slopes of these linear functions constitute the private multidimensional type of the receiver.

The receiver's utility from inaction ($a = 0$) is normalized to zero, whereas his utility from action ($a = 1$) equals the benefit ω less the private cost r . The sender's utility is a weighted sum of the receiver's utility and action. The sender is biased towards the receiver's action but also puts a type-specific weight $\rho(r)$ on the receiver's utility. In particular, if the weight ρ is large, then the sender's and receiver's interests are aligned, but if the weight is zero, then the sender cares only about whether the receiver acts or not.

We assume that the set of states is $\Omega = [0, 1]$ and $\mathbb{E}[\omega] \in (0, 1)$. Because the optimal action of the receiver of any type above 1 or below 0 is independent of the state, we assume that the set of types is $R = [0, 1]$. These assumptions allow for elegance of presentation; relaxing them poses no difficulty.

2.2. Persuasion Mechanisms and Experiments. In order to influence the action taken by the receiver, the sender can design a mechanism that asks the receiver to report his private information and sends a message to the receiver conditional on his report and the realized state.

A *persuasion mechanism* π asks the receiver to report $\hat{r} \in R$ and then recommends him to take one of two actions: for every $\omega \in \Omega$, it recommends to act ($\hat{a} = 1$) with probability $\pi(\omega, \hat{r})$ and not to act ($\hat{a} = 0$) with complementary probability $1 - \pi(\omega, \hat{r})$, where π is a measurable function from $\Omega \times R$ to $[0, 1]$. A persuasion mechanism is *incentive compatible* if the receiver finds it optimal to report his true type and to follow the mechanism's recommendation.

In comparison, an experiment communicates the same messages to each type of the receiver. An *experiment* σ sends to the receiver a random message that depends on the realized state ω . Denote by $\sigma(m|\omega)$ the distribution of message $m \in M = [0, 1]$ conditional on state ω .

For a given experiment σ , each message m induces a posterior belief of the receiver about the state. Because the receiver's utility is monotonic in his type, we can identify every message m with the cutoff type r who is indifferent between the two actions conditional on receiving this message. An experiment is *direct* if its messages are

equal to the cutoff types, $m = \mathbb{E}[\omega|m]$. Without loss of generality, we focus on direct experiments (as in, e.g., Kamenica and Gentzkow, 2011).

A persuasion mechanism can be interpreted as a menu of (non-direct) experiments from which the sender can freely choose one experiment. For a given mechanism π , the corresponding menu of experiments is $\{\sigma_{\hat{r}}\}_{\hat{r} \in R}$, where, conditional on a state ω , an experiment $\sigma_{\hat{r}}$ sends messages 1 and 0 with probabilities $\pi(\omega, \hat{r})$ and $1 - \pi(\omega, \hat{r})$. Conversely, by the revelation principle, any equilibrium outcome induced by any menu of direct or non-direct experiments offered to the receiver can be replicated by a persuasion mechanism. In particular, the restriction to mechanisms that return binary messages is without loss, because the receiver's action is binary.

A persuasion mechanism π is *equivalent to an experiment* if, conditional on each state, the behavior of the receiver as stipulated by π is a best-reply behavior to the posterior beliefs generated by some direct experiment σ ,

$$\pi(\omega, r) \in [1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)] \quad \text{for all } \omega \in \Omega \text{ and all } r \in R, \quad (2)$$

where $\sigma(r_-|\omega)$ denotes the left limit of $\sigma(\cdot|\omega)$ at r . To understand (2), note that, upon receiving a direct message m , every type $r < m$ strictly prefers to choose $a = 1$, and type $r = m$ is indifferent between $a = 0$ and $a = 1$ and may optimally choose any lottery over the two actions. Consequently, the probability that type r acts conditional on state ω can take any value in the interval $[1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)]$.

2.3. Benchmark With Known Types. Consider the benchmark case, studied in Kolotilin (2015), in which the sender knows the receiver's type. In this case, a persuasion mechanism can condition the information disclosed directly on the receiver's type, rather than on the receiver's report about his type. Suppose, for simplicity, that $\rho(r) > 0$ for all r and that F admits a density.

By (1), the sender would like the receiver of type r to act in each state ω that satisfies $1 + \rho(r)(\omega - r) > 0$ or, equivalently, $\omega > r - 1/\rho(r)$. If the receiver who is recommended to act whenever $\omega > r - 1/\rho(r)$ obtains a nonnegative payoff from following the recommendation, then the sender achieves her first best outcome from type r . Otherwise, the mechanism recommends the receiver to act on the largest interval $(x, 1]$ of states such that the receiver is willing to follow the recommendation given his belief that the state is in $(x, 1]$. Thus, the optimal mechanism recommends the receiver to act if $\omega > x^*(r)$ and not to act if $\omega < x^*(r)$, where $x^*(r)$ is the minimum $x \in \Omega$ that satisfies $1 + \rho(r)(x - r) \geq 0$ and $\int_x^1 (\omega - r) dF(\omega) \geq 0$.

2.4. Envelope Characterization of Incentive Compatibility. Denote the expected utility of a receiver of type $r \in R$ who reports $\hat{r} \in R$ and takes actions $a_0 \in \{0, 1\}$ and $a_1 \in \{0, 1\}$ after recommendations $\hat{a} = 0$ and $\hat{a} = 1$, respectively, by

$$U_\pi(r, \hat{r}, a_0, a_1) = \int_\Omega (a_0(1 - \pi(\omega, \hat{r})) + a_1\pi(\omega, \hat{r})) (\omega - r) dF(\omega).$$

The expected utility of the truthful ($\hat{r} = r$) and obedient ($a_0 = 0$ and $a_1 = 1$) receiver is equal to

$$U_\pi(r) = U_\pi(r, r, 0, 1) = \int_{\Omega} \pi(\omega, r)(\omega - r) dF(\omega).$$

We consider mechanisms that satisfy the incentive compatibility constraint

$$U_\pi(r) \geq U_\pi(r, \hat{r}, a_0, a_1) \text{ for all } r, \hat{r} \in R \text{ and } a_0, a_1 \in A. \quad (3)$$

It is convenient to introduce the notation for the expected utility of the *obedient* receiver of type r , who reports \hat{r} and then obeys the recommendation of the mechanism:

$$U_\pi(r, \hat{r}) = U_\pi(r, \hat{r}, 0, 1) = p_\pi(\hat{r}) - q_\pi(\hat{r})r,$$

where

$$q_\pi(\hat{r}) = \int_{\Omega} \pi(\omega, \hat{r}) dF(\omega) \quad \text{and} \quad p_\pi(\hat{r}) = \int_{\Omega} \omega \pi(\omega, \hat{r}) dF(\omega).$$

With this representation of the utility function, we can draw the parallel to the standard mechanism design problem with transfers, where r is a private value, \hat{r} is a reported value, $q_\pi(\hat{r})$ is the probability of transaction, and $p_\pi(\hat{r})$ is the expected monetary transfer. The classical envelope argument yields the following lemma.

Lemma 1. *A mechanism π is incentive compatible if and only if*

$$q_\pi \text{ is nonincreasing,} \quad (4)$$

$$U_\pi(r) = \int_r^1 q_\pi(s) ds, \quad (5)$$

$$U_\pi(0) = \mathbb{E}[\omega]. \quad (6)$$

The proof is in Appendix B.

Interestingly, the obedience constraints for the intermediate types are implied by the boundary conditions, $U_\pi(1) = 0$ and $U_\pi(0) = \mathbb{E}[\omega]$, and truth telling, $U_\pi(r) \geq U_\pi(r, \hat{r})$. To disobey by ignoring the recommendation, that is, to act (not to act) irrespective of what is recommended, is not better than pretending to be the lowest type $\hat{r} = 0$ (the highest type $\hat{r} = 1$, respectively). To disobey by taking the opposite action to the recommended one is never beneficial due to the linearity of the receiver's utility.

In our model, there are no transfers, and there are obedience constraints instead of individual rationality constraints. These differences between our and the standard environment with transfers translate into the following differences in characterization.

First, there are two boundary conditions, $U_\pi(1) = 0$ and $U_\pi(0) = \mathbb{E}[\omega]$:

- (a) We have $\omega - r \leq 0$ for all $\omega \in \Omega = [0, 1]$ and $r \geq 1$. Hence, type 1's utility is maximized by not acting for any belief about the state, implying $U_\pi(1) = 0$. This is (5) evaluated at $r = 1$.
- (b) We have $\omega - r \geq 0$ for all $\omega \in \Omega = [0, 1]$ and $r \leq 0$. Hence, type 0's utility is maximized by acting for any belief about the state, implying $U_\pi(0) = \mathbb{E}[\omega]$. This is (6).

Second, not every pair (q, U) that satisfies conditions (4)–(6) is feasible, that is, a mechanism π that implements such a pair need not exist. For example, if F assigns probability 1 on $\omega = 1/2$, then every nonincreasing q with $\int_0^1 q(r)dr = 1/2$ satisfies (4)–(6). Among these functions, q is feasible if and only if it satisfies $q(r) = 1$ for $r < 1/2$ and $q(r) = 0$ for $r > 1/2$.

3. IMPLEMENTATION EQUIVALENCE

In this section, we characterize the pairs of the interim utility U and action q implementable by persuasion mechanisms, and show that the same pairs of the interim utility and action are implementable by experiments.

This is a step towards solving the sender's optimization problem, because the sender's interim utility $V_\pi(r)$ when the receiver's type is r is a weighted sum of the receiver's interim utility and action,

$$V_\pi(r) = \int_{\Omega} (1 + \rho(r)(\omega - r))\pi(\omega, r)dF(\omega) = q_\pi(r) + \rho(r)U_\pi(r). \quad (7)$$

3.1. Bounds on Receiver's Interim Utility. Consider two simple mechanisms. The *full disclosure* mechanism informs the receiver about the state, so the receiver acts ($\pi(\omega, r) = 1$) if $\omega > r$ and does not act ($\pi(\omega, r) = 0$) if $\omega < r$. Therefore, the interim utility is

$$\bar{U}(r) = \int_r^1 (\omega - r) dF(\omega).$$

The *no disclosure* mechanism does not convey any information to the receiver, so the receiver acts if $\mathbb{E}[\omega] > r$ and does not act if $\mathbb{E}[\omega] < r$. Therefore, the interim utility is

$$\underline{U}(r) = \max\{\mathbb{E}[\omega] - r, 0\}.$$

Thus, $\underline{U}(r)$ is the receiver's interim utility based on the prior information about ω as given by F , whereas $\bar{U}(r)$ is the receiver's interim utility if he observes ω .

Note that every mechanism π must satisfy

$$\underline{U}(r) \leq U_\pi(r) \leq \bar{U}(r) \quad \text{for all } r \in R. \quad (8)$$

The left-hand side inequality of (8) is implied by incentive compatibility: the receiver cannot be better off by ignoring the sender's recommendation. The right-hand side inequality of (8) is the *feasibility constraint*: the receiver's utility cannot exceed the utility attained under full disclosure of ω .

3.2. Implementable Interim Utility and Action. The receiver's interim utility U and action q are *implementable* if there exists an incentive-compatible persuasion mechanism π such that $U(r) = U_\pi(r)$ and $q(r) = q_\pi(r)$ for all $r \in R$. Moreover, U and q are *implementable by an experiment* if π is equivalent to an experiment.

Let \mathcal{U} be the set of all convex functions bounded by \underline{U} and \bar{U} (see Fig. 1).




FIGURE 1. Set \mathcal{U} contains every convex function between \underline{U} and \bar{U} .

Theorem 1. *The following statements are equivalent:*

- (a) U is a convex function between \underline{U} and \bar{U} ;
- (b) U is implementable;
- (c) U is implementable by an experiment.

Proof. Observe that (a) states the necessary conditions for the implementation of U . The incentive-compatibility constraint requires the convexity of U by Lemma 1, and the feasibility constraint (8) requires $\underline{U}(r) \leq U(r) \leq \bar{U}(r)$. Hence, (b) \Rightarrow (a). Also, the implication (c) \Rightarrow (b) is trivial by definition. It remains to show that (a) \Rightarrow (c).

Let $U \in \mathcal{U}$. Define $H(r) = 1 + U'(r_+)$, where $U'(r_+)$ denotes the right-derivative of U at r , so that H is right-continuous. Since $\underline{U}(0) = \bar{U}(0) = \mathbb{E}[\omega]$ and $\underline{U}'(0_+) = -1$, we have $H(0) \geq 0$; since $\underline{U}(1) = \bar{U}(1) = 0$ and $\underline{U}'(1_+) = \bar{U}'(1_+) = 0$, we have $H(1) = 1$. Also, since U is convex, H is nondecreasing. Hence, H is a distribution function. Next, observe that

$$\int_r^1 (1 - H(s)) ds = U(r) \leq \bar{U}(r) = \int_r^1 (1 - F(s)) ds \quad (9)$$

for all r , with equality at $r = 0$, because $\underline{U}(0) = \bar{U}(0) = \mathbb{E}[\omega]$. That is, F is a mean-preserving spread of H . By Blackwell (1953), there exists a joint distribution function $P(\omega, m)$ such that the marginal distribution of ω is F , the marginal distribution of m is H , and $\mathbb{E}_P[\omega|m] = m$ for all m .

Now, consider an experiment $\sigma(m|\omega) = P(m|\omega)$. By construction, $\mathbb{E}[\omega|m] = m$ for all m of σ , and the probability that σ generates message $m \leq x$, for any given value x , is $H(x)$. Thus, σ induces type r to act with probability $q_\sigma(r) \in [1 - H(r), 1 - H(r_-)]$, where indeterminacy arises at each discontinuity point r of H , because type r is indifferent between the two actions and can, therefore, optimally choose any lottery over these actions. Finally, (5) implies that σ implements U . \square

Theorem 1 yields the following characterization of implementable interim actions.

Corollary 1. *The following statements are equivalent:*

(a) *q is a nonincreasing function that satisfies*

$$\int_r^1 q(\omega)d\omega \leq \int_r^1 (1 - F(\omega))d\omega \text{ for all } r, \text{ with equality at } r = 0; \quad (10)$$

(b) *q is implementable;*

(c) *q is implementable by an experiment.*

The heart of these characterization results is that every implementable pair of U and q is implementable by an experiment. This result relies on the following connection between the Mirrlees (1971) representation of incentive compatibility and the Blackwell (1953) representation of garbling. Incentive compatibility (5) and feasibility (8) imply that every implementable interim utility and action must satisfy the requirements in parts (a) of Theorem 1 and Corollary 1. In turn, part (a) of Corollary 1 implies that F is a mean-preserving spread of $1 - q$. Therefore, interim action q can be implemented by an appropriate garbling of the full disclosure experiment. Because every interim utility U is pinned down by interim action q through (5), we obtain the result.

We now highlight a connection to the literature on the equivalence of Bayesian and dominant-strategy incentive compatibility in linear environments with transfers (Manelli and Vincent, 2010, and Gershkov et al., 2013). Using Gutmann et al. (1991), Gershkov et al. (2013) show that for a given monotonic expected allocation (Bayesian incentive-compatible mechanism) there exists a monotonic ex-post allocation (dominant-strategy incentive-compatible mechanism) that delivers the same expected allocation. Relatedly, using Blackwell (1953), we show that for a given nonincreasing q that satisfies (10) there exists an incentive-compatible $\pi(\omega, r)$ that is nonincreasing in r for each $\omega \in \Omega$. Thus, this π is equivalent to an experiment (see Proposition 2 in Section 7), and satisfies

$$\int_{\Omega} \pi(\omega, r)dF(\omega) = q(r) \text{ for each } r \in R.$$

Both Blackwell (1953) and Gutmann et al. (1991) are based on the majorization theory initiated by Hardy et al. (1929).

We will discuss in Sections 6 and 7 how the equivalence between persuasion mechanisms and experiments changes when the receiver has more than two actions and when the receiver's utility is nonlinear in (any transformation of) the state.

4. OPTIMAL MECHANISMS

In this section, we use Theorem 1 to characterize the persuasion mechanisms that are optimal for the sender under additional smoothness assumptions. We assume that the weight ρ in the sender's utility is continuous in the receiver's type and that the distributions F and G admit densities f and g , where g is strictly positive and continuously differentiable.

In addition, we solve a special case. Assuming log-concave (log-convex) g and type-independent ρ , we show that simple censorship persuasion mechanisms, which

include the full and no disclosure mechanisms, maximize the sender's expected utility. We, thus, solve the sender's problem for many commonly-used density functions (see Tables 1 and 3 in Bagnoli and Bergstrom, 2005, for the list of log-concave and log-convex density functions).

4.1. Sender's Problem. The sender seeks an incentive-compatible persuasion mechanism π that maximizes

$$\int_R V_\pi(r) dG(r),$$

where $V_\pi(r)$ is the sender's interim utility when the receiver's type is r , as defined by (7). The following lemma is a useful tool for finding optimal persuasion mechanisms. It expresses the sender's expected utility as a function of the receiver's interim utility.

Lemma 2. *For every incentive-compatible mechanism π ,*

$$\int_R V_\pi(r) dG(r) = g(0)\mathbb{E}[\omega] + \int_R U_\pi(r)I(r)dr,$$

where

$$I(r) = g'(r) + \rho(r)g(r) \text{ for all } r \in R.$$

This formulation of the sender's objective is similar to the Myerson (1981) formulation of the seller's objective in terms of virtual valuations. Both formulations account for the designer's biased preferences and information rents left to the privately informed parties. Lemma 2 relies on the assumption that the sender's utility is a weighted sum of the receiver's utility and action. By incentive compatibility (Lemma 1), the receiver's interim action is the derivative of his interim utility. The lemma then follows by integration by parts.

Proof. Observe that, by (5) and (7), we have

$$\int_R V_\pi(r) dG(r) = \int_R (q_\pi(r) + \rho(r)U_\pi(r)) g(r) dr = \int_R (-U'_\pi(r_+) + \rho(r)U_\pi(r)) g(r) dr.$$

Using integration by parts, we obtain

$$-\int_R U'_\pi(r_+)g(r)dr = -U_\pi(r)g(r)|_0^1 + \int_R U_\pi(r)g'(r)dr.$$

Since $-U_\pi(r)g(r)|_0^1 = \mathbb{E}[\omega]g(0)$ by (6), the lemma follows. \square

By Theorem 1, the receiver's interim utility is implementable by some persuasion mechanism if and only if it is in \mathcal{U} . Hence, the sender's problem can be expressed as

$$\max_{U \in \mathcal{U}} \int_R U(r)I(r)dr. \tag{11}$$

We say that U is *optimal* if it solves the above problem.

4.2. Structure of Optimal Mechanisms. We characterize the structure of optimal mechanisms under the assumption that function I is nonzero almost everywhere and changes sign $n \geq 0$ times.⁶ Let $\{r_1, r_2, \dots, r_n\} \subset (0, 1)$ be the set of types at which I changes its sign, and let $r_0 = 0$ and $r_{n+1} = 1$.




FIGURE 2. The optimal interim utility on the interval where $I(r)$ is positive.

Clearly, as follows from (11), on any interval (r_i, r_{i+1}) where $I(r)$ is positive, for any given values of $U(r_i)$ and $U(r_{i+1})$, the optimality requires that $U(r)$ is pointwise maximized subject to feasibility ($U \leq \bar{U}$) and the convexity of U on $[r_i, r_{i+1}]$. That is, the interim utility U on $[r_i, r_{i+1}]$ is a straight line that passes through the endpoints $U(r_i)$ and $U(r_{i+1})$, unless $U(r) = \bar{U}(r)$ for some $r \in [r_i, r_{i+1}]$, as shown in Fig. 2. Formally,

(P₁) On every interval (r_i, r_{i+1}) where $I(r)$ is positive, U is the greatest convex function that passes through the endpoints $U(r_i)$ and $U(r_{i+1})$ and does not exceed \bar{U} .

Similarly, on any interval (r_i, r_{i+1}) where $I(r)$ is negative, for any given pairs of $(U(r_i), U'(r_i))$ and $(U(r_{i+1}), U'(r_{i+1}))$, the optimality requires that $U(r)$ is pointwise minimized subject to $U \geq \underline{U}$ and the convexity of U on $[r_i, r_{i+1}]$. That is, the interim utility U on $[r_i, r_{i+1}]$ is an upper envelope of two straight lines that pass through the endpoints $U(r_i)$ and $U(r_{i+1})$ and have slopes $U'(r_i)$ and $U'(r_{i+1})$, as shown in Fig. 3. Formally,

(P₂) On every interval (r_i, r_{i+1}) where $I(r)$ is negative, U is piecewise linear with at most one kink and satisfies

$$U(r) = \max \{U(r_i) + U'(r_i)(r - r_i), U(r_{i+1}) + U'(r_{i+1})(r - r_{i+1})\}.$$

Therefore, we can reduce the sender's problem (11) of optimization on the function space \mathcal{U} to an n -variable optimization problem. The optimal U is pinned down by

⁶If there are intervals where $I(r) = 0$, then on those intervals the sender is indifferent about the choice of U . Hence, multiple solutions emerge in this case. Characterization of these solutions is a straightforward but tedious extension of the result in this section.




FIGURE 3. The optimal interim utility on the interval where $I(r)$ is negative.

properties (P_1) and (P_2) within each interval (r_i, r_{i+1}) where the sign of I is constant. Thus, the optimal U is fully defined by the utility values at the points $\{r_1, \dots, r_n\}$ where I changes its sign.

For every vector $\bar{y} = (y_1, \dots, y_n) \in \mathbb{R}^n$, consider the set of all $U \in \mathcal{U}$ that satisfy $U(r_i) = y_i$ for all $i = 1, \dots, n$. If this set is nonempty, the properties (P_1) and (P_2) uniquely determine the optimal U on this set, denoted by $U_{\bar{y}}^*$. For completeness, define $U_{\bar{y}}^* = \underline{U}$ if this set is empty. Thus, we obtain

Theorem 2. *The sender's problem (11) is an n -variable optimization problem,*

$$\max_{\bar{y} \in \mathbb{R}^n} \int_R U_{\bar{y}}^*(r) I(r) dr.$$

Properties (P_1) and (P_2) imply that the optimal U is piecewise linear except for the intervals where $U(r) = \bar{U}(r)$, as shown in Figs. 2–3. For a given optimal U , the corresponding experiment's distribution H of m is given by $H(r) = 1 + U'(r_+)$ for all $r \in [0, 1]$. Therefore, the intervals on which U is linear correspond to constant H (i.e., zero density). The kinks of U correspond to the mass points of the distribution equal to the difference between the right and left derivatives of U at those points. The intervals on which $U(r) = \bar{U}(r)$ correspond to the full disclosure intervals, $H(r) = F(r)$.

4.3. Upper- and Lower-Censorship Mechanisms. An *upper-censorship* mechanism is a mechanism that reveals the states below ω^* and pools the states above ω^* for some threshold $\omega^* \in \Omega$. This mechanism sends the message $m(\omega) = \omega$ if $\omega < \omega^*$ and $m(\omega) = \mathbb{E}[\omega | \omega > \omega^*]$ if $\omega > \omega^*$. A *lower-censorship* mechanism is defined symmetrically. This mechanism sends $m(\omega) = \omega$ if $\omega > \omega^*$ and $m(\omega) = \mathbb{E}[\omega | \omega < \omega^*]$ if $\omega < \omega^*$ for some $\omega^* \in \Omega$.

The receiver's interim utility under an upper-censorship mechanism is shown as a kinked black curve in Fig. 4. A receiver of type $r < \omega^*$ is fully informed whether ω




FIGURE 4. The interim utility under an upper-censorship mechanism.

exceeds r or not, and hence receives the highest feasible utility $\bar{U}(r)$ by choosing to act if $\omega > r$. A receiver of type $r > \omega^*$ is informed only of whether ω exceeds ω^* or not. Hence, he acts if $\mathbb{E}[\omega | \omega > \omega^*] - r > 0$. His utility is, thus, either $\mathbb{E}[\omega | \omega > \omega^*] - r$ or zero, whichever is greater, corresponding to the two linear segments of the utility curve.

Theorem 3. *Let $\rho(r) = \rho \in \mathbb{R}$. An upper-censorship (lower-censorship) mechanism is optimal for all F and all $\rho \in \mathbb{R}$ if and only if g is log-concave (log-convex) on Ω .*

Proof. We prove Theorem 3 using the following lemma proved in Appendix B.

Lemma 3. *An upper-censorship (lower-censorship) mechanism is optimal for all F if and only if I crosses the horizontal axis at most once and from above (from below) on Ω .*

Function $I(r) = g'(r) + \rho g(r)$ crosses the horizontal axis at most once and from above (from below) on Ω for all $\rho \in \mathbb{R}$ if and only if $g'(r)/g(r)$ is nonincreasing (nondecreasing) on Ω by Proposition 1 of Quah and Strulovici (2012). \square

4.4. Comparative Statics. Theorem 3 allows for a sharp comparative statics analysis on the amount of information that is optimally disclosed by the sender.

Let $\rho(r)$ be constant, $\rho(r) = \rho \in \mathbb{R}$. We extend the definition of density function g to the real line and assume that g is log-concave with $g' + \rho g$ being nonzero almost everywhere on \mathbb{R} . Consider a family of densities g_t of the receiver's type

$$g_t(r) = g(r - t),$$

where $t \in \mathbb{R}$ is a parameter. Because g_t is log-concave on Ω for every t , an upper-censorship mechanism is optimal by Theorem 3. Let $\omega^*(\rho, t) \in \Omega$ be the optimal upper-censorship threshold.

We now show that the sender optimally discloses more information when she is less biased relative to the receiver, and when the receiver is more reluctant to act.

Corollary 2. *For all ρ and t such that $\omega^*(\rho, t) \in (0, 1)$,*

- (a) $\omega^*(\rho, t)$ *is strictly increasing in ρ ,*
- (b) $\omega^*(\rho, t)$ *is strictly increasing in t .*

The proof is in Appendix B.

The intuition for part (a) is that for a higher ρ , the sender puts more weight on the receiver's utility, so she optimally endows the receiver with a higher utility by providing more information.

The intuition for part (b) is that for a higher t , each type of the receiver has a greater cost of action, so to persuade the same type of the receiver, the sender needs to increase $\mathbb{E}[\omega | \omega > \omega^*]$ by expanding the full disclosure interval $[0, \omega^*]$.

4.5. Full and No Disclosure Mechanisms. Theorem 3 also allows to characterize the class of distributions of the receiver's type under which the optimal choice of the sender polarizes between the full and no disclosure mechanisms, as, for example, in Lewis and Sappington (1994) and Johnson and Myatt (2006).⁷

Corollary 3. *Let $\rho(r) = \rho \in \mathbb{R}$. Either the full or no disclosure mechanism is optimal for all F and all ρ if and only if there exist λ and c such that $g(r) = ce^{-\lambda r}$ for all $r \in \Omega$.*

Proof. The only two mechanisms that are both upper- and lower-censorship mechanisms are full and no disclosure. By Theorem 3, both upper- and lower-censorship mechanisms are optimal for all F and all ρ if and only if g is both log-concave and log-convex on Ω . That is, $g'(r)/g(r)$ is both nonincreasing and nondecreasing on Ω , which holds if and only if $g'(r)/g(r)$ is equal to a constant, $-\lambda$, on Ω . \square

Using the same argument as in the proof of Lemma 3, we can also characterize the conditions for the optimality of full and no disclosure for a type-dependent ρ .

Corollary 4. *The full (no) disclosure mechanism is optimal for all F if and only if I is positive (negative) on Ω .*

⁷In the case of log-concave density g , the set of upper- and lower-censorship mechanisms is totally ordered by the rotation order of Johnson and Myatt (2006). If we restrict attention to the set of lower-censorship mechanisms (which are generally suboptimal under log-concave g), then the rotation point is decreasing with respect to the rotation order. By Lemma 1 of Johnson and Myatt (2006), the sender's expected utility is quasiconvex and one of the extreme lower-censorship mechanisms (the full or no disclosure mechanism) is optimal for the sender. But, if we restrict attention to the set of upper-censorship mechanisms (which are optimal under log-concave g), then the rotation point is not decreasing with respect to the rotation order. Therefore, Lemma 1 of Johnson and Myatt (2006) does not apply, and an interior upper-censorship mechanism is generally optimal for the sender.

5. APPLICATION: MEDIA CENSORSHIP

In this section, we apply our model to the problem of media control by the government. In the contemporary world, people obtain information about the government state of affairs through various media sources such as television, newspapers, and internet blogs. Without the media, most people would not know what policies and reforms the government pursues and how effective they are. Media outlets have different positions on the political spectrum and differ substantially in how they select and present facts to cover the same news. People choose their sources of information based on their political ideology and socioeconomic status. This information is valuable for significant individual decisions in migration, investment, occupation, and voting, to name a few. Individuals do not fully internalize externalities that their decisions impose on the society. Likewise, the government may not have the society's best interest at heart. To further its goals, the government then wishes to influence individual decisions by manipulating the information through media. In autocracies and countries with weak checks and balances, the government has power to control the media content.⁸

The government's problem of media control can be represented as the persuasion design problem in Section 2. We apply Theorems 1 and 3 to provide conditions for the optimality of simple censorship policies that shut down all media outlets except the most supportive ones. In other words, to act optimally, the government needs no sophisticated instruments of information disclosure other than censorship. An interpretation of our comparative statics results in Corollary 2 is as follows. First, the government increases censorship if influencing society decisions becomes relatively more important than maximizing individual welfare. Second, the government increases censorship if the society experiences an ideology shock in favor of the government.

5.1. Media Censorship Problem. We present a stylized model with a government, media outlets, and readers. Our model has standard ingredients from the theoretical media literature. The novelty of the model is that the government can censor media outlets.

⁸By measuring the frequency with which media outlets use key words and phrases that tend to discredit government actions, a pro-Kremlin news site Politonline.ru on March 31, 2014, published a list of top-20 most disloyal media outlets in Russia. A Russian political scientist and government adviser, Aleksandr Dugin, suggested that this list should be the order in which the government shuts down these media outlets. Using the broadly formulated legislation on combating extremism, the government agency Roskomnadzor, responsible for overseeing the media and mass communications, has taken actions to effectively shut down a number of media outlets at the top of the disloyalty list. An opposition leader and former chess champion, Garry Kasparov, tweeted in response to the crackdown: "These are huge news sites, not political groups. Giant Echo of Moscow site now just gone. Grani, EJ, Navalny's blog, all blocked in Russia." Russia's leading opposition television channel, Dozhd, lost its main cable and satellite providers. In December 2014, the Guardian reported: "Several news companies have had their editors fired while others have lost studio space." Before a Siberian TV channel, TV-2, was forced to shut down in December 2014, its owner said to the Guardian: "We were under constant pressure to change our editorial policies. We're not an opposition channel; we simply give everyone an opportunity to speak."

The government's state of affairs is a random variable ω drawn from $[0, 1]$ according to the distribution F that admits a density f . There is a continuum of media outlets $[0, 1]$. A media outlet $s \in [0, 1]$ has an editorial policy that endorses the government (sends message $m_s = 1$) if $\omega > s$ and criticizes it (sends message $m_s = 0$) if $\omega < s$.^{9,10} The cutoff s can be interpreted as a slant or political bias of the outlet against the government and can be empirically measured as the frequency with which the outlet uses anti-government language.¹¹

There is a continuum of heterogeneous readers indexed by $r \in [0, 1]$ distributed with G that admits a log-concave density g . Each reader observes endorsements of all available media outlets¹² and chooses between action ($a = 1$) and inaction ($a = 0$). A reader's utility is equal to

$$u(\omega, r, a, \bar{a}) = a(\omega - r) + \bar{a}\xi(r),$$

where \bar{a} denotes the average action in the society, and $\bar{a}\xi(r)$ is a type-specific externality term that contributes to the reader's utility but does not affect the reader's optimal action.

The government's utility is equal to

$$\int_R u(\omega, r, a_r, \bar{a})dG(r) + \bar{a}\gamma = \int_R (a_r(\omega - r) + \bar{a}(\xi(r) + \gamma))dG(r),$$

where a_r denotes an action of type r , and $\bar{a}\gamma$ is the government's intrinsic benefit from the average action. We assume that

$$\rho = \left(\int_R (\xi(r) + \gamma)dG(r) \right)^{-1} > 0,$$

meaning that the government is biased towards a greater average action in the society.

The government's censorship policy is a measurable set of the media outlets $S \subset [0, 1]$ that are prohibited to broadcast. Readers observe messages only from the permitted media outlets in $[0, 1] \setminus S$.

The timing is as follows. First, the government chooses a set of prohibited media outlets. Second, the state of affairs is realized, and every permitted media outlet

⁹If the set of media outlets is finite, we can redefine the state to be the posterior mean of ω given the messages of all media outlets. To characterize the optimal censorship policy, we will need to adjust the analysis in Section 4 to allow for discrete distributions of the state.

¹⁰As in Suen (2004), Chan and Suen (2008), and Chiang and Knight (2011), binary media reports that communicate only whether the state of affairs ω is above some standard s can be justified by a cursory reader's preference for simple messages such as positive or negative opinions and yes or no recommendations.

¹¹Gentzkow and Shapiro (2010) construct such a slant index for U.S. newspapers. Empirical findings of their paper suggest that the editorial policies of media outlets are driven by reader preferences, justifying our assumption of the existence of a large variety of editorial policies. Theoretical literature has explored the determinants of media slant of an outlet driven by its readers (Mullainathan and Shleifer, 2005, Gentzkow and Shapiro, 2006, and Chan and Suen, 2008) and its owners (Baron, 2006, and Besley and Prat, 2006).

¹²This will be compared to Chan and Suen (2008) where each reader observes a single media outlet and does not communicate with other readers.

endorses or criticizes the government, according to its editorial policy. Finally, readers observe messages of the permitted media outlets and decide whether to act or not.

There are various interpretations of the reader's action a , such as refraining from emigration, investing in domestic stocks, volunteering for military service, and voting for the government. We can interpret a reader's type r as his ideological position. A reader who is more supportive of the government has a smaller r .¹³

5.2. Results. We now permit the government to use general persuasion mechanisms and show that the optimal persuasion mechanism can be achieved by a censorship policy.

As in Section 2.2, a persuasion mechanism $\pi(\omega, \hat{r})$ maps a realized state ω and a reader's report \hat{r} about his type to a probability with which action is recommended. Each reader can observe only his own recommendation and cannot share information with other readers. Persuasion mechanisms may be hard to implement in practice, but they serve as a benchmark of what the government could possibly achieve if it had full information control.

Consider an incentive-compatible persuasion mechanism π . Let $\bar{a}_\pi(\omega)$ be the average action conditional on ω , and let $q_\pi(r)$ be the interim action of type r ,

$$\bar{a}_\pi(\omega) = \int_R \pi(\omega, r) dG(r) \quad \text{and} \quad q_\pi(r) = \int_\Omega \pi(\omega, r) dF(\omega).$$

Denote by \bar{A}_π the expected average action,

$$\bar{A}_\pi = \int_R q_\pi(r) dG(r).$$

It is convenient to consider the reader's interim utility net of the externality term,

$$U_\pi(r) = \int_\Omega (\omega - r) \pi(\omega, r) dF(\omega).$$

Lemma 1 implies that $U'_\pi(r) = -q_\pi(r)$, and, hence,

$$\bar{A}_\pi = - \int_R U'_\pi(r) dG(r).$$

By integration by parts, the government's expected utility is equal to

$$V_\pi = \int_R \left(U_\pi(r) + \frac{1}{\rho} \bar{A}_\pi \right) dG(r) = \int_R \left(U_\pi(r) - \frac{1}{\rho} U'_\pi(r) \right) dG(r) = \frac{1}{\rho} \int_R I(r) U_\pi(r) dr,$$

where

$$I(r) = g'(r) + \rho g(r).$$

¹³This media censorship problem can be applied to spatial voting models, as in Chiang and Knight (2011). Consider the government party ($p = G$) and opposition party ($p = O$) competing in an election. If party p wins, a voter with ideological position r gets utility $w_p - (r - r_p)^2$, where w_p is the quality or valence of party p , and r_p is the ideology or policy platform of party p . Voters know the parties' ideologies and obtain information about the parties' qualities from all available media outlets. Each voter supports the party that maximizes his expected utility. Our analysis applies, because the voter's utility difference between the government and opposition parties is proportional to $\omega - r$, where $\omega = (w_G - w_O + r_O^2 - r_G^2) / 2(r_O - r_G)$ represents the state.

From Theorem 3, it follows that the optimal persuasion mechanism is an upper-censorship mechanism with some threshold s^* .¹⁴

Corollary 5. *The government's optimal censorship policy is to prohibit all media outlets $s > s^*$ and permit all media outlets $s < s^*$.*

Note that a media outlet with a higher editorial policy cutoff is more disloyal to the government, in the sense that it criticizes the government on a larger set of states. Corollary 5 says that it is optimal for the government to prohibit all sufficiently disloyal media outlets from broadcasting.

This government's censorship policy is optimal among all persuasion mechanisms. In particular, the government would not be better off if it could restrict each reader to follow a single media outlet of his choice and ban readers from communicating with one another, as in Chan and Suen (2008). Nor would the government be better off if it could create more complex mechanisms that aggregate information from multiple media outlets and add noise.

Enikolopov et al. (2011) study the effect of voters' access to NTV, the only independent national TV channel, on the regional results of the 1999 Russian parliamentary elections. They show that local access to NTV substantially decreased the regional aggregate vote for the government party. Entertaining the possibility that access to NTV is controlled by the government, rather than exogenous as argued by Enikolopov et al. (2011), our paper suggests a different interpretation of their findings. By Corollary 2, the government optimally permits access to NTV only in the regions with low initial support of the government. Thus, access to NTV in these regions may be a consequence, not a cause, of low electoral support of the government.

6. MULTIPLE ACTIONS

In this section, we allow the receiver to make a choice among multiple actions. We characterize the implementable receiver's interim utilities and show that the sender can generally implement a strictly larger set of the receiver's interim utilities by persuasion mechanisms than by experiments. We also formulate the sender's optimization problem and show that the sender can achieve a strictly higher expected utility by persuasion mechanisms than by experiments.

6.1. Preferences. Let $A = \{0, 1, \dots, n\}$ be a finite set of actions available to the receiver. The state $\omega \in \Omega = [0, 1]$ and the receiver's type $r \in R = [0, 1]$ are independent and have distributions F and G . We continue to assume that the receiver's utility is linear in the state for every type and every action.

¹⁴This result relies on the linearity of $\bar{a}\xi(r)$ and $\bar{a}\gamma$ in \bar{a} . Suppose instead that the externality term is $\xi(\bar{a}, r)$ and the government's intrinsic benefit is $\gamma(\bar{a})$, where $\xi(\bar{a}, r)$ and $\gamma(\bar{a})$ need not be linear in \bar{a} . Then, for any censorship policy S , we can still express the government's expected utility as $\int_R \tilde{I}(r)U_S(r)dr$, where U_S is the reader's interim utility net of the externality term, and \tilde{I} is a function independent of S . Therefore, we can use the results of Section 4 to characterize the government's optimal censorship policy. In particular, if the government's intrinsic benefit $\gamma(\bar{a})$ is reverse S -shaped, as under a majoritarian voting system, then the government's optimal censorship policy still corresponds to an upper-censorship mechanism.

It is convenient to define the receiver's and sender's utilities, $u(\omega, r, a)$ and $v(\omega, r, a)$, recursively by the utility difference between each two consecutive actions. For each $a \in \{1, \dots, n\}$,

$$\begin{aligned} u(\omega, r, a) - u(\omega, r, a - 1) &= b_a(r)(\omega - x_a(r)), \\ v(\omega, r, a) - v(\omega, r, a - 1) &= z_a(r) + \rho(r)(u(\omega, r, a) - u(\omega, r, a - 1)), \end{aligned}$$

and the utilities from action $a = 0$ are normalized to zero, $u(\omega, r, 0) = v(\omega, r, 0) = 0$ for all ω and all r .

For each $a \in \{1, \dots, n\}$, the receiver's and sender's utilities can be expressed as

$$\begin{aligned} u(\omega, r, a) &= \left(\sum_{i=1}^a b_i(r) \right) \omega - \left(\sum_{i=1}^a b_i(r) x_i(r) \right) \\ v(\omega, r, a) &= \left(\sum_{i=1}^a z_i(r) \right) + \rho(r) u(\omega, r, a). \end{aligned}$$

We assume that $b_a(r) > 0$ for all r and all $a \in \{1, \dots, n\}$. This assumption means that every type r prefers higher actions in higher states. Note that $x_a(r)$ is the cutoff state at which the receiver of type r is indifferent between two consecutive actions $a - 1$ and a . Define $x_0(r) = -\infty$ and $x_{n+1}(r) = \infty$.

Denote by $\bar{x}_a(r)$ the cutoff truncated to the unit interval,

$$\bar{x}_a(r) = \max \{0, \min \{1, x_a(r)\}\}.$$

We assume that the cutoffs are ordered on $[0, 1]$ such that

$$\bar{x}_1(r) \leq \bar{x}_2(r) \leq \dots \leq \bar{x}_n(r) \text{ for all } r \in R.$$

Thus, type r optimally chooses action a on the interval of states $(\bar{x}_a(r), \bar{x}_{a+1}(r))$.¹⁵

6.2. Experiments. Because the receiver's utility is linear in the state for every type and every action, every experiment σ can be equivalently described by the probability that the posterior mean state is below a given value $x \in \Omega$,

$$H_\sigma(x) = \int_\Omega \sigma(x|\omega) dF(\omega).$$

In fact, as in Blackwell (1951), Rothschild and Stiglitz (1970), and Gentzkow and Kamenica (2016d), it is convenient to describe an experiment by a convex function $C_\sigma : \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$C_\sigma(x) = \int_x^\infty (1 - H_\sigma(m)) dm.$$

Observe that by (9), for every experiment σ , we have $C_\sigma(r) = U_\sigma(r)$ for all r , where $U_\sigma(r)$ is the receiver's interim utility under σ in the problem of Section 2, with two actions and $u(\omega, r, a) = a(\omega - r)$. Hence, by Theorem 1, the set of all C_σ is equal to

$$\mathcal{C} = \{C : \underline{C} \leq C \leq \overline{C} \text{ and } C \text{ is convex}\},$$

¹⁵This assumption ensures that the actions that can be optimal for type r are consecutive. If actions $a - 1$ and $a + 1$ are optimal for type r under states ω' and ω'' , then there must be a state between ω' and ω'' where action a is optimal. This assumption simplifies the exposition. Relaxing this assumption poses no difficulty: it will only require us for each type r to omit from the analysis the actions that are never optimal for this type.

where \overline{C} and \underline{C} correspond to the full and no disclosure experiments,

$$\begin{aligned}\overline{C}(x) &= \int_x^\infty (1 - F(m))dm, \\ \underline{C}(x) &= \max\{\mathbb{E}[\omega] - x, 0\}.\end{aligned}$$

6.3. Implementable Interim Utilities. The expected utility of type r under experiment σ is equal to

$$U_\sigma(r) = \int_\Omega \left(\max_{a \in A} u(m, r, a) \right) dH_\sigma(m) \text{ for all } r \in R.$$

Proposition 1. *U is implementable by an experiment if and only if there exists $C \in \mathcal{C}$ such that*

$$U(r) = \sum_{a=1}^n b_a(r)C(x_a(r)) \text{ for all } r \in R. \tag{12}$$

The proof is in Appendix B.

A persuasion mechanism can be described by a (possibly, infinite) menu of experiments, Σ . The receiver of type r chooses one experiment from the menu and then observes messages only from this experiment. Obviously, the receiver chooses the experiment that maximizes his expected utility,

$$U_\Sigma(r) = \max_{\sigma \in \Sigma} U_\sigma(r) \text{ for all } r \in R.$$

By Proposition 1, it is immediate that the receiver's interim utility U is implementable if and only if there exists a menu $\mathcal{C}_\Sigma \subset \mathcal{C}$ such that

$$U(r) = \max_{C \in \mathcal{C}_\Sigma} \left\{ \sum_{a=1}^n b_a(r)C(x_a(r)) \right\} \text{ for all } r \in R.$$


Theorem 1 shows that the sender can implement the same set of receiver's interim utilities by experiments as by persuasion mechanisms. With more than two actions, however, the sender can generally implement a strictly larger set of interim utilities by persuasion mechanisms than by experiments, as shown in Example 1.

Example 1. Let $A = \{0, 1, 2\}$, F admit a strictly positive density, and $u(\omega, r, a)$ be continuous in r for all ω and all a . Furthermore, suppose that there exist two types $r' < r''$, such that for all $r \in (r', r'')$,

$$0 < x_1(r') < x_1(r) < x_1(r'') < x_2(r') < x_2(r) < x_2(r'') < 1.$$

Consider a persuasion mechanism consisting of the menu of two experiments represented by partitions $\{P', P''\}$, where P' and P'' are the first-best partitions for types r' and r'' ,

$$\begin{aligned}P' &= \{[0, x_1(r')), [x_1(r'), x_2(r')), [x_2(r'), 1]\}, \\ P'' &= \{[0, x_1(r'')), [x_1(r''), x_2(r'')), [x_2(r''), 1]\}.\end{aligned}$$

FIGURE 5. The utility of type r^* in Example 1.

Types r' and r'' choose, respectively, P' and P'' and get their maximum possible utilities $\bar{U}(r')$ and $\bar{U}(r'')$. By the continuity of $u(\omega, r, a)$ in r , there exists a type $r^* \in (r', r'')$ who is indifferent between choosing P' and P'' . By this indifference,

$$L_1(P') + L_2(P') = L_1(P'') + L_2(P''),$$

where, for each $a \in \{1, 2\}$,

$$L_a(P') = \int_{x_a(r^*)}^{x_a(r')} \left(u(\omega, r^*, a) - u(\omega, r^*, a-1) \right) dF(\omega)$$

denotes the utility loss of type r^* from using cutoff $x_a(r')$ rather than his first-best cutoff $x_a(r^*)$ to decide between actions $a-1$ and a . Analogously, for each $a \in \{1, 2\}$, we define $L_a(P'')$.

Fig. 5 illustrates this example. The three blue lines depict the utility of the receiver of type r^* from taking action a , $u(\omega, r^*, a)$, for each $a = 0, 1, 2$. The kinked solid blue line is the utility of type r^* from taking the optimal action, $\max_{a \in \{0, 1, 2\}} u(\omega, r^*, a)$. In Fig. 5, the loss of type r^* from experiment P' relative to the first best, $L_1(P') + L_2(P')$, is the total area of the two shaded triangles (assuming that ω is uniformly distributed). Similarly, the loss of type r^* from experiment P'' relative to the first best, $L_1(P'') + L_2(P'')$, is the total area of the two hatched triangles. For type r^* , these shaded and hatched areas are equal, so type r^* is indifferent between the two experiments.

An experiment that gives the maximum possible utilities $\bar{U}(r')$ and $\bar{U}(r'')$ to types r' and r'' must at least communicate the common refinement of partitions P' and P'' .

Therefore, the utility of type r^* under such an experiment is at least

$$\bar{U}(r^*) - \min\{L_1(P'), L_1(P'')\} - \min\{L_2(P'), L_2(P'')\},$$

which is strictly larger than his utility under the persuasion mechanism,

$$\bar{U}(r^*) - L_1(P') - L_2(P'),$$

unless $L_1(P') = L_1(P'')$ and $L_2(P') = L_2(P'')$.

In Fig. 5, for type r^* , the loss $L_1(P'')$ (left hatched triangle) is smaller than the loss $L_1(P')$ (left shaded triangle). Similarly, the loss $L_2(P')$ (right shaded triangle) is smaller than the loss $L_2(P'')$ (right hatched triangle). Hence, the total loss is smaller under the experiment that is the coarsest common refinement of P' and P'' (the area of the smaller shaded and hatched triangles) than under either experiment P' or experiment P'' .

6.4. Sender's Problem. In this section, we impose the following additional assumptions. For each $a \in \{0, 1, \dots, n\}$, function $x_a(r)$ is strictly increasing, and its range contains Ω . Let function $r_a(x)$ be the inverse of function $x_a(r)$. Moreover, for each $a \in \{0, 1, \dots, n\}$, function $z_a(r)$ is differentiable, and function $G(r_a(x))$ is twice differentiable.

For a given experiment σ , the sender's expected utility conditional on the receiver's type being r is¹⁶

$$V_\sigma(r) = \rho(r)U(r) + \sum_{a=1}^n \left(\sum_{i=1}^a z_i(r) \right) (H_\sigma(x_{a+1}(r)) - H_\sigma(x_a(r))).$$

We now express the sender's expected utility as a function of C_σ , and all the model parameters are summarized in function I , the same way as in Lemma 2 in Section 4.

Lemma 4. *For every experiment σ ,*

$$\int_R V_\sigma(r) dG(r) = K + \int_\Omega C_\sigma(x) I(x) dx,$$

where K is a constant independent of σ and

$$I(x) = \sum_{a=1}^n \left(\frac{d}{dx} (z_a(r_a(x))) \frac{d}{dx} G(r_a(x)) + \rho(r_a(x)) b_a(r_a(x)) \frac{d}{dx} G(r_a(x)) \right).$$

The proof is in Appendix B.

The sender's optimal experiment is described by a function $C \in \mathcal{C}$ that solves

$$\max_{C \in \mathcal{C}} \int_\Omega C(x) I(x) dx.$$

The solutions to this problem are characterized by Theorem 2.

¹⁶For each r where $H_\sigma(x_{a+1}(r))$ is discontinuous, this formula assumes that type r breaks the indifference in favor of action a if the posterior mean state is $x_{a+1}(r)$. This assumption is innocuous because G admits a density, and there are at most countably many discontinuities of H_σ .

As shown in Section 6.3, when the receiver has more than two actions, the sender can implement a strictly larger set of receiver's interim utilities by persuasion mechanisms than by experiments. We now show that the set of implementable interim actions is also strictly larger under persuasion mechanisms. Therefore, even if the sender cares only about the receiver's action, and not his utility, $\rho(r) = 0$ for all r , the sender can achieve a strictly larger expected utility under persuasion mechanisms.

Example 1 (Continued). In addition, let there exist $x_2^* \in (x_1(r''), x_2(r'))$ such that $\mathbb{E}[\omega | \omega \geq x_2^*] = x_2(r'')$ and $\mathbb{E}[\omega | \omega < x_2^*] < x_1(r')$.

An experiment that maximizes the probability of action $a = 2$ for type r'' must send message $x_2(r'')$ if and only if $\omega \in [x_2^*, 1]$. Under any such experiment, type r' takes action $a = 2$ if and only if $\omega \in [x_2^*, 1]$, because for $\omega < x_2^*$, this experiment must generate messages distinct from $x_2(r'')$ and, thus, below x_2^* , which is in turn below $x_2(r')$.

Consider now a persuasion mechanism consisting of the menu of two experiments represented by the following partitions:

$$P' = \{[0, x_2^* - \varepsilon) \setminus [x_1(r'), x_1(r'')), [x_1(r'), x_1(r'')), [x_2^* - \varepsilon, 1]\},$$

$$P'' = \{[0, x_2^* - \varepsilon), [x_2^* - \varepsilon, x_2^*], [x_2^*, 1]\}$$

where $\varepsilon > 0$ is sufficiently small. Type r' strictly prefers P' (to P'') because, for a sufficiently small ε , the benefit of taking action $a = 1$ (rather than $a = 0$) on $[x_1(r'), x_1(r''))$ exceeds the cost of taking action $a = 2$ (rather than $a = 1$) on $[x_2^* - \varepsilon, x_2^*]$. Type r'' is indifferent between P'' and P' , because under both partitions he weakly prefers to take action $a = 0$ on $[0, x_2^* - \varepsilon)$ and action $a = 1$ on $[x_2^* - \varepsilon, 1]$. Therefore, under this mechanism, type r'' takes action $a = 2$ if and only if $\omega \in [x_2^*, 1]$, but type r' takes action $a = 2$ if and only if $\omega \in [x_2^* - \varepsilon, 1]$. As shown above, these probabilities of $a = 2$ for types r' and r'' cannot be achieved by any experiment.

Finally, an optimal persuasion mechanism need not be an experiment. Suppose that the sender cares only about action $a = 2$, i.e., $\rho(r) = z_0(r) = z_1(r) = 0$ and $z_2(r) = 1$ for all r . Also suppose that the support of G contains only r' and r'' with r'' being likely enough, so that the sender's optimal experiment maximizes the probability of action $a = 2$ for type r'' . The persuasion mechanism constructed above gives a strictly larger expected utility to the sender than any experiment.

7. NONLINEAR UTILITY

In this section we allow the receiver's utility to be nonlinear in the state ω .

7.1. Preferences. As in Section 2, the receiver has two actions, $A = \{0, 1\}$, the set of states is $\Omega = [0, 1]$, and the set of receiver's types is $R = [0, 1]$. The receiver's utility, however, is

$$u(\omega, r, a) = au(\omega, r),$$

where $u(\omega, r)$ is differentiable, strictly increasing in ω , and strictly decreasing in r . We also normalize the utility such that, for each $\omega \in \Omega$,

$$u(\omega, \omega) = 0.$$

The sender's utility is

$$v(\omega, r, a) = av(\omega, r),$$

where $v(\omega, r)$ is differentiable. State ω and type r are independent and have distributions F and G .¹⁷

7.2. Characterization of Experiments. We start with the characterization of persuasion mechanisms that are equivalent to experiments.

Proposition 2. *An incentive-compatible persuasion mechanism π is equivalent to an experiment if and only if $\pi(\omega, r)$ is nonincreasing in r for every $\omega \in \Omega$.*

The proof is in Appendix B.

Intuitively, because for each experiment σ , the distribution $\sigma(r|\omega)$ of r conditional on each state ω is nondecreasing in r , each $\pi(\omega, r) \in [1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)]$ (see (2)) is nonincreasing in r .

7.3. Cutoff Mechanisms. Consider a special class of persuasion mechanisms called *cutoff mechanisms*. A cutoff mechanism is described by a menu of cutoff experiments, each of which communicates to the receiver whether the state is above some value or not.

Because a cutoff experiment is identified by its cutoff value x , a cutoff mechanism is defined by a compact set of cutoff values $X \subset \Omega = [0, 1]$. Type r chooses a cutoff $x(r) \in X$, so the selected experiment informs the receiver of whether the state is above $x(r)$ or not. Obviously, different types of the receiver may prefer different cutoffs.

Each upper-censorship (lower-censorship) mechanism defined in Section 4.3 is cutoff mechanism. The upper-censorship mechanism with censorship threshold ω^* can be represented as a set of all cutoff values that are below ω^* .

Proposition 2 allows us to make the following statement.

Corollary 6. *Every cutoff mechanism is equivalent to an experiment.*

The proof is in Appendix B.

For a given cutoff mechanism X , the corresponding experiment is the monotone partition of the state space generated by the observation of all experiments X . For example, if X consists of two elements $x', x'' \in (0, 1)$, then the corresponding partition reveals whether the state ω is below x' , between x' and x'' , or above x'' .

7.4. Binary State. Here we apply Proposition 2 to show that if there are only two states in the support of the prior F , then every incentive-compatible mechanism is equivalent to an experiment.

Corollary 7. *Let the support of F consist of two states. Then every incentive-compatible mechanism π is equivalent to an experiment.*

¹⁷Note that if ω and r are correlated, the analysis below applies if we impose strict monotonicity on function $\tilde{u}(\omega, r) = u(\omega, r)g(r|\omega)/g(r)$ rather than on u , where $g(r)$ and $g(r|\omega)$ denote, respectively, the marginal density of r and the conditional density of r for a given ω . This is because the receiver's interim utility under a mechanism π can be written as $U(r) = \int_{\Omega} \tilde{u}(\omega, r)\pi(\omega, r)dF(\omega)$.

The proof is in Appendix B.

Note that if F has a two-point support, then the receiver's utility is linear in the state without loss of generality, and, hence, Theorem 1 applies. However, Corollary 7 makes a stronger statement, because it asserts that every incentive-compatible mechanism is equivalent to an experiment, not just implements the same receiver's interim utility.

7.5. Beyond Binary State. Suppose now that the support of the prior F consists of three states $\omega_1 < \omega_2 < \omega_3$ and let $f_i = \Pr(\omega_i) > 0$ for $i = 1, 2, 3$.

When there are at least three states and the utility of the receiver is nonlinear in (any transformation of) the state, then the posterior distribution of the state induced by an experiment can no longer be parametrized by a one-dimensional variable—such as the posterior mean state in the case of linear utilities, the posterior probability of one of the states in the case of binary-valued state, and the cutoff value in the case of cutoff mechanisms.

As a consequence, the interim action $q(r)$ and, hence, the sender's interim utility $V(r)$ are no longer pinned down by the receiver's interim utility $U(r)$.

Proposition 3. *Let π_1 and π_2 be two mechanisms that are distinct for each $r \in (\omega_1, \omega_3)$ but implement the same differentiable receiver's interim utility U . Then, the interim action q is the same for π_1 and π_2 if and only if there exist functions b , c , and d such that $u(\omega, r) = c(r) + b(r)d(\omega)$ for each $(\omega, r) \in \{\omega_1, \omega_2, \omega_3\} \times (\omega_1, \omega_3)$.*

The proof is in Appendix B.

When the receiver's utility is nonlinear, the sender can implement a strictly larger set of the receiver's interim actions by persuasion mechanisms than by experiments. Therefore, the sender can achieve a strictly higher expected utility by persuasion mechanisms, even if her utility v is state-independent.

Example 2. Let there be three states, $\omega_1 < \omega_2 < \omega_3$, and two types of the receiver, $r' > r''$. Denote $u'_i = u(\omega_i, r')$ and $u''_i = u(\omega_i, r'')$ for $i \in \{1, 2, 3\}$. Assume $u''_2 < 0 < u'_3$. Moreover, assume $u'_1/u''_1 > u'_2/u''_2$, which means that point (ω_1, u'_1) lies above the dashed line in Fig. 6. Finally, assume that the probability masses (f_1, f_2, f_3) on the states satisfy $f_1u'_1 + f_3u'_3 < 0$ and $f_2u'_2 + f_3u'_3 < 0$.

Let Σ' be the set of all experiments that maximize the probability of action for type r' . It is easy to check that any $\sigma' \in \Sigma'$ induces type r' to act with probability 1 if $\omega = \omega_3$, with probability $-f_3u'_3/(f_2u'_2)$ if $\omega = \omega_2$, and with probability 0 if $\omega = \omega_1$.

Observe that by the monotonicity of u in r , each message of $\sigma' \in \Sigma'$ that induces type r' to act also induces type $r'' < r'$ to act. Moreover, by the definition of Σ' , each message of $\sigma' \in \Sigma'$ that induces type r' not to act can be sent only in states ω_1 or ω_2 where the utility of type r'' is negative by assumption, $u''_1 < u''_2 < 0$, so type r'' does not act either. Thus, for each experiment σ under which type r' acts with probability $f_3(1 - u'_3/u'_2)$ (i.e., $\sigma \in \Sigma'$), type r'' acts with the same probability as type r' .

We now construct a persuasion mechanism that also maximizes the probability that type r' acts, but induces type r'' to act with a different probability. Let Σ'' be the set of all experiments σ'' that induce type r' to act with probability 1 if $\omega = \omega_3$,




FIGURE 6. The utilities of types r' (blue) and r'' (black) in Example 2.

with probability 0 if $\omega = \omega_2$, and with probability $-f_3 u'_3 / (f_1 u'_1)$ if $\omega = \omega_1$. Consider a persuasion mechanism that consists of a menu of two experiments $\{\sigma', \sigma''\}$ with $\sigma' \in \Sigma'$ and $\sigma'' \in \Sigma''$. Notice that type r' is indifferent between σ' and σ'' as he obtains zero expected utility in either case. However, type r'' strictly prefers σ'' to σ' , because $u'_1/u''_1 > u'_2/u''_2$ by assumption. Therefore, under this persuasion mechanism, type r' acts with probability $f_3(1 - u'_3/u'_2)$, but type r'' acts with different probability $f_3(1 - u'_3/u'_1)$.

Finally, when the receiver's utility is nonlinear, the set of receiver's interim utilities implementable by persuasion mechanisms (as compared to experiments) can be strictly larger.

Example 2 (Continued). In addition, let $r^* \in (r', r'')$ be such that $u'_1/u_1^* < u'_2/u_2^*$, where $u_i^* = u(\omega_i, r^*)$ for $i \in \{1, 2, 3\}$.

It is easy to check that Σ' and Σ'' are the sets of experiments σ' and σ'' that maximize the utility of types r^* and r'' , respectively, subject to the constraint that type r' gets utility $\underline{U}(r') = 0$. Because Σ' and Σ'' do not intersect, no experiment can achieve the interim utility induced by a persuasion mechanism that consists of the menu of two experiments $\sigma' \in \Sigma'$ and $\sigma'' \in \Sigma''$.

APPENDIX A. LINEAR UTILITIES

This appendix extends our main results to the class of linear utility functions. The receiver's and sender's utilities are normalized to zero if the receiver does not act,

$a = 0$, and are linear in the state if the receiver acts, $a = 1$,

$$\begin{aligned} u(\omega, r, a) &= a \cdot b(\omega - t), \\ v(\omega, r, a) &= a \cdot (c(\omega - t) + d), \end{aligned}$$

where $r = (b, c, d, t) \in \mathbb{R}^4$ denotes the receiver's type. The type has distribution G that admits a differentiable density g , which is strictly positive on a compact set in \mathbb{R}^4 and zero everywhere else. The state $\omega \in \Omega = [0, 1]$ is independent of r and has distribution F .

Let H_σ be the distribution of the posterior mean induced by an experiment σ . As in Section 6.2, it is convenient to describe σ by

$$C_\sigma(t) = \int_t^\infty (1 - H_\sigma(m)) dm.$$

Proposition 4. *For each experiment σ , the receiver's interim utility is*

$$U_\sigma(r) = |b|C_\sigma(t) + \min\{0, b\}(\mathbb{E}[\omega] - t). \quad (13)$$

There exist $K \in \mathbb{R}$ and $I : \mathbb{R} \rightarrow \mathbb{R}$ such that for each σ the sender's expected utility is

$$V_\sigma = K + \int_{t \in \mathbb{R}} C_\sigma(t) I(t) dt. \quad (14)$$

Proposition 4 allows us to extend Theorems 1 and 2 to this setting. Recall from Section 6 that the set of all C_σ is equal to

$$\mathcal{C} = \{C : \underline{C} \leq C \leq \overline{C} \text{ and } C \text{ is convex}\},$$

where \overline{C} and \underline{C} correspond to the full and no disclosure experiments. Each persuasion mechanism can be described by a (possibly, infinite) menu of experiments, Σ , which the receiver chooses from. By (13), for a given menu Σ , the receiver's interim utility is

$$\max_{\sigma \in \Sigma} U_\sigma(r) = |b| \left(\max_{\sigma \in \Sigma} C_\sigma(t) \right) + \min\{0, b\}(\mathbb{E}[\omega] - t).$$

Notice that $\max_{\sigma \in \Sigma} C_\sigma$ is the upper envelope of convex functions $C_\sigma \in \mathcal{C}$ and hence it is in \mathcal{C} . Therefore, by Proposition 4, any implementable pair of the sender's and receiver's expected utilities is implementable by an experiment. Moreover, the sender's problem can be expressed as

$$\max_{C \in \mathcal{C}} \int_{\mathbb{R}} C(t) I(t) dt,$$

and Theorem 2 holds with U replaced by C .

Proof of Proposition 4. Fix a type $r = (b, c, d, t) \in \mathbb{R}^4$ and evaluate

$$U_\sigma(r) = \int_0^1 \max\{0, b(m - t)\} dH_\sigma(m).$$

Clearly, if $b = 0$, then $U_\sigma(r) = 0$. We now consider two cases, $b > 0$ and $b < 0$.

Case 1: $b > 0$. Given a posterior mean m , the receiver acts if and only if $t < m$. By integration by parts, the receiver's interim utility is

$$U_\sigma(r) = \int_t^1 b(m-t)dH_\sigma(m) = bC_\sigma(t).$$

Again, by integration by parts, the sender's interim utility is

$$V_\sigma(r) = \int_t^1 (c(m-t) + d)dH_\sigma(m) = cC_\sigma(t) - dC'_\sigma(t).$$

Case 2: $b < 0$. Given a posterior mean m , the receiver acts if and only if $t \geq m$. By integration by parts, the receiver's interim utility is

$$U_\sigma(r) = \int_0^t b(m-t)dH(m) = -bC(t) + b(\mathbb{E}[\omega] - t).$$

Again, by integration by parts, the sender's interim utility is

$$V_\sigma(r) = \int_0^t (c(m-t) + d)dH(m) = -cC(t) + dC'(t) + d + c(\mathbb{E}[\omega] - t).$$

We, thus, obtain (13).

We now show that the sender's expected utility is given by (14). Let $g(b, c, d|t)$ be the density of (b, c, d) conditional on t , and let $g_t(t)$ be the marginal density of t . Define

$$c_+(t) = \int cg(b, c, d|t)\mathbf{1}_{b>0}d(b, c, d),$$

and

$$c_-(t) = \int cg(b, c, d|t)\mathbf{1}_{b<0}d(b, c, d).$$

Similarly, define $d_+(t)$ and $d_-(t)$.

Fix t and take the expectation with respect to (b, c, d) on the set of $b > 0$:

$$\int_{(b,c,d)} V_\sigma(b, c, d, t)\mathbf{1}_{b>0}g(b, c, d|t)d(b, c, d) = c_+(t)C_\sigma(t) - d_+(t)C'_\sigma(t).$$

Now, integrating with respect to t ,

$$\begin{aligned} \int_{(b,c,d,t)} V_\sigma(b, c, d, t)\mathbf{1}_{b>0}dG(b, c, d, t) &= \int_t (c_+(t)C_\sigma(t) - d_+(t)C'_\sigma(t))g_t(t)dt \\ &= \int_t (c_+(t)g_t(t) + \frac{d}{dt}[d_+(t)g_t(t)])C_\sigma(t)dt. \end{aligned}$$

Similarly, for $b < 0$,

$$\int_{(b,c,d)} V_\sigma(b, c, d, t)\mathbf{1}_{b<0}g(b, c, d|t)d(b, c, d) = -c_-(t)C_\sigma(t) + d_-(t)C'_\sigma(t) + d_-(t) + c_-(t)(\mathbb{E}[\omega] - t).$$

Now, integrating with respect to t ,

$$\begin{aligned} \int_{(b,c,d,t)} V_\sigma(b, c, d, t) \mathbf{1}_{b < 0} dG(b, c, d, t) &= - \int_t (c_-(t)C_\sigma(t) - d_-(t)C'_\sigma(t))g_t(t)dt + K \\ &= - \int_t (c_-(t)g_t(t) + \frac{d}{dt}[d_-(t)g_t(t)])C_\sigma(t)dt + K, \end{aligned}$$

where

$$K = \int_t \left(d_-(t) + c_-(t)(\mathbb{E}[\omega] - t) \right) g_t(t) \mathbf{1}_{b < 0} dt$$

is a constant independent of C_σ .

Since the measure of types with $b = 0$ is zero, we obtain

$$\int_r V_\sigma(r) dG(r) = \int_r V_\sigma(r) (\mathbf{1}_{b > 0} + \mathbf{1}_{b < 0}) dG(r) = \int_t I(t) C_\sigma(t) dt + K,$$

where

$$I(t) = (c_+(t) - c_-(t))g_t(t) + \frac{d}{dt}[(d_+(t) - d_-(t))g_t(t)].$$

□

APPENDIX B. OMITTED PROOFS

Proof of Lemma 1. Necessity. Incentive compatibility (3) requires that for each $\hat{r} > r$, both r and \hat{r} prefer truth telling

$$\begin{aligned} U_\pi(r) &\geq U_\pi(r, \hat{r}) = U_\pi(\hat{r}) + q_\pi(\hat{r})(\hat{r} - r), \\ U_\pi(\hat{r}) &\geq U_\pi(\hat{r}, r) = U_\pi(r) + q_\pi(r)(r - \hat{r}). \end{aligned}$$

Therefore,

$$-q_\pi(r)(\hat{r} - r) \leq U_\pi(\hat{r}) - U_\pi(r) \leq -q_\pi(\hat{r})(\hat{r} - r),$$

which implies (4). Letting $\hat{r} \rightarrow r$ and then integrating from r to 1 gives

$$U_\pi(1) - U_\pi(r) = - \int_r^1 q_\pi(s) ds.$$

Also, observe that type $r = 1$ can secure his maximal attainable utility of 0 by always acting (irrespective of a recommendation); so $U_\pi(1) = 0$ and (5) follows. Finally, the maximal attainable utility of type $r = 1$ is $\mathbb{E}[\omega]$, which can be secured by never acting; so (6) follows.

Sufficiency. It remains to show that (4)–(6) imply (3). If either $\tilde{r} \geq \hat{r} \geq r$ or $\tilde{r} \leq \hat{r} \leq r$, then (4) and (5) imply that

$$\begin{aligned} U_\pi(r, \hat{r}) &= U_\pi(r) + q_\pi(\hat{r})(\hat{r} - r) = \int_{\hat{r}}^1 q_\pi(s) ds + q_\pi(\hat{r})(\hat{r} - r) \\ &\geq \int_{\hat{r}}^{\tilde{r}} q_\pi(\tilde{r}) ds + \int_{\tilde{r}}^1 q_\pi(s) ds + q_\pi(\tilde{r})(\hat{r} - r) \\ &= \int_{\tilde{r}}^1 q_\pi(s) ds + q_\pi(\tilde{r})(\tilde{r} - r) = U_\pi(r, \tilde{r}), \end{aligned}$$

meaning that $U_\pi(r, \hat{r})$ is single-peaked in \hat{r} , with the peak located at $\hat{r} = r$. Therefore,

$$U_\pi(r) \geq U_\pi(r, \hat{r}) = U_\pi(r, \hat{r}, 0, 1) \text{ for all } r, \hat{r} \in R.$$

Moreover, letting $\hat{r} = 1$ and $\hat{r} = 0$ gives

$$\begin{aligned} U_\pi(r) &\geq U_\pi(r, 1) = U_\pi(1) + q_\pi(1)(1 - r) \geq 0 = U_\pi(r, \hat{r}, 0, 0) \text{ for all } r, \hat{r} \in R, \\ U_\pi(r) &\geq U_\pi(r, 0) = U_\pi(0) - q_\pi(0)r \geq \mathbb{E}[\omega] - r = U_\pi(r, \hat{r}, 1, 1) \text{ for all } r, \hat{r} \in R. \end{aligned}$$

Thus, we are left to show that $U_\pi(r) \geq U_\pi(r, \hat{r}, 1, 0)$ for all $r, \hat{r} \in R$. Notice that

$$U_\pi(r, \hat{r}, 1, 0) = \int_{\Omega} (1 - \pi(\omega, r))(\omega - r) dF(\omega) = \mathbb{E}[\omega] - r - U_\pi(r, \hat{r}).$$

Since $U_\pi(r, \hat{r})$ is single-peaked, we have

$$\begin{aligned} U_\pi(r, \hat{r}, 1, 0) &= \mathbb{E}[\omega] - r - \min\{U_\pi(r, 0), U_\pi(r, 1)\} \leq \max\{0, \mathbb{E}[\omega] - r\} \\ &= \max\{U_\pi(r, \hat{r}, 0, 0), U_\pi(r, \hat{r}, 1, 1)\} \leq U_\pi(r) \text{ for all } r, \hat{r} \in R. \end{aligned}$$

Proof of Lemma 3. The ‘if’ part follows from properties (P₁) and (P₂) in Section 4.2.

Suppose that it is *not* the case that I crosses the horizontal axis at most once and from above on Ω . Then there exist $0 \leq r_1 < r_2 < r_3 \leq 1$ such that I is negative on (r_1, r_2) and positive on (r_2, r_3) (since by assumption I is continuous and nonzero almost everywhere). Therefore, by the ‘if’ part of this lemma, for any F that has support only on $[r_1, r_3]$ a lower-censorship mechanism is optimal. Moreover, by (11) every upper-censorship mechanism is strictly suboptimal. The analogous argument applies for I that does not cross the horizontal axis at most once and from below on Ω .

Proof of Corollary 2. Types $r < 0$ always act, and types $r > 1$ never act; so we omit these types from the analysis. The sender’s expected utility under experiment σ can be written as:

$$V_\sigma = K + \int_0^1 J_t(r) dH_\sigma(r), \quad (15)$$

where K is a constant independent of σ , H_σ is the distribution of posterior values $\mathbb{E}_\sigma[\omega|m]$, and

$$J_t(r) = \int_0^r (g_t(s) + \rho G_t(s)) ds.$$

Consider ρ and t such that $\omega^*(\rho, t) \in (0, 1)$. Under an upper-censorship mechanism with threshold ω^* , $H_\sigma(x) = F(x)$ for $x \in [0, \omega^*)$, $H_\sigma(x) = F(\omega^*)$ for $x \in [\omega^*, \omega^{**})$, and $H_\sigma(x) = 1$ for $x \in [\omega^{**}, 1]$, where $\omega^{**} = \mathbb{E}[\omega|\omega > \omega^*]$.

Part (a). The derivative of the sender’s expected utility (15) under an upper-censorship mechanism with respect to threshold ω^* is

$$\begin{aligned} \frac{dV}{d\omega^*} &= f(\omega^*) \int_{\omega^*}^{\omega^{**}} (J'_t(\omega^{**}) - J'_t(s)) ds \\ &= f(\omega^*) \left[\int_{\omega^*}^{\omega^{**}} (g_t(\omega^{**}) - g_t(s)) ds + \rho \int_{\omega^*}^{\omega^{**}} (G_t(\omega^{**}) - G_t(s)) ds \right]. \end{aligned}$$

This derivative is strictly increasing in ρ ; so ω^* is strictly increasing in ρ by Theorem 1 of Edlin and Shannon (1998).

Part (b). Notice that

$$\begin{aligned} \frac{d^2V}{dt d\omega^*} &= -f(\omega^*) \int_{\omega^*}^{\omega^{**}} (J_t''(\omega^{**}) - J_t''(s)) ds \\ &= -f(\omega^*) J_t''(\omega^{**}) (\omega^{**} - \omega^*) + f(\omega^*) (J_t'(\omega^{**}) - J_t'(\omega^*)). \end{aligned}$$

At the optimal interior threshold, we have $dV/d\omega^* = 0$; so

$$J_t'(\omega^{**}) (\omega^{**} - \omega^*) = \int_{\omega^*}^{\omega^{**}} J_t'(s) ds. \quad (16)$$

Since $J_t''(r)$ is nonzero almost everywhere, J_t' is not constant on (ω^*, ω^{**}) . Moreover, by Theorem 3, $I_t(r) = J_t''(r)$ crosses the horizontal axis at most once and from above; so $J_t'(r)$ is quasiconcave. Therefore, (16) implies that $J_t''(\omega^{**}) < 0$ and $J_t'(\omega^{**}) > J_t'(\omega^*)$; so $d^2V/dt d\omega^* > 0$ and ω^* is strictly increasing in t by Theorem 1 of Edlin and Shannon (1998).

Proof of Proposition 1. The receiver's interim utility under experiment σ is

$$U_\sigma(r) = \sum_{a=0}^n \int_{x_a(r)}^{x_{a+1}(r)} u(m, r, a) dH_\sigma(m) = \sum_{a=1}^n \int_{x_a(r)}^{\infty} b_a(r)(m - x_a(r)) dH_\sigma(m),$$

where we used $u(m, r, 0) = 0$. For each $a \in \{1, \dots, n\}$, integration by parts yields

$$\begin{aligned} \int_{x_a(r)}^{\infty} b_a(r)(m - x_a(r)) dH_\sigma(m) &= -b_a(r)(m - x_a(r))(1 - H_\sigma(m)) \Big|_{x_a(r)}^{\infty} \\ &+ b_a(r) \int_{x_a(r)}^{\infty} (1 - H_\sigma(m)) dm = b_a(r) C_\sigma(x_a(r)). \end{aligned}$$

Summing up the above over $a \in \{1, \dots, n\}$ yields (12).

It follows from Section 6.2 that the set of the receiver's interim utilities implementable by experiments is equal to the set of functions that satisfy (12) for every $C \in \mathcal{C}$.

Proof of Lemma 4. We have

$$\begin{aligned} \sum_{a=1}^n \left(\sum_{i=1}^a z_i(r) \right) (H_\sigma(x_{a+1}(r)) - H_\sigma(x_a(r))) &= \sum_{a=1}^n z_a(r) (1 - H_\sigma(x_a(r))) \\ &= - \sum_{a=1}^n z_a(r) C'_\sigma(x_a(r)), \end{aligned}$$

where we used $x_{n+1}(r) = \infty$ (hence, $H_\sigma(x_{n+1}(r)) = 1$) and $1 - H_\sigma(x) = -C'_\sigma(x)$. By Proposition 1, we, thus, obtain

$$\begin{aligned} V(r) &= \rho(r)U(r) + \sum_{a=1}^n \left(\sum_{i=1}^a z_i(r) \right) (H_\sigma(x_a(r)) - H_\sigma(x_{a+1}(r))) \\ &= \rho(r)U(r) - \sum_{a=1}^n z_a(r)C'_\sigma(x_a(r)) = \sum_{a=1}^n \left(\rho(r)b_a(r)C_\sigma(x_a(r)) - z_a(r)C'_\sigma(x_a(r)) \right). \end{aligned}$$

Fix $a \in \{1, \dots, n\}$ and define the variable $x = x_a(r)$. Hence, $r = r_a(x)$. Using this variable change, we have

$$\begin{aligned} &\int_R \left(\rho(r)b_a(r)C_\sigma(x_a(r)) - z_a(r)C'_\sigma(x_a(r)) \right) dG(r) \\ &= \hat{K}_a + \int_\Omega \left(\rho(r_a(x))b_a(r_a(x))C_\sigma(x) - z_a(r_a(x))C'_\sigma(x) \right) dG(r_a(x)), \end{aligned}$$

where \hat{K}_a is a constant independent of σ because $\Omega = [0, 1] \subset [x_a(0), x_a(1)]$, and, for all σ and all $x \notin \Omega$, we have $C_\sigma(x) = \max\{0, \mathbb{E}[\omega] - x\}$. Now we integrate by parts

$$\begin{aligned} \int_\Omega z_a(r_a(x))C'_\sigma(x) dG(r_a(x)) &= \tilde{K}_a - \int_\Omega C_\sigma(x) \frac{d}{dx} \left(z_a(r_a(x)) \frac{d}{dx} G(r_a(x)) \right) dx \\ &= \tilde{K}_a - \int_\Omega C(x) \frac{d}{dx} \left(z_a(r_a(x)) \frac{d}{dx} G(r_a(x)) \right) dx, \end{aligned}$$

where \tilde{K}_a is a constant independent of σ because, for all σ , we have $C_\sigma(0) = \mathbb{E}[\omega]$ and $C_\sigma(1) = 0$. Thus, we obtain

$$\begin{aligned} &\int_R \left(\rho(r)b_a(r)C_\sigma(x_a(r)) - z_a(r)C'_\sigma(x_a(r)) \right) dG(r) \\ &= K_a + \int_\Omega \left(\rho(r_a(x))b_a(r_a(x)) \frac{dG(r_a(x))}{dx} + \frac{d}{dx} \left(z_a(r_a(x)) \frac{d}{dx} G(r_a(x)) \right) \right) C_\sigma(x) dx, \end{aligned}$$

where $K_a = \hat{K}_a + \tilde{K}_a$. Summing the above over $a \in \{1, \dots, n\}$, we obtain $K + \int_\Omega C_\sigma(x)I(x)dx$, where $K = \sum_a K_a$, and I is defined in Lemma 4.

Proof of Proposition 2. Consider a mechanism π that is equivalent to an experiment σ . Since $\sigma(r|\omega)$ is a distribution function of r conditional on ω , it is nondecreasing in r for each ω . Then, by (2), $\pi(\omega, r)$ is nonincreasing in r for each ω .

Conversely, let $\pi(r, \omega)$ be nonincreasing in r for all ω . For every ω and r define $\sigma(r|\omega) = 1 - \pi(r_+, \omega)$, where $\pi(r_+, \omega)$ denotes the right limit of $\pi(\cdot, \omega)$ at r . Since $\pi(r_+, \omega) \in [0, 1]$ is nonincreasing and right-continuous in r , the function $\sigma(r|\omega)$ is a distribution, which describes the distribution of messages for every given state ω . Thus, σ is an experiment. It remains to verify that the constructed experiment is direct and induces the same action by the receiver as mechanism π , i.e., when the experiment sends a message r , then type r is indifferent between the two actions. For

all r ,

$$\begin{aligned} U_\pi(r) &= \int_{\Omega} u(\omega, r) \pi(\omega, r) dF(\omega) = \int_{\Omega} u(\omega, r_+) \pi(\omega, r_+) dF(\omega) \\ &= \int_{\Omega} u(\omega, r) \pi(\omega, r_+) dF(\omega) = \int_{\Omega} u(\omega, r) (1 - \sigma(r|\omega)) dF(\omega) = U_\sigma(r), \end{aligned}$$

where the first equality holds by the definition of U_π , the second by the absolute continuity of U_π (Theorem 1 of Milgrom and Segal, 2002), the third by the continuity of u in r , the fourth by the definition of σ , and the last by the definition of U_σ for direct experiments. There exist left and right derivatives of U_π for all r (Theorem 3 of Milgrom and Segal, 2002) that satisfy:

$$\begin{aligned} U'_\pi(r_+) &= \int_{\Omega} \frac{\partial u(\omega, r)}{\partial r} \pi(r_+, \omega) dF(\omega), \\ U'_\pi(r_-) &= \int_{\Omega} \frac{\partial u(\omega, r)}{\partial r} \pi(r_-, \omega) dF(\omega). \end{aligned}$$

Since $U_\pi(r) = U_\sigma(r)$ and $\sigma(r|\omega) = 1 - \pi(r_+, \omega)$ for all r , we have

$$\begin{aligned} U'_\sigma(r_+) &= \int_{\Omega} \frac{\partial u(\omega, r)}{\partial r} (1 - \sigma(r|\omega)) dF(\omega), \\ U'_\sigma(r_-) &= \int_{\Omega} \frac{\partial u(\omega, r)}{\partial r} (1 - \sigma(r_-|\omega)) dF(\omega), \end{aligned}$$

showing that type r is indifferent between the two actions upon receiving message r .

Proof of Corollary 6. Consider a cutoff mechanism with the set of cutoff values X . The receiver's choice to ignore the mechanism messages and always act (never act) can be described by the auxiliary cutoff $x = -\epsilon$ ($x = 1 + \epsilon$, respectively), where $\epsilon > 0$. Define $\tilde{X} = X \cup \{-\epsilon, 1 + \epsilon\}$. Since the receiver's utility is increasing in ω , for every type r , an optimal cutoff is

$$x^*(r) \in \arg \max_{x \in \tilde{X}} \int_x^1 u(r, \omega) dF(\omega),$$

and the corresponding mechanism is

$$\pi(\omega, r) = \begin{cases} 0, & \text{if } \omega < x^*(r), \\ 1, & \text{if } \omega > x^*(r). \end{cases}$$

Since the receiver's utility is decreasing in r , $x^*(r)$ is nondecreasing in r ; hence $\pi(\omega, r)$ is nonincreasing in r for every ω . By Proposition 2, π is equivalent to an experiment.

Proof of Corollary 7. Consider F whose support consist of two states, without loss of generality, $\{0, 1\}$, and let π be an incentive-compatible persuasion mechanism. By

Proposition 2, it is sufficient to show that π is nonincreasing in r for all $r \in (0, 1)$. Incentive compatibility implies that for all $r, \hat{r} \in (0, 1)$,

$$\sum_{\omega=0,1} u(\omega, r) (\pi(\omega, r) - \pi(\omega, \hat{r})) \Pr(\omega = 1) \geq 0. \quad (17)$$

Rewriting (17) twice, with $(r, \hat{r}) = (r_2, r_1)$ and $(r, \hat{r}) = (r_1, r_2)$, yields the inequalities

$$-\frac{u(0, r_2)}{u(1, r_2)} \delta(r_2, r_1, 0) \leq \delta(r_2, r_1, 1) \leq -\frac{u(0, r_1)}{u(1, r_1)} \delta(r_2, r_1, 0) \quad (18)$$

where $\delta(r_2, r_1, \omega) = (\pi(\omega, r_2) - \pi(\omega, r_1)) \Pr(\omega = 1)$. Because $u(0, r) < 0$ and $u(1, r) > 0$ for $r = r_1, r_2$, the monotonicity of u in r implies that

$$0 < -\frac{u(0, r_2)}{u(1, r_2)} \leq -\frac{u(0, r_1)}{u(1, r_1)} \text{ for } r_2 \leq r_1. \quad (19)$$

Combining (18) and (19) gives $\pi(\omega, r_2) \geq \pi(\omega, r_1)$ if $r_2 \leq r_1$ for each $\omega = 0, 1$.

Proof of Proposition 3. For all $r \in (\omega_1, \omega_3)$ and $j = 1, 2$, we have

$$\begin{aligned} U(r) &= \sum_{i=1}^3 u(\omega_i, r) \pi_j(\omega_i, r) f_i, \\ U'(r) &= \sum_{i=1}^3 \frac{\partial u(\omega_i, r)}{\partial r} \pi_j(\omega_i, r) f_i, \end{aligned}$$

where the first line holds by the definition of U , and the second line by the incentive compatibility of π .

The expected action, $q_{\pi_j}(r) = \sum_{i=1}^3 \pi_j(\omega_i, r) f_i$ is the same across $j = 1, 2$ for each r if and only if the vectors $u(\omega, r)$, $\frac{\partial u(\omega, r)}{\partial r}$, and $\mathbf{1}$ are linearly dependent for each r . That is, for each $(\omega, r) \in \{\omega_1, \omega_2, \omega_3\} \times (\omega_1, \omega_3)$, there exist functions $\gamma(r)$ and $\mu(r)$ such that

$$\frac{\partial u(\omega, r)}{\partial r} + \mu(r) u(\omega, r) = \gamma(r). \quad (20)$$

The solution of differential equation (20) is given by

$$u(\omega, r) = e^{-\int_{\omega_1}^r \mu(x) dx} \left(\eta(\omega) + \int_{\omega_1}^r \gamma(x) e^{\int_{\omega_1}^x \mu(y) dy} dx \right),$$

where function $\eta(\omega)$ satisfies the (initial) normalization condition $u(\omega, \omega) = 0$. This completes the proof with $b(r)$, $c(r)$, and $d(\omega)$ given by

$$(b(r), c(r), d(\omega)) = \left(e^{-\int_{\omega_1}^r \mu(x) dx} \int_{\omega_1}^r \gamma(x) e^{\int_{\omega_1}^x \mu(y) dy} dx, e^{-\int_{\omega_1}^r \mu(x) dx}, \eta(\omega) \right).$$

REFERENCES

- ALONSO, R., AND O. CÂMARA (2016a): “Bayesian Persuasion with Heterogeneous Priors,” *Journal of Economic Theory*, 165, 672–706.
- (2016b): “On the Value of Persuasion by Experts,” mimeo.
- (2016c): “Persuading Voters,” *American Economic Review*, 106, 3590–3605.
- AU, P. H. (2015): “Dynamic Information Disclosure,” *RAND Journal of Economics*, 46, 791–823.
- BAGNOLI, M., AND T. BERGSTROM (2005): “Log-concave Probability and Its Applications,” *Economic Theory*, 26, 445–469.
- BARON, D. P. (2006): “Persistent Media Bias,” *Journal of Public Economics*, 90(1), 1–36.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2016a): “The Design and Price of Information,” mimeo.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- BERGEMANN, D., B. A. BROOKS, AND S. MORRIS (2016b): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, forthcoming.
- BERGEMANN, D., AND S. MORRIS (2013): “Robust Predictions in Games With Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.
- BESLEY, T., AND A. PRAT (2006): “Handcuffs for the Grabbing Hand? Media Capture and Government Accountability,” *American Economic Review*, 96(3), 720–736.
- BLACKWELL, D. (1951): “Comparison of Experiments,” in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, pp. 93–102.
- (1953): “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 24, 265–272.
- BOARD, S., AND J. LU (2016): “Competitive Information Disclosure in Search Markets,” mimeo.
- CHAN, J., AND W. SUEN (2008): “A Spatial Theory of News Consumption and Electoral Competition,” *Review of Economic Studies*, 75, 699–728.
- CHIANG, C.-F., AND B. KNIGHT (2011): “Media Bias and Influence: Evidence from Newspaper Endorsements,” *Review of Economic Studies*, 78, 795–820.
- EDLIN, A. S., AND C. SHANNON (1998): “Strict Monotonicity in Comparative Statics,” *Journal of Economic Theory*, 81, 201–219.
- ELY, J., A. FRANKEL, AND E. KAMENICA (2015): “Suspense and Surprise,” *Journal of Political Economy*, 103, 215–260.
- ELY, J. C. (2016): “Beeps,” *American Economic Review*, forthcoming.
- ENIKOLOPOV, R., M. PETROVA, AND E. ZHURAVSKAYA (2011): “Media and Political Persuasion: Evidence from Russia,” *American Economic Review*, 101, 3253–3285.
- ESÖ, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705–731.
- GENTZKOW, M., AND E. KAMENICA (2014): “Costly Persuasion,” *American Economic Review, Papers & Proceedings*, 104, 457–462.
- (2016a): “Bayesian Persuasion with Multiple Senders and Rich Signal Spaces,” mimeo.
- (2016b): “Competition in Persuasion,” *Review of Economic Studies*, forthcoming.
- (2016c): “Disclosure of Endogenous Information,” *Economic Theory Bulletin*, forthcoming.
- (2016d): “A Rothschild-Stiglitz Approach to Bayesian Persuasion,” *American Economic Review, Papers & Proceedings*, 106, 597–601.
- GENTZKOW, M., AND J. SHAPIRO (2006): “Media Bias and Reputation,” *Journal of Political Economy*, 114(2), 280–316.

- (2010): “What Drives Media Slant? Evidence from US Daily Newspapers,” *Econometrica*, 78(1), 35–71.
- GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On the Equivalence of Bayesian and Dominant Strategy Implementation,” *Econometrica*, 81, 197–220.
- GILL, D., AND D. SGROI (2008): “Sequential Decisions with Tests,” *Games and Economic Behavior*, 63, 663–678.
- (2012): “The Optimal Choice of Pre-launch Reviewer,” *Journal of Economic Theory*, 147, 1247–1260.
- GUTMANN, S., J. H. B. KEMPERMAN, J. A. REEDS, AND L. A. SHEPP (1991): “Existence of Probability Measures with Given Marginals,” *Annals of Probability*, 19, 1781–1797.
- HARDY, G. H., J. E. LITTLEWOOD, AND G. PÓLYA (1929): “Some Simple Inequalities Satisfied by Convex Functions,” *Messenger Math*, 58, 145–152.
- HÖRNER, J., AND A. SKRZYPACZ (2016): “Selling Information,” *Journal of Political Economy*, forthcoming.
- JOHNSON, J. P., AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96, 756–784.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KOLOTILIN, A. (2015): “Experimental Design to Persuade,” *Games and Economic Behavior*, 90, 215–226.
- (2016): “Optimal Information Disclosure: A Linear Programming Approach,” mimeo.
- LEWIS, T. R., AND D. SAPPINGTON (1994): “Supplying Information to Facilitate Price Discrimination,” *International Economic Review*, 35, 309–327.
- LI, F., AND P. NORMAN (2016): “On Persuasion with Multiple Senders,” mimeo.
- LI, H., AND X. SHI (2016): “Discriminatory Information Disclosure,” mimeo.
- MANELLI, A., AND D. VINCENT (2010): “Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model,” *Econometrica*, 78, 1905–1938.
- MATHEVET, L., J. PEREGO, AND I. TANEVA (2016): “Information Design: The Belief Hierarchy Approach,” mimeo.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MIRRLEES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175–208.
- MULLAINATHAN, S., AND A. SHLEIFER (2005): “The Market for News,” *American Economic Review*, 95(4), 1031–1053.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- OSTROVSKY, M., AND M. SCHWARZ (2010): “Information Disclosure and Unraveling in Matching Markets,” *American Economic Journal: Microeconomics*, 2, 34–63.
- PEREZ-RICHET, E. (2014): “Interim Bayesian Persuasion: First Steps,” *American Economic Review, Papers & Proceedings*, 104, 469–474.
- QUAH, J., AND B. STRULOVICI (2012): “Aggregating the Single Crossing Property,” *Econometrica*, 80, 2333–2348.
- RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949 – 987.
- ROTHSCHILD, M., AND J. STIGLITZ (1970): “Increasing Risk: I. A definition,” *Journal of Economic Theory*, 2, 225–243.
- SUEN, W. (2004): “The Self-Perpetuation of Biased Beliefs,” *Economic Journal*, 114, 377–396.
- TANEVA, I. (2016): “Information Design,” mimeo.