ON THE EQUIVALENCE OF INFORMATION DESIGN BY UNINFORMED AND INFORMED PRINCIPALS

ANDRIY ZAPECHELNYUK

ABSTRACT. We compare information design, or Bayesian persuasion, by an uninformed principal (who has no information about the state of the world when making her choice) and by an informed principal (who has private information and can condition her choice on that information). We show that, under the assumptions of monotone preferences of the principal and nondegenerate information structures, a Pareto undominated outcome is implementable by the uninformed principal if and only if it is implementable by the informed principal.

JEL Classification: D82, D83

Keywords: Information design, Bayesian persuasion, informed principal, implementation, sequential equilibrium

---

Date: February 24, 2023.

Zapechelnyuk: School of Economics, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK. *E-mail:* azapech@gmail.com.

The author thanks Dirk Bergemann, Simon Board, Anton Kolotilin, Eduardo Perez-Richet, and anonymous referees for very helpful comments.

1. Introduction

The literature on information design, or Bayesian persuasion, traditionally assumes that the principal, or information designer, commits to an information disclosure protocol without any private knowledge about what she is about to disclose. While this feature is plausible in a variety of contexts, it is often more realistic to consider an alternative, where the principal may possess private information and use it to her advantage when deciding how the information should be disclosed.[1] We will refer to the former and latter settings as, respectively, the uninformed and informed principal models.

In general, the informed principal can implement fewer outcomes than the uninformed one. This is because the informed principal has to make sequentially rational choices given her private information, whereas the uninformed principal has no such constraint. We are interested in the conditions when this sequential rationality constraint entails no loss of generality, namely, when the uninformed principal's information disclosure protocol can be sustained as a sequentially rational play for the informed principal.

To illustrate the central idea of this paper, consider an example with a plaintiff (principal, she) and a judge (agent, he), as in Kamenica and Gentzkow (2011). Based on presented evidence, the judge chooses whether to rule in favor or against the plaintiff's case. In litigation, especially in civil lawsuits, usually there is a specific and detailed procedure, or practice direction, that explains the conduct and sets out the steps the court normally expects the plaintiff to follow. It is plausible that the plaintiff is privately informed about the evidence before presenting it to the judge. So the plaintiff might have an incentive to alter the procedure in some way, depending on her information. What can be done to deter such deviations? A reasonable answer is that a deviation from the procedure may raise the judge's suspicion that he is being manipulated. Provided there is no incriminating evidence, so there remains some uncertainty about the truth no matter what the plaintiff discloses, any alteration of the procedure could predispose the judge's against the plaintiff's case, so much that this change in the disposition dominates the informational benefit for the plaintiff. As a result, the plaintiff's sequentially rational choice is to adhere to the procedure irrespectively of her private information.

[1]For literature surveys see Bergemann and Morris (2019) and Kamenica (2019), and for an outline of questions of interest in Bayesian persuasion, including the theme of private information of a sender, see Kamenica, Kim, and Zapechelnyuk (2021).

In this paper, we show that, under certain assumptions, every implementable Pareto undominated outcome in the uninformed principal model is implementable as a sequential equilibrium in the informed principal model. This equilibrium is pooling, in the sense that the informed principal chooses the same information disclosure protocol irrespective of her private information. Coupled with the observation that every outcome implementable in the informed principal model is also implementable in the uninformed principal model[2], we draw the conclusion about the equivalence of implementation of Pareto undominated outcomes by means of information design in these two models. A notable consequence of this result is the optimality equivalence: the mechanism that induces the optimal sequential equilibrium in the informed principal's problem can be found by solving the uninformed principal's problem.[3]

Our result holds under two assumptions. The first assumption states that the principal has *monotone* preferences over the agent's actions. Specifically, there exists an order over the agent's actions along which the principal's utility is increasing irrespective of the state of the world. This assumption includes state-independent preferences of the principal as a special case. The consequence of this assumption is that there exists an agent's belief that leads to a state-independent "punishment" action. This is an action that, regardless of the state, is inferior to every action inducible in Pareto undominated information disclosure in the uninformed principal model. This punishment is used to deter the principal's deviations conditional on learning the state.

The second assumption states that information structures, referred to as *tests*, that are available to the principal cannot be absolutely accurate. That is, no test can make the agent absolutely certain about the state of the world, although it can be arbitrarily close to providing this certainty. We use this assumption to ensure that the agent's posterior beliefs conditional on tests and their messages are defined by Bayes' rule for any prior. This assumption entails no loss of generality in the uninformed principal model, because we define outcomes to be implementable if they can be induced in the limit by a convergent sequence of tests. However, this assumption is substantive in the informed principal model, because it prevents the principal to deviate to absolutely accurate tests, thus imposing a refinement on the set of equilibria.

[2]This observation is reminiscent of the inscrutability principle of Myerson (1983). In the context of information design by informed principal, this observation was first made by Perez-Richet (2014) in a setting with two states and two actions.

[3]Note that the informed principal model is a signaling game that generally has multiple sequential equilibria. So our argument of optimality equivalence presumes that the equilibrium selection can be made in favor of the principal.

Related Literature. The closest paper in the literature to our paper is Koessler and Skreta (2022), thereafter, KS. Like our paper, KS compare the problems of uninformed and informed information designer, but they reach a different conclusion. KS show that an equilibrium in the informed principal model that implements the optimal outcome for the uninformed principal need not exist. This is because KS impose a specific, albeit natural and commonly accepted, constraint on the agent's out-of-equilibrium beliefs. Under this constraint, the beliefs that induce the "punishment" action, which is used to deter the principal's deviations in our setting, need not be feasible. Our paper adopts a complementary approach. We do not impose constraints on the agent's out-of-equilibrium beliefs. Instead, we assume that tests are never absolutely accurate. This difference between KS and our paper is illustrated by example in Section 4. Another difference is that our model has more structure due to the assumption on the principal's preferences. KS make no such assumption, so a state-independent punishment that deters the principal's deviations need not exist in KS's setting.

A few other papers study the model of information design by an informed principal. Perez-Richet (2014) and Degan and Li (2021) consider a more specialized setting with two states and two actions. Hedlund (2017) analyzes the setting with two states and multiple actions, where the principal is partially informed about the state. Applying the D1 equilibrium refinement criterion, Hedlund (2017) shows that the resulting outcome either fully reveals the principal's private information about the state, or fully reveals the state itself. Chen and Zhang (2020) consider an interaction between a privately informed seller and a potential buyer. They allow the seller to communicate her type to the buyer via two channels, information disclosure and pricing, and show that a credible type separation is generally impossible via one channel alone. Bizzotto and Vigier (2021) study Bayesian persuasion over multiple periods with exogenous news, where the sender is unable to commit to the information that she will supply in future periods. Lastly, Serena (2022) study a model with an informed principal designs information disclosure to maximizes the aggregate effort of two contestants who compete in a Tullock contest.[4]

The papers mentioned above allow the principal to choose arbitrary information structures. Alonso and Câmara (2018) restrict the principal's choice of tests to a given set and characterize the conditions when the uninformed principal can benefit from accessing additional information about the state. There is also a substantial literature

[4]In a related paper, Antsygina and Teteryatnikova (2022) study information design in contests with uninformed principal.

on information disclosure with evidence by an informed principal that restricts the use of information structures, in particular, focusing on deterministic and partitional ones. Prominent papers in this literature include Okuno-Fujiwara, Postlewaite, and Suzumura (1990), Seidmann and Winter (1997), Hagenbach, Koessler, and Perez-Richet (2014), Hart, Kremer, and Perry (2017), Ben-Porath, Dekel, and Lipman (2019), and Ivanov (2021) to name a few.

Our paper is also is related to the literature on the informed principal in the standard mechanism design setting that was set in motion by the seminal paper of Myerson (1983). Some of this literature touches upon the question of information disclosure of the principal's private type to agents, particularly focusing on when full disclosure does not hurt the principal (e.g., Maskin and Tirole, 1990; Yilankaya, 1999; Skreta, 2011; Mylovanov and Tröger, 2014; Bedard, 2017; Mekonnen, 2021).

One can interpret the principal's ability to change her mind after observing private information as her lack of commitment. In this sense, our paper is related to the literature that investigates information design where the principal chooses an information structure without full commitment to its messages (Guo and Shmaya, 2021; Lipnowski, Ravid, and Shishkin, 2022; Min, 2021; Eilat and Neeman, 2023). In these models, failure to commit means cheap talk.[5] In contrast, in our informed principal model, the principal can still credibly communicate information. This is because the principal does not freely choose a message after learning the state, instead she publicly commits to a state contingent disclosure mechanism that communicates messages on behalf of the principal.

2. Model

2.1. **Preliminaries.** Consider a setting with two players, a principal (she) and an agent (he), whose utilities depend on the agent's action a and the state of the world θ, and are given by $u_P(a,\theta)$ and $u_A(a,\theta)$, respectively. The set of actions A and the set of states Θ are finite. There is a common prior $q_0 \in \Delta(\Theta)$ about the state.

We assume that the principal has *monotone* preferences over the agent's actions. Specifically, the principal's ordinal comparison of any pair of actions does not change with the state, so $u_P(a,\theta)$ satisfies

$$u_P(a',\theta') \geq u_P(a'',\theta') \iff u_P(a',\theta'') \geq u_P(a'',\theta'') \\ \text{for all } a',a'' \in A \text{ and all } \theta',\theta'' \in \Theta. \tag{A$_1$}$$

[5]Credible communication by cheap talk when the sender has state-independent preferences is characterized by Lipnowski and Ravid (2020).

Under this assumption, without loss of generality, actions in A can be ordered so that the principal's utility is increasing in the action. Assumption (A_1) includes a special case where the principal's utility is state-independent, so $u_P(a,\theta) = u_P(a)$.[6]

In addition, to simplify the exposition, we assume that in each state the agent has a single optimal action, so

$$\text{for each } \theta \in \Theta \text{ there exists } a^*_\theta \in A \text{ s.t. } u_A(a^*_\theta, \theta) > u_A(a,\theta) \text{ for all } a \neq a^*_\theta. \tag{1}$$

The agent is initially uninformed about the state. He receives information about the state via a test designed by the principal. Let M be a set of messages, with at least as many messages as actions in A. A *test* t is a conditional probability distribution that sends each message $m \in M$ with probability $t(m|\theta)$ when the realized state is $\theta \in \Theta$.

We consider *nondegenerate tests*. Formally, the set of nondegenerate tests, denoted by T, is the set of all conditional probability distributions $t(\cdot|\theta)$ over M that satisfy

$$\left(\sum_{\theta\in\Theta} t(m|\theta) > 0 \implies t(m|\theta) > 0 \;\text{ for each } \theta \in \Theta \right) \;\text{ for each } m \in M. \tag{A_2}$$

Assumption (A_2) means that every message m either cannot occur under t at all, or it occurs with a positive probability in every state. This assumption captures the idea that no test can make the agent absolutely certain about the state, although tests can be arbitrarily close to providing this certainty. It is also necessary and sufficient for the agent's posterior beliefs conditional on messages of the test to be defined by Bayes' rule for all priors.

After having observed a test $t \in T$ and a message m generated by that test, the agent forms a posterior belief about the state, denoted by $\beta(\cdot|t,m) \in \Delta(\Theta)$, according to Bayes' rule whenever possible.[7] Given a posterior, the agent chooses an action that maximizes his expected utility using an exogenously given decision rule d that satisfies

$$d(q) \in \arg\max_{a\in A} \sum_{\theta\in\Theta} u_A(a,\theta)q(\theta) \;\text{ for each } q \in \Delta(\Theta). \tag{2}$$

[6] Assumption (A_1) can be relaxed. It is sufficient to assume that there exists a "punishment" action that is worst for the principal among those actions that are optimal for the agent in some state. We further comment on this in Remark 2 (Section 3).

[7] In fact, the agent updates his prior twice, first after seeing test t, and second after seeing message m generated by the test. But because the agent only acts after seeing both t and m, all that matters is the final posterior $\beta(\cdot|t,m)$.

Rule d is used for tie breaking whenever a utility maximizing action is not unique. For example, as often assumed in the literature, the ties can be resolved in favor of the principal.

We now describe the principal's information and behavior. We consider two settings, one where the principal is uninformed about the state and one where the principal is informed.

2.2. **Uninformed Principal.** We first consider the standard Bayesian persuasion setting as in Kamenica and Gentzkow (2011). In this setting, the principal is uninformed about the state when choosing a test $t \in T$.

Given a test $t \in T$, the interaction proceeds as follows. First, Nature draws a state θ from Θ according to the prior q_0. Then, test t produces a message $m \in M$ according to the conditional probability distribution $t(\cdot|\theta)$. Finally, the agent observes the test t and its message m, forms posterior belief $\beta(\cdot|t, m)$, and chooses action $a = d(\beta(\cdot|t, m))$.

Every test $t \in T$ induces a conditional probability distribution over the agent's choices of actions, denoted by λ_t. The probability $\lambda_t(a|\theta)$ of action a conditional on state θ is given by

$$\lambda_t(a|\theta) = \sum_{m \in M} t(m|\theta)\mathbf{1}_{\{a=d(\beta(\cdot|t,m))\}}, \tag{3}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. We will refer to λ_t as the *outcome* induced by t.

An outcome is sufficient to describe the interim utilities of the principal and the agent. Given an outcome λ, for each $\theta \in \Theta$, these interim utilities are given by

$$U_i(\theta; \lambda) = \sum_{a \in A} u_i(a, \theta)\lambda(a|\theta), \quad i = P, A. \tag{4}$$

An outcome is implementable by the uninformed principal if it is approachable by outcomes of tests in T.

Definition 1. An outcome $\lambda \in (\Delta(A))^{|\Theta|}$ is *implementable by the uninformed principal* if there exists a sequence of tests $(t_k)_{k \in \mathbb{N}}$ such that $t_k \in T$ for each $k \in \mathbb{N}$, and $\lim_{k\to\infty} \lambda_{t_k} = \lambda$.[8]

2.3. **Informed Principal.** We now consider the informed principal setting. In this setting, the principal is privately informed about the state and designs a test that

[8]Here and elsewhere in the paper, the convergence is pointwise.

depends on her information. The principal's *mechanism* $\tau : \Theta \to \Delta(T)$ specifies a probability distribution $\tau(\cdot|\theta)$ over the set of tests conditional on each state θ.

Consider a pair (τ, β) of the principal's mechanism τ and the agent's system of posterior beliefs β that specifies a posterior belief $\beta(\cdot|\hat{t}, m) \in \Delta(\Theta)$ for each test $\hat{t} \in T$ and each message $m \in M$. Given a pair (τ, β), the interaction proceeds as follows. First, Nature draws a state θ from Θ according to the prior q_0. Then, the mechanism draws a test $t \in T$ according to the conditional distribution $\tau(\cdot|\theta)$. Next, the test produces a message $m \in M$ according to the conditional distribution $t(\cdot|\theta)$. Finally, the agent observes the realized test t and its message m, forms posterior belief $\beta(\cdot|t, m)$, and chooses action $a = d(\beta(\cdot|t, m))$.

Analogously to the uninformed principal setting, every pair (τ, β) induces a conditional probability distribution over the agent's choices of actions, denoted by $\lambda_{(\tau,\beta)}$. The probability $\lambda_{(\tau,\beta)}(a|\theta)$ of action a conditional on state θ is given by

$$\lambda_{(\tau,\beta)}(a|\theta) = \int_{t\in T} \left(\sum_{m\in M} t(m|\theta) \mathbf{1}_{\{a=d(\beta(\cdot|t,m))\}} \right) \tau(\mathrm{d}t|\theta). \tag{5}$$

We will refer to $\lambda_{(\tau,\beta)}$ as the *outcome* induced by (τ, β). As in the uninformed principal setting, an outcome is sufficient to describe the interim utilities of the principal and agent, which are given by (4).

Our notion of implementation by the informed principal is based on sequential equilibrium (Kreps and Wilson, 1982). We rule out the agent's out-of-equilibrium beliefs that cannot withstand small perturbations of the principal's mechanism.

Let $\varepsilon > 0$. Given a belief system β, a mechanism τ is called *ε-sequentially rational under β* if the principal cannot improve her interim expected utility by more than ε in any state, so

$$U_P(\theta; \lambda_{(\tau,\beta)}) \geq \sup_{t\in T} \left(\sum_{m\in M} u_P(d(\beta(\cdot|t, m), \theta) t(m|\theta) \right) - \varepsilon \quad \text{for each } \theta \in \Theta. \tag{6}$$

A mechanism τ is called a *full-support mechanism* if conditional distribution $\tau(\cdot|\theta)$ has full support on T for each $\theta \in \Theta$.

Given a mechanism τ, a belief system β is called *consistent with* τ if there exists a sequence of full-support mechanisms $(\tau_k)_{k\in\mathbb{N}}$ and the corresponding sequence of belief systems $(\beta_k)_{k\in\mathbb{N}}$ derived by Bayes' rule such that $\lim_{k\to\infty}(\tau_k, \beta_k) = (\tau, \beta)$.

Definition 2. A pair (τ, β) is an ε-sequential equilibrium if τ is ε-sequentially rational under β, and β is consistent with τ.[9]

An outcome is implementable by the informed principal if it is approachable by outcomes of ε-sequential equilibria with arbitrarily small ε.

Definition 3. An outcome $\lambda \in (\Delta(A))^{|\Theta|}$ is *implementable by the informed principal* if there exists a sequence $(\varepsilon_k, \tau_k, \beta_k)_{k\in\mathbb{N}}$ such that (i) for each $k \in \mathbb{N}$, $\varepsilon_k > 0$ and (τ_k, β_k) is ε_k-sequential equilibrium, and (ii) $\lim_{k\to\infty} \varepsilon_k = 0$ and $\lim_{k\to\infty} \lambda_{(\tau_k,\beta_k)} = \lambda$.

3. Result

Before presenting our result, we introduce the notion of Pareto dominance.

Definition 4. Consider the uninformed principal model. An implementable outcome λ is *Pareto undominated* if there is no implementable outcome that is weakly preferred to λ by both principal and agent in each state, and strictly so by at least one of them in some state.

Theorem 1. *(a) If an outcome is implementable by the informed principal, then it is implementable by the uninformed principal.*[10]

(b) Let Assumptions (A_1) *and* (A_2) *hold. If an outcome is implementable by the uninformed principal and Pareto undominated, then it is implementable by the informed principal.*

The proof is in the Appendix. Here we provide the intuition for the result, and in the next section we will illustrate it by an example.

Part (a) follows from an application of the sure-thing principle.[11] If the agent chooses the same optimal action in two distinct events, then he should choose the same action without knowing which of those events has occurred. So one can bundle together all pairs (t, m) that lead to the same action a, and identify all these pairs with a single message that "recommends" action a. We would like to point out that Part (a) is not a trivial statement that always holds. It relies on the richness of the set of tests

[9]This solution concept corresponds to Myerson and Reny's (2020) *perfect conditional ε-equilibrium* who extend sequential equilibrium (Kreps and Wilson, 1982) to infinite sets of signals and actions.

[10]Note that we do not impose Assumptions (A_1) and (A_2) in part (a). Assumption (A_1) plays no role for this result. If we restrict attention to the tests that satisfy Assumption (A_2), part (a) continues to hold as we show in Remark 3 (see the Appendix).

[11]The idea behind the statement in Part (a) is not novel. A version of this result for the model with two states and two actions appears in Perez-Richet (2014).

available to the principal and need not be true if the set of tests is restricted (see Alonso and Câmara, 2018).

The intuition for Part (b) is as follows. Let λ^* be a Pareto undominated outcome implementable by the uninformed principal. Suppose that there is a test, t^*, that implements λ^*. This test can be replicated within the informed principal setting by the mechanism $\bar{\tau}$ that prescribes the same test t^* independently of θ. The problem is that $\bar{\tau}$ might not be sequentially rational, nor even ε-sequentially rational for a small enough ε. For example, the principal might have an incentive to deviate by choosing a highly accurate test that nearly reveals the state.

As the key part of the proof, we construct a sequential equilibrium with out-of-equilibrium beliefs for the agent that deter the principal's deviations from the prescribed test. Specifically, whenever the principal deviates from the prescribed test t^* to any different test $\hat{t}$, the agent (who observes this deviation) becomes "skeptical" and forms a posterior belief that assigns probability one to a specific "punishment" state. This "punishment" state induces an action of the agent that hurts the principal no matter what the state is, referred to as the "punishment" action. Assumption (A_1) ensures that the same action is the worst for the principal in all states. Assumption (A_2) ensures that messages of the deviation test $\hat{t}$ cannot alter the agent's degenerate posterior belief that the state is equal to the "punishment" state. The condition that λ^* is Pareto undominated ensures that there actually exists such a degenerate posterior belief under which the agent optimally chooses the principal's "punishment" action (which is, by definition, is weakly inferior to λ^* for the principal).

Lastly, the described belief system satisfies the consistency requirement of sequential equilibrium, because it is obtained as the limit of a sequence of perturbed mechanisms constructed as follows. With a probability that approaches zero, instead of choosing t^*, the perturbed mechanism chooses a full-support lottery over tests. The probability of choosing this lottery is by the order of magnitude larger in the "punishment" state than in all other states. Thus, whenever the principal deviates from the prescribed test t^* to any different test $\hat{t}$, the agent's posterior probability of the "punishment" state approaches one as the perturbation vanishes.

Remark 1. The restriction to Pareto undominated outcomes stems from the requirement of consistency of out-of-equilibrium beliefs in sequential equilibrium. If the solution concept was prefect Bayesian equilibrium (PBE), so that out-of-equilibrium beliefs could be arbitrary, then, under assumptions (A_1) and (A_2), *all* outcomes implementable by the uninformed principal would be sustainable in PBE of the informed principal game.

Remark 2. The proof of Theorem 1(b) relies on the existence of the agent's action that punishes the principal's deviations uniformly in all states, thus allowing to sustain pooling equilibria, where the principal chooses the same test in all states. Assumption (A_1) is sufficient for the existence of such a uniform punishment, but not necessary. It can be relaxed as follows. Let $\tilde{A}(\theta)$ be the set of actions that are optimal for the agent in state θ, and let $\tilde{A} = \bigcup_{\theta\in\Theta} \tilde{A}(\theta)$.[12] For the proof of Theorem 1(b), we only need to assume that there exists an action $a_* \in \tilde{A}$ that is the worst for the principal among all actions in $\tilde{A}$ in each state, so $a_* \in \arg\min_{a\in\tilde{A}} u_P(a,\theta)$ for each $\theta \in \Theta$.

4. Example

In this section, we present an example borrowed from Koessler and Skreta (2022). The role of this example is twofold. First, it illustrates Theorem 1 and shows how an equivalent pooling equilibrium in the informed principal setting is constructed for a given optimal information design by the uninformed principal. Second, it shows why the assumption that tests cannot be perfectly accurate (Assumption (A_2)) is crucial for this construction, and why relaxing this assumption (and thus enlarging the set of tests that the principal can deviate to) can destroy the pooling equilibrium. It also highlights the difference between Koessler and Skreta (2022) and this paper.

Consider the following example. There are three actions, $A = \{l, m, h\}$, and two states, $\Theta = \{L, H\}$. The prior probability of state H is denoted by q_0 and is given by $q_0 = 1/6$. Table 1 shows the players' utilities, where each pair of numbers presents the utilities of the principal and agent, respectively.

	l	m	h
L	$0,3$	$2,2$	$3,0$
H	$0,0$	$2,2$	$3,3$

Table 1. Utilities of Principal and Agent

Note that the principal's utility is state-independent, $u_P(a,\theta) = u_P(a)$. Her preference over actions is $l \prec m \prec h$. Given a posterior probability q of state H, the agent's

[12] Note that $\tilde{A}$ is not the same as the set of actions that are rationalizable, as there can be actions that are rationalizable, i.e., optimal under some belief about state, and yet not optimal in any state.

preferred action $d(q)$ is given by

$$d(q) = \begin{cases} l, & \text{if } q < 1/3, \\ m, & \text{if } 1/3 \leq q < 2/3, \\ h, & \text{if } q \geq 2/3, \end{cases}$$

with ties resolved in favor of the principal.

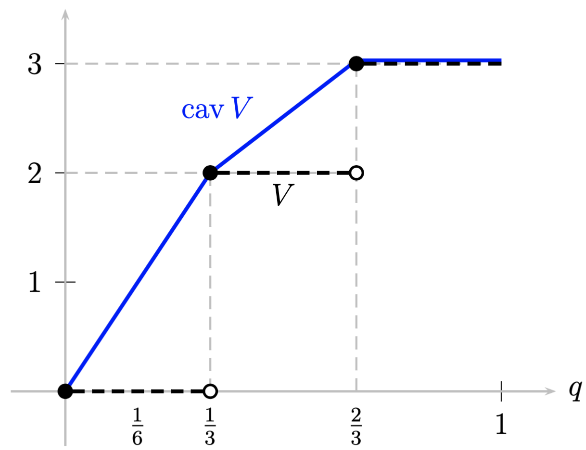

FIGURE 1. The principal's expected utility V and its concavification cav V.

The concavification method of Kamenica and Gentzkow (2011) allows us to find the optimal test for the uninformed principal. Provided the agent chooses actions according to d, the principal's expected utility $V(q)$ as a function of the agent's posterior probability q of state H is given by

$$V(q) = u_P(d(q)) = \begin{cases} 0, & \text{if } q < 1/3, \\ 2, & \text{if } 1/3 \leq q < 2/3, \\ 3, & \text{if } q \geq 2/3. \end{cases}$$

Figure 1 illustrates V (step function depicted by dashed lines) and its concavification[13] cav V (piecewise linear function depicted by solid lines). The value of the optimal test t^* for the uninformed principal is cav V evaluated at the prior $q_0 = 1/6$. As apparent from Figure 1, the value cav $V(1/6)$ is a convex combination of $V(q)$ at two posteriors, $q = 0$ and $q = 1/3$. The optimal test is given by $t^*(\cdot|H) = (0, 1, 0)$ and $t^*(\cdot|L) = (3/5, 2/5, 0)$. In words, when the state is H, the test recommends action m; when the state is L, the test recommends actions l and m with probabilities $3/5$ and

[13]We write cav V for the smallest concave function that weakly exceeds V.

$2/5$, respectively. The agent's posterior is $q = 0$ after observing l, and it is $q = 1/3$ after observing m. Message h is never sent by the test.

Note that the optimal test t^* is degenerate, in the sense that it does not satisfy Assumption (A_2). However, it can be represented as the limit of a sequence of nondegenerate tests. So the outcome of this optimal test is implementable by the uninformed principal.

Koessler and Skreta (2022) argue that the outcome of the test t^* is not implementable by the informed principal. The reason is that if it was implementable, the informed principal would have to choose t^* in both states, H and L. But when the state was H, the principal would prefer to deviate by revealing the state. The agent would then optimally choose h instead of m, which would be a strict improvement for the principal. The conclusion is that the sequential rationality of the informed principal poses a substantive constraint that prevents the attainment of the outcome that is optimal for the uninformed principal.

The above argument has a potential caveat. Imagine that after observing the principal's deviation, the agent becomes "skeptical" and believes that the state is L with certainty. But then, the test reveals that the state is H with certainty. The agent's posterior belief is indeterminate under Bayes' rule. It could be that the agent doubts his conviction that $\theta = L$ and thus believes the result of the test. Alternatively, it could be that the agent doubts the accuracy of the test, and thus remains convinced that $\theta = L$.

Perez-Richet (2014) and Koessler and Skreta (2022) deal with the above belief indeterminacy problem using to the principle of the preeminence of tests. They impose the constraint that every out-of-equilibrium posterior belief must assign probability one to each event that is revealed as certain by the test. For example, if the test reveals that the state is H with certainty, the agent's posterior belief must be that the state is H with certainty, irrespective of the prior.

This paper takes a complementary approach. We do not impose constraints on the agent's out-of-equilibrium beliefs. Instead, we assume that tests are never absolutely accurate, as captured by Assumption (A_2). With such tests, no events are certain, so the principle of the preeminence of tests has no substance.

We argue that the outcome of test t^* is implementable by the informed principal, provided her deviations conditional on learning the state are restricted to nondegenerate tests. Consider the pooling mechanism τ^* that chooses test t^* with certainty in every state. This mechanism is sequentially rational under the following agent's belief system β^*. If the principal deviates from t^* to a different test t, the agent forms the belief that the state is $\theta = L$ with certainty. Because t is nondegenerate, after every message of t the agent remains certain that $\theta = L$, and thus chooses action l, which is the principal's least preferred action. Moreover, (τ^*, β^*) is a sequential equilibrium, since the above degenerate out-of-equilibrium belief can be sustained as an outcome of Bayes' rule by slightly perturbing the mechanism τ^*. Let f be a full-support distribution over the set of tests T. For a small $\tilde{\varepsilon} > 0$, consider the following perturbed mechanism $\tau_{\tilde{\varepsilon}}$:

(i) when $\theta = L$, $\tau_{\tilde{\varepsilon}}(\cdot|\theta)$ chooses test t^* with probability $1 - \tilde{\varepsilon}$, and with the complementary probability it draws a random test from T according to distribution f;

(ii) when $\theta = H$, $\tau_{\tilde{\varepsilon}}(\cdot|\theta)$ chooses test t^* with probability $1 - \tilde{\varepsilon}^2$, and with the complementary probability it draws a random test from T according to distribution f.

As $\tilde{\varepsilon}$ vanishes, $\tau_{\tilde{\varepsilon}}$ approaches τ^*. At the same time, conditional on observing a deviation $t \neq t^*$ and a message of t, the agent's posterior beliefs approach the degenerate belief that $\theta = L$.

Our conclusion is that, under the assumptions of our setting, the sequential rationality of the informed principal does not pose a substantive constraint.

Appendix. Proof of Theorem 1

Recall our assumption that the set M has at least as many messages as actions in A. Let us identify actions with messages, so assume that

$$A \subseteq M.$$

We then interpret each message $m \in A$ as a recommendation to choose action m.

A test t is called *obedient* if
(i) it induces the agent's choice equal to the recommended action, so $d(\beta(\cdot|t, m)) = m$ for each $m \in A$, and
(ii) it never sends messages outside of A, so $t(m|\theta) = 0$ for each $m \in M \backslash A$.

By (3), the outcome λ_t of an obedient test t in the uninformed principal setting is given by

$$\lambda_t(a|\theta) = \sum_{m \in M} t(m|\theta) \mathbf{1}_{\{a = d(\beta(\cdot|t,m))\}} = t(a|\theta). \tag{7}$$

By the revelation principle, in the uninformed principal model, any test can be replaced by an obedient test without changing the outcome.[14]

Proof of Theorem 1(a). Let (τ, β) be an ε-sequential equilibrium in the informed principal model for some $\varepsilon > 0$. For each action $a \in A$ let $Y_\beta(a)$ be the set of all pairs (t, m) of a test t and a message m such that the agent's optimal choice is a,

$$Y_\beta(a) = \{(t, m) \in T \times M : d(\beta(\cdot|t, m)) = a\}.$$

Using this notation, the outcome $\lambda_{(\tau,\beta)}$, which is given by (5), can be rewritten as

$$\lambda_{(\tau,\beta)}(a|\theta) = \int_{t\in T} \left(\sum_{m\in M} t(m|\theta) \mathbf{1}_{\{a=d(\beta(\cdot|t,m))\}} \right) \tau(\mathrm{d}t|\theta) = \int_{(t,m)\in Y_\beta(a)} t(m|\theta)\tau(\mathrm{d}t|\theta).$$

Given (τ, β), construct a test $\tilde{t}$ as follows. For each $\theta \in \Theta$ let

$$\tilde{t}(m|\theta) = \begin{cases} \lambda_{(\tau,\beta)}(m|\theta) & \text{for each } m \in A, \\ 0 & \text{for each } m \in M\backslash A. \end{cases} \tag{8}$$

By the definition of $Y_\beta(a)$, the test $\tilde{t}$ is obedient. By (7) and (8), the outcome of this test in the uninformed principal setting, $\lambda_{\tilde{t}}$, is given by

$$\lambda_{\tilde{t}}(a|\theta) = \tilde{t}(a|\theta) = \lambda_{(\tau,\beta)}(a|\theta), \quad a \in A, \ \theta \in \Theta.$$

Let λ be an outcome implementable by the informed principal. Then there is a sequence $(\varepsilon_k, \tau_k, \beta_k)_{k\in\mathbb{N}}$ such that $\lim_{k\to\infty} \varepsilon_k = 0$ and $\lim_{k\to\infty} \lambda_{(\tau_k,\beta_k)} = \lambda$. For each $k \in \mathbb{N}$, using the construction (8), the pair (τ_k, β_k) is replaced by the test $\tilde{t}_k$ with the same outcome, $\lambda_{\tilde{t}_k} = \lambda_{(\tau_k,\beta_k)}$. Thus, $\lim_{k\to\infty} \lambda_{\tilde{t}_k} = \lim_{k\to\infty} \lambda_{(\tau_k,\beta_k)} = \lambda$, which means that λ is implementable by the uninformed principal.

Remark 3. We do not impose Assumption (A_2) in part (a) of Theorem 1. However, this result continues to hold if we make this assumption, specifically, if every test in the support of τ satisfies (A_2). Observe that $\tilde{t}$ constructed in (8) need not satisfy (A_2). Nevertheless, the outcome $\lambda_{\tilde{t}}$ of test $\tilde{t}$ is still implementable by the uninformed principal, because $\tilde{t}$ can be approximated by a sequence of tests $(t_k)_{k\in\mathbb{N}}$ such that t_k satisfies (A_2) for each $k \in \mathbb{N}$, and $\lim_{k\to\infty} \lambda_{t_k} = \lambda_{\tilde{t}}$. In other words, ($A_2$) does not pose a substantive restriction for the uninformed principal when it is imposed on the informed principal.

[14]Note that in the informed principal model restricting to obedient tests entails loss of generality. This is because tests can be used as signals of information, so the principal can potentially use two tests with the same outcome as distinct signals.

Proof of Theorem 1(b). Using Assumption (A_1), let actions in A be ordered so that the principal's utility is increasing in a. Let $\prec$ denote this order.

For each state $\theta \in \Theta$ let δ_θ be the degenerate belief that puts probability one on θ. Let a_* be the worst action for the principal among the actions that can be induced by degenerate beliefs, so

$$a_* = d(\delta_{\theta^*}) = \min_{\theta \in \Theta} d(\delta_\theta), \tag{9}$$

We refer to a_* and δ_{θ_*} as the *punishment action* and *punishment belief*, respectively.

We prove the following statement. In the uninformed principal model, if an outcome is Pareto undominated, then every action induced in this outcome is at least as good for the principal as the punishment action a_*.

Lemma 1. *Consider the uninformed principal model. Let λ be implementable and Pareto undominated. Then for each $a' \in A$,*

$$a' \prec a_* \implies \sum_{\theta \in \Theta} \lambda(a'|\theta) = 0.$$

Proof. By contradiction, suppose that there exists $a' \in A$ such that

$$a' \prec a_* \text{ and } \sum_{\theta \in \Theta} \lambda(a'|\theta) > 0. \tag{10}$$

Let $(\varepsilon_k)_{k \in N}$ be a sequence of positive numbers with $\lim_{k\to\infty} \varepsilon_k = 0$. Because λ is implementable, there exists a sequence of tests $(t_k)_{k\in\mathbb{N}}$ in T such that $\lim_{k\to\infty} \lambda_{t_k} = \lambda$.

Consider $k \in \mathbb{N}$. By the revelation principle, without loss of generality, let t_k be obedient. Consequently, by (7),

$$\lambda_{t_k}(a|\theta) = t_k(a|\theta) \text{ for each } a \in A \text{ and each } \theta \in \Theta.$$

Moreover, by (10) and by $t_k \in T$ (so t_k satisfies (A_2)),

$$\lambda_{t_k}(a'|\theta) = t_k(a'|\theta) > 0 \text{ for each } \theta \in \Theta \text{ and each } k \in \mathbb{N}. \tag{11}$$

Let a^*_θ be the agent's preferred action in state θ, and let A^* be the set of such actions,

$$a^*_\theta = d(\delta_\theta) \text{ and } A^* = \{a^*_\theta\}_{\theta\in\Theta}.$$

Construct a test $\tilde{t}_k$ as follows. For each θ, with probability ε_k let $\tilde{t}_k(\cdot|\theta)$ send a random message with uniform distribution over the set A^*. With the complementary probability, $1-\varepsilon_k$, whenever $t_k(\cdot|\theta)$ sends message a', let $\tilde{t}_k(\cdot|\theta)$ send instead message a^*_θ, and whenever $t_k(\cdot|\theta)$ sends message $m \neq a'$, let $\tilde{t}_k(\cdot|\theta)$ send the same message as

$t_k(\cdot|\theta)$. In summary,

$$\tilde{t}_k(m|\theta) = \begin{cases} \frac{\varepsilon_k}{|A^*|} + (1-\varepsilon_k)(t_k(m|\theta) + t_k(a'|\theta)) & \text{if } m = a^*_\theta, \\ \frac{\varepsilon_k}{|A^*|} + (1-\varepsilon_k)t_k(m|\theta) & \text{if } m \in A^* \backslash \{a^*_\theta\}, \\ 0, & \text{if } m = a', \\ (1-\varepsilon_k)t_k(m|\theta) & \text{if } m \notin A^* \cup \{a'\}, \end{cases}$$

The following observations are in order.

First, $\tilde{t}_k \in T$. This is because t_k is in T, so the posteriors of the messages that are sent under t_k with positive probability have full support in both t_k and $\tilde{t}_k$. Furthermore, the messages a^*_θ, which may or may not be sent by t_k, are sent by $\tilde{t}_k$ with a positive probability in each state by construction of $\tilde{t}_k$.

Second, $\tilde{t}_k$ is obedient for every large enough k. If $\tilde{t}_k$ sends message $m \notin A^* \cup \{a'\}$, then the posterior under $\tilde{t}_k$ is the same as under t_k. Alternatively, if $\tilde{t}_k$ sends message $a^*_\theta \in A^*$, then the posterior $\beta(\cdot|\tilde{t}_k, a^*_\theta)$ is a perturbed mixture of two beliefs. The first belief is $\beta(\cdot|t_k, a^*_\theta)$ induced in t_k (provided t_k generates a^*_θ with a positive probability), in which case we know that a^*_θ must be optimal for the agent by the obedience of t_k. The second belief is δ_θ, in which case a^*_θ is uniquely optimal for the agent. The mixture of these two beliefs is perturbed, with the magnitude of the perturbation proportional to ε_k. When k is sufficiently large, so that the perturbation ε_k is small enough, a^*_θ is uniquely optimal for the agent under the posterior $\beta(\cdot|t_k, a^*_\theta)$.

Finally, for every large enough k, the agent is better off and the principal is strictly better off under $\tilde{t}_k$, as compared to t_k. This is because in each state θ, whenever test t_k sends a' and test $\tilde{t}_k$ sends a^*_θ, the agent prefers a^*_θ because it is the agent's uniquely optimal action in state θ. For the principal, by Assumption (A_1) and by (9), we have $u_P(a', \theta') < u_P(a_*, \theta') \le u_P(a^*_\theta, \theta')$ for every $\theta' \in \Theta$. Therefore, there is a constant $c > 0$ such that

$$u_P(a^*_\theta, \theta') - u_P(a', \theta') \ge c \quad \text{for all } \theta, \theta' \in \Theta.$$

In addition, by (10) and (11), there exists a probability $p > 0$ such that a' is played with probability at least p in test t_k for each sufficiently large k. Consequently, the principal's utility increment under $\tilde{t}_k$ as compared to t_k is at least $pc > 0$ for every sufficiently large k. We conclude that λ is Pareto dominated by $\tilde{\lambda} = \lim_{k\to\infty} \lambda_{\tilde{t}_k}$. □

Equipped with Lemma 1, we return to the proof of Part (b) of Theorem 1. Let λ^* be implementable in the uninformed principal model, and Pareto undominated. Then there exists a sequence $(\varepsilon_k, t_k)_{k\in\mathbb{N}}$ with $\varepsilon_k > 0$ and $t_k \in T$ such that

$$\lim_{k\to\infty} \varepsilon_k = 0 \quad \text{and} \quad \lim_{k\to\infty} \lambda_{t_k} = \lambda^*, \tag{12}$$

and, by Lemma 1,

$$\sum_{a'\in A: a' \prec a_*} \lambda_{t_k}(a'|\theta) \leq \varepsilon_k \quad \text{for each } \theta \in \Theta, \text{ and each } k \in \mathbb{N}. \tag{13}$$

That is, test t_k induces actions that are weakly superior to a_* for the principal with the probability at least $1 - \varepsilon_k$ in each state.

Fix $k \in \mathbb{N}$. Consider the following pair (τ_k, β_k) in the informed principal model. Let mechanism τ_k choose test t_k in all states, so $\tau_k(\cdot|\theta)$ assigns probability one to t_k for each $\theta \in \Theta$. Let β_k satisfy

$$\beta_k(\cdot|t, m) = \begin{cases} \beta(\cdot|t_k, m) & \text{if } t = t_k, \\ \delta_{\theta_*} & \text{if } t \in T\backslash\{t_k\} \end{cases}$$

for each $t \in T$ and each m that has a positive probability under t. In words, after observing test t_k, the agent forms an interim belief (i.e., the belief given the test but before observing the message of the test) equal to the prior, and then, given a message of t_k, the agent forms the posterior belief according to Bayes' rule. However, after observing a deviation $t \neq t_k$, the agent forms an interim belief equal to the punishment belief δ_{θ_*}. As this belief is degenerate but test t is nondegenerate by Assumption (A_2), messages of t do not affect the belief, leading to the same posterior belief $\beta_k(\cdot|t, m) = \delta_{\theta_*}$ for any message m. Recall that δ_{θ_*} induces the punishment action $a_* = d(\delta_{\theta_*})$. By (13), with probability at least $1 - \varepsilon_k$, test t_k generates one of the actions that are at least as good as a_* for the principal. Thus, the principal's utility increment from the deviation to $t \neq t_k$ is bounded by $\bar{u}_P\varepsilon_k$, where

$$\bar{u} = \max_{a',a''\in A,\theta\in\Theta} |u_P(a', \theta) - u_P(a'', \theta)|.$$

We thus conclude that τ_k is $(\bar{u}\varepsilon_k)$-sequentially rational under β_k. Moreover, by construction,

$$\lambda_{(\tau_k,\beta_k)} = \lambda_{t_k}, \quad k \in \mathbb{N}. \tag{14}$$

We now show that β_k is consistent with τ_k. Let θ_* be the "punishment" state, so it satisfies $a_* = d(\delta_{\theta_*})$. Let f be an arbitrary atomless full-support distribution over T.

Consider a sequence $(\varepsilon_n, \tau_{kn})_{n\in\mathbb{N}}$, where $\varepsilon_n > 0$ and $\lim_{n\to\infty} \varepsilon_n = 0$, and, for each $n \in \mathbb{N}$, the mechanism τ_{kn} is defined as follows. For each $\theta \neq \theta_*$, with probability $1 - (\varepsilon_n)^2$ let $\tau_{kn}(\cdot|\theta)$ choose the test t_k, and with probability $(\varepsilon_n)^2$ let it choose a random test according to distribution f. When $\theta = \theta_*$, with probability $1 - \varepsilon_n$ let $\tau_{kn}(\cdot|\theta)$ choose the test t_k, and with probability ε_n let it choose a random test according to distribution f. Thus, when observing t_k and message m, the agent's posterior belief $\beta_{kn}(\cdot|t_k, m)$ under τ_{kn} is given by

$$\beta_{kn}(\theta|t_k, m) = \begin{cases} \frac{t_k(m|\theta_*)q_0(\theta_*)}{t_k(m|\theta_*)q_0(\theta_*)+(1+\varepsilon_n)\sum_{\theta'\neq\theta_*} t(m|\theta')q_0(\theta')} & \text{if } \theta = \theta_*, \\ \frac{(1+\varepsilon_n)t(m|\theta)q_0(\theta)}{t_k(m|\theta_*)q_0(\theta_*)+(1+\varepsilon_n)\sum_{\theta'\neq\theta_*} t(m|\theta')q_0(\theta')} & \text{if } \theta \neq \theta_*. \end{cases}$$

When observing a test $t \neq t_k$ and message m, the posterior belief is given by

$$\beta_{kn}(\theta|t, m) = \begin{cases} \frac{t(m|\theta_*)q_0(\theta_*)\varepsilon_n}{t(m|\theta_*)q_0(\theta_*)\varepsilon_n+\sum_{\theta'\neq\theta_*} t(m|\theta')q_0(\theta')(\varepsilon_n)^2} & \text{if } \theta = \theta_*, \\ \frac{t(m|\theta)q_0(\theta)(\varepsilon_n)^2}{t(m|\theta_*)q_0(\theta_*)\varepsilon_n+\sum_{\theta'\neq\theta_*} t(m|\theta')q_0(\theta')(\varepsilon_n)^2} & \text{if } \theta \neq \theta_*. \end{cases}$$

As $n \to \infty$, $\beta_{kn}(\cdot|t_k, m) \to \beta_k(\cdot|t_k, m)$ and $\beta_{kn}(\cdot|t, m) \to \delta_{\theta^*}$ for each $t \neq t_k$ pointwise. We thus obtain that $\lim_{n\to\infty}(\tau_{kn}, \beta_{kn}) = (\tau_k, \beta_k)$, so β_k is consistent with τ_k.

To summarize, (τ_k, β_k) is $(\bar{u}\varepsilon_k)$-sequential equilibrium for each $k \in \mathbb{N}$. By (12) and (14), we conclude that λ^* is implementable by the informed principal. □

References

Alonso, R., and O. Câmara (2018): "On the value of persuasion by experts," *Journal of Economic Theory*, 174, 103–123. https://doi.org/10.1016/j.jet.2017.12.001

Antsygina, A., and M. Teteryatnikova (2022): "Optimal information disclosure in contests with stochastic prize valuations," *Economic Theory*. https://doi.org/10.1007/s00199-022-01422-8

Bedard, N. C. (2017): "Contracts in informed-principal problems with moral hazard," *Economic Theory Bulletin*, 5, 21–34. https://doi.org/10.1007/s40505-016-0093-0

Ben-Porath, E., E. Dekel, and B. L. Lipman (2019): "Mechanisms with evidence: Commitment and robustness," *Econometrica*, 87, 529–566. https://doi.org/10.3982/ECTA14991

Bergemann, D., and S. Morris (2019): "Information design: A unified perspective," *Journal of Economic Literature*, 57, 44–95. , https://doi.org/10.1257/jel.20181489

Bizzotto, J., and A. Vigier (2021): "Can a better informed listener be easier to persuade?," *Economic Theory*, 72, 705–721. https://doi.org/10.1007/s00199-020-01321-w

Chen, Y., and J. Zhang (2020): "Signalling by bayesian persuasion and pricing strategy," *Economic Journal*, 130, 976–1007. https://doi.org/10.1093/ej/ueaa002

Degan, A., and M. Li (2021): "Persuasion with costly precision," *Economic Theory*, 72, 869–908. https://doi.org/10.1007/s00199-021-01346-9

Eilat, R., and Z. Neeman (2023): "Communication with endogenous deception costs," *Journal of Economic Theory*, 207, 105572. https://doi.org/10.1016/j.jet.2022.105572

Guo, Y., and E. Shmaya (2021): "Costly miscalibration," *Theoretical Economics*, 16, 477–506. https://doi.org/10.3982/TE3991

Hagenbach, J., F. Koessler, and E. Perez-Richet (2014): "Certifiable pre-play communication: Full disclosure," *Econometrica*, 82, 1093–1131. https://doi.org/10.3982/ECTA11070

Hart, S., I. Kremer, and M. Perry (2017): "Evidence games: Truth and commitment," *American Economic Review*, 107(3), 690–713. https://doi.org/10.1257/aer.20150913

Hedlund, J. (2017): "Bayesian persuasion by a privately informed sender," *Journal of Economic Theory*, 167, 229–268. https://doi.org/10.1016/j.jet.2016.11.003

Ivanov, M. (2021): "Optimal monotone signals in Bayesian persuasion mechanisms," *Economic Theory*, 72, 955–1000. https://doi.org/10.1007/s00199-020-01277-x

Kamenica, E. (2019): "Bayesian persuasion and information design," *Annual Review of Economics*, 11, 249–272. https://doi.org/10.1146/annurev-economics-080218-025739

Kamenica, E., and M. Gentzkow (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615. https://doi.org/10.1257/aer.101.6.2590

Kamenica, E., K. Kim, and A. Zapechelnyuk (2021): "Bayesian persuasion and information design: Perspectives and open issues," *Economic Theory*, 72, 701–704. https://doi.org/10.1007/s00199-021-01383-4

Koessler, F., and V. Skreta (2022): "Informed information design," CEPR Discussion Paper No. DP17028, available at SSRN https://ssrn.com/abstract=4069816.

Kreps, D. M., and R. Wilson (1982): "Sequential equilibria," *Econometrica*, 50, 863–894. https://doi.org/10.2307/1912767

Lipnowski, E., and D. Ravid (2020): "Cheap talk with transparent motives," *Econometrica*, 88, 1631–1660. https://doi.org/10.3982/ECTA15674

Lipnowski, E., D. Ravid, and D. Shishkin (2022): "Persuasion via weak institutions," *Journal of Political Economy*, 130, 2705–2730. https://doi.org/10.1086/720462

Maskin, E., and J. Tirole (1990): "The principal-agent relationship with an informed principal: The case of private values," *Econometrica*, 58, 379–409. https://doi.org/10.2307/2938208

Mekonnen, T. (2021): "Informed principal, moral hazard, and limited liability," *Economic Theory Bulletin*, 9, 119–142. https://doi.org/10.1007/s40505-021-00201-3

Min, D. (2021): "Bayesian persuasion under partial commitment," *Economic Theory*, 72, 743–764. https://doi.org/10.1007/s00199-021-01386-1

Myerson, R. B. (1983): "Mechanism design by an informed principal," *Econometrica*, 51, 1767–1797. https://doi.org/10.2307/1912116

Myerson, R. B., and P. J. Reny (2020): "Perfect Conditional ε-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions," *Econometrica*, 88, 495–531. https://doi.org/10.3982/ECTA13426

Mylovanov, T., and T. Tröger (2014): "Mechanism design by an informed principal: Private values with transferable utility," *Review of Economic Studies*, 81, 1668–1707. https://doi.org/10.1093/restud/rdu019

Okuno-Fujiwara, M., A. Postlewaite, and K. Suzumura (1990): "Strategic information revelation," *Review of Economic Studies*, 57, 25–47. https://doi.org/10.2307/2297541

Perez-Richet, E. (2014): "Interim bayesian persuasion: First steps," *American Economic Review*, 104, 469–74. https://doi.org/10.1257/aer.104.5.469

Seidmann, D. J., and E. Winter (1997): "Strategic information transmission with verifiable messages," *Econometrica*, 65, 163–169. https://doi.org/10.2307/2171817

Serena, M. (2022): "Harnessing beliefs to optimally disclose contestants' types," *Economic Theory*, 74, 763–792. https://doi.org/10.1007/s00199-021-01378-1

Skreta, V. (2011): "On the informed seller problem: Optimal information disclosure," *Review of Economic Design*, 15, 1–36. https://doi.org/10.1007/s10058-009-0099-2

Yilankaya, O. (1999): "A note on the seller's optimal mechanism in bilateral trade with two-sided incomplete information," *Journal of Economic Theory*, 87, 267–271. https://doi.org/10.1006/jeth.1999.2532