

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 85, No. 6 (November, 2017), 1949–1964

PERSUASION OF A PRIVATELY INFORMED RECEIVER

ANTON KOLOTILIN

School of Economics, UNSW Business School

TYMOFIY MYLOVANOV

Department of Economics, University of Pittsburgh and Kyiv School of Economics

ANDRIY ZAPECHELNYUK

School of Economics and Finance, University of St Andrews

MING LI

Department of Economics, Concordia University and CIREQ

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

NOTES AND COMMENTS

PERSUASION OF A PRIVATELY INFORMED RECEIVER

ANTON KOLOTILIN

School of Economics, UNSW Business School

TYMOFIY MYLOVANOV

Department of Economics, University of Pittsburgh and Kyiv School of Economics

ANDRIY ZAPECHELNYUK

School of Economics and Finance, University of St Andrews

MING LI

Department of Economics, Concordia University and CIREQ

We study persuasion mechanisms in linear environments. A receiver has a private type and chooses between two actions. A sender designs a persuasion mechanism or an experiment to disclose information about a payoff-relevant state. A persuasion mechanism conditions information disclosure on the receiver's report about his type, whereas an experiment discloses information independent of the receiver's type. We establish the equivalence of implementation by persuasion mechanisms and by experiments, and characterize optimal persuasion mechanisms.

KEYWORDS: Bayesian persuasion, information disclosure, information design, mechanism design without transfers, experiments, persuasion mechanisms.

1. INTRODUCTION

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

Anton Kolotilin: akolotilin@gmail.com

Tymofiy Mylovanov: mylovanov@gmail.com

Andriy Zapechelnyuk: azapech@gmail.com

Ming Li: ming.li@concordia.ca

The authors would like to express their gratitude to anonymous referees for their very helpful comments and to Navik Kartik for a thoughtful discussion at the Boston ASSA meetings. They are also grateful to Ricardo Alonso, Dirk Bergemann, Patrick Bolton, Alessandro Bonatti, Rahul Deb, Péter Esö, Johannes Hörner, Florian Hoffman, Roman Inderst, Emir Kamenica, Daniel Kräehmer, Marco Ottaviani, Mallesh Pai, Larry Samuelson, and Ilya Segal for their excellent comments. Kolotilin acknowledges financial support from the Australian Research Council. Zapechelnyuk acknowledges financial support from the Economic and Social Research Council Grant ES/N01829X/1.

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

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

We now outline the argument behind this result. Because the utilities are linear in the state, the only relevant content of a message about the state is the posterior mean. Therefore, an experiment can be fully described by the distribution of the posterior mean state H that it generates. By Blackwell (1953), there exists an experiment that generates H if and only if the prior distribution of the state is a mean-preserving spread of H . By linearity, the receiver's interim utility U can be represented as an appropriate integral of H . Using the integral form of the mean-preserving spread condition as in Blackwell (1951), we show that the prior is a mean-preserving spread of H if and only if U is convex and lies between the upper and lower bounds generated by full and no disclosure. Therefore, any U that satisfies the necessary conditions that we identified above is implementable by an experiment.

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

Related Literature

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

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

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

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) imposed more structure on the receiver's utility and the distribution of the receiver's type than our paper. At the same time, they allowed the sender's utility to be nonlinear in the state. Rayo and Segal (2010) partially characterized optimal experiments and represented 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 (2017) allowed for nonlinear sender's and receiver's utility functions and an arbitrary joint distribution of the state and the receiver's type. The linear-programming approach in Kolotilin (2017) permits to verify whether a candidate experiment is optimal. But this approach has limited applicability because it does not allow to directly characterize the optimal experiments.

The three papers above study experiments.³ In contrast, we consider persuasion mechanisms, which can tailor information disclosure to the receiver's report.⁴ Linear utility functions with two possible actions of the receiver enable us to use the envelope representation of incentive compatibility as in Myerson (1981). However, the characterization of implementable mechanisms in our setting differs from Myerson (1981) because there are no transfers between the sender and receiver, and there are obedience constraints instead of participation constraints.⁵ The equivalence between persuasion mechanisms and experiments relies on the majorization theory of Hardy, Littlewood, and Pólya (1929) and, thus, relates to the equivalence between Bayesian and dominant-strategy implementation in Manelli and Vincent (2010) and Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013).

2. MODEL

2.1. Setup

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

Let the receiver's and sender's utilities be

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

³There is a rapidly growing Bayesian persuasion literature that studies optimal experiments. Bayesian persuasion with a privately informed sender was considered in Gill and SgROI (2008, 2012), Perez-Richet (2014), Alonso and Câmara (2017), and Hedlund (2017). Bayesian persuasion with multiple senders was analyzed in Ostrovsky and Schwarz (2010), Board and Lu (2017), Gentzkow and Kamenica (2017a, 2017b), and Li and Norman (2017). Dynamic Bayesian persuasion was examined in Au (2015), Ely, Frankel, and Kamenica (2015), and Ely (2017). Information acquisition and the value of information in Bayesian persuasion were explored in Gentzkow and Kamenica (2014, 2017c), Kolotilin (2015), Alonso and Câmara (2016a), and Bergemann and Morris (2016).

⁴Focusing on experiments is *with* loss of generality in settings with multiple receivers in Bergemann and Morris (2013), Alonso and Câmara (2016b), Mathevet, Perego, and Taneva (2017), and Taneva (2016).

⁵Bayesian persuasion with monetary transfers was investigated in Bergemann and Pesendorfer (2007), Eső and Szentés (2007), Bergemann, Brooks, and Morris (2015, 2017), Bergemann, Bonatti, and Smolin (2017), Li and Shi (2017), and Hörner and Skrzypacz (2016).

where ρ is a bounded measurable function.⁶

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, whereas if the weight is zero, then the sender cares only about whether the receiver acts or not.

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

2.2. Persuasion Mechanisms

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

A *persuasion mechanism* π asks the receiver to report $\hat{r} \in R$ and then recommends him to take one of two actions: for every $\omega \in \Omega$, it recommends to act ($\hat{a} = 1$) with probability $\pi(\omega, \hat{r})$ and not to act ($\hat{a} = 0$) with complementary probability $1 - \pi(\omega, \hat{r})$, where $\pi : \Omega \times R \rightarrow [0, 1]$ is a measurable function.

A persuasion mechanism is *incentive compatible* if the receiver finds it optimal to report his true type and to follow the mechanism's recommendation. Denote the interim utility of a receiver of type $r \in R$ who reports $\hat{r} \in R$ and takes actions $a_0 \in \{0, 1\}$ and $a_1 \in \{0, 1\}$ after recommendations $\hat{a} = 0$ and $\hat{a} = 1$, respectively, by

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

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

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

We consider mechanisms that satisfy the incentive-compatibility constraint

$$U_\pi(r) \geq U_\pi(r, \hat{r}, a_0, a_1) \quad \text{for all } r, \hat{r} \in R \text{ and } a_0, a_1 \in A. \quad (2)$$

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

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

⁶For our analysis, it is sufficient to assume that $u(\omega, r, a)$ and $v(\omega, r, a)$ are linear in (a transformation of) the state ω . The unidimensionality of the type r and the linearity of the utilities in the type are assumed only for elegance of exposition. We further comment on this in Section 3.3, and refer an interested reader to Appendix C of the Supplemental Material (Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017)) for technical details.

where $q_\pi(\hat{r})$ is the interim action and $p_\pi(\hat{r})$ is the expected state after action conditional on a report \hat{r} :

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

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

LEMMA 1: *A mechanism π is incentive compatible if and only if*

$$q_\pi \text{ is nonincreasing,} \tag{3}$$

$$U_\pi(r) = \int_r^1 q_\pi(s) ds, \tag{4}$$

$$U_\pi(0) = \mathbb{E}[\omega]. \tag{5}$$

The proof is in the [Appendix](#).

Interestingly, the obedience constraints for the intermediate types are implied by the boundary conditions, $U_\pi(1) = 0$ and $U_\pi(0) = \mathbb{E}[\omega]$, and truth telling, $U_\pi(r) \geq U_\pi(r, \hat{r})$. To disobey by ignoring the recommendation, that is, to act (not to act) irrespective of what is recommended, is not better than pretending to be the lowest type $\hat{r} = 0$ (the highest type $\hat{r} = 1$, respectively). Due to the linearity of the receiver’s utility, to misreport and disobey by taking the opposite action to the recommended one is always inferior to reporting either the lowest type, $\hat{r} = 0$, or the highest type, $\hat{r} = 1$, and then following the recommendation.

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

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

(a) We have $\omega - r \leq 0$ for all $\omega \in \Omega = [0, 1]$ and $r = 1$. Hence, type 1’s utility is maximized by not acting for any belief about the state, implying $U_\pi(1) = 0$. This is (4) evaluated at $r = 1$.

(b) We have $\omega - r \geq 0$ for all $\omega \in \Omega = [0, 1]$ and $r = 0$. Hence, type 0’s utility is maximized by acting for any belief about the state, implying $U_\pi(0) = \mathbb{E}[\omega]$. This is (5).

Second, not every pair (q, U) that satisfies conditions (3)–(5) is feasible, that is, a mechanism π that implements such a pair need not exist. For example, if F assigns probability 1 on $\omega = 1/2$, then every nonincreasing q with $\int_0^1 q(r) dr = 1/2$ satisfies (3)–(5). Among these functions, q is feasible if and only if it satisfies $q(r) = 1$ for $r < 1/2$ and $q(r) = 0$ for $r > 1/2$.

2.3. Experiments

An experiment communicates a one-way message to the receiver and does not require the receiver to report his type. Formally, an *experiment* σ sends to the receiver a random message m from a message space $M = [0, 1]$ that depends on a realized state ω . Denote by $\sigma(m|\omega)$ the distribution of message m conditional on state ω .

A persuasion mechanism can be interpreted as a menu of experiments from which the receiver can freely choose one experiment. For a given mechanism π , the corresponding menu of experiments is $\{\sigma_{\hat{r}}\}_{\hat{r} \in R}$, where, conditional on a state ω , an experiment $\sigma_{\hat{r}}$ sends messages 1 and 0 with probabilities $\pi(\omega, \hat{r})$ and $1 - \pi(\omega, \hat{r})$. Conversely, by the revelation principle, any equilibrium outcome induced by any menu of experiments offered to the receiver can be replicated by a persuasion mechanism.

For a given experiment σ , each message m induces a posterior belief of the receiver about the state. Because the receiver’s utility is monotonic in his type, we can identify every message m with the cutoff type r who is indifferent between the two actions conditional on receiving this message. An experiment is *direct* if its messages are equal to the cutoff types, $m = \mathbb{E}[\omega|m]$. Without loss of generality, we focus on direct experiments (as in, e.g., [Kamenica and Gentzkow \(2011\)](#)).

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

$$\pi(\omega, r) \in [1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)] \quad \text{for all } \omega \in \Omega \text{ and all } r \in R, \tag{6}$$

where $\sigma(r_-|\omega)$ denotes the left limit of $\sigma(\cdot|\omega)$ at r . To understand (6), 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 $[1 - \sigma(r|\omega), 1 - \sigma(r_-|\omega)]$.

3. IMPLEMENTATION EQUIVALENCE

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

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

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

3.1. Bounds on Receiver’s Interim Utility

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

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

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

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

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

3.2. Implementable Interim Utility and Action

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

Let \mathcal{U} be the set of all convex functions bounded by \underline{U} and \overline{U} (see Figure 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 the implementation of U . The incentive-compatibility constraint requires the convexity of U by Lemma 1, and the feasibility constraint (8) requires $\underline{U}(r) \leq U(r) \leq \overline{U}(r)$. Hence, (b) \Rightarrow (a). Also, the implication (c) \Rightarrow (b) is trivial by definition. It remains to show that (a) \Rightarrow (c).

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




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

U is convex, H is nondecreasing. Hence, H is a distribution function. Next, observe that

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

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

Now, consider an experiment $\sigma(m|\omega) = P(m|\omega)$. By construction, $\mathbb{E}[\omega|m] = m$ for all m of σ , and the probability that σ generates message $m \leq x$, for any given value x , is $H(x)$. Thus, σ induces type r to act with probability $q_\sigma(r) \in [1 - H(r), 1 - H(r_-)]$, where indeterminacy arises at each discontinuity point r of H , because type r is indifferent between the two actions and can, therefore, optimally choose any lottery over these actions. Finally, (4) implies that σ implements U . Q.E.D.

Theorem 1 yields the following characterization of implementable interim actions.

COROLLARY 1: *The following statements are equivalent:*

(a) q is a nonincreasing function that satisfies

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

(b) q is implementable;

(c) q is implementable by an experiment.

The heart of the characterization result, Theorem 1, is that every implementable pair of U and q is implementable by an experiment. This result relies on the following connection between the Mirrlees (1971) representation of incentive compatibility and the Blackwell (1953) representation of garbling. Incentive compatibility (4) and feasibility (8) imply that every implementable interim utility U and action q must satisfy the requirements in parts (a) of Theorem 1 and Corollary 1. Part (a) of Corollary 1 states that F is a mean-preserving spread of $1 - q$. But any such q can be implemented by an appropriate garbling of the full disclosure experiment.⁷ Because each interim utility U is pinned down by some interim action q through (4), we obtain the result.

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

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

⁷Suppose that an experiment induces the distribution of posterior means equal to $H = 1 - q$. It is optimal for a type r to act if the posterior mean is at least r , which occurs with probability $1 - H(r) = q(r)$. Therefore, this experiment implements q .

Since $\pi(r, \omega) \in [0, 1]$ is nonincreasing and, w.l.o.g., right-continuous in r , we can interpret the function $\sigma(r|\omega) = 1 - \pi(r, \omega)$ as a distribution of messages for every given state ω . That is, σ is an experiment, and mechanism π is equivalent to experiment σ . Both Blackwell (1953) and Gutmann et al. (1991) are based on the majorization theory initiated by Hardy, Littlewood, and Pólya (1929).

3.3. Discussion

We now discuss extensions and limitations of Theorem 1, when the sender’s and receiver’s utility functions are nonlinear and the receiver has more than two actions.

Linearity in Type

The result in Theorem 1 can be easily extended to the case where the utilities are linear in (a transformation of) the state ω , but are arbitrary functions of the type r . Suppose that

$$u(\omega, r, a) = a \cdot (c(r) + b(r)d(\omega)) \quad \text{and} \quad v(\omega, r, a) = a + \rho(r)u(\omega, r, a),$$

where b, c, d , and ρ are bounded measurable functions. In addition, assume that $b(r)$ is strictly positive, and $c(r)/b(r)$ is strictly monotone.⁸ All other details of the model remain as in Section 2; in particular, the receiver chooses one of two actions, $a \in \{0, 1\}$. We now demonstrate how a change of variables gives back the problem formulated in Section 2. Define new variables

$$\tilde{\omega} = d(\omega) \quad \text{and} \quad \tilde{r} = -c(r)/b(r),$$

and let

$$\tilde{u}(\tilde{\omega}, \tilde{r}, a) = \frac{u(\omega, r, a)}{b(r)} = a \cdot \left(\frac{c(r)}{b(r)} + d(\omega) \right) = a \cdot (\tilde{\omega} - \tilde{r}).$$

The utility of each type r is now multiplied by the factor $1/b(r)$, which has no effect on the optimal choice of action $a \in \{0, 1\}$ for this type. Next, define $\tilde{\rho}(\tilde{r}) = \rho(r)b(r)$ and let

$$\tilde{v}(\tilde{\omega}, \tilde{r}, a) = v(\omega, r, a) = a + \rho(r)b(r) \frac{u(\omega, r, a)}{b(r)} = a + \tilde{\rho}(\tilde{r})\tilde{u}(\tilde{\omega}, \tilde{r}, a).$$

Thus, after this change of variables, the sender’s and receiver’s utilities have the functional forms identical to (1).

Linearity in State

The linearity of the utilities in (a transformation of) the state implies that a posterior distribution of the state can be summarized into a scalar, the posterior mean. For any message of any experiment, the receiver acts if and only if his type is below the cutoff equal to the posterior mean state conditional on that message. An experiment induces a probability distribution over these cutoffs. A persuasion mechanism, which is a menu of experiments, induces a compound distribution over cutoffs, which is itself an experiment.

Without the linearity in the state, a posterior distribution of the state cannot be summarized into a scalar. Even if the receiver’s utility is increasing in the state and decreasing

⁸These assumptions are unnecessary for our results and are made for elegance of exposition. A more general case of utility functions linear in ω is analyzed in Appendix C of the Supplemental Material.

in the type, different types of the receiver may have different attitudes to risk. Thus, there may be a value of screening types by offering the receiver a menu of experiments.

For example, consider an experiment that sends one message if $\omega \geq \omega'$ and another message otherwise. For a type r' that is exactly indifferent between $a = 0$ and $a = 1$ upon learning that $\omega \geq \omega'$, this experiment leaves type r' with zero surplus and uniquely maximizes the probability that type r' acts. Under this experiment, a nearby type $r'' < r'$ also acts if and only if $\omega \geq \omega'$. Without linearity, it is possible to design another experiment, preferred by r'' but not by r' , which induces type r'' to act with a different probability. Thus, the persuasion mechanism corresponding to the menu of these two experiments implements the interim actions that cannot be implemented by any experiment.

In Appendix B of the Supplemental Material (Kolotilin et al. (2017)), we formulate the model in which the utilities are nonlinear in the state, characterize persuasion mechanisms that are equivalent to experiments, and demonstrate that the equivalence of implementation by persuasion mechanisms and by experiments generally fails.

Binary Action Space

The assumption that the receiver chooses one of *two* actions is crucial for Theorem 1 to hold. Relaxing this assumption will affect the result in two ways.

First, the number of actions affects the complexity of the obedience constraints. With two actions, there are only two relevant ways to disobey a mechanism, by ignoring the mechanism's message and either always act or never act. However, the outcomes of these deviations can be achieved by reporting the extreme types, $\hat{r} = 0$ and $\hat{r} = 1$, which always act and never act, respectively. Thus, the incentive to report the type truthfully automatically implies obedience. With more than two actions, there are more ways to disobey than just ignore the messages. The obedience constraints are nontrivial and are not implied by the truth telling. Consequently, the implication (a) \Rightarrow (b) of Theorem 1 (that every convex function between \underline{U} and \bar{U} is implementable) is false, because the obedience constraints impose additional necessary conditions for implementability.

Second, with more than two actions, there exist outcomes implementable by persuasion mechanisms but not by experiments. For example, suppose that there are three actions available to the receiver, $A = \{0, 1, 2\}$, and consider a persuasion mechanism described by a menu of experiments that send binary messages, "high" or "low". Under this mechanism, each type of the receiver will choose a preferred experiment from the menu, and then choose one of two best-response actions to the messages. Generally, this mechanism is not implementable by an experiment. If an experiment sends a binary message (the same for all types), then, generally, there will be types who would have chosen a different experiment from the menu and would have acted differently under the original persuasion mechanism. If, instead, an experiment sends multiple messages, then, generally, there will be types that have more than two best-response actions to those messages, thus behaving differently than under the original mechanism. Consequently, the implication (b) \Rightarrow (c) of Theorem 1 (that every implementable U is implementable by an experiment) does not hold.

In Appendix A of the Supplemental Material, we formulate the model with multiple actions of the receiver, characterize the implementable receiver's interim utilities, and show that the sender can generally implement a strictly larger set of the receiver's interim utilities and actions by persuasion mechanisms than by experiments.

4. OPTIMAL MECHANISMS

In this section, we use Theorem 1 to characterize the persuasion mechanisms that are optimal for the sender under additional smoothness assumptions. We assume that the weight ρ in the sender’s utility is continuous in the receiver’s type and that the distribution of types G admits a continuously differentiable density g .

4.1. Sender’s Problem

The sender seeks an incentive-compatible persuasion mechanism π that maximizes

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

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

LEMMA 2: For every incentive-compatible mechanism π ,

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

where

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

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

PROOF: Observe that, by (4) and (7), we have

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

Using integration by parts, we obtain

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

Since $-U_\pi(r)g(r)|_0^1 = \mathbb{E}[\omega]g(0)$ by (5), the lemma follows.

Q.E.D.

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

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

We say that U is *optimal* if it solves the above problem. For a given optimal U , the corresponding optimal experiment is defined by the distribution H of the posterior mean state,

$$H(r) = 1 + U'(r_+) \quad \text{for all } r \in [0, 1].$$

4.2. Structure of Optimal Mechanisms

We characterize the structure of optimal mechanisms under the assumption that function I is nonzero almost everywhere and changes sign $n \geq 0$ times.⁹ Let $\{r_1, r_2, \dots, r_n\} \subset (0, 1)$ be the set of types at which I changes its sign, and let $r_0 = 0$ and $r_{n+1} = 1$.

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

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

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

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

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

Therefore, we can reduce the sender's problem (11) of optimization on the function space \mathcal{U} to an n -variable optimization problem. The optimal U is pinned down by properties (P₁) and (P₂) within each interval (r_i, r_{i+1}) where the sign of I is constant. Thus, the




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

⁹If there are intervals where $I(r) = 0$, then on those intervals the sender is indifferent about the choice of U . Hence, multiple solutions emerge in this case. Characterization of these solutions is a straightforward but tedious extension of the result in this section.




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

optimal U is fully defined by the utility values at the points $\{r_1, \dots, r_n\}$ where I changes its sign.

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

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

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

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

EXAMPLE 1: Suppose that $I(r) = g'(r) + \rho(r)g(r)$ crosses the horizontal axis at most once and from above. This is the case, for instance, if the density $g(r)$ is single-peaked and $\rho(r) = 0$. Then, Theorem 2 implies that there exists a threshold ω^* such that the optimal experiment reveals the state if it is below ω^* and sends the same pooling message for all states above ω^* .

5. CONCLUSION

In this paper, we study the model of persuasion of a privately informed receiver in linear environments, and show the equivalence of persuasion by experiments and by persuasion mechanisms. A privately informed receiver can be interpreted as a random draw from a heterogeneous population of receivers. Within this interpretation, a persuasion

mechanism establishes a private communication channel with each individual in the population; it asks each individual to report his type, and returns a private message about the state. In contrast, an experiment broadcasts the same message to the entire population. Our equivalence result can be viewed as the equivalence between persuasion via private communication to receivers and by broadcasting information to the whole population. In particular, so long as the receivers' types are unobserved by the sender, there is no loss of optimality in publicity of messages and free information sharing among the receivers.

Within the above interpretation, it would be natural to assume that the receivers' payoffs depend not only on their own actions, but also on collective actions in the population, thus being affected by network externalities. This extension is beyond the scope of this paper and it is left open for future research.

APPENDIX: PROOF OF LEMMA 1

Necessity. Incentive compatibility (2) requires that for each $\hat{r} > r$, both r and \hat{r} prefer truth telling

$$U_\pi(r) \geq U_\pi(r, \hat{r}) = U_\pi(\hat{r}) + q_\pi(\hat{r})(\hat{r} - r),$$

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

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

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

Also, observe that type $r = 1$ can secure his maximal attainable utility of 0 by always acting (irrespective of a recommendation); so $U_\pi(1) = 0$ and (4) follows. Finally, the maximal attainable utility of type $r = 0$ is $\mathbb{E}[\omega]$, which can be secured by never acting; so (5) follows.

Sufficiency. It remains to show that (3)–(5) imply (2). If either $\tilde{r} \geq \hat{r} \geq r$ or $\tilde{r} \leq \hat{r} \leq r$, then (3) and (4) imply that

$$U_\pi(r, \hat{r}) = U_\pi(r) + q_\pi(\hat{r})(\hat{r} - r) = \int_{\hat{r}}^1 q_\pi(s) ds + q_\pi(\hat{r})(\hat{r} - r)$$

$$\geq \int_{\hat{r}}^{\tilde{r}} q_\pi(\tilde{r}) ds + \int_{\tilde{r}}^1 q_\pi(s) ds + q_\pi(\tilde{r})(\hat{r} - r)$$

$$= \int_{\tilde{r}}^1 q_\pi(s) ds + q_\pi(\tilde{r})(\tilde{r} - r) = U_\pi(r, \tilde{r}),$$

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

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

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

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

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

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

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

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

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

REFERENCES

- ALONSO, R., AND O. CÂMARA (2016a): “Bayesian Persuasion With Heterogeneous Priors,” *Journal of Economic Theory*, 165, 672–706. [1951]
- (2016b): “Persuading Voters,” *American Economic Review*, 106, 3590–3605. [1951]
- (2017): “On the Value of Persuasion by Experts,” Working Paper, University of Southern California. [1951]
- AU, P. H. (2015): “Dynamic Information Disclosure,” *RAND Journal of Economics*, 46, 791–823. [1951]
- BERGEMANN, D., AND S. MORRIS (2013): “Robust Predictions in Games With Incomplete Information,” *Econometrica*, 81, 1251–1308. [1951]
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522. [1951]
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609. [1951]
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2017): “The Design and Price of Information,” *American Economic Review* (forthcoming). [1951]
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957. [1951]
- BERGEMANN, D., B. A. BROOKS, AND S. MORRIS (2017): “First Price Auctions With General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143. [1951]
- BLACKWELL, D. (1951): “Comparison of Experiments,” in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, 93–102. [1950]
- (1953): “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 24, 265–272. [1950,1956,1957]
- BOARD, S., AND J. LU (2017): “Competitive Information Disclosure in Search Markets,” *Journal of Political Economy* (forthcoming). [1951]
- ELY, J., A. FRANKEL, AND E. KAMENICA (2015): “Suspense and Surprise,” *Journal of Political Economy*, 103, 215–260. [1951]
- ELY, J. C. (2017): “Beeps,” *American Economic Review*, 107 (1), 31–53. [1951]
- ESÖ, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705–731. [1951]
- GENTZKOW, M., AND E. KAMENICA (2014): “Costly Persuasion,” *American Economic Review, Papers & Proceedings*, 104, 457–462. [1951]
- (2016): “A Rothschild–Stiglitz Approach to Bayesian Persuasion,” *American Economic Review, Papers & Proceedings*, 106, 597–601. [1950]
- (2017a): “Bayesian Persuasion With Multiple Senders and Rich Signal Spaces,” *Games and Economic Behavior*, 104, 411–429. [1951]
- (2017b): “Competition in Persuasion,” *Review of Economic Studies*, 84, 300–322. [1951]
- (2017c): “Disclosure of Endogenous Information,” *Economic Theory Bulletin*, 5, 47–56. [1951]
- 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. [1951,1956]
- GILL, D., AND D. SGROI (2008): “Sequential Decisions With Tests,” *Games and Economic Behavior*, 63, 663–678. [1951]
- (2012): “The Optimal Choice of Pre-Launch Reviewer,” *Journal of Economic Theory*, 147, 1247–1260. [1951]

- 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. [1956,1957]
- HARDY, G. H., J. E. LITTLEWOOD, AND G. PÓLYA (1929): "Some Simple Inequalities Satisfied by Convex Functions," *Messenger of Mathematics*, 58, 145–152. [1951,1957]
- HEDLUND, J. (2017): "Bayesian Persuasion by a Privately Informed Sender," *Journal of Economic Theory*, 167, 229–268. [1951]
- HÖRNER, J., AND A. SKRZYPACZ (2016): "Selling Information," *Journal of Political Economy*, 124 (6), 1515–1562. [1951]
- KAMENICA, E., AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615. [1950,1954]
- KOLOTILIN, A. (2015): "Experimental Design to Persuade," *Games and Economic Behavior*, 90, 215–226. [1951]
- (2017): "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics* (forthcoming). [1951]
- KOLOTILIN, A., T. MYLOVANOV A. ZAPECHELNYUK, M. LI (2017): "Supplement to 'Persuasion of a Privately Informed Receiver'," *Econometrica Supplemental Material*, 85, <http://dx.doi.org/10.3982/ECTA13251>. [1952, 1958]
- LI, F., AND P. NORMAN (2017): "On Bayesian Persuasion With Multiple Senders," Working Paper, University of North Carolina at Chapel Hill. [1951]
- LI, H., AND X. SHI (2017): "Discriminatory Information Disclosure," *American Economic Review*, 107, 3363–3385. [1951]
- MANELLI, A., AND D. VINCENT (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 78, 1905–1938. [1951,1956]
- MATHEVET, L., J. PEREGO, AND I. TANEVA (2017): "On Information Design in Games," Working Paper, New York University. [1951]
- MIRRLEES, J. A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175–208. [1956]
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. [1951, 1959]
- OSTROVSKY, M., AND M. SCHWARZ (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2, 34–63. [1951]
- PEREZ-RICHET, E. (2014): "Interim Bayesian Persuasion: First Steps," *American Economic Review, Papers & Proceedings*, 104, 469–474. [1951]
- RAYO, L., AND I. SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987. [1951]
- TANEVA, I. (2016): "Information Design," Working Paper, University of Edinburgh. [1951]

Co-editor Joel Sobel handled this manuscript.

Manuscript received 23 February, 2015; final version accepted 26 June, 2017; available online 28 June, 2017.