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 characterize 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 it 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

---

Kolotilin: School of Economics, UNSW Australia, Sydney, NSW 2052, Australia. E-mail: akolotilin@gmail.com.
Li: Concordia University and CIREQ, 1455 Boulevard de Maisonneuve Ouest, Department of Economics, Concordia University, Montréal, Québec, Canada H3G 1M8. Email: 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.

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.

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 have targeted different types of the receiver with different information. 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 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 utility profiles 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. Since each type chooses an optimal mapping, the receiver's interim utility profile is an upper envelope of linear functions and, hence, is convex. To sum up, an implementable interim utility profile 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 profile can be attained by an experiment. Moreover, in our model the receiver's interim utility profile 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. Since 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

[1] For 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.

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

spread of H. By linearity, the receiver's interim utility profile U can be represented as an 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 to 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, optimal mechanisms generally take a simple form of censorship mechanisms. A persuasion mechanism is called upper-censorship if it reveals the state below some threshold and pools the states above this threshold. We show that the optimal persuasion mechanism is upper-censorship, 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 upper-censorship 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 all types of the receiver experience a preference shock detrimental to the sender.

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 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 non-linear 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 it has limited applicability because it does not allow to directly characterize 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).

2. Model

2.1. **Setup.** There are two players: the 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 independently distributed with c.d.f.s 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

[3] There 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 Câmara (2016b). Bayesian persuasion with multiple senders is analyzed in Ostrovsky and Schwarz (2010), Board and Lu (2015), Au and Kawai (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 Câmara (2016a), and Bergemann and Morris (2016).

[4] Focusing on experiments is with loss of generality in settings with multiple receivers studied in Bergemann and Morris (2013), Alonso and Câmara (2016c), Mathevet et al. (2016), and Taneva (2016).

[5] Bayesian persuasion with monetary transfers is investigated in Bergemann and Pesendorfer (2007), Esö and Szentes (2007), Bergemann et al. (2015, 2016), Li and Shi (2015), and Hörner and Skrzypacz (2016).

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.

For convenience, we assume that the set of types is $R = \mathbb{R}$, the set of states is $\Omega = [0, 1]$, and the expected state is in the interior, $\mathbb{E}[\omega] \in (0, 1)$. These assumptions allow for elegance of presentation; relaxing them poses no difficulty.

2.2. **Persuasion mechanisms.** 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 $\pi : \Omega \times R \to [0, 1]$ asks the receiver to report $\hat{r} \in R$ 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 persuasion mechanism is *incentive compatible* if the receiver finds it optimal to report his true type and to follow the mechanism's recommendation.

By the revelation principle, the focus on incentive-compatible persuasion 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 persuasion mechanism. In particular, the restriction that the mechanism returns a binary recommendation about action instead of messages about the state is without loss, because the receiver's action is binary.

2.3. **Experiments.** Experiments communicate the same messages to all types 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 c.d.f. over message space $M = \mathbb{R}$ conditional on state ω.

For a given experiment σ, each message m induces a posterior belief of the receiver about the state. Since the receiver's utility is monotonic, 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 called *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 π is *equivalent to an experiment* if, conditional on each state, the behavior of the receiver stipulated by π is a best-reply behavior to the posterior beliefs generated by some experiment σ,

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

where $\sigma(r_-|\omega)$ denotes the left limit of $\sigma(.|\omega)$ 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 $[1 - \sigma(r|\omega), 1 - \sigma_-(r|\omega)]$.

2.4. **Envelope characterization of incentive compatibility.** Denote the expected utility of a receiver of type r who reports $\hat{r}$ and takes actions a_0 and a_1 in $A = \{0, 1\}$ after recommendations $\hat{a} = 0$ and $\hat{a} = 1$, respectively, by

$$U_\pi(r, \hat{r}, a_0, a_1) = \int_\Omega \left(a_0(1 - \pi(\omega, \hat{r})) + a_1\pi(\omega, \hat{r})\right)(\omega - r)\mathrm{d}F(\omega).$$

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

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

We consider mechanisms that satisfy the incentive compatibility constraint

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

It is convenient to introduce the notation for the expected utility 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

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

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, $q_\pi(\hat{r})$ 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*

$$q_\pi \quad \textit{is} \quad \textit{non-increasing}, \tag{4}$$

$$U_\pi(r) = \int_r^1 q_\pi(s)\mathrm{d}s, \tag{5}$$

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

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 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 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 in 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 $\omega = 1/2$ with certainty, then every non-increasing q with $\int_0^1 q(r)\mathrm{d}r = 1/2$ 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

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

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

$$V_\pi(r) = \int_\Omega v(\omega, r, 1)\pi(\omega, r)\mathrm{d}F(\omega) = q_\pi(r) + \rho(r)U_\pi(r). \tag{7}$$

3.2. **Bounds on the receiver's utility.** Consider two simple mechanisms. The *full disclosure* mechanism informs the receiver about the state, so the receiver acts if and only if his type is below the realized state ω, i.e., $\pi(\omega, r) = 1$ iff $r \leq \omega$, and the interim utility is

$$\overline{U}\left(r\right) = \int_r^1 \left(\omega - r\right) \mathrm{d}F\left(\omega\right).$$

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 interim utility is

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

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

Note that every mechanism π must satisfy

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

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.

3.3. **Implementable utility profiles.** The receiver's interim utility profile U and interim action profile 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.

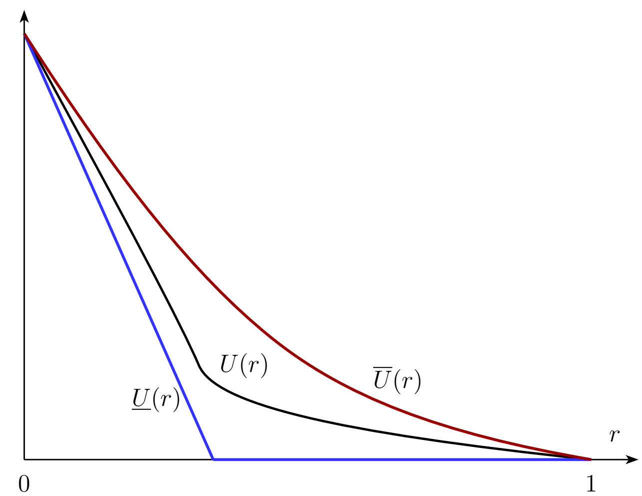

FIGURE 1. Set $\mathcal{U}$ contains every convex function between $\underline{U}$ and $\overline{U}$.

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

Theorem 1. *The following statements are equivalent:*
(a) U is a convex function between $\underline{U}$ and $\overline{U}$;
(b) U is implementable;
(c) U is implementable by an experiment.

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

By (2), every U implementable by an experiment is implementable. By Lemma 1 and (8), every implementable U belongs to $\mathcal{U}$. It remains to prove that every $U \in \mathcal{U}$ is implementable by an experiment.

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

$$\int_r^1 (1 - H(s))\mathrm{d}s = U(r) \leq \overline{U}(r) = \int_r^1 (1 - F(s))\mathrm{d}s \tag{9}$$

for all r, with equality at $r = 0$, because $\underline{U}(0) = \overline{U}(0) = \mathbb{E}[\omega]$. 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 σ given by $\sigma(m|\omega) = P(m|\omega)$. By construction, $\mathbb{E}[\omega|m] = m$ for all m of σ, and the probability that σ generates a message below any given value x is $H(x)$; so σ induces type r to act with probability $q_\sigma(r) \in [1 - H(r), 1 - H(r_-)]$, where indeterminacy arizes 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 utility profile U. □

Theorem 1 yields the following characterization of implementable action profiles.

Corollary 1. *The following statements are equivalent:*
(a) q is nonincreasing function that satisfies

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

(b) q is implementable;
(c) q is implementable by an experiment.

The heart of these characterization results is that every implementable U and q are implementable by an experiment. This result relies on the following connection between Mirrlees (1971) representation of incentive compatibility and Blackwell (1953) representation of garbling. Incentive compatibility (5) and feasibility (8) imply that every implementable utility and action profiles must satisfy 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, action profile q can be implemented by an appropriate garbling of the fully informative signal. Since every utility profile U is pinned down by an action profile q through (5), we obtain the result.

We now highlight a connection to the literature on 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), they 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 $\pi(\omega, r)$ which is nonincreasing in r for each $\omega \in \Omega$, thus being equivalent to an experiment (see Proposition 2), and satisfies

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

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

We will discuss 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 in Sections 6 and 7.

4. Optimal Mechanisms

In this section we use Theorem 1 to characterize 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 c.d.f.s F and G of the sender's and receiver's types 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 full-disclosure and no-disclosure mechanisms, maximize the sender's expected utility. We thus solve the problem for many commonly-used density functions (see Tables 1 and 3 of log-concave or log-convex density functions in Bagnoli and Bergstrom, 2005).

4.1. **Sender's problem.** The sender seeks a persuasion mechanism π that maximizes

$$\int_R V_\pi(r)\mathrm{d}G(r),$$

where $V_\pi(r)$ is the sender's expected 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 utility profile.

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

$$\int_R V_\pi(r)\mathrm{d}G(r) = \int_R U_\pi(r)I(r)\mathrm{d}r,$$

where

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

Proof. Observe that, by (5) and (7), $V_\pi(r) = q_\pi(r) + \rho(r)U_\pi(r) = -U'_\pi(r) + \rho(r)U_\pi(r)$. The result is immediate by integration by parts and by the assumption of compact support of G. □

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:

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

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 almost everywhere nonzero and changes sign $n \geq 0$ times.[6]

Let $\{r_1, r_2, \ldots, r_n\} \subset (0,1)$ be the set of types at which I changes its sign, and let $r_0 = 0$ and $r_{n+1} = 1$.[7]

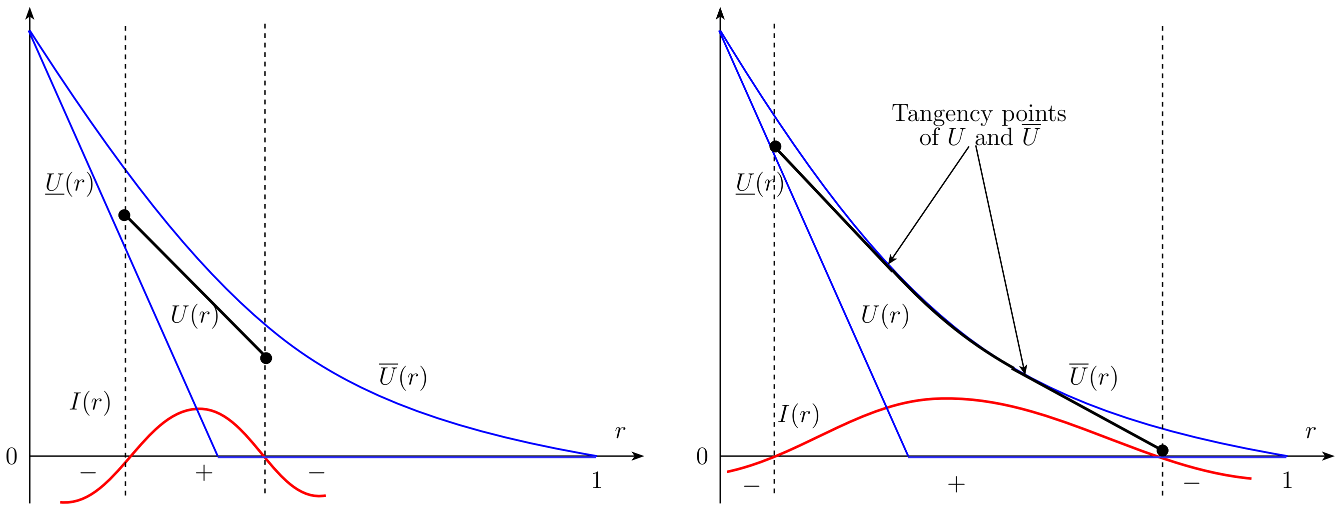

FIGURE 2. Optimal utility profile on the interval of where $I(r)$ is positive.

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 ($U \leq \overline{U}$) and convexity of U. That is, for any given values of $U(r_i)$ and $U(r_{i+1})$ at boundary points of interval (r_i, r_{i+1}), the utility profile U on the interior of (r_i, r_{i+1}) is a straight line unless $U(r) = \overline{U}(r)$, as illustrated by Fig. 2. Formally:

(P_1) 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 whose slopes are the derivatives at the endpoints, $U'(r_i)$ and $U'(r_{i+1})$, as illustrated by Fig. 3. Formally:

(P_2) On every interval (r_i, r_{i+1}) on which I is negative, U is piecewise linear with at most one kink, and satisfies

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

[6]If there are intervals where $I(r) = 0$, 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.

[7]Recall that $\underline{U}$ and $\overline{U}$ coincide outside of interval $(0,1)$, thus $I(r)$ plays no role on $\mathbb{R}\backslash(0,1)$.

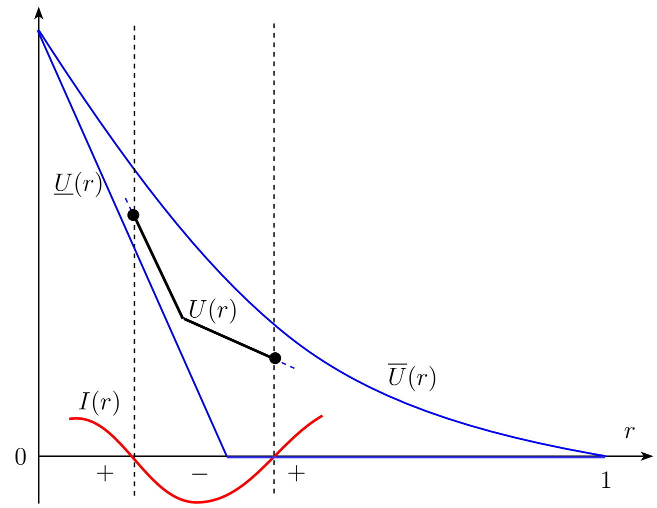

FIGURE 3. Optimal utility profile on the interval of where $I(r)$ is negative.

In effect, we can reduce the sender's problem (11) of optimization on the function space $\mathcal{U}$ to an n-variable optimization problem. An optimal utility profile is pinned down by properties (P_1) and (P_2) within each interval (r_i, r_{i+1}) where the sign of I is constant, and thus fully defined by the utility values at the points $(r_1, ..., r_n)$ where I changes its sign.

Consider the Cartesian product of n intervals

$$Y = \prod\nolimits_{i=1}^{n} \left[\underline{U}(r_i), \overline{U}(r_i)\right].$$

For every vector $\bar{y} = (y_1, ..., y_n) \in Y$, consider the set of all utility profiles in $\mathcal{U}$ that satisfy $U(r_i) = u_i$ for all $i = 1, .., n$. If this set is nonempty, properties (P_1) and (P_2) uniquely determine the seller's preferred utility profile on this set, denoted by $U^*_{\bar{y}}$. For completeness, define $U^*_{\bar{y}} = \underline{U}$ whenever that set is empty. Thus we obtain:

Theorem 2. *The sender's problem* (11) *is an n-variable optimization problem on domain $Y \subset \mathbb{R}^n$,*

$$\max_{\bar{y} \in Y} \int_R U^*_{\bar{y}}(r) I(r) \mathrm{d}r.$$

Properties (P_1) and (P_2) imply that the optimal U is piecewise linear except for the intervals where $U(x) = \overline{U}(x)$, as shown on Fig. 2–3. For a given optimal U, the corresponding experiment's c.d.f. H is constructed as follows. The linear intervals of U correspond to constant H (i.e., zero density). The kinks of U correspond to the atom points of the distribution equal to the difference between the right and left derivatives of U at those points. The intervals of $U(x) = \overline{U}(x)$ correspond to full-disclosure intervals, $H(x) = F(x)$.

4.3. **Upper- and lower-censorship.** An experiment 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

$$m(\omega) = \begin{cases} \omega, & \text{if } \omega < \omega^*, \\ \mathbb{E}[\omega|\omega \geq \omega^*], & \text{if } \omega \geq \omega^*. \end{cases}$$

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

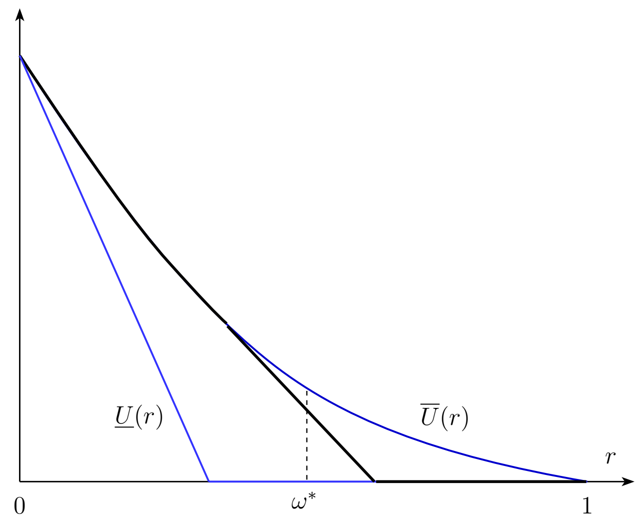

FIGURE 4. Utility under upper-censorship mechanism.

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. 4. A receiver with type $r < \omega^*$ is fully informed whether ω exceeds r or not, and hence receives the highest feasible utility $\overline{U}(r)$ by acting if $\omega > r$. A receiver with type $r > \omega^*$ is only informed whether ω exceeds ω^* or not, hence he acts if $\mathbb{E}[\omega|\omega \geq \omega^*] - r > 0$. His utility is thus either $\mathbb{E}[\omega|\omega \geq \omega^*] - r$ 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. *Upper-censorship (lower-censorship) 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). □

4.4. **Comparative statics.** Theorem 3 allows for sharp comparative statics analysis on the amount of information that is optimally disclosed by the sender. Let $\rho(r) = \rho \in \mathbb{R}$ and let g be a log-concave density with $g' + \rho g$ being almost everywhere nonzero on $\mathbb{R}$.[8] 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. Since g_t is log-concave on Ω for every t, upper censorship is optimal by Theorem 3. Let $\omega^*\,(\rho, t) \in \Omega$ be the optimal upper-censorship cutoff.

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 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 same types of the receiver, the sender needs to increase $\mathbb{E}[\omega|\omega \geq \omega^*]$ by expanding the full disclosure interval $[0, \omega^*)$.

4.5. **Full-disclosure and no-disclosure mechanisms.** Theorem 3 also allows to characterize the class of c.d.f.s of the receiver's type under which the optimal choice of the sender polarizes between full-disclosure and no-disclosure mechanisms, as, for example, in Lewis and Sappington (1994) and Johnson and Myatt (2006).[9]

Corollary 3. *Let $\rho(r) = \rho \in \mathbb{R}$. Either full disclosure or no disclosure is optimal for all F and all ρ if and only if there exist $\lambda \in \mathbb{R}$ and $c \in (0, \lambda/(1 - e^{-\lambda}))$ such that $g(r) = ce^{-\lambda r}$ for all $r \in \Omega$.*

Proof. The only two mechanisms that are both upper-censorship and lower-censorship are full disclosure and no disclosure. By Theorem 3, both upper-censorship and lower-censorship 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 non-increasing and non-decreasing on Ω, which holds if and only if $g'(r)/g(r)$ is equal to a constant, $-\lambda$, on Ω. Constraint $c \in (0, \lambda/(1 - e^{-\lambda}))$ ensures that $g(r)$ is a probability density function. $\square$

[8] For log-convex g the results are symmetric.

[9] In 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 (full disclosure or no disclosure) 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.

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

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

5. Application: Media Censorship

In this section, we apply our model to the question 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 their information through media. In autocracies and countries with weak checks and balances, the government has power to control the media content.[10]

The government 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, the government needs no sophisticated instruments of information disclosure other than censorship to act optimally. 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.

[10] A pro-Kremlin website politonline.ru has conducted a quantitative study that ranks Russian media outlets according to disloyalty and published a list of the top-20 most "negative" news sources on the Russian Internet (*GlobalVoices*, April 2014). 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."

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.

The government's state of affairs is a random variable ω drawn from $[0,1]$ according to c.d.f. F that admits 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 $\omega \geq s$ and criticizes it (sends message $m_s = 0$) if $\omega < s$.[11,12] The threshold 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 the anti-government language.[13]

There is a continuum of heterogeneous readers indexed by $r \in [0,1]$ distributed according to c.d.f. G that admits a log-concave density g. Each reader observes endorsements of all available media outlets[14] 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

$$\int_R u(\omega, r, a_r, \bar{a})\mathrm{d}G(r) + \bar{a}\gamma = \int_R \Big(a_r(\omega - r) + \bar{a}(\xi(r) + \gamma)\Big)\mathrm{d}G(r),$$

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

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

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 $[0,1]\backslash S$.

[11]If 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.

[12]As in Suen (2004), Chan and Suen (2008), and Chiang and Knight (2011), binary media reports that only communicate 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.

[13]Gentzkow and Shapiro (2010) construct such a slant index for U.S. newspapers. The theoretical literature has explored 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).

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

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 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 into domestic stocks, volunteering for the military service, and voting for the government.

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 only observe 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,

$$\bar{a}_\pi(\omega) = \int_R \pi(\omega, r)\mathrm{d}G(r) \quad \text{and} \quad q_\pi(r) = \int_\Omega \pi(\omega, r)\mathrm{d}F(\omega).$$

Denote by $\bar{A}_\pi$ the expected average action,

$$\bar{A}_\pi = \int_R q_\pi(r)\mathrm{d}G(r).$$

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

$$U_\pi(r) = \int_\Omega (\omega - r)\pi(\omega, r)\mathrm{d}F(\omega).$$

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

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

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

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

where

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

From Theorem 3 it follows that the optimal persuasion mechanism is an upper-censorship with some threshold s^*.[15] Notice that the government can achieve this mechanism by prohibiting each media outlet s to broadcast if and only if $s \geq s^*$.

Corollary 5. *The government's optimal censorship policy is* $[s^*, 1]$ *for some* $s^* \in [0, 1]$.

Note that a media outlet with a higher editorial policy 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 to broadcast.

The government's censorship policy $[s^*, 1]$ 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 the government would be better off if it could create more complex signals that aggregate information from multiple media outlets and add noise.

6. Multiple Actions

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

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

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 every two consecutive actions. For each $a = 1, ..., n$,

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

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.

[15] This 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 (a linear transformation of) $\int_R \tilde{I}(r) U_S(r) \mathrm{d}r$, where U_S is the receiver'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.

For each $a = 1, ..., n$, the receiver's and sender's utility can be expressed as

$$u(\omega, r, a) = \left(\sum\nolimits_{i=1}^{a} b_i(r)\right)\omega - \left(\sum\nolimits_{i=1}^{a} b_i(r)x_i(r)\right)$$
$$v(\omega, r, a) = \left(\sum\nolimits_{i=1}^{a} z_i(r)\right) + \rho(r)u(\omega, r, a).$$

We assume that $b_a(r) > 0$ for all r and all $a = 1, ..., n$. 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,

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

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

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

Thus, the receiver of type r optimally chooses action a on the interval of states $(\bar{x}_a(r), \bar{x}_{a+1}(r))$.[16]

6.2. **Experiments.** Since 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 $x \in \Omega$,

$$H_\sigma(x) = \int_\Omega \sigma(x|\omega)\mathrm{d}F(\omega).$$

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

$$C_\sigma(x) = \int_x^\infty (1 - H_\sigma(m))\mathrm{d}m.$$

Observe that by equation (9), for every experiment σ, we have $C_\sigma(r) = U_\sigma(r)$ for all r, where $U_\sigma(r)$ is the receiver's utility profile 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

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

where $\overline{C}$ and $\underline{C}$ correspond to the full-disclosure and no-disclosure signals,

$$\overline{C}(x) = \int_x^\infty (1 - F(m))\mathrm{d}m,$$
$$\underline{C}(x) = \max\{\mathbb{E}[\omega] - x, 0\}.$$

[16]This 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 exposition. Relaxing this assumption poses no difficulty; it will only require for each type r to omit from the analysis the actions that are never optimal.

6.3. **Implementable utility profiles.** The expected utility of type r under signal σ is equal to

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

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

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

A persuasion mechanism can be described by a (possibly, infinite) menu of signals, Σ. The receiver of type r chooses a signal from the menu, and then observes messages only from that signal. Obviously, the receiver chooses the signal 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 a utility profile U of the receiver is implementable if and only if there exists a menu $\mathcal{C}_\Sigma \subset \mathcal{C}$ such that

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

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

[Figure here]

Figure 5. Three-action example: Utility of type r^*

Example 1. Let $A = \{0, 1, 2\}$, F admit a strictly positive density, and $u(\omega, r, a)$ be continuous in r. 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 signals represented by partitions $\{P', P''\}$, where P' and P'' are the first-best partitions for types r' and r'',

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

Types r' and r'' choose, respectively, P' and P'' and get their maximum possible utilities $\overline{U}(r')$ and $\overline{U}(r'')$. By continuity of $u(\omega, r, a)$ in r, there exists 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 $a \in \{1, 2\}$,

$$L_a(P') = \int_{x_a(r^*)}^{x_a(r')} \Big(u(\omega, r^*, a) - u(\omega, r^*, a-1) \Big) \mathrm{d}F(\omega)$$

denotes the loss of type r^* from using cutoff $x_a(r')$ rather than her 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 type r^* from taking action a, $u(\omega, r^*, a)$, for each $a = 0, 1, 2$. The kinked solid blue line is the "envelope" utility of r^* from taking the optimal action, $\max_{a \in \{0,1,2\}} u(\omega, r^*, a)$. On Fig. 5, the loss of r^* from signal 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 r^* from signal 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 r^* is indifferent between the two signals.

An experiment that gives the maximum possible utilities $\overline{U}(r')$ and $\overline{U}(r'')$ 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

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

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

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

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

On 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 common refinement of P' and P'' (the area of the smaller shaded and hatched triangles) than under either signal P' or signal P''.

6.4. **Sender's problem.** In this section we impose the following additional assumptions. For each $a \in \{0, 1, \dots, n\}$, the function $z_a(r)$ is differentiable and the function $x_a(r)$ has an inverse $r_a(x)$ on $\mathbb{R}$. Moreover, $G(r_a(x))$ is twice differentiable in x.

For a given signal σ, the sender's expected utility conditional on the receiver's type being r is[17]

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

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

Lemma 4. *The sender's expected utility under experiment σ is*

$$\int_R V_\sigma(r)\mathrm{d}G(r) = \int_{\mathbb{R}} C_\sigma(x)I(x)\mathrm{d}x,$$

where

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

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

$$\max_{C\in\mathcal{C}} \int_{\mathbb{R}} C(x)I(x)\mathrm{d}x.$$

The solutions to this problem are characterized by Theorem 2. The sender's optimal persuasion mechanism is described by the set of functions $\mathcal{C}_\Sigma \subset \mathcal{C}$ that solves

$$\max_{\mathcal{C}_\Sigma\subset\mathcal{C}} \int_{\mathbb{R}} C(x)I(x)\mathrm{d}x$$

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

We already know that, when the receiver has more than two actions, the set of implementable receiver's utility profiles is strictly larger under persuasion mechanisms than under experiments. We now show that the set of implementable action profiles 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 signal, type r' takes action $a = 2$ if and only if $\omega \in [x_2^*, 1]$, because for $\omega < x_2^*$, this signal must generate messages distinct from $x_2(r'')$ and thus below x_2^*, which is in turn below $x_2(r')$.

[17]For each r where $H_\sigma(x_{a+1}(r))$ is discontinuous this formula assumes that type r breaks the indifference in favour 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.

Consider now a persuasion mechanism consisting of the menu of two signals 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 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, this action profile can not be achieved by any experiment.

Finally, an optimal persuasion mechanism need not be an experiment. Suppose that the sender only cares 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

So far we have considered the receiver's utility which is linear in the state ω. In this section we relax this assumption.

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 types is $R = \mathbb{R}$. The receiver's utility is

$$u(\omega, r, a) = au(\omega, r), \quad a \in \{0, 1\},\ \omega \in \Omega,\ r \in R,$$

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

$$u(\omega, \omega) = 0, \quad \omega \in \Omega.$$

The sender's utility is

$$v(\omega, r, a) = av(\omega, r), \quad a \in \{0, 1\},\ \omega \in \Omega,\ r \in R,$$

where $v(\omega, r)$ is differentiable. State ω and type r are independently distributed according to c.d.f.s F and G, respectively.[18]

[18] Note that if 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, 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 $U(r) = \int_\Omega \tilde{u}(\omega, r)\pi(\omega, r)\mathrm{d}F(\omega)$.

7.2. **Characterisation of experiments.** We start with the characterisation of persuasion mechanisms that are equivalent to experiments.

Recall that an experiment sends (stochastic) messages as a function of realized states. We denote by $\sigma(\cdot|\omega)$ the c.d.f. of messages conditional on state ω. As in Section 2.3, without loss of generality we can consider *direct experiments* whose messages are identified with cutoff types who are indifferent between acting and not acting. This is because the receiver's utility is monotonic, and hence for every posterior belief about state there is a cutoff type such that the receiver acts if and only if his type is below that cutoff.

Also recall that a persuasion mechanism π is equivalent to an experiment if the ex-post (state-contingent) behavior of the receiver stipulated by π is a best-reply behavior to the posterior beliefs generated by some experiment.

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 $\omega \in \Omega$.*

Intuitively, since for each experiment σ, the c.d.f. $\sigma(r|\omega)$ of r conditional on each state ω is nondecreasing in r, each $\pi(\omega, r) \in \left[1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)\right]$ (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 signals, each of which communicates to the receiver whether the state is above some value or not.

Since a cutoff signal 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 signal will inform the receiver whether the state is above $x(r)$ or not. Obviously, different types of the receiver may prefer different cutoffs.

Upper-censorship and lower-censorship mechanisms defined in Section 4.3 are cutoff mechanisms. For example, an upper-censorship mechanism with censorship cutoff w^* can be represented as the menu of cutoff signals $X = [0, \omega^*]$.

Proposition 2 allows us to make the following statement.

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

For a given cutoff mechanism X, the correspondent experiment is the monotone partition of the state space generated by observation of all experiments X. For example, interval $(x', x'']$, with $x', x'' \in (0, 1)$, is the element of that partition if and only if x' and x'' are two consecutive cutoffs in 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.*

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

7.5. **Beyond the binary state case.** 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 posterior distributions of the state induced by signal realizations can no longer be parametrized by a one-dimensional variable, such as the posterior mean state in case of linear utilities, the posterior probability of one of the states in case of binary-valued state, and the cutoff value in case of cutoff mechanisms.

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

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 utility profile U. Then the expected action profile 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 $(\omega, r) \in \{\omega_1, \omega_2, \omega_3\} \times (\omega_1, \omega_3)$.*

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

Example 2. Let $r' < r''$ be two types of the receiver and denote $u'_i = u(\omega_i, r')$ and $u''_i = u(\omega_i, r'')$ for $i \in \{1, 2, 3\}$. Let $u'_2 < 0 < u''_3$, $u'_1/u''_1 < u'_2/u''_2$, $f_1 u'_1 + f_3 u'_3 < 0$, and $f_2 u'_2 + f_3 u'_3 < 0$.

Let Σ'' be the set of all signals σ'' that induce type r'' to act with probability 1 if $\omega = \omega_3$, with probability $-f_3 u''_3/(f_2 u''_2)$ if $\omega = \omega_2$, and with probability 0 if $\omega = \omega_1$. It is easy to check that Σ'' is the set of signals that maximize the probability of $a = 1$ for type r''. By monotonicity of u in r, for each message of $\sigma'' \in \Sigma''$, if type r'' acts, type $r' < r''$ also acts. Moreover, by definition of Σ'', for each message of $\sigma'' \in \Sigma''$, if type r'' does not act, the state can only be ω_1 or ω_2; so type r' also does not act, because $u'_2 < 0$. Therefore, for any experiment that induces type r'' to act with probability $f_3\big(1 - u''_3/u''_2\big)$, type r' acts with the same probability.

Let Σ' be the set of all signals σ' that induce type r' to act with probability 1 if $\omega = \omega_3$, with probability 0 if $\omega = \omega_2$, and with probability $-f_3 u''_3/(f_1 u''_1)$ if $\omega = \omega_1$. By the same argument as above, for each message of $\sigma' \in \Sigma'$, type r' acts if and only if type r'' acts. Notice that type r'' is indifferent between $\sigma' \in \Sigma'$ and $\sigma'' \in \Sigma''$, but type r' strictly prefers σ' to σ'', because $u'_1/u''_1 < u'_2/u''_2$. Therefore, under a persuasion mechanism that consists of the menu of two signals $\sigma' \in \Sigma'$ and $\sigma'' \in \Sigma''$, type r'' acts with probability $f_3 - f_3 u''_3/u''_2$, but type r' acts with (different) probability $f_3 - f_3 u''_3/u''_1$.

Finally, when the receiver's utility is nonlinear, the set of receiver's utility profiles 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 signals σ' and σ'' that maximize the utility of types r' and r^*, respectively, subject to the constraint that type r'' gets utility $\underline{U}(r'') = 0$. Since Σ' and Σ'' do not intersect, no experiment can achieve the utility profile induced by a persuasion mechanism that consists of the menu of two signals $\sigma' \in \Sigma'$ and $\sigma'' \in \Sigma''$.

Appendix A. Linear Utilities

This section 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, $a = 1$,

$$u(\omega, r, a) = a \cdot b(\omega - t),$$
$$v(\omega, r, a) = a \cdot (c(\omega - t) + d),$$

where $r = (b, c, d, t) \in \mathbb{R}^4$ denotes the receiver's type. The type is distributed according to a c.d.f. 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 distributed according to c.d.f. F independently of r.

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

$$C_\sigma(t) = \int_t^1 (1 - H_\sigma(m))\mathrm{d}m.$$

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

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

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

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

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

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

where $\overline{C}$ and $\underline{C}$ correspond to the full-disclosure and no-disclosure signals. 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 the sender's and receiver's expected utilities is implementable by an experiment. Moreover, the sender's problem can be expressed as

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

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

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

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

$$U_\sigma(r) = \int_t^1 b(m-t)\mathrm{d}H_\sigma(m) = bC_\sigma(t).$$

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

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

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

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

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

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

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

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

and

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

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

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

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

Now, integrating with respect to t,

$$\int_{(b,c,d,t)} V_\sigma(b,c,d,t)\mathbf{1}_{b>0}\mathrm{d}G(b,c,d,t) = \int_t (c_+(t)C_\sigma(t) - d_+(t)C'_\sigma(t))g_t(t)\mathrm{d}t$$
$$= \int_t (c_+(t)g_t(t) + \tfrac{\mathrm{d}}{\mathrm{d}t}[d_+(t)g_t(t)])C_\sigma(t)\mathrm{d}t.$$

Similarly, for $b < 0$,

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

Now, integrating w.r.t. t,

$$\int_{(b,c,d,t)} V_\sigma(b,c,d,t)\mathbf{1}_{b<0}\mathrm{d}G(b,c,d,t) = -\int_t (c_-(t)C_\sigma(t) - d_-(t)C'_\sigma(t))g_t(t)\mathrm{d}t + K$$
$$= -\int_t (c_-(t)g_t(t) + \tfrac{\mathrm{d}}{\mathrm{d}t}[d_-(t)g_t(t)])C_\sigma(t)\mathrm{d}t + K,$$

where

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

is a constant independent of C_σ.

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

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

where

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

□

Appendix B. Omitted Proofs

B.1. **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 $\hat{r} \to r$ and then integrating from r to 1 gives

$$U_\pi(1) - U_\pi(r) = -\int_r^1 q_\pi(s)\mathrm{d}s.$$

Also, observe that type $r = 1$ can secure his maximal attainable utility of 0 by always acting (irrespective of 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

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

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

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

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

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

B.2. **Proof of Lemma 3.** The 'if' part follows from Properties (P_1) and (P_2) 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 almost everywhere nonzero). 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 argument for I that does not cross the horizontal axis at most once and from below on Ω is symmetric.

B.3. **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:

$$V_\sigma = C + \int_0^1 J_t(r)\, \mathrm{d}H_\sigma(r), \tag{15}$$

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

$$J_t(r) = \int_0^r (g_t(s) + \rho G_t(s))\,\mathrm{d}s.$$

Consider ρ and t such that $\omega^*(\rho, t) \in (0,1)$. Under upper censorship with cutoff 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 $\omega^{**} = \mathbb{E}[\omega|\omega \geq \omega^*]$.

Part (a). The derivative of the sender's expected utility (15) under upper censorship with respect to cutoff ω^* is:

$$\begin{aligned}\frac{\mathrm{d}V}{\mathrm{d}\omega^*} &= f(\omega^*) \int_{\omega^*}^{\omega^{**}} (J_t'(\omega^{**}) - J_t'(s))\,\mathrm{d}s \\ &= f(\omega^*) \left[\int_{\omega^*}^{\omega^{**}} (g_t(\omega^{**}) - g_t(s))\,\mathrm{d}s + \rho \int_{\omega^*}^{\omega^{**}} (G_t(\omega^{**}) - G_t(s))\,\mathrm{d}s\right].\end{aligned}$$

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

$$\begin{aligned}\frac{\mathrm{d}^2V}{\mathrm{d}t\,\mathrm{d}\omega^*} &= -f(\omega^*) \int_{\omega^*}^{\omega^{**}} (J_t''(\omega^{**}) - J_t''(s))\,\mathrm{d}s \\ &= -f(\omega^*)\,J_t''(\omega^{**})(\omega^{**} - \omega^*) + f(\omega^*)(J_t'(\omega^{**}) - J_t'(\omega^*)).\end{aligned}$$

At the optimal interior cutoff, we have $\mathrm{d}V/\mathrm{d}\omega^* = 0$; so

$$J_t'(\omega^{**})(\omega^{**} - \omega^*) = \int_{\omega^*}^{\omega^{**}} J_t'(s)\,\mathrm{d}s. \tag{16}$$

Since $J_t''(r)$ is almost everywhere nonzero, 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 $\mathrm{d}^2V/\mathrm{d}t\mathrm{d}\omega^* > 0$ and ω^* is strictly increasing in t by Theorem 1 of Edlin and Shannon (1998).

B.4. **Proof of Proposition 1.** The expected utility of type r under signal σ is

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

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

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

Summing up the above over $a = 1, ..., n$ yields (12).

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

B.5. **Proof of Lemma 4.** We have

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

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

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

Fix $a = 1, ..., n$ and define the variable $x = x_a(r)$, and hence $r = r_a(x)$. Using this variable change, we have

$$\begin{aligned}&\int_R \Big(\rho(r)b_a(r)C_\sigma(x_a(r)) - z_a(r)C'_\sigma(x_a(r))\Big)\mathrm{d}G(r) \\ &= \int_{\mathbb{R}} \Big(\rho(r_a(x))b_a(r_a(x))C_\sigma(x) - z_a(r)C'_\sigma(x)\Big)\mathrm{d}G(r_a(x)).\end{aligned}$$

Now we integrate by parts

$$\begin{aligned}\int_{\mathbb{R}} z_a(r)C'_\sigma(x)\mathrm{d}G(r_a(x)) &= C(x)z_a(r)\frac{\mathrm{d}G(r_a(x))}{\mathrm{d}x}\Big|_{-\infty}^{\infty} - \int_{\mathbb{R}} C(x)\tfrac{\mathrm{d}}{\mathrm{d}x}\left(z_a(r)\tfrac{\mathrm{d}}{\mathrm{d}x}G(r_a(x))\right)\mathrm{d}x \\ &= -\int_{\mathbb{R}} C(x)\tfrac{\mathrm{d}}{\mathrm{d}x}\left(z_a(r)\tfrac{\mathrm{d}}{\mathrm{d}x}G(r_a(x))\right)\mathrm{d}x.\end{aligned}$$

Thus we obtain

$$\begin{aligned}&\int_R \Big(\rho(r)b_a(r)C_\sigma(x_a(r)) - z_a(r)C'_\sigma(x_a(r))\Big)\mathrm{d}G(r) \\ &= \int_{\mathbb{R}} \Big(\rho(r_a(x))b_a(r_a(x))\frac{\mathrm{d}G(r_a(x))}{\mathrm{d}x} + \tfrac{\mathrm{d}}{\mathrm{d}x}\left(z_a(r)\tfrac{\mathrm{d}}{\mathrm{d}x}G(r_a(x))\right)\Big)C_\sigma(x)\mathrm{d}x.\end{aligned}$$

Summing the above over $a = 1, ..., n$ we obtain $\int_{\mathbb{R}} C_\sigma(x)I(x)\mathrm{d}x$, where I is defined in Lemma 4.

B.6. **Proof of Proposition 2.** Consider a mechanism π which is equivalent to an experiment σ. Since $\sigma(r|\omega)$ is a c.d.f. of r conditional on ω, it is nondecreasing in r for every ω. Then, by (2), $\pi(\omega, r)$, is nonincreasing in r for every 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 $\pi(.,\omega)$ at r. Since $\pi(r_+,\omega) \in [0,1]$ nonincreasing and right-continuous in r, function $\sigma(r|\omega)$ is a c.d.f., which describes the distribution of messages for every given state ω. Thus, σ is a signal. It remains to verify that the constructed signal is direct and induces the same action by the receiver as mechanism π, i.e., when the signal sends a message r, then type r is indifferent between the two actions. For all r,

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

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

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

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

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

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

B.7. **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 auxiliary cutoffs $x = -\epsilon$ ($x = 1+\epsilon$, 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

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

and the corresponding mechanism is

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

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.

B.8. **Proof of Corollary 7.** Consider F whose support consist of two states, w.l.o.g., $\{0,1\}$, 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 $r, \hat{r} \in (0,1)$,

$$\sum_{\omega=0,1} u(\omega, r)\left(\pi(\omega, r) - \pi(\omega, \widehat{r})\right) \Pr(\omega = 1) \geq 0. \tag{17}$$

Writing (17) for $(r, \widehat{r}) = (r_2, r_1)$ and $(r, \widehat{r}) = (r_1, r_2)$ yields:

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

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

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

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

B.9. **Proof of Proposition 3.** For all $r \in (\omega_1, \omega_3)$ and $j = 1, 2$, we have

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

where the first line holds by 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

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

The solution of differential equation (20) is given by

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

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

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

8. Attic

(1) Receiver's indifference is measure zero event
(2) the sender's and receiver's
(3) sender (she) and receiver (he)
(4) interim utility profiles $V(r)$ and $U(r)$
(5) Check Tex console for Warnings
(6) We need compact support of G from the start (Sec 2).
(7) Proof refer to Appendix B.

References

Alonso, R., and O. Câmara (2016a): "Bayesian Persuasion with Heterogeneous Priors," *Journal of Economic Theory*, forthcoming.

——— (2016b): "On the Value of Persuasion by Experts," mimeo.

——— (2016c): "Persuading Voters," *American Economic Review*, forthcoming.

Au, P. H. (2015): "Dynamic Information Disclosure," *RAND Journal of Economics*, 46, 791–823.

Au, P. H., and K. Kawai (2016): "Competition in Information Disclosure," mimeo.

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 (2016): "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., 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," *The 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 (2015): "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.

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: Quantity vs. Quality," 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 (2015): "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," *The 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," *The Economic Journal*, 114, 377–396.
Taneva, I. (2016): "Information Design," mimeo.