

PERSUASION OF A PRIVATELY INFORMED RECEIVER

ANTON KOLOTILIN, MING LI, TYMOFIY MYLOVANOV, ANDRIY
ZAPECHELNYUK

ABSTRACT. We study Bayesian persuasion in linear environments with a privately informed receiver. We allow the sender to condition information provided to the receiver on the receiver's report about his type (private persuasion). We describe implementable outcomes, establish equivalence between public persuasion (studied in the prior literature) and private persuasion, and draw connections with the standard quasi-linear mechanism design with transfers. We also characterize optimal persuasion rules, establish monotone comparative statics, and consider several applications, such as a competitive market for a good with consumption externalities (e.g., cigarettes) in which a benevolent government designs an educational campaign about a payoff-relevant characteristic (e.g., the health risks of smoking).

JEL Classification: D81, D82, D83

Keywords: information disclosure, Bayesian persuasion, screening, mechanism design without transfers, advertising, certification, mass media

Date: February 12, 2015.

Kolotilin: School of Economics, University of New South Wales, Sydney, NSW 2052, Australia. E-mail: akolotilin@gmail.com.

Li: Department of Economics, Concordia University and CIREQ. E-mail: ming.li@concordia.ca.

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.

1. INTRODUCTION

A sender wishes to manipulate a receiver's beliefs about an optimal action, but there is uncertainty about the receiver's preferences. This is a Bayesian persuasion environment à la Kamenica and Gentzkow (2011) and Rayo and Segal (2010),¹ with a novelty that the receiver has private information (type) about his tastes. Numerous natural applications fit this environment: public educational campaigns, optimal certification policies of financial and consumer products, media censorship regulations, informational lobbying, etc. (Some of the applications are discussed in detail in Section 6.)

By way of motivation, consider, for example, a government that would like to persuade the public to reduce the amount of public smoking by commissioning an educational campaign about the health risks of cigarettes. What is the optimal way to conduct the educational campaign? Should the government target different consumers by providing them with different information or is it optimal to provide the same information to everyone? When should the government withhold information from the public? How does the distribution of tastes (and, possibly biases) among the consumers affect the amount of information released by the government?

In our model, the receiver has a binary choice: to act or not. His utility difference between action and inaction depends on his private type and an uncertain state, and it is additively separable. The sender is biased and might wish the receiver to act in situations where a fully informed receiver would have not. The sender chooses any mechanism that sends messages depending on the type reported by the receiver and the realized state. The receiver makes a report, observes a message from the mechanism, and then takes an optimal action based on his updated beliefs about the state. Thus, the sender can screen different types of receivers by targeting information disclosure to the receiver's reports. We call this model *private* persuasion.

The important benchmark is *public* persuasion in which the sender designs a public signal about the state that is independent of the receiver's report. The literature on Bayesian persuasion focuses on the environments in which the receiver does not have private information and, hence, private persuasion is not relevant. Public persuasion with private information is considered in Rayo and Segal (2010), Kamenica and Gentzkow (2011), and Kolotilin (2014b).

The first result of the paper is equivalence of private and public persuasion. The result is based on an elegant connection between Mirlees integral representation of incentive-compatibility (Mirlees 1971) and Blackwell integral representation of garbling (Blackwell 1953). First, using the integral envelope representation of the incentive compatibility we characterize the set of receiver's utility profiles implementable by private persuasion. It consists of every convex function bounded by the utility profiles obtained if the sender reveals the state (full information) and if the sender reveals nothing (no information). Second, using Blackwell integral representation of garbling we observe that every convex function bounded by full information and no

¹We study the environment of Kamenica and Gentzkow (2011); our results do not apply to the environment in Rayo and Segal (2010).

information utility profiles can be obtained by a public signal that is a garbling of the fully revealing signal, establishing equivalence.

The equivalence result tells us that screening the receiver’s type by the provided conditioning information on the receiver’s report does not expand the set of implementable utility profiles. The result is non-obvious because different types of the receiver have different preferences over disclosed information: The ideal signal for the receiver reveals whether his utility difference between action and inaction is positive or negative, which varies with the receiver’s type. The equivalence result will fail if in addition to information disclosure a mechanism can use transfers (c.f. Bergemann, Bonatti, and Smolin (2014) and Krähmer and Mylovanov (2014)).

The equivalence has implications for optimality of public persuasion for the sender in our model. We consider the sender’s payoff to be a weighted sum of the probability of acting and the utility of the receiver. By incentive compatibility, the probability of acting is equal to receiver’s marginal utility. Hence, the sender’s payoff can be expressed in terms of the receiver’s utility profile, which together with the equivalence result establishes that public persuasion is optimal.

Our second result is the characterization of optimal persuasion mechanisms and comparative statics. A public signal is called *upper censorship* if it sends the message equal to the realized value of the state whenever it is below some threshold, and sends no message otherwise. In the latter case, the receiver infers only that the state is above the threshold. Upper censorship is optimal if the probability density of the receiver’s types is logconcave—the condition which is satisfied for most densities used in applications.²

The structure of the optimal mechanism underscores the effect of privacy of the receiver’s information. By way of example, assume that the sender would like to maximize the probability of action. If the receiver’s preferences were public, the optimal (public) mechanism would be a signal that reveals whether the state is above or below some cutoff. The cutoff is designed to make the receiver indifferent to act if the news is good (i.e., the state is above the cutoff) and strictly prefers not to act otherwise. In this mechanism, the sender offers no value to the receiver, raising the question why the receiver should employ the sender in the first place.³ By contrast, in our model privacy of the type generates information rents for the receiver.

Upper censorship provides more information than cutoff mechanisms (which are optimal in the benchmark environment with publicly known type of the receiver). It also resembles some of the persuasion mechanisms observed in reality. Consider, for example, the question of optimal certification of financial products.⁴ Upper censorship can be approximated by a coarse ranking, with one high grade that bunches together

²Lower censorship, defined symmetrically, is optimal if the probability density of the receiver’s types is logconvex.

³Since the receiver weakly prefers not to act regardless of the signal realization, he obtains the same payoff as from ignoring the signals and taking no action. For this setting, Kolotilin (2014a) derives the comparative statics over the distributions of states of nature and public types of the receiver

⁴For optimal certification see Lizzeri (1999) and the subsequent literature.

good securities (the states above the threshold) and multiple lower grades (the states below the threshold). The high grade is rather uninformative, but is nevertheless a signal of good enough quality that makes cautious investors buy the security. The other grades signal negative news, and are pretty informative, inducing purchase from a smaller set of less risk-averse investors.

We establish a comparative statics result that the sender optimally discloses more information if she puts a higher weight on the receiver’s utility and if economic fundamentals are more biased against the sender. As extreme cases, revealing full information is optimal if the density of the receiver’s type is increasing and revealing no information can be optimal if the receiver type’s density is decreasing. Thus, privacy of the receiver’s tastes can have an ambiguous result on the amount of information revealed by the sender. Nevertheless, a relatively higher probability of a more skeptical receiver will tend to force the sender to reveal more information.

In Section 6 we consider three applications to demonstrate versatility of the model. We start with the design of optimal educational campaigns about the value of a product in competitive markets with positive or negative consumption externalities and consumer biases. Our second application is the government’s control of the editorial policies of media outlets. Finally, we consider optimal informational lobbying policy of a politician with uncertain stand on an issue. In all three applications, we provide conditions under which upper censorship policy is optimal and show the relationship between the economic fundamentals and the amount of information revealed by the sender.

We conclude the paper with comparison of our results with those in Johnson and Myatt (2006). The sets of upper-censorship (and lower-censorship) mechanisms are totally ordered by cutoff points and, thus, by rotation order (Johnson and Myatt 2006). The set of lower-censorship mechanisms satisfies conditions of Lemma 1 of Johnson and Myatt (2006), which entails that in the environments without uncertainty about the receiver’s type one of the extremes (full information or no information) must be optimal. In Section 7, we discuss why this result does not hold in our environment with a privately informed receiver.

Two classes of results are relegated to the Appendix. We provide structural characterization of optimal public signals for general environments, beyond logconcave and logconvex density, in Appendix A. We explore the limits of equivalence between public and private persuasion outside of linear environments in Appendix B. Some proofs are deferred to Appendix C.

Related literature. Our Bayesian persuasion model is a variation of Kamenica and Gentzkow (2011), with more structure on the preferences and the action space, and a novelty that the receiver is privately informed.⁵ We study private persuasion in which information revealed to the receiver depends on her reported private information. Public persuasion, in which the information revealed must be identical for all

⁵For recent work on persuasion by disclosure, see, for example, Che, Dessein, and Kartik (2013), Koessler and Renault (2012), Hagenbach, Koessler, and Perez-Richet (2014), and Hoffmann, Inderst, and Ottaviani (2014).

receiver types, has been covered in Kamenica and Gentzkow (2011, Section VI.A), Rayo and Segal (2010), and Kolotilin (2014b). Kamenica and Gentzkow (2011) provide a methodological contribution; Rayo and Segal (2010) consider a sufficiently different model, with a more complex preference structure; and Kolotilin (2014b) provides an algebraic criterion of optimality for public persuasion mechanisms.

Bayesian persuasion with a privately informed sender is considered in Rayo and Segal (2010), Perez-Richet (2014) and Alonso and Câmara (2014a).⁶ Competition in Bayesian persuasion is analyzed in Gentzkow and Kamenica (2012),⁷ Bayesian persuasion with costly signals is studied in Gentzkow and Kamenica (2014a), and endogenous acquisition of information in Bayesian persuasion is introduced in Gentzkow and Kamenica (2014b). Dynamic Bayesian persuasion environments appear in Ely, Frankel, and Kamenica (2015) and Ely (2014).⁸ Variations of Bayesian persuasion with monetary transfers and selling information are studied in Horner and Skrzypacz (2009), Bergemann and Bonatti (2013), Bergemann, Bonatti, and Smolin (2014), and Krähmer and Mylovannov (2014). Alonso and Câmara (2014b) explore Bayesian persuasion without common priors.

Bayesian persuasion is a single-agent (decision) problem. Informational structures in multi-agent environments are studied in Bergemann and Morris (2013a) and Bergemann and Morris (2013b). Bergemann, Brooks, and Morris (2013) analyze limits of price discrimination on the set of all information structures and Bergemann, Brooks, and Morris (2012) focus on information structures in first price auctions.⁹ Alonso and Câmara (2014c) study Bayesian persuasion in voting environments. Bayesian persuasion with multiple receivers is also considered in Taneva (2014) and Wang (2013).

2. MODEL

2.1. Setup. There are two players: the sender (she) and the receiver (he). The receiver takes a binary action: to act or not to act. There are two payoff-relevant random variables: state of nature $\omega \in \Omega = [0, 1]$ and the receiver's type $r \in R = [0, 1]$. The state of nature is uncertain, but the receiver privately observes his type. Random variables ω and r are independently distributed with c.d.f. F and G , respectively, where G admits a strictly positive differentiable density g .

The receiver's payoff is $u(\omega, r) = \omega - r$ if he acts and is normalised to zero if he does not act. The sender's payoff is $v(\omega, r) = 1 + \rho(r)u(\omega, r)$ if the receiver acts and is zero if the receiver does not act, where $\rho(r) \in \mathbb{R}$. That is, the sender is biased towards $a = 1$, but she also puts a (type-specific) weight $\rho(r)$ on the receiver's payoff. In particular, if the weight is very large, then the sender's and receiver's interests

⁶Perez-Richet and Prady (2012) is a related model with a restricted set of persuasion rules.

⁷Boleslavsky and Cotton (2015a, 2015b) study related models with competing persuaders.

⁸See also Kremer, Mansour, and Perry (2014) for an optimal disclosure policy in a dynamic environment.

⁹The literature on optimal informational structures in auctions has been initiated in Bergemann and Pesendorfer (2007). See Eso and Szentes (2007), Bergemann and Wambach (2013), and Hao and Shi (2015) for recent contributions.

are aligned, but if the weight is zero, then the sender cares only about whether the receiver acts or not. The utility is not transferrable; there are no monetary payments.

The assumption that $u(\omega, r) = \omega - r$ is without loss of generality if we consider the class of the receiver's payoff functions of the form $u(\omega, r) = b(r) + c(r)d(\omega)$ for some functions b , c , and d , with c having a constant sign, as we note in Remark 2. The boundary conditions $u(\omega, 0) \geq 0$ and $u(\omega, 1) \leq 0$ for all $\omega \in \Omega$ allow for elegance of presentation; relaxing these assumptions poses no difficulty.

2.2. Private persuasion. In order to influence the decision made by the receiver, the sender can design a test that asks the receiver to report his private information and sends a message to the receiver conditional on his report and the realized state. This is an environment of private persuasion because the test can depend non-trivially on the information revealed by the receiver.

We are interested in optimal tests and adopt the mechanism design approach. A private persuasion mechanism (test) $\pi : \Omega \times R \rightarrow [0, 1]$ asks the receiver to report $\hat{r} \in [0, 1]$ and then provides him with a binary message: 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 the complementary probability. A private mechanism is incentive-compatible if the receiver finds it optimal to report his true type and follow the mechanism's recommendation.

By the revelation principle, the focus on private incentive compatible-mechanisms is without loss of generality in that any equilibrium outcome of any game between the sender and the receiver, in which the value of ω is disclosed in some way to the receiver, can be replicated by an incentive compatible private mechanism. In particular, the restriction that the mechanism returns a binary recommendation about action instead of messages about the state is without loss since action is binary.

2.3. Public persuasion. Under public persuasion, the messages are independent of reports of the receiver, so all types of the receiver are informed identically. A public persuasion mechanism $\sigma : \Omega \rightarrow \Delta(M)$ sends to the receiver a randomized message distributed according to measure $\sigma(\omega)$ for each realized state ω , where M denotes a message space.

For a given public signal σ , each message m induces a posterior belief of the receiver about the distribution of states, and hence the posterior value of state $\mathbb{E}[\omega|m]$. As in Kamenica and Gentzkow (2011), without loss assume that the set of messages is $M = \Omega = [0, 1]$ and that messages are *direct*, in the sense that they inform the receiver about the posterior value of the state, $m = \mathbb{E}[\omega|m]$.

Observe that every public signal σ is identical to mechanism π defined by

$$\pi(\omega, r) = \Pr \left[m \geq r \mid m \in \text{Supp}(\sigma(\omega)) \right]. \quad (1)$$

2.4. Interpretation: Continuum of receivers and ex-post implementation.

In some applications, it can be useful to think about a population of heterogeneous receivers of mass one parametrized by the privately known preference parameter r distributed according to G . Under private persuasion, the receivers report their types

to the sender and then receive private recommendations from the mechanism. Under public persuasion, all receivers observe the same public message about the state.

By definition (c.f. (1)), an outcome of public persuasion can be achieved through a private mechanism. Furthermore, in this private mechanism, the incentive compatibility and obedience constraints are satisfied ex-post, even after the receiver learns the entire profile of recommendations that would be made to other types. The reverse statement is true as well: if a private mechanism is ex-post incentive compatible and obedient, then it can be implemented through public persuasion.

3. ENVELOPE CHARACTERIZATION OF INCENTIVE COMPATIBILITY

Denote the expected payoff of a receiver of type r who reports \hat{r} and takes actions a_0 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 payoff 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

$$U_\pi(r) \geq U_\pi(r, \hat{r}, a_0, a_1) \text{ for all } r, \hat{r} \in R, a_0, a_1 \in \{0, 1\}. \quad (\text{IC})$$

It is convenient to introduce the notation for the expected payoff of the *obedient* receiver, who makes report \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 payoff function we can draw the parallel to the standard linear mechanism design problem with transfers, where r is a private value, $q_\pi(\hat{r})$ is the probability of transaction and $p_\pi(\hat{r})$ is the expected monetary transfer that depend on report \hat{r} . The classical envelope argument yields the following lemma:

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

$$U_\pi(r) = \int_r^1 q_\pi(s) ds, \quad (2)$$

$$U_\pi(0) = \mathbb{E}[\omega], \quad (3)$$

$$q_\pi \text{ is non-increasing.} \quad (4)$$

Proof. The proof is in the Appendix. □

Interestingly, the obedience constraints for the intermediate types are implied by the obedience constraints for the boundary types and incentive compatibility, $U_\pi(r) \geq U_\pi(r, \hat{r})$. To disobey by ignoring the recommendation, that is, to act (not to act) irrespective of what is recommended, is the same as to pretend 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 monotonicity 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 in the following differences in characterization.

First, there are two boundary conditions, $q(0) = 1$ and $q(1) = 0$:

- (a) Since $\omega \geq 0$, we have $\omega - r \geq 0$ for all ω if $r = 0$. Hence, type 0's payoff is maximized by taking action for any belief about the state, implying $q(0) = 1$ and the payoff equal to the expected state. This is (3).
- (b) Since $\omega \leq 1$, we have $\omega - r \leq 0$ for all ω if $r = 1$. Hence, type 1 never acts, implying $q(1) = 0$ and $U(1) = 0$. This is (2), with $r = 1$.

Second, monotonicity of q 's and boundary conditions $q(0) = 1$ and $q(1) = 0$ are *not* sufficient for incentive compatibility. Conditions (2) and (3) imply that the area under q must be equal to the expected state, $\int_0^1 q_\pi(s) ds = \mathbb{E}[\omega]$.

Finally, not every pair (q, U) that satisfies conditions (2)–(4) is feasible, that is, a mechanism π that implements such a pair need not exist. For example, if $\omega = 1/2$ with certainty, then every monotonic q with $q(1) = 0$, $q(0) = 1$, and $\int_0^1 q(r) dr = \frac{1}{2}$ satisfies (2)–(4). Among these functions, the unique feasible q is $q(r) = 1$ if $r \leq 1/2$ and $q(r) = 0$ otherwise.

4. EQUIVALENCE OF PRIVATE AND PUBLIC PERSUASION

4.1. Bounds on the receiver's payoff. Consider two simple mechanisms. The *full disclosure* mechanism informs the receiver about the state, so the receiver acts iff his type is below the realized state ω , i.e., $\pi(\omega, r) = 1$ iff $r \leq \omega$, and the expected payoff 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 and only if his type r is below the ex-ante expected value of the state, i.e., $\pi(r, \omega) = 1$ iff $r \leq \mathbb{E}[\omega]$, and the expected payoff is

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

Thus, $\underline{U}(r)$ is the optimal payoff of the receiver based on prior information about ω as given by F , while $\bar{U}(r)$ is the receiver's expected payoff if he observes ω .

Note that every mechanism π must satisfy

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

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

4.2. Implementable utility profiles. We say that receiver’s utility profile U is implementable if there exists a private persuasion mechanism π such that $U(r) = U_\pi(r)$ for all $r \in R$. We say that U is implementable by public signal if there exists a public persuasion mechanism σ such that the receiver’s utility in this mechanism U_σ satisfies $U(r) = U_\sigma(r)$ for all $r \in R$.




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

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

Theorem 1. *The following statements are equivalent:*

- (a) U is a convex function bounded by \underline{U} and \bar{U} ;
- (b) U is implementable;
- (c) U is implementable by public signal.

Proof. By (1), every U implementable by public signal is implementable. By Lemma 1 and (B), every implementable U belongs to \mathcal{U} . It remains to prove that every $U \in \mathcal{U}$ is implementable by a public signal.

Every public signal σ can be equivalently described by the c.d.f. of the probability that type r will *not act*:

$$H(r) = \int_{\Omega} \Pr \left[u(m, r) < 0 \mid m \in \text{Supp}(\sigma(\omega)) \right] dF(\omega). \quad (5)$$

In words, $H(r)$ is the ex-ante probability that the posterior expected payoff $\mathbb{E}[u(\omega, r)|\cdot]$ conditional on σ will be negative. Since $u(\omega, r) = \omega - r$, $H(r)$ can be interpreted as the probability that the posterior value $\mathbb{E}[\omega|\cdot]$ is below r .

Observe that the c.d.f. of the posterior value of the full disclosure signal is given by $F(r)$.

Let $U \in \mathcal{U}$. Define $H(r) = 1 + U'(r)$, where U' denotes the right-derivative. Observe that $\underline{U}'(0) = \bar{U}'(0) = -1$ and $\underline{U}'(1) = \bar{U}'(1) = 0$, hence $H(0) = 0$ and $H(1) = 1$. Also, H is increasing, since U is convex. Hence, H is a c.d.f. Next, observe that

$$\int_r^1 (1 - H(s))ds = U(r) \leq \bar{U}(r) = \int_r^1 (1 - F(s))ds. \quad (6)$$

Furthermore, $\underline{U}(0) = \bar{U}(0) = \mathbb{E}[\omega]$ implies that (6) is satisfied with equality for $r = 0$. Thus, F is a mean-preserving spread of H . It follows that signal σ that induces utility profile U can be obtained by an appropriate garbling of full disclosure signal $\bar{\sigma}$ (Blackwell 1953). \square

Remark 1. Theorem 1 indirectly characterizes the set of implementable triples (q, U, V) , where $q : R \rightarrow [0, 1]$ is the receiver's action profile and $V : R \rightarrow \mathbb{R}$ is the sender's utility profile. Indeed, profile U uniquely pins down profiles q and V :

$$q(r) = -U'(r), \quad (7)$$

$$V(r) = q(r) + \rho(r)U(r), \quad (8)$$

where the first equality holds by Lemma 1 and the second by definitions of q , U , and V . In particular, action profile q is implementable by a persuasion mechanism if and only if q is non-increasing and satisfies

$$\int_r^1 q(s)ds \leq \bar{U}(r) \quad \text{for all } r \in R. \quad (9)$$

Remark 2. For Theorem 1 to hold, it suffices to assume that the receiver's payoff $u(\omega, r)$ is $b(r) + c(r)d(\omega)$ for some functions b , c , and d , with c positive. In addition, the receiver's utility profile U pins down the sender's utility profile V if $v(r, \omega) = 1 + \rho(r)u(\omega, r)$ for some function ρ .¹⁰ Otherwise, (8) does not hold. Hence, a public persuasion mechanism and a private persuasion mechanism that implement the same utility profile U for the receiver may yield different payoffs to the sender. As a consequence, optimal persuasion mechanisms may be outside of the set of public persuasion mechanisms.

¹⁰The receiver's incentive compatibility constraints do not change if his payoff function is divided by $c(r)$. After this normalization, without loss of generality, we can rename the types, $\tilde{r} = -b(r)/c(r)$, and the states, $\tilde{\omega} = d(\omega)$. The sender's payoff can also be expressed as $v(\tilde{\omega}, \tilde{r}) = 1 + \tilde{\rho}(\tilde{r})(\tilde{\omega} - \tilde{r})$ where $\tilde{\rho}(\tilde{r}) = \rho(r)c(r)$.

5. OPTIMAL MECHANISMS

Denote by $V_\pi(r)$ the sender's expected payoff under mechanism π when the receiver's type is r :

$$V_\pi(r) = \int_{\Omega} v(\omega, r) \pi(\omega, r) dF(\omega).$$

The sender seeks a persuasion mechanism π that maximises $\int_R V_\pi(r) dG(r)$. The following lemma is a useful tool for finding optimal persuasion mechanisms. It expresses the sender's expected payoff as a function of the receiver's utility profile.

Lemma 2. *For every mechanism π ,*

$$\int_R V_\pi(r) dG(r) = C + \int_R U_\pi(r) I(r) dr, \quad (10)$$

where

$$I(r) = g'(r) + \rho(r)g(r) \quad \text{and} \quad C = g(0)\mathbb{E}[\omega].$$

This result follows from the observation that V_π is a function of U_π for every π , as we note in Remark 1. Integration by parts yields (10). The proof is in the Appendix.

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

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

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

5.1. Upper- and lower-censorship. A public signal is an *upper-censorship* if there exists a cutoff $\omega^* \in \Omega$ such that the signal truthfully reveals the state whenever it is below the cutoff and does not reveal the state whenever it is above the cutoff, so

$$\sigma(\omega) = \begin{cases} \omega, & \text{if } \omega \leq \omega^*, \\ \mathbb{E}[\omega | \omega > \omega^*], & \text{if } \omega > \omega^*. \end{cases}$$

A *lower-censorship* signal is defined symmetrically, $\sigma(\omega) = \omega$ whenever $\omega > \omega^*$ and $\sigma(\omega) = \mathbb{E}[\omega | \omega \leq \omega^*]$ whenever $\omega \leq \omega^*$ for some $\omega^* \in \Omega$.

The utility profile under an upper-censorship mechanism (which pools all states on interval $[\omega^*, 1]$) is shown as a black kinked curve on Fig. 2. Since $\omega < \omega^*$ is revealed to the receiver, if receiver's type r is below ω^* , he is informed whether ω exceeds r or not and hence receives the highest feasible utility $\bar{U}(r)$ by acting iff $\omega > r$. On the other hand, a receiver with type r above ω^* is essentially informed whether ω exceeds ω^* or not, hence he acts iff $\mathbb{E}[\omega | \omega > \omega^*] - r > 0$. His utility is thus either $\mathbb{E}[\omega | \omega > \omega^*] - r$ or zero, whichever is greater, corresponding to the two straight segments of the utility curve.

The next result provides conditions for optimality of upper- (lower-) censorship.

Lemma 3. *Upper-censorship (lower-censorship) is optimal for all F iff I crosses the horizontal axis at most once and from above (from below).*




FIGURE 2. Utility under upper-censorship mechanism.




FIGURE 3. Construction of upper-censorship mechanism under quasiconcave I .

Proof. Consider a function I that crosses the horizontal axis once and from above at some r^* , so $I(r)$ is positive for $r < r^*$ and negative for $r > r^*$. By (11) the sender wishes to design a mechanism that maximizes $U(r)$ for all $r < r^*$ and minimizes it for

all $r > r^*$, subject to the incentive and feasibility constraints. To prove that upper-censorship is optimal, given U , we now construct an upper-censorship function that has exactly the same utility at r^* , while higher utility for all $r < r^*$ and lower utility for all $r > r^*$, thus being preferable for the sender. The construction is illustrated by Fig. 3. Function U is depicted by the solid dark red curve; the constructed upper-censorship function is depicted by the kinked black curve. This function coincides with $\bar{U}(r)$ on interval $[0, \omega^*]$ and then connects points $(\omega^*, \bar{U}(\omega^*))$ and $(r^*, U(r^*))$ by a straight line, where ω^* is the tangency point. This straight line continues until it hits zero utility and forms a kink at zero. This utility function is induced by an upper-censorship mechanism with cutoff point ω^* , and it is superior to U by (11). The argument for lower-censorship is symmetric.

Conversely, suppose that I does not cross the horizontal axis at most once and from above. Then there exist $r_1 < r_2 < r_3$ such that I is negative on (r_1, r_2) and positive on (r_2, r_3) , because I is continuous by assumption. Therefore, lower-censorship is optimal for any F that has support only on $[r_1, r_3]$. In fact, by (11), every upper-censorship mechanism is strictly inferior. The argument for I that does not cross the horizontal axis at most once and from below is symmetric. \square

Full disclosure of ω and no disclosure of any information about ω are two natural policies that can be available to the sender even if she cannot use the entire set of persuasion mechanisms. The privacy of the receiver's information can protect the receiver and force the sender to fully disclose her information. However, disclosing no information about the state can also be optimal. From (11), full disclosure is optimal among all persuasion mechanisms if $I \geq 0$ and no disclosure is optimal if $I \leq 0$. By the same argument as in the proof of Lemma 3, the converse is also true.

Corollary 1. *Full disclosure (no disclosure) is optimal for all F iff I is positive (negative).*

Optimality of full disclosure or no disclosure in Corollary 1 is determined only by the distribution of type of the receiver. In particular, full disclosure can be optimal even when the sender is indifferent about the payoff of the receiver ($\rho(r) = 0$, so $v(\omega, r) = 1$), and the ex-ante expected state $\mathbb{E}[\omega]$ is so high that no disclosure induces action 1 with probability close to 1.

Under the assumption of type-independent weight ρ we can obtain a stronger result for optimality of upper- (lower-) censorship:

Theorem 2. *Let $\rho(r) = \rho \in \mathbb{R}$. Upper censorship (lower censorship) is optimal for all F and all ρ iff density g of the receiver's type is logconcave (logconvex).*

Proof. By Lemma 3, upper censorship (lower censorship) is optimal iff function $I = g'(r) + \rho g(r)$ crosses the horizontal axis at most once and from above (from below). This holds for all $\rho \in \mathbb{R}$ iff $\frac{g'(r)}{g(r)}$ is nonincreasing (nondecreasing) by Proposition 1 of Quah and Strulovici (2012). \square

This result is useful, because many commonly-used density functions are logconcave or logconvex (Tables 1 and 3 of Bagnoli and Bergstrom, 2005).

Under the assumption of type-independent weight ρ , we can also obtain a stronger result for optimality of full disclosure and no disclosure. We say that G is a *maximum entropy distribution with parameter* $\lambda \in \mathbb{R}$ if

$$g(r) = ce^{\lambda r}, \quad r \in [0, 1],$$

where $c > 0$ is a constant such that $\int_0^1 g(r)dr = 1$.¹¹ Note that the uniform distribution is the maximum entropy distribution with $\lambda = 0$.

Corollary 2. *Let $\rho(r) = \rho \in \mathbb{R}$. Either full disclosure or no disclosure is optimal for all F and all ρ iff G is the maximum entropy distribution function.*

Proof. The only two mechanisms that are both upper-censorship and lower-censorship are full disclosure and no disclosure. By Theorem 2, both upper-censorship and lower-censorship are optimal for all F and all ρ iff g is both logconcave and logconvex. That is, $\frac{g'(r)}{g(r)}$ is both non-increasing and non-decreasing, which holds iff $\frac{g'(r)}{g(r)}$ is equal to some constant λ . \square

The class of maximum entropy distributions is a very special class under which the optimal choice of the sender polarizes, in the sense that either full disclosure or no disclosure is optimal. So, if $g(r) = ce^{\lambda r}$ for some $\lambda \in \mathbb{R}$, then full disclosure is optimal whenever $\rho \geq -\lambda$ and no disclosure is optimal whenever $\rho \leq -\lambda$ (and anything is optimal when $\rho = -\lambda$).

As an illustration, let $\rho = 0$ and let G be the maximum entropy distribution with mean r_0 . Then, no disclosure is optimal if and only if $r_0 \leq \frac{1}{2}$ and full disclosure is optimal if and only if $r_0 \geq \frac{1}{2}$. Indeed, with $\rho = 0$, the optimal persuasion mechanism is fully determined by the sign of λ , which is in turn determined by whether mean r_0 is greater or smaller than $\frac{1}{2}$.

5.2. Comparative statics. Theorem 2 allows for sharp comparative statics analysis on the amount of information that is optimally disclosed by the sender. Specifically, if g is logconcave and if a change in primitives increases optimal cutoff ω^* , then the optimal mechanism discloses more information. We now show that the sender optimally discloses more information when she is less biased relative to the receiver, or when receivers are more reluctant to act.

Let $\rho(r) = \rho \in \mathbb{R}$. Recall that ρ is the alignment parameter between the sender's and the receiver's preferences: a greater ρ means a smaller sender's bias.

Consider a family of c.d.f.s G_t of the receiver's type

$$G_t(r) \equiv G(r - t),$$

where $t \in \mathbb{R}$ is a parameter and G admits a strictly positive, differential, and logconcave density g on \mathbb{R} with $g' + \rho g$ being almost everywhere nonzero. Parameter t is a parallel shift of the distribution of types. A greater t means a more "conservative" (reluctant to act) distribution, since every type of the receiver has a greater cost of action.

¹¹That is, G is the maximum entropy distribution on the class of probability distributions with support on $[0, 1]$ and a given mean.

Since G_t is logconcave for every t , upper censorship is optimal by Theorem 2. Let $\omega^*(\rho, t) \in \Omega$ be the optimal upper-censorship cutoff.

Proposition 1. *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 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. To get the intuition for part (b), consider the case of $\rho = 0$. When t increases, the peak of the density function g shifts to the right. For a large enough t density g is increasing on $[0, 1]$, so full disclosure is optimal. Symmetrically, when t is sufficiently small, density g is decreasing on $[0, 1]$, so no disclosure is optimal.

6. APPLICATIONS

6.1. Public education campaign. The government would like to persuade the public to reduce the amount of smoking by commissioning a public education campaign about the health risks of cigarettes.

There is a competitive market for cigarettes. The supply is perfectly elastic: the market clearing price is equal to the cost parameter, $p = c$. The demand is induced by a continuum of consumers. Each consumer values cigarettes at $\omega - \varepsilon$, where ω is an uncertain common value of cigarettes and ε is an idiosyncratic preference shock; $\omega \in [0, 1]$ and $\varepsilon \in \mathbb{R}$ are distributed according to F and G , respectively, where density g is log-concave. Thus, a consumer would like to purchase cigarettes if and only if her value $\omega - \varepsilon$ exceeds price p .

Denote $r = \varepsilon + p$. Then a buyer's net payoff is $u(\omega, r) = \omega - r$ and the c.d.f. of r is $G_p(r) \equiv G(r - p)$. The total consumed quantity will depend on consumers' beliefs about the expected value of ω conditional on the information they have. If this expected value is m , then a consumer would like to purchase cigarettes if and only if $m \geq r$. The demand curve is thus equal to $G(m)$. Thus, the consumer surplus is equal to $CS(m) = \int_0^m (m - r) dG_p(r)$.

Smoking imposes a negative externality on the entire population proportional to the amount of cigarette consumption, $\alpha G_p(m)$, where $\alpha > 0$. The social planner would like to maximize the social welfare by designing a public education campaign (persuasion mechanism) that reveals information about ω and thereby alters its expected value, m . For each ω and r social planner's preferences are captured by payoff

$$v(\omega, r) = u(\omega, r) - \alpha$$

if consumer with type r buys cigarettes, and zero if he does not buy. The social welfare is equal to $V(m) = CS(m) - \alpha G_p(m)$.

As in Section 2, let π be a persuasion mechanism. The social planner maximizes

$$\begin{aligned} V_\pi(r) &= \int \int v(\omega, r) \pi(\omega, r) dF(\omega) dG_p(r) = \int (U_\pi(r) - \alpha q_\pi(r)) dG_p(r) \\ &= C + \int \hat{I}_p(r) U_\pi(r) dr \end{aligned}$$

where $\hat{I}_p(r) = g_p(r) - \alpha g'_p(r)$ and C is a constant. Our results from Section 5 apply, allowing us to characterize the optimal educational campaign. In particular, full information disclosure is optimal if g is decreasing (so a large mass of consumers has high value for cigarettes), irrespective of the magnitude of the negative externality, α . A lower-censorship that fully informs consumers only if quality ω is above a cutoff is optimal if g is log-concave (Theorem 2).

Naturally, absent negative externality ($\alpha = 0$), full disclosure is always optimal: the social planner only increases the consumer surplus by letting the consumers make informed choices. But in presence of the externality, the social planner may wish to hide information. A larger negative externality and a higher price of cigarettes lead to (weakly) less disclosure (Proposition 1). The intuition is that if the price is sufficiently large, then g is increasing (so a large mass of consumers has low value for cigarettes) and a substantial fraction of consumers would not buy cigarettes if uninformed. So, a better information is more likely to increase cigarette consumption than to decrease it. Even though there is a positive effect of information due to consumers making better informed choices, this is offset by the externality effect if α is large.

Finally, our analysis can easily handle environments with positive externalities (negative α), such as educational campaigns about benefits of vaccination. In such environments, upper-censorship that fully informs consumers only if quality ω is below a cutoff is optimal. Moreover, a larger positive externality and a lower price of vaccines lead to less disclosure.

6.2. Lobbying. We consider a modification of the lobbying example from Kamenica and Gentzkow (2011) adapted for the case of two actions of the receiver and uncertainty about the receiver's preferences. A lobbying group (sender) commissions a study with the goal of influencing a politician's decision about adoption of a policy. The politician (receiver) chooses whether to adopt the policy ($\hat{a} = 1$) or not ($\hat{a} = -1$). The social value of the policy is determined by state $\hat{\omega} \in [-1, 1]$ and politician's type $\hat{r} \in [-1, 1]$. The politician's payoff is equal to $\tilde{u}(\hat{a}, \hat{\omega}, \hat{r}) = -(a - a_p(\hat{\omega}, \hat{r}))^2$, where the politician's bliss point is given by $a_p(\hat{\omega}, \hat{r}) = \hat{\omega} - \hat{r}$. The lobbyist's payoff is equal to $\tilde{v}(\hat{a}, \hat{\omega}) = -(a - a_l(\hat{\omega}, \hat{r}))^2$, where the lobbyist's bliss point is biased towards c : $a_l(\hat{\omega}, \hat{r}) = \alpha a_p(\hat{\omega}, \hat{r}) + (1 - \alpha)c$ with $\alpha \in [0, 1]$ and $c > 0$.

Note that

$$\frac{1}{4} (\tilde{u}(1, \hat{\omega}, \hat{r}) - \tilde{u}(-1, \hat{\omega}, \hat{r})) = \frac{1}{4} ((1 + (\hat{\omega} - \hat{r}))^2 - (1 - (\hat{\omega} - \hat{r}))^2) = a_p(\hat{\omega}, \hat{r})$$

and

$$\frac{1}{4} (\tilde{v}(1, \hat{\omega}) - \tilde{v}(-1, \hat{\omega})) = a_l(\hat{\omega}, \hat{r}).$$

Making the substitution $\omega = \frac{\hat{\omega}+1}{2}$, $r = \frac{\hat{r}+1}{2}$, and $\rho = \frac{\alpha}{2(1-\alpha)c}$, and changing the notation to $a = 0$ for $\hat{a} = -1$ and $a = 1$ for $\hat{a} = 1$, we obtain

$$\begin{aligned} u(\omega, r) &= \frac{1}{2} (\tilde{u}(1, \hat{\omega}, \hat{r}) - \tilde{u}(-1, \hat{\omega}, \hat{r})) = \omega - r, \\ v(\omega, r) &= \frac{(1-\alpha)c}{2} (\tilde{v}(1, \hat{\omega}) - \tilde{v}(-1, \hat{\omega})) = 1 + \rho(\omega - r) \end{aligned}$$

with $\omega, r \in [0, 1]$ and $\rho \geq 0$. Thus, our analysis applies. In particular, if density g of \hat{r} is logconcave, then upper-censorship is optimal with more information disclosed if the lobbyist is less biased (α closer to 1 and c closer to 0).

There are two noteworthy differences with the results in Kamenica and Gentzkow (2011) for this environment. First, the optimal persuasion mechanism is very much driven by the shape of the distribution of the private information of the receiver and much less by the degree of the conflict of preferences between the parties expressed in terms of the effective bias. In Kamenica and Gentzkow (2011), the value of α determines the structure of the optimal mechanism, whereas c played no role. In our environment, both of them interact. Second, in Kamenica and Gentzkow (2011) the lobbyist either commissions a fully revealing study or no study at all. As Kamenica and Gentzkow (2011) remark: ‘‘This contrasts with the observation that industry-funded studies often seem to produce results more favorable to the industry than independent studies.’’ In our environment, more complex mechanisms can be optimal and the expected probability of approval of the policy can be higher than the ex-ante probabilities of approval under full or no information.

6.3. Media censorship. The following is an adaptation of the model of media endorsement of competing political candidates (parties) in Chiang and Knight (2011) augmented by the possibility of censorship by a partial government.¹²

There is a government, two political parties $c \in \{G, O\}$ (pro-Government and Opposition) competing in a parliamentary election, a continuum of voters and a continuum of newspapers, each of unit measure. Parties are characterized by a pair of platform and quality (i_c, ω_c) , where we assume that $i_c, \omega_c \in [0, 1]$, and without loss of generality set $i_G < i_O$. If party c wins, the voter’s payoff is equal to

$$u(c, r) = \omega_c - \frac{\beta}{2}(i_c - r)^2,$$

where $r \in [0, 1]$ is the voter’s ideal platform (ideology) and $\beta > 0$ represents the relative utility weight placed on the party’s platform. The value of r is distributed across the population of voters according to c.d.f. G . The voters know the platforms of the parties but are uncertain about their relative quality differential $\hat{\omega} = \omega_G - \omega_O$ of the candidates on the party ballots. We assume that $\hat{\omega}$ is distributed on \mathbb{R} according to c.d.f. \hat{F} . Voters support the party that maximizes their expected utility.

The newspapers receive information about the quality of the candidates and choose to endorse one of the parties. Each newspaper n observes signal $\omega_n = \hat{\omega} + \epsilon_n$, where ϵ_n is an error term with zero mean and iid distribution across newspapers. We denote the

¹²Another model with similar structure is Chan and Suen (2008).

distribution of ω_n by F_n . A newspaper endorses the pro-Government party iff $\omega_n > \gamma_n$, where γ_n describes the newspaper's editorial policy. Voters know the editorial policies and choose a single newspaper to read.

In Chiang and Knight (2011), the newspaper editorial policies are exogenous. We extend the model by allowing the government to determine the set of editorial policies available in the market by shutting down newspapers with undersirable editorial policies and subsidizing creation of additional newspapers if some desired editorial policies are lacking. Thus, we allow the government to determine the set of the editorial policies $\Gamma = \{\gamma_n | n \in N\}$. The objective of the government is to maximize the share of the votes for the pro-Government party.

The timing of the game is as follows. The government chooses the set of permitted editorial policies.¹³ The voters choose their preferred newspapers. The newspapers observe the information about the party quality and make their endorsements. The votes are cast. The solution concept is perfect Bayesian equilibrium.

Proposition 2. *Assume that the density of the distribution of the voter ideology is log-concave. Then, $\Gamma = \{\gamma | \gamma \leq \omega^*\}$ for some $\omega^* \in [0, 1]$, that is, the government censors all newspapers with the editorial policies above ω^* and ensures that all editorial policies with $\gamma \leq \omega^*$ are represented in the market.*

7. DISCUSSION

In this section we discuss side issues and extensions that connect this paper to the literature beyond the topic of Bayesian persuasion.

7.1. Rotation order. The set of upper-censorship (lower-censorship) mechanisms is totally ordered by cutoff points ω^* in Ω . The extremes, $\omega^* = 0$ and $\omega^* = 1$, correspond to the two extreme mechanisms, full disclosure and no disclosure.

Interestingly, the set of c.d.f.s of the posterior expected utility induced by upper-censorship (lower-censorship) mechanisms, as defined by (5), is also totally ordered by the *rotation order* \succeq_R (Johnson and Myatt 2006): for every two c.d.f.s H and \hat{H} one can find a rotation point r^* such that $H \succeq_R \hat{H}$ iff $H(r) \geq \hat{H}(r)$ for all $r \leq r^*$ and $H(r) \leq \hat{H}(r)$ for all $r \geq r^*$. Roughly speaking, a higher c.d.f. in the rotation order has a higher dispersion. For upper-censorship mechanisms, the rotation point r^* is $\mathbb{E}[\omega | \omega > \omega^*]$ and it is decreasing w.r.t. \succeq_R ; for lower-censorship mechanisms, the rotation point r^* is $\mathbb{E}[\omega | \omega \leq \omega^*]$ and it is increasing w.r.t. \succeq_R .

To compare to Johnson and Myatt (2006), consider the case of no uncertainty about the receiver's type (as in Kamenica and Gentzkow 2011), so all the uncertainty about whether the receiver acts or not is due to the information about ω . The set of lower-censorship mechanisms satisfies conditions of Lemma 1 of Johnson and Myatt (2006), which entails that one of the extremes (full disclosure or no disclosure) must be optimal on this set.¹⁴ For illustration, imagine that the rotation point r^* is constant,

¹³The newspapers are not strategic, they only follow their editorial policies (unless censored by the government).

¹⁴Of course, neither may be optimal on a more general domain of mechanisms.

and let the receiver's type r be below r^* . Then, by (11), depending whether $I(r) > 0$ or $I(r) < 0$, the sender wants to choose a c.d.f. H that maximizes or minimizes $H(r)$. But $H(r)$ is increasing w.r.t. \succeq_R , hence the optimum is at one of the extremes. More generally, if r^* is increasing w.r.t. \succeq_R , then $H(r)$ is quasiconvex w.r.t. \succeq_R , leading to the same conclusion. However, if there is uncertainty about r , and there are positive measures of types both to the left and to the right of the rotation point, then the effect of rotation of c.d.f. on the sender's payoff is ambiguous.

Note that the set of upper-censorship mechanisms does not satisfy the condition of Johnson and Myatt's (2006) Lemma 1 that the rotation point should be increasing w.r.t. \succeq_R . In this case, $H(r)$ is not quasiconvex w.r.t. \succeq_R , so the sender's payoff can be maximized at an interior upper-censorship mechanism, even if there is no uncertainty about the receiver's type.

7.2. Cutoff mechanisms and optimal delegation. A mechanism π is a *cutoff mechanism* if it recommends the receiver to act if and only if ω is above some cutoff $\varphi(\hat{r})$ in $[0, 1]$:

$$\pi(\omega, \hat{r}) = \begin{cases} 1, & \text{if } \omega \geq \varphi(\hat{r}), \\ 0, & \text{otherwise.} \end{cases}$$

Upper-censorship and lower-censorship mechanisms are cutoff mechanisms.

A cutoff mechanism can be equivalently described as a menu $T \subset [0, 1]$ of cutoff tests for the receiver to choose, such that $T = \{\varphi(r)\}_{r \in R}$. A *cutoff test* $t \in T$ is a binary signal that recommends the receiver to act if and only if $\omega \geq t$.

Observe that test $t = 0$ ($t = 1$) yields the uninformative recommendation to always act (never act), irrespective of the state. Since type $r = 0$ always acts and type $r = 1$ never acts, for every incentive compatible cutoff mechanism the set of cutoff tests T must contain 0 and 1. There are no further restrictions: every $T \subset [0, 1]$ such that $\{0, 1\} \subset T$ corresponds to an incentive compatible cutoff mechanism.

The problem of optimal persuasion on the class of cutoff mechanisms is equivalent to the problem of optimal delegation of choice from set $T \subset [0, 1]$, with the constraint that T must contain the endpoints, 0 and 1. In this delegation problem, the receiver must pick a cutoff test and then follow the recommendation of the test. The sender restricts the receiver's choice to a subset $T \subset [0, 1]$ such that $\{0, 1\} \subset T$ and lets the receiver choose t from T .

Restriction to the class of cutoff mechanisms is *not* without loss of generality.

Example 1. Let $\rho(r) = \bar{\rho} = 0$ (so the sender wishes to maximize the probability that the receiver acts). Let distribution F be uniform and let distribution G be bimodal as shown on Fig. 4.

Under this assumptions, G is the value function for the sender: conditional on posterior expected value of the state $\bar{\omega}$, the receiver acts iff $r \leq \bar{\omega}$, thus the sender's payoff is $G(\bar{\omega})$, the probability that the receiver acts. The sender's unconditional payoff is a convex combination of these values.

The upper bound on the sender's expected payoff, V^* , is given by the concave closure of G evaluated at $E[\omega]$ (cf. Kamenica and Gentzkow, 2011). This upper




FIGURE 4. Utility under upper-censorship mechanism.

bound can be achieved by the binary public signal that sends message $\bar{\omega} = 1/3$ if $\omega \in (1/12, 7/12)$ and $\bar{\omega} = 2/3$ otherwise. However, this upper bound cannot be achieved by any cutoff mechanism.

APPENDIX A. STRUCTURE OF OPTIMAL MECHANISMS: GENERAL CASE

In Section 5 we studied optimal persuasion mechanisms under the assumption that function I changes sign at most once. Here we describe the structure of the optimal mechanism for a more general case.

Assume that I is almost everywhere nonzero and changes sign n times.¹⁵ Let $\{r_1, r_2, \dots, r_n\}$ be the set of types at which I changes its sign.

Theorem 3. *Every optimal utility profile U is the convex envelope of the graph of \bar{U} and a set of at most $\frac{n+1}{2}$ points. Each point lies between \underline{U} and \bar{U} and corresponds to a distinct interval (r_i, r_{i+1}) on which I is negative.*

Theorem 3 implies that optimal U is piecewise linear unless $U(r) = \bar{U}(r)$, with kinks on intervals where I is negative. Fig. 5 illustrates the structure of the optimal utility profile for an example where $I(r)$ (the red curve) changes its sign three times. The optimal utility profile $U(r)$ (the black kinked curve) is the lower contour of the convex hull of points A and B and the graph of \bar{U} (the solid blue curve). Points A

¹⁵On the intervals where $I(r) = 0$ the sender is indifferent about the choice of U , hence multiple solutions emerge in this case. The characterization of the solutions is a straightforward but tedious, and thus omitted.




FIGURE 5. Optimal utility profile for the case where $I(r)$ changes sign three times.

and B are kinks of $U(r)$, and located over the intervals where $I(r)$ is negative. On the first interval of positive $I(r)$, $U(r)$ is the straight line. On the second interval of positive $I(r)$, $U(r)$ is the straight line that smoothly merges into $\bar{U}(r)$. The segment of $U(r)$ to the right of the tangency point coincides with \bar{U} .

Moreover, an optimal U is fully determined by a nonincreasing profile (u_1, \dots, u_n) , where each u_i defines the value of U at point r_i , so $u_i \in [\underline{U}(r_i), \bar{U}(r_i)]$ for all $i = 1, \dots, n$. This follows from the two properties below.

Clearly, as follows from (11), on any interval (r_i, r_{i+1}) where $I(r)$ is positive, the optimality requires that $U(r)$ is pointwise maximized subject to feasibility ($U \leq \bar{U}$) and convexity of U . That is, for any given values of U at boundary points r_i and r_{i+1} , the utility profile U on the interior of (r_i, r_{i+1}) is a straight line unless $U(r) = \bar{U}(r)$, as illustrated by Fig. 6 and Fig. 7. Formally:

- (P₁) On every interval (r_i, r_{i+1}) on which I 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, the optimality requires that $U(r)$ is pointwise minimized subject to $U \geq \underline{U}$ and convexity of U . That is, on the interior of (r_i, r_{i+1}) the utility profile U is the minimum convex function that passes through endpoints $U(r_i)$ and $U(r_{i+1})$. It is an upper envelope of two straight lines




FIGURE 6. Optimal utility profile on the interval of where $I(r)$ is positive.

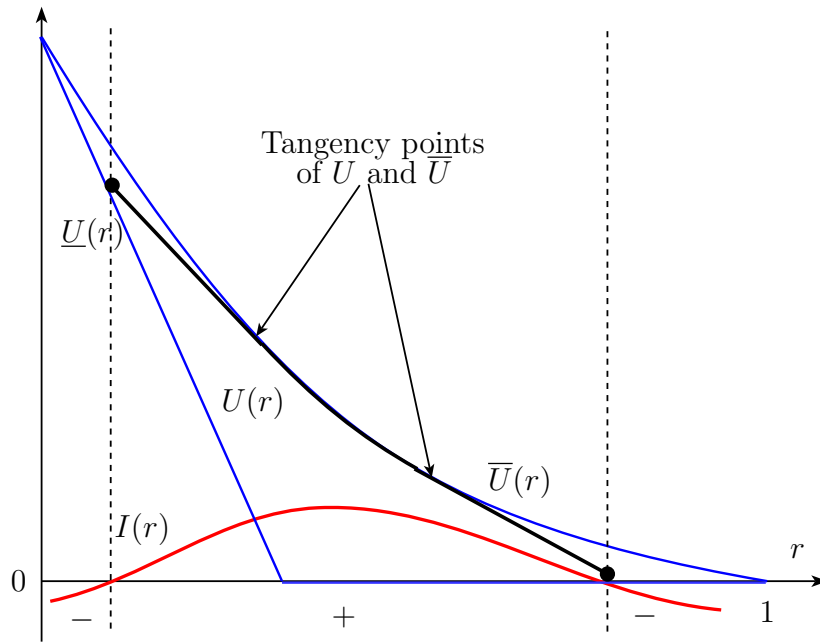


FIGURE 7. Optimal utility profile on the interval of where $I(r)$ is positive.

whose slopes are the derivatives at the endpoints, $U'(r_i)$ and $U''(r_{i+1})$, as illustrated by Fig. 8. Formally:




FIGURE 8. Optimal utility profile on the interval of where $I(r)$ is negative.

(P₂) On every interval (r_i, r_{i+1}) on which I is negative, U is piecewise linear with at most one kink, and thus must satisfy

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

Consider a nonincreasing sequence (u_1, \dots, u_n) that satisfies $u_i \in [\underline{U}(r_i), \bar{U}(r_i)]$ for all $i = 1, \dots, n$. Let u_i determine the value of U at point r_i , $U(r_i) = u_i$. Observe that the interior of every interval (r_i, r_{i+1}) is uniquely determined by properties (P₁) and (P₂). Thus, optimal U can be fully described by the profile (u_1, \dots, u_n) . Finding an optimal U is thus an n -variable optimization problem.

Proof of Theorem 3. The proof is immediate by necessary conditions (P₁) and (P₂) that every optimal U must satisfy. \square

APPENDIX B. NONLINEAR UTILITY OF THE RECEIVER

In this section we discuss the complications of handling the case of nonlinear utility of the receiver. We relax the linearity assumption and only assume that the receiver's payoff from acting, $u(\omega, r)$, is continuous and strictly monotonic (w.l.o.g., increasing in ω and decreasing in r), and satisfies normalization $u(r, r) = 0$ for all r .

Note that if types ω and r are correlated, the analysis below carries over if we impose strict monotonicity on function $\tilde{u}(\omega, r) = u(\omega, r)g(r|\omega)/g(r)$ rather than on u , because the receiver's interim expected payoff under mechanism π can be written as $U(r) = \int_{\Omega} \tilde{u}(\omega, r) \pi(\omega, r) dF(\omega)$.

B.1. Monotonicity. Monotonicity of u implies the same ordinal preferences over the states by all types of the receiver. Without this assumption, optimal persuasion mechanisms need not be public, as the following example demonstrates.

Example 2. Let $\omega = (\omega_1, \omega_2) \in \{0, 1\} \times \{0, 1\}$ and $r \in \{r_1, r_2, r_3\}$. Let all pairs (ω, r) be equally likely. Define $u(\omega, r_1) = \omega_1 - 1$, $u(\omega, r_2) = \omega_2 - 1$, and $u(\omega, r_3) = 3/4 - \omega_1\omega_2$. Suppose that the sender wishes to maximize the probability that the receiver acts.

Consider the following mechanism:

$$\pi(\omega, r) = \begin{cases} 1 & \text{if } r = r_1 \text{ and } \omega_1 = 1, \text{ or if } r = r_2 \text{ and } \omega_2 = 1, \text{ or if } r = r_3, \\ 0 & \text{otherwise.} \end{cases}$$

Essentially, this mechanism allows the receiver to learn at most one component of ω . Clearly, under this mechanism, it is incentive compatible for the receiver to truthfully report r and follow the recommendation. Thus, the probability that the receiver acts ($a = 1$) under π is:

$$\Pr_{\pi}(a = 1) = \Pr(r = r_1) \Pr(\omega_1 = 1) + \Pr(r = r_2) \Pr(\omega_2 = 1) + \Pr(r = r_3) = \frac{2}{3}.$$

However, the sender cannot induce the receiver to choose $a = 1$ with probability $2/3$ using a public signal. To see this, note first that the receiver r_1 chooses $a = 1$ only if he is certain that $\omega_1 = 1$, and the receiver r_2 chooses $a = 1$ only if he is certain that $\omega_2 = 1$. Thus, under any public signal, the probability that the receiver acts cannot exceed $2/3$. The only possibility of how the sender could achieve this probability would be to reveal both ω_1 and ω_2 , but in that case the receiver r_3 would not act when $(\omega_1, \omega_2) = (1, 1)$.

B.2. Additive-multiplicative separability. Recall that Theorem 1 only proves equivalence of implementation of the receiver's utility profile under private and public persuasion. To establish equivalence of the sender's payoff, we used two assumptions, that the sender's utility is a function of the receiver's utility and the receiver's action, and that the receiver's action profile is pinned down by her utility profile (Remark 1). The latter property is a consequence of linear u . More generally, this property obtains if and only if u is additively-multiplicatively separable in the sense of Remark 2 (see Proposition 3 below). In other words, when u is not additively-multiplicatively separable, it may be the case that two persuasion mechanisms implement the same utility profile for the receiver, and yet the sender's expected payoff is different.

Proposition 3. *Let π_1 and π_2 be two mechanisms that are distinct for every r but implement the same receiver's utility profile U . The probability of $a = 1$ is the same for π_1 and π_2 if and only if there exist functions b , c , and d such that $u(\omega, r) = b(r) + c(r) d(\omega)$ for every (ω, r) .*

Proof. Consider two mechanisms π_1 and π_2 that implement the same utility profile U for the receiver. We have for each $i = 1, 2$ and for all r

$$\begin{aligned} \int u(\omega, r) \pi_i(\omega, r) dF(\omega) &= U(r), \\ \int \frac{\partial u(\omega, r)}{\partial r} \pi_i(\omega, r) dF(\omega) &= U'(r), \end{aligned}$$

where the first line holds by definition of U and the second line is the local incentive compatibility condition.

The expected action, $q_{\pi_i}(r) = \int \pi_i(\omega, r) dF(\omega)$ is the same across $i = 1, 2$ for all r if and only if the vectors $u(\omega, r)$, $\frac{\partial u(\omega, r)}{\partial r}$, and $\mathbf{1}$ are linearly dependent for all r :

$$\gamma(r) \frac{\partial u(\omega, r)}{\partial r} + \mu(r) u(\omega, r) = \mathbf{1} \text{ for some } \mu \text{ and } \gamma. \quad (12)$$

Note that $\gamma(r) \neq 0$, because otherwise $u(\omega, r)$ would not depend on ω . For every ω , the solution of differential equation (12) is given by

$$u(\omega, r) = e^{-\int \frac{\mu(r)}{\gamma(r)} dr} \left(Const(\omega) + \int \frac{1}{\gamma(r)} e^{\int \frac{\mu(s)}{\gamma(s)} ds} dr \right),$$

which completes the proof with $b(r)$, $c(r)$, and $d(\omega)$ defined by

$$(b(r), c(r), d(\omega)) = \left(e^{-\int \frac{\mu(r)}{\gamma(r)} dr} \int \frac{1}{\gamma(r)} e^{\int \frac{\mu(s)}{\gamma(s)} ds} dr, e^{-\int \frac{\mu(r)}{\gamma(r)} dr}, Const(\omega) \right).$$

□

B.3. Ex-post equivalence. Theorem 1 establishes equivalence of private and public persuasion mechanisms in terms of the interim expected payoff of the receiver. This definition is in line with recent papers on standard mechanism design with transfers (Manelli and Vincent, 2010 and Gerskov et al., 2013).

A stronger notion of equivalence requires private and public persuasion mechanisms implement the same ex-post outcomes (actions by the receiver). Notice that equivalence in terms of ex-post outcomes implies equivalence in terms of interim expected payoffs, without any assumption (such as linearity) on the sender's and receiver's payoff functions. Under monotonic u , it is easy to verify whether a given private mechanism π is implementable by public signal:

Proposition 4. *An incentive compatible persuasion mechanism π is ex-post equivalent to a public signal if and only if $\pi(\omega, r)$ is non-increasing in r for every ω .*

Proof. Consider an arbitrary public signal σ . By monotonicity of u , for every message m if type r acts, then every type $r' < r$ will act as well. Consequently, we can focus on direct signals where every message m is equal to the type who is indifferent to act or not after receiving this message, $\mathbb{E}_\omega[u(\omega, m)|m] = 0$. Thus, the mechanism¹⁶

$$\pi(\omega, r) = \sum_m \Pr[m \geq r] \sigma_m(\omega).$$

is non-increasing in r for every ω .

¹⁶We write $\sigma_m(\omega)$ for the probability of message m conditional on state ω .

Consider a mechanism π that is non-increasing in r . Thus, the probability $q_\pi(r) = \int_0^1 \pi(\omega, r) dF(\omega)$ that the receiver of type r will act is non-increasing in r . The c.d.f. $H(r) = 1 - q_\pi(r)$ defines a public signal as described in the proof of Theorem 1. \square

Observe that when u is nonlinear, an incentive compatible mechanism π need not have a monotonic q_π , and hence it may be impossible to construct an equivalent public signal.

B.4. Binary state. If the state is binary, then there is ex-post equivalence between private and public mechanisms under the assumption of monotonic u .¹⁷

Proposition 5. *Let the support of F consist of $\{0, 1\}$. Then every incentive compatible mechanism π is ex-post equivalent to a public signal.*

Proof. It is sufficient to show that π is non-increasing in r for $r \in (0, 1)$. The receiver r prefers to report $n = r$ rather than r' only if:

$$\sum_{\omega=0,1} u(\omega, r) (\pi(\omega, r) - \pi(\omega, \hat{r})) \Pr(\omega) \geq 0. \quad (13)$$

Writing (13) for $(r, \hat{r}) = (r_2, r_1)$ and $(r, \hat{r}) = (r_1, r_2)$ yields:

$$-\frac{u(0, r_2)}{u(1, r_2)} \delta(r_2, r_1, 0) \leq \delta(r_2, r_1, 1) \leq -\frac{u(0, r_1)}{u(1, r_1)} \delta(r_2, r_1, 0) \quad (14)$$

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

$$0 < -\frac{u(0, r_2)}{u(1, r_2)} \leq -\frac{u(0, r_1)}{u(1, r_1)} \text{ for } r_2 \leq r_1. \quad (15)$$

Combining (14) and (15) gives $\pi(\omega, r_2) \geq \pi(\omega, r_1)$ for $\omega = 0, 1$, and $r_2 \leq r_1$. \square

APPENDIX C. OMITTED PROOFS

C.1. Proof of Lemma 1. Incentive compatibility (IC) implies truth-telling,

$$U_\pi(r) \geq U_\pi(r, \hat{r}) \text{ for all } r, \hat{r} \in R. \quad (16)$$

Also, observe that type $r = 0$ can secure the maximal attainable payoff $\mathbb{E}[\omega]$ by always acting (irrespective of recommendation). On the other hand, the maximal attainable payoff of type $r = 1$ is zero, which can be secured by never acting. Together with (IC), this implies

$$U_\pi(0) = \mathbb{E}[\omega] \text{ and } U_\pi(1) = 0. \quad (17)$$

By the standard envelope argument, (16)–(17) is equivalent to (2)–(4).

It remains to show that (16)–(17) imply (IC). By contradiction, suppose that (16)–(17) hold and there exist r and \hat{r} such that

$$U_\pi(r) < \max_{a_0, a_1 \in \{0,1\}} \{U_\pi(r, \hat{r}, a_0, a_1)\}.$$

¹⁷In fact, this result is tight: if the state has at least three values, then there is no ex-post equivalence between private and public mechanisms, even if the receiver's utility is linear.

By (16) we have $U_\pi(r, \hat{r}, 0, 1) \leq U_\pi(r)$. Also, it can be easily verified that $U_\pi(r, \hat{r}, 0, 1) \leq U_\pi(r, \hat{r}, 1, 0)$ implies

$$U_\pi(r, \hat{r}, 1, 0) \leq \max \{U_\pi(r, \hat{r}, 0, 0), U_\pi(r, \hat{r}, 1, 1)\}.$$

Thus we are left with

$$U_\pi(r) < \max \{U_\pi(r, \hat{r}, 0, 0), U_\pi(r, \hat{r}, 1, 1)\}.$$

Note that $a_0 = a_1 = 1$ ($a_0 = a_1 = 0$) means that the receiver chooses to act (not to act) regardless of the recommendation, in which case his expected payoff is $\mathbb{E}[\omega]$ (zero). By (17) this is also the payoff of the obedient receiver who reports $\hat{r} = 0$ ($\hat{r} = 1$, respectively). Hence

$$U_\pi(r, \hat{r}, 0, 0) = U_\pi(r, 0) = \mathbb{E}[u(\omega)] \quad \text{and} \quad U_\pi(r, \hat{r}, 1, 1) = U_\pi(r, 1) = 0.$$

But (16) implies $U_\pi(r) \geq \max \{U_\pi(r, 0), U_\pi(r, 1)\}$, a contradiction.

C.2. Proof of Lemma 2. We have

$$V_\pi(r) = \int_{\Omega} \left(1 + \rho(r)u(\omega, r)\right) \pi(\omega, r) dF(\omega) = q_\pi(r) + \rho(r)U_\pi(r).$$

Since $\int_r^1 q_\pi(s) ds = U_\pi(r)$ by Lemma 1, integrating by parts yields

$$\int_R U_\pi(r) dG(r) = g(0)U_\pi(0) + \int_R g'(r)U_\pi(r) dr.$$

Hence

$$\int_R V_\pi(r) dG(r) = U_\pi(0)g(0) + \int_R U_\pi(r) \left((g'(r) + \rho(r)g(r)) \right) dr.$$

Substituting $C = U_\pi(0)g(0) = \mathbb{E}[\omega]g(0)$ and $I(r) = g'(r) + \rho(r)g(r)$ yields (10).

C.3. Proof of Proposition 1. Types $r < 0$ always act and types $r > 1$ never act; so we omit these types from the analysis. The sender's expected payoff under public mechanism σ can be written as:

$$V_\sigma = C + \int_0^1 J_t(r) dH_\sigma(r), \quad (18)$$

where C is a constant that does not depend on σ , H_σ is c.d.f. 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 $\omega^*(\rho, t) \in (0, 1)$. Denote $\omega^{**} = \mathbb{E}[\omega|\omega \geq \omega^*]$. The derivative of the sender's expected payoff (18) under upper censorship with respect to cutoff 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).

Consider t at which $\omega^*(\rho, t) \in (0, 1)$. Notice that

$$J'_t(r) = g(r-t) + \rho G(r-t) = J'(r-t);$$

so

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

At optimal interior cutoff, we have $\frac{dV}{d\omega^*} = 0$; so

$$J'_t(\omega^{**}) (\omega^{**} - \omega^*) = \int_{\omega^*}^{\omega^{**}} J'_t(s) ds. \quad (19)$$

Since $J''_t(r)$ is almost everywhere nonzero, J'_t is not constant on (ω^*, ω^{**}) . Moreover, by Theorem 2, $I_t(r) = J''_t(r)$ crosses the horizontal axis at most once and from above; so $J'_t(r)$ is quasiconcave. Therefore, (19) implies that $J''_t(\omega^{**}) < 0$ and $J'_t(\omega^{**}) > J'_t(\omega^*)$; so $\frac{d^2V}{dt d\omega^*} > 0$ and ω^* is strictly increasing in t by Theorem 1 of Edlin and Shannon (1998).

C.4. Proof of Proposition 2. Given a newspaper's signal about quality differential $\omega_n = \hat{\omega} + \epsilon_n$, the voter's expected net payoff from choosing the pro-Government party versus the Opposition party is equal to

$$\begin{aligned} \tilde{u}(\hat{\omega}, r) &= \omega_n - \frac{\beta}{2}(i_G - r)^2 + \frac{\beta}{2}(i_O - r)^2 \\ &= \omega_n - \beta(i_O - i_G)r + \frac{\beta}{2}(i_O^2 - i_G^2) \end{aligned}$$

Define $\omega = \frac{\omega_n + \frac{\beta}{2}(i_O^2 - i_G^2)}{\beta(i_O - i_G)}$ and let

$$u(\omega, r) = \frac{1}{\beta(i_O - i_G)} \hat{u}(\omega_n, r) = \omega - r.$$

Let F be the distribution of ω . Let $\gamma(r)$ be the choice of the editorial policy by voter r and let $\hat{\pi}(\omega_n, r)$ be the expected probability of the vote for the pro-Government party if the newspaper to which voter r subscribes observes signal ω_n . By independence of ϵ_n

$$V = \int_R \int_{\epsilon} \int_{\hat{\omega}} \hat{\pi}(\hat{\omega} + \epsilon, r) dH(\epsilon) d\hat{F}(\omega) dG(r) = \int_R \int_{\omega} \hat{\pi}(\omega, r) dF(\omega) dG(r)$$

where H is the c.d.f of ϵ . Now, consider the auxiliary problem in which the government maximizes V over the set of persuasion mechanisms. Then, by Proposition 2, the optimal persuasion policy is upper-censorship. We now show that $\Gamma = [0, \omega^*]$ implements the outcome of the optimal persuasion policy, which establishes the claim of the proposition.

Observe that the voter's payoff from choosing the newspaper with editorial policy γ net of his full information payoff $\bar{U}(r)$ is equal to

$$v(\gamma, r) = \int_{\gamma}^r (\omega - r) dF(\omega)$$

If $\Gamma = [0, \omega^*]$, this payoff is maximized by choosing the newspaper with the editorial policy $\gamma = r$ if $r \leq \omega^*$ and $\gamma = \omega^*$ otherwise. It follows that the voter will choose to vote for the pro-Government party iff $\omega \geq r$ for $r \leq \omega^*$, iff $\omega > \omega^*$ if $\mathbb{E}[\omega | \omega > \omega^*] \geq r$ and $r > \omega^*$, and never otherwise. Observe that the same behavior is induced by the optimal persuasion policy of upper-censorship at ω^* .

REFERENCES

- ALONSO, R., AND O. CÂMARA (2014a): "On the value of persuasion by experts," Mimeo.
 ——— (2014b): "Persuading skeptics and reaffirming believers," Mimeo.
 ——— (2014c): "Persuading voters," Mimeo.
- BAGNOLI, M., AND T. BERGSTROM (2005): "Log-concave probability and its applications," *Economic Theory*, 26(2), 445–469.
- BERGEMANN, D., AND A. BONATTI (2013): "Selling Cookies," Cowles Foundation Discussion Papers 1920, Cowles Foundation for Research in Economics, Yale University.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2014): "Selling Experiments: Menu Pricing of Information," *working paper*.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2012): "Extremal Information Structures in First Price Auctions," Cowles Foundation Discussion Papers 1299, Cowles Foundation for Research in Economics, Yale University.
 ——— (2013): "The Limits of Price Discrimination," Cowles Foundation Discussion Papers 1896, Cowles Foundation for Research in Economics, Yale University.
- BERGEMANN, D., AND S. MORRIS (2013a): "Bayes Correlated Equilibrium and Comparison of Information Structures," *mimeo*.
 ——— (2013b): "Robust Predictions in Games With Incomplete Information," *Econometrica*, 81(4), 1251–1308.
- BERGEMANN, D., AND M. PESENDORFER (2007): "Information structures in optimal auctions," *Journal of Economic Theory*, 137, 580–609.
- BERGEMANN, D., AND A. WAMBACH (2013): "Sequential Information Disclosure in Auctions," Cowles Foundation Discussion Papers 1900, Cowles Foundation for Research in Economics, Yale University.
- BLACKWELL, D. (1953): "Equivalent comparisons of experiments," *Annals of Mathematical Statistics*, 24, 265–272.
- BOLES LAVSKY, R., AND C. COTTON (2015a): "Grading Standards and Education Quality," *American Economic Journal: Microeconomics*, forthcoming.
 ——— (2015b): "Information and Extremism in Elections," *American Economic Journal: Microeconomics*, forthcoming.
- CHAN, J., AND W. SUEN (2008): "A Spatial Theory of News Consumption and Electoral Competition," *Review of Economic Studies*, 75(3), 699–728.
- CHE, Y.-K., W. DESSEIN, AND N. KARTIK (2013): "Pandering to Persuade," *American Economic Review*, 103(1), 47–79.
- CHIANG, C.-F., AND B. KNIGHT (2011): "Media Bias and Influence: Evidence from Newspaper Endorsements," *Review of Economic Studies*, 78(3), 795–820.
- EDLIN, A. S., AND C. SHANNON (1998): "Strict Monotonicity in Comparative Statics," *Journal of Economic Theory*, 81, 201–219.

- ELY, J. (2014): “Beeps,” *mimeo*.
- ELY, J., A. FRANKEL, AND E. KAMENICA (2015): “Suspense and Surprise,” *Journal of Political Economy*, forthcoming.
- ESO, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705–731.
- GENTZKOW, M., AND E. KAMENICA (2012): “Competition in Persuasion,” Discussion paper, University of Chicago.
- (2014a): “Costly Persuasion,” *American Economic Review*, 104(5), 457–62.
- (2014b): “Disclosure of Endogenous Information,” Discussion paper, University of Chicago.
- HAGENBACH, J., F. KOESSLER, AND E. PEREZ-RICHET (2014): “Certifiable Pre-Play Communication: Full Disclosure,” *Econometrica*, 83(3), 1093–1131.
- HAO, L., AND X. SHI (2015): “Discriminatory Information Disclosure,” *mimeo*.
- HOFFMANN, F., R. INDERST, AND M. OTTAVIANI (2014): “Persuasion through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation,” *Mimeo*.
- HORNER, J., AND A. SKRZYPACZ (2009): “Selling Information,” Cowles Foundation Discussion Papers 1743, Cowles Foundation for Research in Economics, Yale University.
- JOHNSON, J. P., AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96(3), 756–784.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615.
- KOESSLER, F., AND R. RENAULT (2012): “When does a firm disclose product information?,” *RAND Journal of Economics*, 43(4), 630–649.
- KOLOTILIN, A. (2014a): “Experimental Design to Persuade,” *mimeo*.
- (2014b): “Optimal Information Disclosure: Quantity vs. Quality,” *mimeo*.
- KRÄHMER, D., AND T. MYLOVANOV (2014): “Selling Information,” *working paper*.
- KREMER, I., Y. MANSOUR, AND M. PERRY (2014): “Implementing the ?Wisdom of the Crowd?,” *Journal of Political Economy*, 122(5), 988 – 1012.
- LIZZERI, A. (1999): “Information Revelation and Certification Intermediaries,” *RAND Journal of Economics*, 30(2), 214–231.
- MIRRLEES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175–208.
- PEREZ-RICHET, E. (2014): “Interim Bayesian Persuasion: First Steps,” *American Economic Review*, 104(5), 469–74.
- PEREZ-RICHET, E., AND D. PRADY (2012): “Complicating to Persuade?,” *Working papers*.
- RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118(5), 949 – 987.
- TANEVA, I. (2014): “Information design,” *mimeo*.
- WANG, Y. (2013): “Bayesian Persuasion with Multiple Receivers,” *Mimeo*.