Dynamic Informational Control*

Maxim Ivanov[†]
McMaster University

March 2011

Abstract

This paper investigates a multi-stage model of informational control, i.e., cheap-talk communication between an informed expert and an uninformed principal by Crawford and Sobel (1982), such that the principal can affect the quality of expert's private information without learning its content. We construct the two-stage procedure of dynamic updating of expert's information that allows the principal to elicit perfect information from the expert about an unknown single- or multi-dimensional state and reach his first-best outcome if the bias in preferences is not too large relative to the size of the state space. If the state space is unbounded, full information extraction is possible for an arbitrarily large bias under some regularity conditions.

JEL classification: C72, D81, D82, D83

Keywords: Communication, Information, Cheap Talk

1 Introduction

This paper focuses on the standard problem of facilitating cheap talk communication between the parties, who are asymmetrically informed and have different decision-making powers. In many situations, principals do not possess important information about the economic, political, or military consequences of their decisions. Moreover, their possibilities of obtaining relevant information and gaining expertise in necessary fields are severely restricted by the large volume, diverse range, and complexity of main responsibilities. As a result, high costs of information acquisition and opportunity costs force the decision makers to consult the experts, who have either more expertise in particular areas or significantly lower costs of

*This paper supersedes my other paper, "Dynamic Information Revelation in Cheap Talk." I am grateful to Dirk Bergemann, Oliver Board, Kalyan Chatterjee, Paul Fischer, Seungjin Han, Vijay Krishna, Bart Lipman, Robert Marshall, Tymofiy Mylovanov, Gregory Pavlov, Marek Pycia, and Neil Wallace for their valuable suggestions and discussions on different versions of the paper. All mistakes are mine.

[†]Department of Economics, McMaster University, 1280 Main Street West, Hamilton, ON, Canada L8S 4M4. E-mail: mivanov@mcmaster.ca.

acquiring and processing new information.[1] Legislators consult lobbyists about the costs and benefits of proposed laws, patients seek advice from doctors about the impact of medical procedures, and headquarters consult managers about numerous corporate issues.

Communication in such environments is typically characterized by two intrinsic features. The first one is the conflict of interest. Because the expert is not disinterested, she might misrepresent information in an attempt to manipulate the decisions of the principal in her favor. These manipulations reduce the quality of disclosed information and, therefore, induce inefficient decisions. The second feature is the imperfect primary information of the informed party. Even the most knowledgeable expert may have noisy or insufficient information.

While the conflict of interest between the involved parties is generally exogenous, the quality of expert's information can be sometimes endogenized by the principal. That is, though the expert observes her information privately, the precision of the information is affected by the principal. For instance, if obtaining new information requires conducting a complex test or experiment, the principal may impose the restrictions on the testing procedures performed by the expert. Also, if the parties interact through multiple stages, the principal may influence the precision of the expert's information before every round, whereas the expert can update her report afterwards. In other words, the principal allows the expert to acquire new information and send another message instead of making a decision on the basis of a single report from the expert.

The major contribution of this paper is that it demonstrates how the principal can use these factors—affecting the quality of expert's private information (hereafter, informational control) and dynamic interaction—in order to extract perfect information from the expert about an unknown single- or multi-dimensional issue in only two rounds of cheap talk conversation. We construct the procedure of updating expert's information over time, which sustains truthful communication in both stages. According to this procedure, the expert precisely learns the state of nature in the second stage and reports it truthfully to the principal. The procedure does not require any commitment on the side of the principal, who can change the quality of expert's information and/or his decision at any moment. In particular, this implies that the principal can influence how much new information the expert can acquire in the second period depending on her previous report.

Dynamic informational control can be interpreted as follows. It is a situation in which the expert conducts a sequence of experiments and reports their results to the principal while the principal determines the precision of each experiment and makes a final decision. Applications of informational control are restrictions on the procedures of acquiring information, which is a subject of communication between parties with conflicting interests. Consider, for example, the communication problem in the defense procurement.[2] While the military is

[1] A natural question of the principal's limited possibilities of obtaining information without involving outside parties has been addressed in multiple works. For instance, Radner (1991) emphasizes: "The decentralization of information-processing is dictated by the large scale of modern enterprises, which makes it impossible for any single person to manage everything." Similarly, Krishna and Morgan (2001a) note: "The diverse range of problems confronted by decision makers, such as corporate CEOs or political leaders, almost precludes the possibility that they themselves are experts in all relevant fields, and hence, the need for outside experts naturally arises." Also, Austen-Smith (1994) states: "when the information has to be acquired, it is natural to suppose the acquisition is costly; for otherwise, there is no reason why such information is asymmetrically distributed."

[2] I am especially thankful to Robert Marshall for sharing this example.

an expert in evaluating the weapons characteristics, the budget for production is determined by the Congress (principal). Moreover, the parties' interests are not necessarily aligned—it is argued that the military can be biased toward weapons with excessive costs (Rogerson, 1990). Remarkably, the Department of Defense has received multiple accusations about manipulating tests results to yield the most favorable interpretation.[3]

In this work, we offer a potential solution to such communication problems. In the context of our model, the principal can impose proper restrictions on the sequence of testing procedures, for example, by approving only particular certified equipment, which must be used by the expert to perform each test. We show that the communication protocol, which elicits full information from the expert is quite simple. It consists of two experiments, such that the precision of the second experiment depends on the expert's report on the first experiment. That is, the expert has to submit a separate report upon acquiring new information instead of conveying all information simultaneously. The experiments are designed as follows. The first experiment returns two signals—the true state and some *complement* state, which is sufficiently distinct from the true one—but does not reveal which signal contains relevant information. The second experiment allows the expert to distinguish between her *reported* state and the complement state to the reported one. Otherwise, the outcome of the experiment is identical to that in the previous stage. Thus, the expert's information is updated in the second period only if she tells the truth in the previous stage. If the divergence in preferences is not large, then the expert prefers to report her information truthfully in both stages.

Another application of dynamic informational control is communication through a channel with strategic noise. This is a situation, in which two parties are interested in truthful communication, but have to interact via a mediator with conflicting objectives. Consider, for example, sharing information in the U.S. Congress. Due to the high volume and complexity of its work, the Congress divides its tasks among committees and subcommittees, which consider bills and issues and recommend measures for consideration by their respective chambers. Thus, if a lobbyist (expert) needs to convey some information with legislators (decision maker), she has to pass it to the specialized committee (mediator). In reality, however, mediators are likely to have their own agenda (e.g., the composition of Congress committees is generally determined by party leaders). In addition, though the activity of the mediator can be often monitored by the expert, e.g., because the mediator's reports are publicly observable, the expert may not have instruments of punishing the mediator for lying. We suggest a solution to the problem of communication via an interested mediator, who can be monitored by the privately informed expert, but is not liable for distorting information. In this case, the two-stage communication protocol that sustains truthful information transmission via the mediator can be organized as follows. At the first stage, the expert submits two reports to the mediator, one of which contains a piece of true information. Then, upon observing that the mediator truthfully conveys obtained information to the decision maker, the expert reveals which report is correct, and so does the mediator.

Intuitively, the efficiency of the constructed algorithm of updating expert's information is driven by a combination of two crucial components. First, upon acquiring imprecise information in the first round, the quality of expert's information in round two and, therefore, the benefits of updating information depend on the current report of the expert. Second, the

[3] See, for example, the U.S. General Accounting Office (1988, 1992).

principal uses the first-stage report in order to (partially) verify the expert's report in the next period and makes only those decisions that are consistent with received information in both stages. This significantly restricts the expert's possibilities of manipulating her information in round two, because any first-stage message shrinks the set of feasible actions. Together, these factors imply that before sending the first message, the expert faces the trade-off between the future informational benefits received in the case of truthful reporting and flexibility over available actions, which can be induced by properly distorting her current information. If the bias in player's interests is not large, the informational benefits dominate the benefits of inducing any decision by the partially informed expert. In addition, the smaller set of feasible actions in stage two helps sustain truthful communication in that stage. In fact, the expert can induce only the principal's best responses to either the true state or the complement state. If the latter option is unfavorable, the expert prefers to tell the truth. An important feature of the introduced procedure of updating expert's information is decomposing communication about the continuous variable into the continuum of cheap-talk conversations about discrete posterior realizations. This decomposition plays a dual role. First, it introduces substantial uncertainty in the first-period information of the expert and, thus, motivates her to learn the state precisely in the second round. Precise learning, however, occurs only after truthful reporting of the interim information. Second, the discretization of the state space reduces the expert's possibilities of manipulating her information by claiming a different state in the second period. Thus, though the number of expert's types is smaller, each type fully separates itself. In short, informational control gives a lot of freedom over decisions to the expert, because the principal can be easily manipulated by the expert's messages. However, the value of such freedom is low if the expert is not well informed about the true situation.

We build the analysis on the extension of the classical model of Crawford and Sobel (1982) (hereafter, CS), which incorporates the discussed ingredients: imperfect information of the expert and communication in multiple stages. In this light, our work is related to two areas in the literature on strategic communication that deal with each of these features. Regarding the first area, Austen-Smith (1994) considers the situation, in which the expert can observe information at some privately known cost and is able to prove the fact of information acquisition. Though the information costs imply that the expert might prefer to remain uninformed, they affect the principal's reaction to all expert's messages and facilitate information transmission from the informed expert. This increases the overall quality of conveyed information. Green and Stockey (2007) first recognized the fact that the principal's expected payoff is not necessarily monotone in the quality of the expert's information. Fischer and Stocken (2001) demonstrate this result for the CS model. Ivanov (2010a) extends this result by showing that informational control can provide the higher payoff to the principal than optimally delegating authority to the expert. Our paper extends this literature by allowing the expert to collect and report information over time.

Another way to improve information transmission between conflicting parties is to organize communication through multiple stages. Aumann and Hart (2003) consider two-person games with two-sided cheap talk, in which one side is better informed than the other and the players communicate without time constraints. For the class of games with discrete types and the bimatrix structure of strategies and payoffs, they demonstrate that the set of equilibrium outcomes can be significantly expanded. Krishna and Morgan (2004) investigate multi-stage communication with active participation of the principal in the communication process and show the existence of equilibria that almost always ex-ante Pareto dominate all direct-talk equilibria. They also recognize that a simple extension of the CS model to the setup with multiple reports of the expert does not affect the set of equilibrium outcomes and, therefore, cannot improve communication.[4] The argument by Krishna and Morgan (2004) can be directly applied to the case of the imperfectly informed expert without information updating. In contrast, if the precision of expert's information can be improved over time, the consequences of multi-stage communication are qualitatively different from outcomes in the single-stage setup.

Our work also complements a few papers on full disclosure of information in communication games. These include cheap talk with multiple experts investigated by Krishna and Morgan (2001a, 2001b), Battaglini (2002), and Eso and Fong (2008). Kartik et al. (2007) consider the settings with lying costs of the expert. Golosov et al. (2010) analyze dynamic cheap talk with multiple rounds of the expert's reports and the principal's actions. Seidmann and Winter (1997) and Mathis (2008) consider the case of verifiable expert's messages. The main distinction of our paper is the endogenous and dynamic quality of expert's information, whereas the information structure of the informed party(s) in the aforementioned works is exogenously given at the beginning of the game.

The rest of the paper is structured as follows. Section 2 presents the formal model. Section 3 highlights an illustrative example. The general analysis of the model is performed in Section 4. Section 5 concludes the paper.

2 The model

We consider the model of two-stage communication game with two players, the expert and the principal, who interact about some issue of communication. It is represented by the *state* θ, which is distributed on the interval $\Theta = [\underline{\theta}, \bar{\theta}], -\infty \leq \underline{\theta} < \bar{\theta} \leq \infty$, according to a distribution function $F(\theta)$ with a positive and continuous density $f(\theta)$. The expert can privately observe some information about θ, whereas the principal makes a decision a that affects the payoffs of both players. The players' payoff functions are of the form $U(a, \theta, b_i), i \in \{E, P\}$ where the inherent bias parameter $b_i \geq 0$ reflects the divergence in players' interests. The principal's bias is normalized to be 0, whereas the expert's bias is $b > 0$. The function $U(a, \theta, b)$ is twice-differentiable and satisfies $U''_{aa} < 0, U''_{a\theta} > 0$, and $U''_{ab} > 0$. Due to these conditions, U has a unique *ideal* decision $a^*(\theta, b) = \arg\max_a U(a, \theta, b)$, which is strictly increasing in θ and b. We assume that $U'_b(a^*(\theta, b), \theta, b) \leq 0$ for all θ and b.[5] Hereafter, we write $U(a, \theta) = U(a, \theta, 0)$ and $V(a, \theta, b) = U(a, \theta, b)$ as the principal's and the expert's payoff functions, respectively. We also write $a_p(\theta) = a^*(\theta, 0)$ and $y(\theta, b) = a^*(\theta, b)$ as the ideal decisions of the principal and the expert, respectively. Since $U''_{ab} > 0$, then $b > 0$ implies $y(\theta, b) > a_p(\theta)$, i.e., the expert always prefers a higher action than is ideal for the principal.

Actions. At the beginning of each period t, the principal determines the expert's *information structure* $I_t = \{F_t(s_t|\theta), \theta \in \Theta\} \in \mathcal{I}, t = 1, 2$, which consists of a family of

[4] Because the expert possesses all available information before communication starts, she sends a sequence of messages that induce the most preferable decision. As a result, the principal infers the same information about the state as in the one-stage case. Thus, the set of induced decisions is also not affected, and any equilibrium in the multi-stage game is outcome-equivalent to that in the one-stage game.

[5] That is, there are no private benefits to the expert for simply being more biased.

conditional distributions of signals. The expert then privately observes a *signal* $s_t \in \mathcal{S}$ drawn from an associated distribution $F_t(s_t|\theta)$. At the end of the period t, the expert sends a *message* $m_t \in \mathcal{M}$ to the principal. Finally, upon receiving messages $\{m_1, m_2\}$, the principal makes a *decision* $a \in \mathbb{R}$.

Strategies. A behavioral pure strategy of the principal is: (1) a pair $\{I_1, I_2(m_1, I_1)\}$, where the second-period information structure $I_2 : \mathcal{M} \times \mathcal{I} \to \mathcal{I}$ is a function of the expert's message m_1 and the first-period information structure I_1; and (2) a decision $a : \mathcal{M}^2 \times \mathcal{I}^2 \to \mathbb{R}$, which is a function of the principal's history $\{m_1, m_2, I_1, I_2\}$.[6] Hereafter, we call the pair of the information structures $\{I_1, I_2(m_1, I_1)\}$ *communication schedule*. The behavioral strategy of the expert is a pair of functions $\{\sigma_1, \sigma_2\}$, where $\sigma_1 : \mathcal{I} \times \mathcal{S} \to \Delta\mathcal{M}$ and $\sigma_2 : \mathcal{I}^2 \times \mathcal{S}^2 \times \mathcal{M} \to \Delta\mathcal{M}$, which map the expert's private histories $h_1 = \{I_1, s_1\} \in \mathcal{I} \times \mathcal{S}$ and $h_2 = \{h_1, I_2, s_2, m_1\} \in \mathcal{I}^2 \times \mathcal{S}^2 \times \mathcal{M}$ in the first and the second periods, respectively, into the space of probability distributions on the message set $\mathcal{M}$.

Beliefs. The principal's belief $\mu : \mathcal{M}^2 \times \mathcal{I}^2 \to \Delta\Theta$ determines the probability distribution of θ given the principal's history $\{m_1, m_2, I_1, I_2\}$. We call the belief system *consistent*, if it is derived from the player's strategies on the basis of Bayes' rule where applicable.[7] A decision a is called *rationalizable* if there is some belief, for which a is a best response. Clearly, the set of rationalizable decisions is $\mathcal{A} = [a_p(0), a_p(1)]$.

Equilibrium. For a fixed communication schedule $\{I_1, I_2(m_1, I_1)\}$, a perfect Bayesian equilibrium (hereafter, equilibrium) is the belief μ and a pair of the player's strategies, $a(.)$ and $\{\sigma_1(.), \sigma_2(.)\}$, such that μ is consistent with the strategies and the strategies satisfy the following conditions.

(1) given $\{I_1, I_2, a(.), s_1, s_2, m_1\}$, $\sigma_2(.)$ maximizes the expert's 2nd-period payoff. That is, if $m_2 \in \operatorname{supp} \sigma_2(.)$, then

$$m_2 \in \arg\max_x E_\theta [V(a(m_1, x, I_1, I_2), \theta, b) | I_1, I_2, s_1, s_2].$$

(2) given $\{I_1, I_2(m_1, I_1), a(.), s_1\}$ and $\sigma_2(.)$, σ_1 maximizes the expert's 1st-period payoff. That is, if $m_1 \in \operatorname{supp} \sigma_1(.)$, then

$$m_1 \in \arg\max_x E_{\theta, s_2} [V(a(x, m_2(I_1, I_2(x, I_1), x, s_1, s_2), I_1, I_2(x, I_1)), \theta, b) | I_1, s_1].$$

(3) given m_1, m_2, I_1, I_2, and μ, a maximizes the principal's payoff:

$$a = \arg\max_y E_\theta [U(y, \theta) | \mu(m_1, m_2, I_1, I_2)].$$

[6] Since a mixture of distribution functions is a distribution function, we do not need to consider randomizing over information structures. Also, because of the strict concavity of the principal's preferences, he never mixes over decisions.

[7] For all experts's messages $m'_t \notin \mathcal{M}, t = 1, 2$, we define the receiver's beliefs in such a way that he interprets them as some $m_t \in \mathcal{M}$.

3 Example: the uniform-quadratic case

We start with an illustrative example that outlines the two-stage procedure of updating information of the expert, which induces her to fully reveal the state at the end of communication. For that purpose, we consider the uniform-quadratic specification of the CS model. It is known for its tractability and the possibility to obtain closed-form solutions in various modifications of the basic CS model and broadly used in the literature on strategic communication.[8] In this setup, the distribution of θ is uniform on $[0,1]$ and the preferences are of the quadratic form:

$$U(a,\theta,b_i) = -(a-\theta-b_i)^2, i \in \{E,P\}.$$

Because $b_P = 0$ and $b_E = b$, the ideal decisions of the principal and the expert are $a_p(\theta) = \theta$ and $y(\theta,b) = \theta + b$, respectively. Suppose that $b \le \bar{b} = 1/4$, where $\bar{b}$ is the largest bias, which sustains informative communication in the CS model.

Define the communication schedule as follows. If the state $\theta < 1/2$, the expert privately observes the signal $s = \theta$. If $1/2 \le \theta < 1$, the expert observes $s = \theta - 1/2$. For example, if θ is equal to $1/5$ or $7/10$, the expert observes the signal $s = 1/5$. Because each pair of states $\{\theta,\theta'\} = \{\theta,\theta+1/2\}, \theta < 1/2$ maps into the same signal $s = \theta$, the expert is not able to distinguish between the state $\theta = s$ and the *complement state* $\theta' = \varphi(s) = s + 1/2$ for all $s \in [0,1/2)$. Such a partitioning of the state space into two-point sets generates the binary posterior distribution of θ conditional on s, $\Pr\{\theta = s|s\} = \Pr\{\theta = s+1/2|s\} = 1/2$.

Depending on the report of the expert in the first stage, her information can be updated in the next period. In particular, given the first-stage message $m_1 < 1/2$, the expert can distinguish between $\theta = m_1$ and the complement state $\theta' = m_1 + 1/2$. For all other θ, the information structure remains unchanged. Thus, the expert learns θ in the second stage *if and only if* she truthfully reports her information in the first stage. Reporting $m_1 \neq s$ does not update the expert's information in the second stage, since she observes the signal s again. Upon receiving the expert's message $m_2 \in \{m_1, m_1+1/2\}$ in the second period, the principal interprets it as truthful and makes a decision $a(m_2) = m_2$.

According to such a communication schedule, the expert in the first stage faces the trade-off between the informational benefits and the flexibility over available actions. On the one hand, the expert can induce any action y in the set of rationalizable decisions $\mathcal{A} = [0,1]$ by sending either $m_1 = y$ or $m_1 = y - 1/2$ in the first period, depending on whether y is below or above $1/2$, respectively, and then reporting $m_2 = y$ in the second period. However, if the expert distorts her information by sending $m_1 \neq s$, she remains imperfectly informed in the next round. On the other hand, truthful reporting $m_1 = s$ in the first stage allows the expert to learn θ precisely, but shrinks the set of feasible actions to two. These are the principal's best-responses to the posterior realizations s and $\varphi(s)$: $a_p(s) = s$ and $a_p(s+1/2) = s + 1/2$.

From the expert's prospective in round one, distorting information and inducing the

[8] See, for example, Blume et al. (2007), Gilligan and Krehbiel (1987, 1989), Goltsman et al. (2007), Ivanov (2010b), Krishna and Morgan (2001a, 2001b, 2004, 2008), Melumad and Shibano (1991), and Ottaviani and Squintani (2006).

decision $y \in \mathcal{A}$ results in the payoff:

$$E\left[V\left(y,\theta,b\right)|s\right] = \frac{1}{2}V\left(y,s,b\right) + \frac{1}{2}V\left(y,\varphi\left(s\right),b\right) = -\frac{1}{2}\left(y-s-b\right)^2 - \frac{1}{2}\left(y-s-\frac{1}{2}-b\right)^2,$$

which is maximized at $y_1\left(s,b\right) = \min\left\{s+1/4+b,1\right\}$. Since $b \leq 1/4$, it follows that $y_1\left(s,b\right) = s+1/4+b < a_p\left(1\right) = 1$ if $s < 1/2$. This decision provides the payoff:

$$E\left[V\left(y_1\left(s,b\right),\theta\right)|s\right] = -\frac{1}{4}\left(\varphi\left(s\right)-s\right)^2 = -\frac{1}{16}, s < 1/2.$$

In the second case, truthful reporting in both stages results in the expert's payoff:

$$\begin{aligned} E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] &= \frac{1}{2}V\left(a_p\left(s\right),s,b\right) + \frac{1}{2}V\left(a_p\left(\varphi\left(s\right)\right),\varphi\left(s\right),b\right) \\ &= -\frac{1}{2}\left(s-s-b\right)^2 - \frac{1}{2}\left(\varphi\left(s\right)-\varphi\left(s\right)-b\right)^2 = -b^2, s < 1/2. \end{aligned}$$

Finally, the second-stage incentive compatibility constraints must prevent the positively biased expert from distorting information in period two. That is, inducing $a_p\left(\theta\right) = \theta$ upon learning $\theta < 1/2$ must be more beneficial than inducing $a_p\left(\varphi\left(\theta\right)\right) = \varphi\left(\theta\right) > a_p\left(\theta\right)$, or

$$V\left(s,s,b\right) \geq V\left(\varphi\left(s\right),s,b\right), \tag{1}$$

for all $s < 1/2$. These constraints are satisfied if and only if $\varphi\left(s\right) - s = s + 1/2 - s \geq 2b$, or $b \leq 1/4$. Therefore, if $b \leq 1/4$, then the inequalities (1) and

$$E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] = -b^2 \geq -\frac{1}{16} = E_\theta\left[V\left(y_1\left(s,b\right),\theta\right)|s\right],$$

imply that the expert conveys her information truthfully in both periods and, hence, reveals θ in the second stage.

It is worth noting that it is impossible to expand the range of biases that sustain full information extraction by choosing another complement state function $\varphi\left(\theta\right)$. This is because the minimum distance between the points in set $\left\{s,\varphi\left(s\right)\right\}$ does not exceed $1/2$.[9] That is, the maximum of $E\left[V\left(y_1\left(s,b\right),\theta\right)|s\right]$ over s is not less than $-\frac{1}{16}$ that is above $E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] = -b^2$ if $b > 1/4$. In addition, the second stage incentive-compatibility constraints (1) are violated for signals s, such that $\varphi\left(s\right) - s < 2b$.

While the detailed discussion of dynamic informational control is left to the remaining sections, a brief intuition is as follows. The first-stage information structure decomposes cheap talk communication about the continuous state into the continuum of single-stage communication subgames over the binary posterior realizations that are sufficiently distant from each other. If the expert is willing to convey truthfully her imprecise information in the first period, the second-period information structure rewards her by perfectly informing about the state. By that moment, however, the principal knows the particular binary distribution of posteriors, which sustains truthful conversation if the bias is not large.

[9] For any partition of $[0,1)$ into two-point sets $\{\theta',\theta\}$, we have $|\theta'-\theta| \leq 1/2$ for $\theta = 1/2$ and all $\theta' \in [0,1)$.

4 Dynamic information extraction

As shown by Fischer and Stocken (2001) and Ivanov (2010a), the principal can improve communication by restricting the quality of the expert's primary information. It is achieved by partitioning the state space into a collection of intervals $\Theta_k = [\theta_k, \theta_{k+1}], k = 1, 2, ..., N$, such that the expert privately observes the subinterval Θ_k that contains the state. The partitional information structure significantly limits the expert's possibilities of distorting information in her favor. For example, a claim that the posterior valuation $E[\theta|s]$ is $m \notin \{E[\theta|\theta \in \Theta_k]\}_{k=1}^N$ is not credible, because the principal knows that this information is not available to the expert. As a result, though the expert knows less, she is willing to report her information truthfully. Moreover, the overall effect of expert's coarse information on the principal's ex-ante payoff is positive.

In this section, we extend the example above to the general settings and show that the principal can extract full information about the state in only two stages by properly defining the expert's communication schedule. Our main focus is the class of *truthtelling* equilibria, in which the expert truthfully reports her signal in each period, that is, $m_t(h_t) = s_t, \forall h_t, t = 1, 2$. Also, we say that the communication schedule sustains the *fully informative* equilibrium, if the expert truthfully reveals the state upon perfectly learning it in the second stage.

4.1 Information structures

Consider first the case of the bounded state space $\Theta = [0, 1]$. Hereafter, we restrict the analysis to a particular class of communication schedules. In this class, the expert's information structure in the first period is determined by the partition of Θ into the continuum of two-point sets and a singleton at $\theta = 1$:

$$I_1 = \left\{ \phi(\theta)_{\theta \in [0,1/2)}, 1 \right\}, \tag{2}$$

where $\phi(\theta) = \{\theta, \varphi(\theta)\}$ is the *joint-state set*, and $\varphi : [0, 1/2] \to [1/2, 1]$ is the *complement state function*, which is continuous, strictly increasing, $\varphi(0) = 1/2$ and $\varphi(1/2) = 1$. We show below that the choice of the function $\varphi(\theta)$ does not affect the results qualitatively.

According to the information structure (2), if $\theta = 1$, the expert learns it perfectly. If $\theta < 1$, the expert observes the signal $s = \theta$ if $\theta < 1/2$, and $s = \varphi^{-1}(\theta)$ if $1/2 \leq \theta < 1$. That is, both $\theta < 1/2$ and the complement state $\theta' = \varphi(\theta) \geq 1/2$ map into the same signal $s = \theta$, so the expert cannot distinguish between s and $\varphi(s)$ upon observing s. Hereafter, we call a pair of posterior realizations θ and $\theta' = \varphi(\theta)$ the *joint states*. The information structure (2) generates the family of the first-stage posterior distributions $\{F(\theta|s), s \in [0, 1/2), \theta \in \{s, \varphi(s)\}\}$ and the degenerated distribution at $\theta = 1$, such that each $F(\theta|s)$ is the two-point distribution on $\{s, \varphi(s)\}$ with the probabilities:[10]

$$p_s = \Pr\{\theta = s|s\} = \frac{f(s)}{f(s) + f(\varphi(s))}, s \in [0, 1/2), \text{ and}$$
$$p_s^c = \Pr\{\theta = \varphi(s)|s\} = 1 - p_s = \frac{f(\varphi(s))}{f(s) + f(\varphi(s))}.$$

[10] Distributions $F(\theta|s), s \in [0, 1/2)$ are non-degenerated, since $f(s) > 0, s \in \Theta$ means $p_s > 0$ and $p_s^c > 0$.

After conveying the message $m_1 \in [0, 1/2)$ to the principal, the expert's second-stage information structure is modelled as follows. It refines the joint-state set $\{m_1, \varphi(m_1)\}$ into separate points, m_1 and $\varphi(m_1)$, but preserves the other sets unrefined. In other words, sending message m_1 allows the expert to distinguish between the states $\theta = m_1$ and $\theta' = \varphi(m_1)$ only and does not update her information about other states. That is,

$$I_2(m_1, I_1) = \begin{cases} \{m_1, \varphi(m_1), I_1 \backslash \{m_1, \varphi(m_1)\}\} \text{ if } m_1 < 1/2, \\ I_1 \text{ if } m_1 = 1. \end{cases} \quad (3)$$

In order to see how the communication schedule determined by (2) and (3) influences the expert's motives to convey information in both stages, consider first the procedure of updating expert's information over time. The particular choice of the expert's information structure in the first stage is driven by two factors. The first factor is related to the quality of expert's information in the first stage—it is very imprecise. In order to demonstrate this, fix an interval $[\theta_1, \theta_2]$ and consider the class of all distributions of a random variable X with the support $L \subset [\theta_1, \theta_2]$ and a given mean value $E[X] \in (\theta_1, \theta_2)$. In this class, the two-point distribution on $\{\theta_1, \theta_2\}$ with the probabilities $p_{\theta_1} = \Pr\{X = \theta_1\} = \frac{\theta_2 - E[X]}{\theta_2 - \theta_1}$ and $p_{\theta_2} = 1 - p_{\theta_1}$, respectively, is dominated by any other distribution by the second-order stochastic dominance.[11] By putting $X = [\theta|s]$, this dominance implies that the risk-averse expert is the most uncertain about the posterior value of θ if it takes one of the boundary points, θ_1 or θ_2. Applying this intuition to our context, the principal is interested in keeping expert's information sufficiently imprecise in round one. This increases the expert's benefits from perfectly learning the state in the next round and, thus, forces her to communicate truthfully in the first period. At the same time, the principal is interested in limiting the set of posterior values of θ for each signal s. This is necessary in order to prevent the expert from claiming a distinct posterior value $\theta' \neq \theta$ after perfectly learning θ in round two.

The other factor is related to both the quality of expert's information in the first stage and her incentives to communicate truthfully in the second stage. In particular, it is the distance between the joint states s and $\varphi(s)$. A larger distance $\varphi(s) - s$ between the posterior realizations leads to two effects. First, it increases expert's uncertainty about the state after observing the signal in the first stage, and thus, her benefits from updating information in the future. Second, it suppresses the expert's incentives to distort information in the second stage. The positively biased expert may have the incentive to distort her information by claiming the higher complement state $\varphi(\theta) > \theta$ after learning $\theta < 1/2$. If $\varphi(\theta)$ is distant from θ, then the induced decision $a_p(\varphi(\theta))$ is substantially different from the expert's ideal decision $y(\theta, b)$, and, therefore, is unfavorable. However, the principal is limited in maximizing the distance $\varphi(s) - s$ for all s. This is because for any partition of $[0, 1)$ consisting of the joint-state sets, the minimum distance between the joint states is bounded by $1/2$.

4.2 Bounded state space

We show now how the constructed communication schedule allows the principal to elicit full information from the expert. Suppose that the expert acquires information according to the information structures (2)–(3) and sends the sequence of messages $\{m_1, m_2\}$. If $m_1 \in [0, 1/2)$

[11] The proof follows, for instance, from (3.A.8) in Shaked and Shanthikumar (2007).

and $m_2 \in \{m_1, \varphi(m_1)\}$, the principal interprets them as truthful and implements the decision $a_p(m_2)$. If $m_1 = 1$, the principal makes a decision $a_p(1)$ unconditionally on m_2.[12]

Let $y_1(s,b) \in \mathcal{A}$ be the expert's optimal decision given her interim first-stage information:

$$y_1(s,b) = \arg\max_{a \in \mathcal{A}} \; p_s V(a,s,b) + (1-p_s) V(a, \varphi(s), b), s \in [0,1/2). \tag{4}$$

That is, $y_1(s,b)$ is the most profitable decision in the case of manipulation of first-stage information. The expert can induce $y = y_1(s,b)$ by sending the message:

$$m_1 = \begin{cases} a_p^{-1}(y) \text{ if } y < a_p(1/2) \text{ or } y = a_p(1), \text{ and} \\ \varphi^{-1}\left(a_p^{-1}(y)\right) \text{ if } a_p(1/2) \leq y < a_p(1), \end{cases}$$

in the first stage and $m_2 = a_p^{-1}(y)$ in the second stage. Because her information will not be updated in the second stage, inducing $y_1(s,b)$ results in the payoff:

$$\begin{aligned} E[V(y_1(s,b), \theta, b) | s] &= \max_{a \in \mathcal{A}} p_s V(a,s,b) + (1-p_s) V(a, \varphi(s), b) \\ &= p_s V(y_1(s,b), s, b) + (1-p_s) V(y_1(s,b), \varphi(s), b), s \in [0,1/2). \end{aligned} \tag{5}$$

If the expert reveals her information truthfully by sending $m_1 = s$, her interim information will be updated in the second stage, but the set of feasible actions shrinks to $\{a_p(s), a_p(\varphi(s))\}$. In this case, reporting the truth in both periods results in the payoff:

$$E[V(a_p(\theta), \theta, b) | s] = p_s V(a_p(s), s, b) + (1-p_s) V(a_p(\varphi(s)), \varphi(s), b). \tag{6}$$

Because the expert can induce only $a_p(s)$ or $a_p(\varphi(s))$ in the second period, her second-stage incentive-compatibility constraints are given by:

$$V(a_p(s), s, b) \geq V(a_p(\varphi(s)), s, b), s \in [0,1/2). \tag{7}$$

That is, fully informative communication is sustainable if (7) holds and

$$E[V(a_p(\theta), \theta, b) | s] \geq E[V(y_1(s,b), \theta, b) | s], s \in [0,1/2).$$

The following theorem demonstrates that the trade-off between the informational benefits and the flexibility over available actions is in favor of the former option if the bias in preferences is not large. All proofs are collected in the Appendix.

Theorem 1 *There exists $\bar{b}$ such that if $b \leq \bar{b}$, there is the fully informative equilibrium in the game with the communication schedule determined by (2) and (3).*

The proof follows directly from the construction of the communication schedule. If the players' interests do not differ significantly, the ideal decisions of the expert and the principal, $y(\theta,b)$ and $a_p(\theta)$, respectively, are sufficiently close for each realization of the state. Since expert's information in the first stage is substantially imprecise, she is more interested in

[12] If the expert reports $m_1 \notin [0,1/2) \cup \{1\}$ and/or $m_2 \notin \{m_1, \varphi(m_1)\}$, the principal interprets such messages as, for instance, $m_1 = m_0 \in [0,1/2)$ and $m_2 = m_1$, respectively.

learning the state perfectly at the cost of inducing the principal's ideal decision than inducing the optimal decision $y_1(s, b)$ at the cost of losing new information. Because the expert can learn the state only by revealing her first-stage information, the principal knows the distribution of posteriors at the beginning of the second period. This discrete distribution sustains truthtelling communication at that period if the bias is not large.

A few comments are necessary. First, there are two key differences between the truthtelling equilibria in the dynamic and single-shot versions of the model. The main distinction of the multi-stage setup is that the principal can use the expert's message in the first period in order to refine her information in the second one. However, the information structure in the second period refines only the set of posterior values of the state, which have been reported by the expert. It is clear that such a procedure of updating the expert's information over time cannot be applied to the static model. Also, any message of the expert in period one reduces the set of feasible decisions in the second period.[13] This is because the principal uses the first-stage message in order to partially verify the second-stage report, and makes a decision consistent with the messages in both stages. In the case of "inconsistent" messages, the principal relies on the first-period report only. Because of such out-of-equilibrium beliefs, the expert cannot rebut the previous report in the second round.[14]

Second, dynamic informational control demonstrates some similarity to other instruments facilitating information transmission in cheap talk environments. Comparing it to multi-stage communication investigated by Krishna and Morgan (2004), one may notice that the expert's incentives to provide more information in both models are driven by active participation of the principal in the communication process. This active interaction between the players entails some uncertainty for the risk-averse expert, which is, nevertheless, might be affected by the expert's messages in the first round of communication. The nature of uncertainty in these models, however, is completely different. In Krishna and Morgan's model, the expert is imperfectly informed about future decisions of the principal because of a random binary outcome of the simultaneous dialog in the first round. In particular, if the outcome of the dialog is "success", the expert may update her information in the next stage. Thus, even though the expert might not be allowed to improve her report in the second round, the uncertainty about future interaction affects the expert's incentives in the first stage. Moreover, in the case of "success", the expert actually updates her first-round message by revealing more precise information in the second round. Together, these factors provide an overall improvement over single-stage direct communication. In the model of informational control, uncertainty arises directly from expert's imprecise information about the true realization of the state. The principal may influence this uncertainty by amending the

[13] Otherwise, if the set of feasible decisions did not change over time, the multi-period game would be equivalent to the single-stage game.

[14] This intuition is somewhat similar to that used in simultaneous communication with multiple experts by Krishna and Morgan (2001a). In the case of consulting two experts with opposing biases, the principal uses each of the two messages in order to verify the other one and believes one of the experts only after receiving "inconsistent" messages. In particular, the credible expert is the one who benefits less from sending her message under the assumption that the second message is truthful. This makes profitable rebuttal of the message of the rival expert impossible. In our model, there is a rivalry between the imperfectly informed expert in the first stage, who is interested in learning more information, and the expert with updated information in the second stage, who wants to manipulate the principal's decision. These factors make the expert's report in period one more credible.

second-stage information structure, which depends on the first-period report of the expert.

Dynamic informational control also inherits some features of the more general class of mediated communication protocols introduced by Myerson (1991) and studied by Blume et al. (2007), Goltsman et al. (2009), and Ivanov (2010b) in the cheap talk context. In mediated communication, the expert faces the message-contingent lottery over induced actions controlled by the mediator, who privately requests information from the expert and gives recommendations to the decision maker. In our model, the expert faces the message-contingent lottery over her own posterior types generated by the information structure in period one. The natural question to ask is why such a lottery over the expert's types is more effective than the lottery over the principal's decisions. The intuition is based on two factors. First, because the set of types is a continuum, the mediated protocol of communication has to be delicate in punishing the expert for lying. Suppose that the mediator aims to punish the expert of type θ for placing some spin on information and claiming type θ' close to θ by inducing some unfavorable distribution over actions. By the continuity of the expert's payoff function, such a lottery also punishes type θ' for telling the truth and motivates this type to distort information. In contrast, the lottery over posterior types is able to separate the magnitudes of expert's punishment for distorting information across first-stage types (signals). As a result, each type of the expert can be punished even for slight distortions of information without damaging the incentives of arbitrarily close types to communicate their information. Reporting a message s' different from the observed signal s in the first period does not provide new information to the particular type s only. These informational losses may lead to the lower payoff to the expert even in the case of inducing the most favorable decision $y_1(s,b)$. This effect, however, does not influence the expert's type s', who can still get high informational benefits after telling the truth.

Second, the incentives of the risk-averse expert to provide more information to the mediator are purely driven by the randomness in the mediator's recommendations to the principal. Because each expert's message maps into several actions, which cannot be the ideal principal's decisions together, inducing the ideal decision for each state is not feasible. In the case of dynamic informational control, there is no need to introduce randomness over actions after the expert learns the state precisely in the second period. By that time, the principal knows the distribution of posterior types, which are distant enough to separate themselves without additional incentives.

In the light of Theorem 1, an important question to examine is whether full information extraction is monotone in b. In other words, what are the conditions on the primitives, i.e., the players' payoff functions and the distribution of states, which sustain full information extraction if and only if the magnitude of the expert's bias is below some cut-off? To address this issue, consider the difference between the payoffs in the cases of truthful communication and inducing an arbitrary decision $y \in \mathbb{R}$:

$$\begin{aligned}\Delta V(y,s,b) &= E\left[V\left(a_p(\theta),\theta,b\right)|s\right] - E\left[V(y,\theta,b)|s\right] \\ &= p_s V\left(a_p(s),s,b\right) + (1-p_s)\,V\left(a_p(\varphi(s)),\varphi(s),b\right) \\ &\quad - p_s V(y,s,b) - (1-p_s)\,V(y,\varphi(s),b)\,.\end{aligned}$$

Suppose that $y_1^*(s,b) \in \mathbb{R}$ is the expert's optimal, but possibly non-feasible decision given her first-stage information:

$$y_1^*(s,b) = \arg\max_{a\in\mathbb{R}} \; E\left[V(a,\theta,b)\,|s\right]$$
$$= \arg\max_{a\in\mathbb{R}} \; p_s V(a,s,b) + (1-p_s)\,V(a,\varphi(s),b)\,, s\in[0,1/2).$$

Then, if the difference in the payoffs in the cases of truthful communication and inducing the (non-feasible) decision $y_1^*(s,b)$ is not increasing in b, then the fully informative equilibrium exists if and only if the bias is below some cut-off.

Lemma 1 *If $\Delta V_b'(y=y_1^*(s,b),s,b) \le 0, s < 1/2$, then there is the fully informative equilibrium in the game with the communication schedule determined by (2) and (3) if and only if $b\le\bar{b}$.*

The main reason for using $y_1^*(s,b)$ instead of $y_1(s,b)$ is that it is not necessary to check if the boundary condition $y_1(s,b)\le a_p(1)$ is binding for all s and b. As a result, verifying the sign of $\Delta V_b'(y=y_1^*(s,b),s,b)$ is a significantly less complicated exercise than checking it for $\Delta V_b'(y=y_1(s,b),s,b)$. The following corollary is a straightforward implication of Lemma 1 for a special class of preferences.

Lemma 2 *If $V(a,\theta,b) = V(a-\beta(b,\theta),\theta)$, where $\beta(0,\theta)=0, \beta_b'\ge 0$, and $\beta_{b\theta}''\le 0$, then there is the fully informative equilibrium in the game with the communication schedule determined by (2) and (3) if and only if $b\le\bar{b}$.*

The payoff functions of type $V(a,\theta,b) = V(a-\beta(b,\theta),\theta)$ are broadly used in the literature on strategic communication and the mechanism design.[15] The first pair of conditions implies that the expert's interests coincide with those of the principal if the bias parameter $b=0$, and that the magnitude of the expert's bias is increasing in b. The third condition implies that the marginal magnitude of the bias is not increasing in θ.

In general, the relationship between the marginal benefits of inducing $y_1(s,b)$ or reporting truthfully is not monotone in the expert's bias. As b falls, the difference in the payoffs $E\left[V(a_p(\theta),\theta,b)\,|s\right]$ and $E\left[V(y_1(s,b),\theta,b)\,|s\right]$ may either increase or decrease. On many occasions, however, $\Delta V(y_1(s,b),s,b)$ satisfies a weaker condition. The following lemma gives the details.

Lemma 3 *If $\Delta V(y_1(s,b),s,b)$ is quasi-monotone in b on $[0,b_0(s)]\,, s\in[0,1/2)$, where $b_0(s)$ is uniquely given by $a_p(\varphi(s)) = y_1(s,b_0(s))$, then there is the fully informative equilibrium in the game with the communication schedule determined by (2) and (3) if and only if $b\le\bar{b}$.*[16]

[15] Besides the quadratic payoff function with a constant bias, this class includes, for example, the generalized quadratic functions $V(a,\theta,b) = -(a-y(\theta,b))^2$ used by Alonso and Matouschek (2008) and Kovác and Mylovanov (2009), and the symmetric functions $V(a,\theta,b) = V(|a-(\theta+b)|)$ exploited by Dessein (2002). Krishna and Morgan (2004) use a special case of the symmetric functions, $V(a,\theta,b) = -|a-(\theta+b)|^\rho$.

[16] The function $H(b)$ is *quasi-monotone* in b on $[0,b_0]$ if $H(b_1)\ge 0$ for some $b_1\in[0,b_0]$ implies $H(b)\ge 0, b\in[0,b_1)$.

According to Lemma 3, the quasi-monotonicity of $\Delta V(y_1(s,b), s, b)$ is required only for relatively small values of the bias. Otherwise, we show that if either the bias or the probability of the higher posterior $\varphi(s)$ is so large that the most profitable "deviating" decision $y_1(s,b)$ exceeds $a_p(\varphi(s))$, then $\Delta V(y_1(s,b), s, b)$ is monotone in b. Also, note that the conditions in Lemmas 1 and 3 are related to the monotonicity of ranking of the expert's payoffs $E[V(a_p(\theta), \theta, b)|s]$ and $E[V(y_1(s,b), \theta, b)|s]$ only. In fact, the monotonicity of the expert's second-stage incentive-compatibility constraints does not require extra conditions on the payoff function $V(a, \theta, b)$ or the distribution $F(\theta)$. The following example highlights the differences in the conditions on $\Delta V(y_1(s,b), s, b)$ used in Lemmas 1 and 3.

Example 2. Suppose $V(a,\theta,b) = -(a - \theta - b\theta)^2$ and $\varphi(\theta) = \theta + 1/2$. Then, $a_p(s) = s$ and $y_1^*(s,b) = (1+b)(p_s s + (1-p_s)(s+1/2)), s \leq 1/2$. If $b \leq b_a(s) = \frac{p_s}{2-p_s}$, then $y_1^*(s,b) \leq a_p(1) = 1$. That is, $y_1(s,b) = y_1^*(s,b)$ and

$$\begin{aligned}\Delta V(y_1(s,b), s, b) &= -b^2\left(p_s s^2 + (1-p_s)\left(s + \frac{1}{2}\right)^2\right) + \frac{p_s(1-p_s)(1+b)^2}{4} \\ &= -\frac{(1+2s-p_s)^2}{4}b^2 + \frac{p_s(1-p_s)}{2}b + \frac{p_s(1-p_s)}{4}.\end{aligned}$$

This function is not monotone in $b \geq 0$. It is, however, concave in b. This property, along with $\Delta V(y_1(s,0), s, 0) = \frac{p_s(1-p_s)}{4} > 0$, implies that $\Delta V(y_1(s,b), s, b)$ is quasi-monotone in b for all s, and $\Delta V(y_1(s,b), s, b) \geq 0$ if and only if

$$b \leq b_y(s) = \frac{p_s(1-p_s) + (p_s(1-p_s)(4s - p_s + 4s^2 - 4p_s s + 1))^{1/2}}{(1+2s-p_s)^2}.$$

Let $f(\theta) = \frac{\alpha e^{-\alpha\theta}}{1-e^{-\alpha}}, \theta \in [0,1]$, where $\alpha = \ln 4 \simeq 1.386$. Then, $p_s = \left(1 + \exp\left(-\frac{\ln 4}{2}\right)\right)^{-1} = 2/3, s \leq 1/2$. This results in $b_a(s) = 1/2$ and

$$b_y(s) = \frac{2 + (6 + 24s + 72s^2)^{1/2}}{(1+6s)^2},$$

which is decreasing in s and $b_y(1/2) = 1/2$.

Finally, the second-stage incentive-compatibility constraints require

$$a_p(\varphi(s)) - a_p(s) = \varphi(s) - s = \frac{1}{2} \geq 2sb \text{ for all } s < 1/2, \text{ or } b \leq \frac{1}{2}.$$

Therefore, the fully informative equilibrium is sustainable if and only if $b \leq 1/2$.

4.3 Unbounded state space

If the state space is unbounded, the expert's first-stage information structure can be partitioned into two-point sets, such that the distance between any joint states θ and $\theta' = \varphi(\theta)$ is arbitrarily large. This might suppress the expert's incentive to distort her information because of two factors. First, a large distance between the joint states implies that the optimal decision y_1 in the case of observing the signal s and deviating from truthtelling in period one would differ significantly from at least one of the ideal decisions $y(s,b)$ and $y(\varphi(s),b)$. As a result, the expected losses from distorting information should increase. Second, upon communicating truthfully in the first stage and observing $s < s' = \varphi(s)$ in the next period, inducing $a_p(s')$ results in high losses of the expert if s' is sufficiently distant from s. These observations suggest that the principal can potentially extract all information from the expert for an arbitrary magnitude of the expert's bias.

The intuition above, however, misses one subtlety. The value of updated information in the second stage is high if the joint states are quite distinct *and* approximately equally likely. For a fixed state θ, if the distance between the joint states θ and $\theta' > \theta$ increases, then the density $f(\theta')$ eventually converges to 0 and the posterior probability $p_\theta = \frac{f(\theta)}{f(\theta)+f(\theta')}$ converges to 1. Thus, the expert infers that the posterior realization $\theta = s$ is more likely than $\theta' = \varphi(s)$. This decreases her benefits of being perfectly informed in the second stage and, as a result, the incentives to truthfully convey information in the first stage. These arguments imply that truthtelling can be sustained only if the magnitude of $f(\theta)$ is comparable to that of $f(\theta')$. The example below highlights this logic.

Example 3. Suppose $f(\theta) = \alpha e^{-\alpha\theta}, \theta \geq 0$ and the payoff functions are quadratic, $U(a,\theta,b_i) = -(a-\theta-b_i)^2$, $a_p(\theta) = \theta$, and $y(\theta,b) = \theta + b$.

We design the algorithm of updating expert's information by replicating the communication schedule in the case of the bounded state space. First, split $\Theta = \mathbb{R}_+$ into intervals $[kd, kd+d), k \in \mathbb{N}_0 = \{0,1,...\}$ of a fixed length $d > 0$. Then, partition each interval $[kd, kd+d)$ into the joint-state sets $\{\theta, \varphi(\theta)\} = \left\{\theta, \theta + \frac{d}{2}\right\}, \theta \in [kd, kd+\frac{d}{2})$. For each pair of θ and $\theta' = \varphi(\theta)$, the expert observes the signal $s = \theta$. The second-stage information structure is constructed in a similar fashion. After sending the first-period message $m_1 \in [kd, kd+\frac{d}{2})$ for some $k \in \mathbb{N}_0$, the expert can separate m_1 from $\varphi(m_1)$, but does not acquire new information if the state is not in the set $\{m_1, \varphi(m_1)\}$.

Given the first-stage signal $s \in [kd, kd+\frac{d}{2})$ for some $k \in \mathbb{N}_0$, reporting the truth in both stages results in the payoff to the expert $E[V(a_p(\theta),\theta,b)|s] = -b^2$. At the same time, inducing the expert's optimal decision conditional on the first-stage information:

$$\begin{aligned} y_1(s,b) &= \arg\max_{y\in\mathbb{R}_+} p_s(a-s-b)^2 - (1-p_s)(a-\varphi(s)-b)^2 \\ &= p_s s + (1-p_s)\varphi(s) + b = s + b + \frac{d}{2}\frac{e^{-\alpha\frac{d}{2}}}{1+e^{-\alpha\frac{d}{2}}}, \end{aligned}$$

where $p_s = \left(1+e^{-\alpha\frac{d}{2}}\right)^{-1}$, provides the payoff:

$$E[V(y_1(s,b),\theta,b)|s] = -p_s(1-p_s)(\varphi(s)-s)^2 = -\frac{e^{-\alpha\frac{d}{2}}}{1+e^{-\alpha\frac{d}{2}}}\frac{d^2}{4} = -\frac{z^2e^{-z}}{\alpha^2(1+e^{-z})^2}, \quad (8)$$

where $z = \frac{\alpha d}{2} > 0$. Because $e^{-z} \in (0,1), z > 0$, it easily follows from (8) that

$$\phi(z^*) \leq \phi(z) < E[V(y_1(s,b),\theta,b)|s] < \frac{\phi(z)}{4},$$

where $\phi(z) = -\frac{z^2 e^{-z}}{\alpha^2}$ and $z^* = \underset{z \geq 0}{\arg\min}\ \phi(z) = 2$. If $\alpha b \geq 4/e$, then

$$E\left[V\left(y_1\left(s,b\right),\theta,b\right)|s\right] > \phi\left(z^*\right) = -\frac{4}{\alpha^2 e^2} \geq -b^2 = E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right]. \tag{9}$$

In contrast, if $\alpha b \leq 1/e$, then taking $d = 2z^*/\alpha = 4/\alpha$ results in

$$E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] = -b^2 \geq \frac{\phi\left(z^*\right)}{4} = -\frac{1}{\alpha^2 e^2} > E\left[V\left(y_1\left(s,b\right),\theta,b\right)|s\right]. \tag{10}$$

Finally, the second-stage incentive-compatibility constraints (1) require

$$\varphi\left(s\right) - s = s + \frac{d}{2} - s \geq 2b,$$

or equivalently, $d \geq 4b$. This inequality holds if $\alpha b \leq 1/e$, because $d = 4/\alpha \geq 4be > 4b$.

The intuition behind (9) is gained from the following observation. For a fixed bias b, if α is sufficiently large, the density $f\left(\theta'\right)$ falls quickly as the distance between the joint states θ and $\theta' > \theta$ increases. Because of that effect, the expert knows almost surely that the true realization of the state is $\theta = s$ rather than $\theta' = \varphi\left(s\right)$ after observing signal s. In this case, her benefits of learning θ perfectly and inducing the principal's ideal action in stage two are relatively low. It is more profitable for the expert to remain imperfectly informed, but induce the optimal decision $y_1\left(\theta,b\right)$, which is close to her ideal point $y\left(\theta,b\right)$. On the other hand, (10) means that if the prior distribution is smooth enough (i.e., $\alpha \leq 1/be$), moving the joint states sufficiently apart relatively to the magnitude of the bias preserves the expert's uncertainty about each posterior realization. As a result, full information extraction is sustainable for an arbitrary bias if the density is sufficiently smooth.

In order to formalize the logic above, consider the distribution of states with a positive density $f\left(\theta\right) > 0$ on $\Theta = \mathbb{R}_+$ and the payoff function $U\left(a,\theta\right)$, such that the set of rationalizable decisions is unbounded, $\mathcal{A} = \left[a_p\left(0\right),\infty\right)$. (The case of $\Theta = \mathbb{R}$ is identical). Let $y_1^*\left(\theta,\theta',b\right)$ be the expert's optimal decision upon learning that the state is either θ or θ' with equal probabilities $p_\theta = p_{\theta'} = 1/2$:

$$y_1^*\left(\theta,\theta',b\right) = \underset{a \in \mathcal{A}}{\arg\max}\ V\left(a,\theta,b\right) + V\left(a,\theta',b\right). \tag{11}$$

We suppose that $V\left(a,\theta,b\right)$ satisfies the following three conditions.

A1: For any b, there exists $\delta\left(b\right) > 0$, such that $y\left(\theta,b\right) - a_p\left(\theta\right) \leq \delta\left(b\right), \forall\theta$.

A2: For any b and $\delta > 0$, there exists $\tilde{\delta}\left(b,\delta\right) > 0$, such that $V\left(y\left(\theta,b\right) - \delta,\theta,b\right) > V\left(y\left(\theta,b\right) + \tilde{\delta}\left(b,\delta\right),\theta,b\right), \forall\theta$.

A3: For any b and $\delta > 0$, there exists $d\left(b,\delta\right) > 0$, such that $y\left(\theta,b\right) + \delta < y_1^*\left(\theta,\theta + d\left(b,\delta\right),b\right) < y\left(\theta + d\left(b,\delta\right),b\right) - \delta, \forall\theta$.

The first condition states that there is a uniform bound on the expert's bias. The second one requires that the expert is not infinitely more sensitive to the actions below her ideal decision than to the actions above it. Finally, (A3) is the "betweenness" condition, which establishes that moving the joint states θ and θ' sufficiently apart from each other guarantees that the expert's optimal decision given her information that the true state is either θ or θ' is not too close to the ideal decisions for either realization. These conditions do not impose strong restrictions on the shape of the expert's payoff function. In fact, imposing only (A1) might be sufficient for satisfying the other two assumptions.[17] Given these preliminaries, the following theorem characterizes the scenarios, in which the principal can extract full information from the expert with an arbitrarily large bias.

Theorem 2 *If $V(a,\theta,b)$ satisfies conditions A1–A3, then for any b, there exist $d(b)>0$ and $\varepsilon(\theta,b)\in(0,1)$, such that if $\varepsilon(\theta,b)\leq\frac{f(\theta+d(b))}{f(\theta)}\leq\frac{1}{\varepsilon(\theta,b)},\forall\theta$, then there is the communication schedule that sustains the fully informative equilibrium.*

In order to prove this result, we employ the communication schedule similar to that used in Example 3. First, the state space is divided into the sequence of half-open intervals of a fixed length, which depends on the expert's bias. Then, each interval is partitioned into joint-state sets, such that the distance between the joint states is equal to a half of the interval's length. After observing the joint-state set that contains the state in the first stage, the expert can either reveal it truthfully and learn the state precisely in the second stage or induce any rationalizable decision by manipulating her information. The conditions (A1)–(A3) imply that if the probabilities of posterior realizations in each joint-state set are equal, then moving the joint states sufficiently apart from each other by increasing the length of intervals forces the expert to communicate truthfully. Because it is impossible to keep the identical posterior probabilities for any distribution with an unbounded support, the condition in Theorem 2 means that the likelihood ratio between any states θ and $\theta+d$ must not vary significantly, where d depends on the magnitude of the bias. In this case, the signal that the state is in some joint-state set $\{\theta,\theta+d\}$ does not provide much information to the expert, because she is sufficiently uncertain about both posterior realizations. This guarantees that the expert's benefits from perfectly learning the state in the second period still exceed the benefits of optimal manipulation of imprecise first-stage information.

To demonstrate how Theorem 2 can be applied to the particular payoff functions, consider now the family of generalized quadratic functions $V(a,\theta,b)=-(a-\theta-\beta(b,\theta))^2$, and suppose that the distribution function $F(\theta)$ satisfies the following regularity condition.

A4: There exist $\alpha>0$ and $d>0$, such that

$$e^{-\alpha d}\leq\frac{f(\theta+d)}{f(\theta)}\leq e^{\alpha d},\theta\geq 0.$$

The regularity condition (A4) states that the likelihood ratio in densities $f(\theta)$ and $f(\theta+d)$ varies less rapidly than exponentially with the rate α for all states. That is, for any pair of states θ and $\theta+d$, the prior distribution varies more "smoothly" than the exponential distribution with the parameter α. Clearly, if (A4) holds for some d and α_0,

[17] Consider the symmetric payoff function $V(a,\theta,b)=V(|a-\theta-\beta(b,\theta)|)$, where $V'(x)<0,V''(x)<0,x\geq 0,\beta(0,\theta)\geq 0$, and $\beta_b'(b,\theta)>0$. That is, $a_p(\theta)=\theta,y(\theta,b)=\theta+\beta(b,\theta)$, and $y_1^*(\theta,\theta+d,b)=\theta+\frac{d}{2}+\frac{\beta(b,\theta)+\beta(b,\theta+d)}{2}$. For a fixed b, (A1) requires that $\beta(b,\theta)\leq\delta(b),\forall\theta$. This implies $V(y(\theta,b)-\delta,\theta,b)=V(\delta)>V\left(\tilde{\delta}\right)=V\left(y(\theta,b)+\tilde{\delta},\theta,b\right)$ if $\tilde{\delta}>\delta$, so (A2) holds. Also, $y(\theta,b)+\delta<y_1^*(\theta,\theta+d,b)<y(\theta+d,b)-\delta$ if $d>\delta(b)+2\delta$, so (A3) holds as well.

then it also holds for d and $\alpha > \alpha_0$. The following lemma establishes the sufficient conditions for full information extraction under the above assumptions.

Lemma 4 *If $V(a,\theta,b) = -(a-\theta-\beta(b,\theta))^2, \beta(b,\theta) \leq \delta$, and $F(\theta)$ satisfies (A4) for $d = 2\delta(2/z^*+1)$ and $\alpha = z^*/\delta$, where $z^* = \exp(-z^*/2-1) \simeq 0.314$, then there exists the communication schedule, which sustains the fully informative equilibrium.*

This lemma highlights the relationship between the magnitudes of the expert's bias and the smoothness in the density. In particular, for any upper bound on the bias, there exists an upper bound on the absolute value of the variation in the likelihood ratio between any two states, such that the distance between these states is fixed and proportional to the bias. These conditions are required to move the joint states apart from each other while keeping the expert sufficiently uncertain about each state after receiving the first-period signal.

4.4 Multi-dimensional state space

In this section, we extend the above results to the N–dimensional state space. Suppose that the state space and the decision set are $N-$dimensional: $a = (a_1, ..., a_N) \in \mathbb{R}^N$ and $\theta = (\theta_1, ..., \theta_N) \in \Theta = [0,1]^N, N \geq 2$. The density function of the state $f(\theta)$ is bounded and positive on Θ. The players' payoff functions are of the form $U(a,\theta,b^i), i \in \{P,E\}$, where $b^P = 0$, and $b^E = b > 0$. Similarly to the one-dimensional case, the function $U(a,\theta,b)$ is twice-differentiable in (a,θ,b), strictly concave in a for all (θ,b), and has a unique ideal decision $a^*(\theta,b) = \arg\max_a U(a,\theta,b)$ for each (θ,b). Also, $U(a,\theta) = U(a,\theta,0)$ and $V(a,\theta,b) = U(a,\theta,b)$ stand for the principal's and the expert's payoff functions, and $a_p(\theta) = a^*(\theta,0)$ and $y(\theta,b) = a^*(\theta,b)$ stand for the ideal decision of the principal and the expert, respectively. We assume that each ideal decision of the principal corresponds to a single state, or $a_p(\theta) \neq a_p(\theta'), \theta' \neq \theta$.[18] Finally, the expert of the highest type in each dimension $\theta_1 = \{1\}^N$ has no incentives to distort her information: $V(a_p(\theta_1), \theta_1, b) \geq V(a, \theta_1, b), a \in \mathcal{A}$.

We construct the communication schedule by extending the expert's information structures (2) and (3) to the multi-dimensional state space. Define the sequence of the "low-state" subsets of Θ:

$$Q_k = \{\theta \in \Theta | \theta_1 = ... = \theta_{k-1} = 1, 0 \leq \theta_k < 1/2\}, k = 1, ..., N,$$

and the sequence of the complement state functions:

$$\varphi_k(\theta) = (\theta_1, ..., \theta_{k-1}, \varphi(\theta_k), \theta_{k+1}, ..., \theta_N), \theta \in Q_k, k = 1, ..., N,$$

where function $\varphi(\theta)$ is similar to the one-dimensional case: it is continuous, strictly increasing, $\varphi(0) = 1/2$ and $\varphi(1/2) = 1$. Then, the union of sets $\phi_1(\theta) = \{\theta, \varphi_1(\theta)\}$ over Q_1:

$$\Theta_1 = \bigcup_{\theta \in Q_1} \phi_1(\theta),$$

[18] In the one-dimensional case, it is guaranteed by the strict monotonicity of $a_p(\theta)$ in θ.

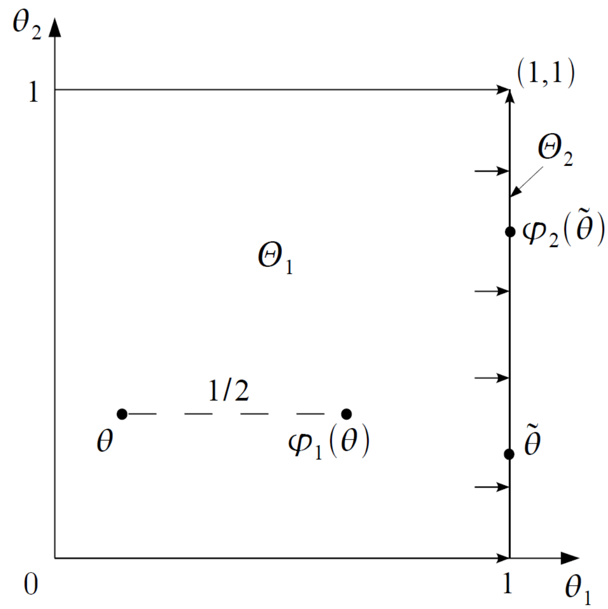

Figure 1: Two-dimensional communication schedule

fills the N–dimensional cube $\Theta = [0,1]^N$ up to the zero-measure set

$$\bar{\Theta}_1 = \Theta \backslash \Theta_1 = \{\theta \in \Theta | \theta_1 = 1\},$$

which is the $(N-1)$–dimensional cube. Then, we can apply the same algorithm of partitioning $\bar{\Theta}_1$ into the joint-state sets $\phi_2(\theta) = \{\theta, \varphi_2(\theta)\}, \theta \in Q_2$. The set $\Theta_2 = \bigcup_{\theta \in Q_2} \phi_2(\theta)$ fills $\bar{\Theta}_1$ up to the set

$$\bar{\Theta}_2 = \bar{\Theta}_1 \backslash \Theta_2 = \{\theta \in \Theta | \theta_1 = \theta_2 = 1\},$$

which is the $(N-2)$-dimensional set. Repeat this procedure until it converges to $\bar{\Theta}_N = \{1\}^N$. By construction, it follows that

$$\Theta = \Theta_1 \cup \bar{\Theta}_1 = \Theta_1 \cup \Theta_2 \cup \bar{\Theta}_2 = ... = \Theta_1 \cup ... \cup \Theta_N \cup \bar{\Theta}_N = \bigcup_{k=1,...,N} \Theta_k \cup \{1\}^N.$$

Define the expert' first-stage information structure as follows. Partition Θ into the collection of the joint-state sets

$$I_1 = \left\{ \{\phi_k(\theta), \theta \in Q_k\}_{k=1}^N, \{1\}^N \right\}. \tag{12}$$

If $\theta \in \phi_k(s)$ for some k and $s \in Q_k$, the expert observes signal s. That is, she cannot distinguish between the joint states $\theta = s$ and $\theta' = \varphi_k(s)$. (Fig. 1 illustrates the communication schedule for $N = 2$ and $\varphi(\theta) = \theta + 1/2$.) By sending a message $m \in Q_k$ in the first stage, the expert is able to distinguish only between the states m and $\varphi_k(m)$ in the second period:

$$I_2(m, I_1) = \{m, \varphi_k(m), I_1 \backslash \phi_k(m)\}, m \in Q_k. \tag{13}$$

If $\theta = \{1\}^N$, then the expert perfectly knows θ in the first period. The following theorem demonstrates that the communication schedule characterized by (12) and (13) sustains the fully informative equilibrium if the expert's bias is not large.

Theorem 3 *There exists $\bar{b} > 0$ such that if $b \leq \bar{b}$, there is the communication schedule determined by (12) and (13), which sustains the fully informative equilibrium.*

The logic behind the proof of Theorem 3 is a reminiscent of the result for the single-dimensional case. Because the principal's ideal decisions are different across the states, partitioning Θ into two-point joint-state sets guarantees that there is no pair of states, which map into the same ideal decision of the expert as long as the bias in preferences is small. From the expert's perspective, this implies that she is interested in learning the particular realization of the state upon observing the joint-state set. In other words, telling the truth in the first round provides her with the informational benefits in the second stage. If the bias is not large, these informational gains outweigh the benefits of inducing any rationalizable decision under imperfect first-stage information.

5 Conclusion

In this paper, we have demonstrated how simple dynamic management of information assessed by a privately informed agent can completely remove the informational asymmetry between the agent and an uninformed decision maker. Though the parties have conflicting interests and interact by means of cheap-talk communication, providing proper access to information allows the principal to elicit perfect information about an unknown issue of communication and reach his first-best decision after only two rounds of communication.

A few comments are in order. First, the analysis and the results are equally applicable to the situation, in which the expert's information structure in the first stage is given by the partition, consisting of a collection of the two-interval sets $\phi_i = \{(\theta_1^i, \theta_2^i] \cup (z_1^i, z_2^i]\}_{i \in \mathcal{J}}, z_i^1 > \theta_i^2$ indexed by i. This partition can sustain the (approximately) fully informative equilibria if the density is zero for some states or significantly varies across the states. Suppose, for example, that the density is concentrated at high and/or intermediate states and virtually disappears if the state is near the lower bound of the support. In this case, partitioning the state into two-point sets collapses truthful communication, because the imbalance between the probabilities of states reduces the expert's uncertainty about posterior realizations in the first period. However, by locally partitioning the subinterval of states around points with zero or rapidly varying density into two-interval sets, it is possible to keep the expert sufficiently uncertain. This provides her the incentive to communicate truthfully in order to learn the particular interval, which contains the state, in the next round (see Appendix B for an example).

In some circumstances, full information extraction can be achieved by expanding the class of the complement state functions $\varphi(\theta)$. This is especially effective in the case of symmetric distributions with zero density at some states. Consider, for example, the symmetric triangle distribution supported on the unit interval. If the class of the complement state functions is restricted to increasing functions that map from the lower half into the upper half of the unit interval, then there is no fully informative equilibrium for any positive bias because of the aforementioned imbalance between the probabilities of posterior realizations. In contrast, if the first-stage information structure partitions the boundary subintervals $[0, \theta_1)$ and $(1-\theta_1, 1]$ into the joint-state sets $\{\theta, 1-\theta\}$, then the posterior realizations of the state in each set are equally likely. Because the density is positive on the subinterval $[\theta_1, 1-\theta_1]$, it is possible to fully extract information for some range of biases by using an increasing complement state function for partitioning this subinterval.

Appendix A

In this section we provide the proofs of the results.

Proof of Theorem 1. Because $f(s) > 0, s \in [0,1]$, we have $p_s = \frac{f(s)}{f(s)+f(\varphi(s))} > 0$ and $1 - p_s > 0$. The continuity of $f(\theta)$, the differentiability of $V(a, \theta, b)$ and $y(\theta, b)$ and the continuity of $\varphi(\theta)$ in (a, θ, b) imply that $V(a_p(s), s, b) - V(a_p(\varphi(s)), s, b), E[V(y_1(s,b), \theta, b)|s]$, and $E[V(a_p(\theta), \theta, b)|s]$ are continuous in (s, b). Also, $V''_{a\theta} > 0$ and $\varphi(s) - s > 0$ imply $y(s,b) < y(\varphi(s), b), s \in [0, 1/2]$.

Consider the case of $s = 1$, so the expert infers that $\theta = 1$. Thus, her information will not be updated in the second stage unconditionally on the message m_1. Sending $m_1 = 1$ induces the principal's decision $a_p(1)$ for any message m_2 in $t = 2$. Sending $m_1 < 1/2$ results in the principal's decision $a \in \{a_p(m_1), a_p(\varphi(m_1))\}$ depending on m_2. Because $a_p(m_1) < a_p(\varphi(m_1)) < a_p(1) < y(1, b)$ and $V'_a(a, \theta, b) > 0, a < y(\theta, b)$, we have $V(a_p(1), 1, b) > V(a_p(\varphi(m_1)), 1, b) > V(a_p(m_1), 1, b), m_1 < 1/2$, so the expert reveals $\theta = 1$ truthfully.

Given $s \in [0, 1/2)$ in $t = 1$, the expert has two options. First, she can induce the optimal decision $y_1(s, b)$ given her current information. Second, she can learn θ perfectly in the second stage by reporting s truthfully. In the latter case, the expert can induce only $a \in \{a_p(s), a_p(\varphi(s))\}$ in the second period.

Consider the expert's second-stage incentive compatibility constraints (7) given truthful reporting in $t = 1$. If $b = 0$, then $a_p(s) = y(s, 0), \forall s$. Because $V(y(s,b), s, b) > V(a, s, b), a \neq y(s, b)$, and $U''_{a\theta} > 0$ implies $a_p(\varphi(s)) > a_p(s)$, it follows that

$$V(a_p(s), s, 0) - V(a_p(\varphi(s)), s, 0) = V(y(s, 0), s, 0) - V(y(\varphi(s), 0), s, 0) > 0, s \in [0, 1/2].$$

That is, (7) is strictly satisfied for all $s \in [0, 1/2]$. By the continuity of $V(a_p(s), s, b) - V(a_p(\varphi(s)), s, b)$ in (s, b), it follows that there is $b_c > 0$, such that

$$V(a_p(s), s, b) \geq V(a_p(\varphi(s)), s, b), s \in [0, 1/2], b \leq b_c.$$

In the case of inducing $y_1(s, b)$, suppose that $b = 0$. Then, $y(\varphi(s), 0) = a_p(\varphi(s)) > a_p(s) = y(s, 0)$ implies that $y_1(s, 0) \neq y(s, 0)$ and/or $y_1(s, 0) \neq y(\varphi(s), 0)$. Thus,

$$\begin{aligned} E[V(y_1(s,0), \theta, 0)|s] &= p_s V(y_1(s,0), s, 0) + (1-p_s) V(y_1(s,0), \varphi(s), 0) \\ &< p_s V(y(s,0), s, 0) + (1-p_s) V(y(\varphi(s), 0), \varphi(s), 0) \\ &= p_s V(a_p(s), s, 0) + (1-p_s) V(a_p(\varphi(s), 0), \varphi(s), 0) \\ &= E[V(a_p(\theta), \theta, 0)|s], s \in [0, 1/2], \end{aligned}$$

where $E[V(a_p(\theta), \theta, b)|s]$ is the expert's payoff in the case of truthful reporting in both stages. Because $E[V(y_1(s,b), \theta, b)|s]$ and $E[V(a_p(\theta), \theta, b)|s]$ are continuous in (s,b), there exists $b_y > 0$, such that

$$E[V(a_p(\theta), \theta, b)|s] \geq E[V(y_1(s,b), \theta, b)|s], s \in [0, 1/2], b \leq b_y.$$

Thus, taking $b \leq \bar{b} = \min\{b_c, b_y\}$ sustains the fully informative equilibrium. ■

Proof of Lemma 1. Suppose that there exists the fully informative equilibrium for some $\bar{b} > 0$. That is:

$$V\left(a_p(s), s, \bar{b}\right) - V\left(a_p(\varphi(s)), s, \bar{b}\right) \geq 0, \tag{14}$$

and

$$\begin{aligned}\Delta V\left(y_1\left(s, \bar{b}\right), s, \bar{b}\right) &= E\left[V\left(a_p(\theta), \theta, \bar{b}\right)|s\right] - E\left[V\left(y_1\left(s, \bar{b}\right), \theta, \bar{b}\right)|s\right] \\ &= p_s V\left(a_p(s), s, \bar{b}\right) + (1-p_s) V\left(a_p(\varphi(s)), \varphi(s), \bar{b}\right) \\ &- p_s V\left(y_1\left(s, \bar{b}\right), s, \bar{b}\right) - (1-p_s) V\left(y_1\left(s, \bar{b}\right), \varphi(s), \bar{b}\right) \geq 0,\end{aligned}$$

for all $s \in [0, 1/2)$. We prove that there exists the fully informative equilibrium for any $b < \bar{b}$.

Because $a_p(\varphi(s)) > a_p(s)$ and $V''_{ab} > 0$, then $V'_b(a_p(s), s, b) - V'_b(a_p(\varphi(s)), s, b) < 0, s \in [0, 1/2)$. This implies that the second-stage incentive-compatibility constraints are satisfied for $b < \bar{b}$:

$$V(a_p(s), s, b) - V(a_p(\varphi(s)), s, b) > V\left(a_p(s), s, \bar{b}\right) - V\left(a_p(\varphi(s)), s, \bar{b}\right) \geq 0, b < \bar{b}. \tag{15}$$

This also means $a_p(\varphi(s)) > y(s,b), s \in [0, 1/2), b < \bar{b}$. (Otherwise, if $a_p(\varphi(s)) \leq y(s,b)$, then $a_p(s) < a_p(\varphi(s)) \leq y(s,b)$ and $V'_a(a, \theta, b) > 0, a < y(\theta, b)$ result in $V(a_p(s), s, b) < V(a_p(\varphi(s)), s, b)$).

Since $V''_{ab} > 0$, then $V'_b(a, s, b)$ is increasing in a. Because $a_p(s) < y(s,b), \forall s, b > 0$, we have:

$$V'_b(a_p(s), s, b) < V'_b(y(s,b), s, b) \leq 0.$$

This leads to the inequality:

$$E\left[V'_b(a_p(\theta), \theta, b)|s\right] = p_s V(a_p(s), s, b) + (1-p_s) V(a_p(\varphi(s)), \varphi(s), b) < 0. \tag{16}$$

Because $y_1(s,b) > a_p(0)$, we need to consider two cases: $y_1(s,b) < a_p(1)$ and $y_1(s,b) = a_p(1)$. First, let $y_1(s,b) < a_p(1)$. In this case, $y_1(s,b) = y_1^*(s,b)$ and $\Delta V(y_1(s,b), s, b) = \Delta V(y_1^*(s,b), s, b)$. This implies $\Delta V'_b(y_1(s,b), s, b) = \Delta V'_b(y_1^*(s,b), s, b) \leq 0$. Second, let $y_1(s,b) = a_p(1)$, which means $y_1(s,b) > a_p(\varphi(s)) > a_p(s)$. Since $V''_{ab} > 0$, then $V'_b(a_p(s), s, b) < V'_b(a_p(1), s, b)$ and $V'_b(a_p(\varphi(s)), \varphi(s), b) < V'_b(a_p(1), \varphi(s), b)$. By the envelope theorem, $\frac{d}{db} E[V(y_1(s,b), \theta, b)|s] = E[V'_b(y_1(s,b), \theta, b)|s]$ results in:

$$\begin{aligned}&\frac{d}{db}\Delta V(y_1(s,b), s, b) = p_s V'_b(a_p(s), s, b) + (1-p_s) V'_b(a_p(\varphi(s)), \varphi(s), b) \\ &- p_s V'_b(y_1(s,b), s, b) - (1-p_s) V'_b(y_1(s,b), \varphi(s), b) = p_s\left(V'_b(a_p(s), s, b) - V'_b(a_p(1), s, b)\right) \\ &+ (1-p_s)\left(V'_b(a_p(\varphi(s)), \varphi(s), b) - V'_b(a_p(1), \varphi(s), b)\right) < 0.\end{aligned}$$

Thus, $\Delta V'_b(y_1(s,b), s, b) \leq 0$ and $\Delta V(y_1(s,b), s, b) \geq \Delta V\left(y_1\left(s, \bar{b}\right), s, \bar{b}\right) \geq 0, b < \bar{b}$. ■

Proof of Lemma 2. If $V(a,\theta,b) = V(a-\beta(b,\theta),\theta)$, then $y_1^*(s,b)$ is given by the FOC:

$$p_s V_a'(y_1^*(s,b) - \beta(b,s), s) + (1-p_s) V_a'(y_1^*(s,b) - \beta(b,\varphi(s)), \varphi(s)) = 0,$$

where $V_a'(y_1^*(s,b) - \beta(b,s), s) < 0 < V_a'(y_1^*(s,b) - \beta(b,\varphi(s)), \varphi(s))$ due to $V_{a\theta}'' > 0$.

Because $V_b'(a-\beta(b,\theta),\theta) = -V_a'(a-\beta(b,\theta),\theta)\beta_b'(b,\theta)$ and $\beta_b'(b,s) \geq 0$, this implies $-V_a'(y_1^*(s,b) - \beta(b,s), s)\beta_b'(b,s) \geq 0$, and

$$\begin{aligned} E\left[V_b'(y_1^*(s,b),\theta,b)\,|s\right] &= -p_s V_a'(y_1^*(s,b) - \beta(b,s), s)\,\beta_b'(b,s) \\ &\quad - (1-p_s) V_a'(y_1^*(s,b) - \beta(b,\varphi(s)), \varphi(s))\,\beta_b'(b,\varphi(s)) \\ &\geq -\beta_b'(b,\varphi(s))\left(p_s V_a'(y_1^*(s,b) - \beta(b,s), s) + (1-p_s) V_a'(y_1^*(s,b) - \beta(b,\varphi(s)), \varphi(s))\right) = 0, \end{aligned}$$

where the inequality follows from $\beta_b'(b,s) \geq \beta_b'(b,\varphi(s))$ due to $\beta_{b\theta}''(b,\theta) \leq 0$. Thus, $E[V_b'(a_p(\theta),\theta,b)\,|s] < 0$ and $E[V_b'(y_1(s,b),\theta,b)\,|s] \geq 0$, which gives

$$\Delta V_b'(y_1(s,b),s,b) = E\left[V_b'(a_p(\theta),\theta,b)\,|s\right] - E\left[V_b'(y_1^*(s,b),\theta,b)\,|s\right] < 0.$$

■

Proof of Lemma 3. Suppose that there is the fully informative equilibrium for some $\bar{b}$. We show now that there exists the fully informative equilibrium for any $b < \bar{b}$. First, according to (15), the second-stage incentive-compatibility constraints hold for $b < \bar{b}$. Compare now the expert's payoffs in the case of inducing $y_1(s,b)$ and being truthful in both stages. The differentiability of $V(a,\theta,b)$ in (a,b) and $y(\theta,b)$ in b implies the differentiability of $\Delta V(y_1(s,b),s,b)$ in b. By the envelope theorem, $\frac{d}{db}E[V(y_1(s,b),\theta,b)\,|s] = E[V_b'(y_1(s,b),\theta,b)\,|s]$, and

$$\begin{aligned} \frac{d}{db}\Delta V(y_1(s,b),s,b) &= p_s\left(V_b'(a_p(s),s,b) - V_b'(y_1(s,b),s,b)\right) \\ &\quad + (1-p_s)\left(V_b'(a_p(\varphi(s)),\varphi(s),b) - V_b'(y_1(s,b),\varphi(s),b)\right). \end{aligned}$$

Note that $y_1(s,b) > a_p(s)$. Otherwise, if $y_1(s,b) \leq a_p(s) \leq y(s,b) < y(\varphi(s),b)$, then $V_a'(a,s,b) > 0, a < y(s,b)$ implies $V_a'(y_1(s,b),s,b) \leq 0, V_a'(y_1(s,b),\varphi(s),b) < 0$, and

$$E\left[V_a'(a,\theta,b)\,|s\right]|_{a=y_1(s,b)} = p_s V_a'(y_1(s,b),s,b) + (1-p_s) V_a'(y_1(s,b),\varphi(s),b) > 0,$$

which contradicts that $y_1(s,b) < a_p(1)$ is the solution to (5). By the same argument, $y_1(s,b) < a_p(1)$ implies $y_1(s,b) > y(s,b)$. Therefore, $y_1(s,b) < a_p(1)$ is given by the FOC:

$$\Phi(a,b)\,|_{a=y_1(s,b)} = p_s V_a'(a,s,b) + (1-p_s) V_a'(a,\varphi(s),b) = 0.$$

Because $V_{ab}'' > 0$ and $V_{aa}'' < 0$, the implicit function theorem leads to:

$$\frac{dy_1(s,b)}{db} = -\frac{\Phi_b'(a,b)}{\Phi_a'(a,b)}|_{a=y_1(s,b)} = -\frac{p_s V_{ab}''(a,s,b) + (1-p_s) V_{ab}''(a,\varphi(s),b)}{p_s V_{aa}''(a,s,b) + (1-p_s) V_{aa}''(a,\varphi(s),b)}|_{a=y_1(s,b)} > 0.$$

Suppose that $y_1(s,b_0(s)) = a_p(\varphi(s)) < a_p(1)$ for some $b_0(s)$. Because $y_1(s,b) < y_1(s,b_0(s)), b < b_0(s)$, then $b_0(s)$ is uniquely determined. Consider two cases: $y_1(s,b) \geq a_p(\varphi(s))$ or, equivalently, $b \geq b_0(s)$ and $y_1(s,b) < a_p(\varphi(s))$, or $b < b_0(s)$. If $y_1(s,b) \geq a_p(\varphi(s)) > a_p(s)$, then $V''_{ab} > 0$ means

$$V'_b\left(a_p\left(s\right), s, b\right) - V'_b\left(y_1\left(s, b\right), s, b\right) < 0, \text{ and}$$
$$V'_b\left(a_p\left(\varphi\left(s\right)\right), \varphi\left(s\right), b\right) - V'_b\left(y_1\left(s, b\right), \varphi\left(s\right), b\right) \leq 0.$$

This leads to:

$$\begin{aligned}\Delta V'_b\left(y_1\left(s, b\right), s, b\right) = & \; p_s\left(V'_b\left(a_p\left(s\right), s, b\right) - V'_b\left(y_1\left(s, b\right), s, b\right)\right) \\ & + \left(1 - p_s\right)\left(V'_b\left(a_p\left(\varphi\left(s\right)\right), \varphi\left(s\right), b\right) - V'_b\left(y_1\left(s, b\right), \varphi\left(s\right), b\right)\right) < 0, b \geq b_0\left(s\right).\end{aligned}$$

Because $\Delta V\left(y_1\left(s, b\right), s, b\right)$ is quasi-monotone in b on $\left[0, b_0\left(s\right)\right]$, it follows that $\Delta V\left(y_1\left(s, b\right), s, b\right) \geq 0$ is quasi-monotone for all $b > 0$. That is, if $\Delta V\left(y_1\left(s, \bar{b}\right), s, \bar{b}\right) \geq 0$, then $\Delta V\left(y_1\left(s, b\right), s, b\right) \geq 0, b < \bar{b}$. ■

Proof of Theorem 2. According to (A1), there exists $\delta\left(b\right) > 0$, such that $a_p\left(s\right) \geq y\left(s, b\right) - \delta\left(b\right), \forall s$. By (A2), there exists $\tilde{\delta}\left(b, \delta\left(b\right)\right) > 0$, such that $V\left(y\left(s, b\right) - \delta\left(b\right), s, b\right) > V\left(y\left(s, b\right) + \tilde{\delta}\left(b, \delta\left(b\right)\right), s, b\right), \forall s$. Take $\delta^*\left(b\right) = \max\left\{\delta\left(b\right), \tilde{\delta}\left(b, \delta\left(b\right)\right)\right\}$. By (A3), there exists $\frac{d\left(b, \delta^*\left(b\right)\right)}{2} > 0$, such that $y\left(s, b\right) + \delta^*\left(b\right) < y_1^*\left(s, s + \frac{d\left(b, \delta^*\left(b\right)\right)}{2}, b\right) < y\left(s + \frac{d\left(b, \delta^*\left(b\right)\right)}{2}, b\right) - \delta^*\left(b\right), \forall s$.

Take $d = d\left(b, \delta^*\left(b\right)\right)$ and construct the communication schedule as follows:

$$I_1 = \left\{\theta, \theta + \frac{d}{2}\right\}_{\theta \in \left[kd, kd + \frac{d}{2}\right), k = 0, 1, \ldots}, \text{ and}$$

$$I_2\left(I_1, m_1\right) = \left\{m_1, \varphi\left(m_1\right), I_1 \backslash \left\{m_1, \varphi\left(m_1\right)\right\}\right\},$$

where

$$\varphi\left(m_1\right) = m_1 + \frac{d}{2}.$$

Consider the expert's second-stage incentive compatibility constraints given truthful reporting of s in the first period and observing $\theta = s < \varphi\left(s\right)$ in the second period. The condition (A1) implies that $y\left(s, b\right) - \delta\left(b\right) \leq a_p\left(s\right) \leq y\left(s, b\right)$. Also, we have

$$\begin{aligned}a_p\left(s + \frac{d}{2}\right) \geq y\left(s + \frac{d}{2}, b\right) - \delta\left(b\right) \geq y\left(s + \frac{d}{2}, b\right) - \delta^*\left(b\right) \\ > y\left(s, b\right) + \delta^*\left(b\right) \geq y\left(s, b\right) + \tilde{\delta}\left(b, \delta\left(b\right)\right), \forall s,\end{aligned}$$

where the first inequality follows from (A1) and the third inequality follows from (A3). This gives

$$\begin{aligned}V\left(a_p\left(s\right), s, b\right) & \geq V\left(y\left(s, b\right) - \delta\left(b\right), s, b\right) \\ & > V\left(y\left(s, b\right) + \tilde{\delta}\left(b, \delta\left(b\right)\right), s, b\right) > V\left(a_p\left(s + \frac{d}{2}\right), s, b\right), \forall s,\end{aligned} \tag{17}$$

where the first and the last inequalities follow from $V'_a\left(a, s, b\right) \gtrless 0, a \lesseqgtr y\left(s, b\right)$, and the second inequality follows from (A2).

Consider now the expert's payoff in the first stage. Note that (A3) implies

$$y\left(b, s\right) + \tilde{\delta}\left(b, \delta\left(b\right)\right) \leq y\left(b, s\right) + \tilde{\delta}^*\left(b\right) < y_1^*\left(s, s + \frac{d}{2}, b\right), \forall s,$$

that results in

$$V\left(a_p\left(s\right),s,b\right) > V\left(y\left(s,b\right)+\tilde{\delta}\left(b,\delta\left(b\right)\right),s,b\right)$$
$$\geq V\left(y\left(s,b\right)+\delta^*\left(b\right),s,b\right) > V\left(y_1^*\left(s,s+\frac{d}{2},b\right),s,b\right), \forall s,$$

where the first inequality follows from (17), and the other inequalities follow from $V_a'\left(a,s,b\right) < 0, a > y\left(s,b\right)$. Also, (A1) and (A3) imply

$$y_1^*\left(s,s+\frac{d}{2},b\right) < y\left(s+\frac{d}{2},b\right)-\delta^*\left(b\right) \leq y\left(s+\frac{d}{2},b\right)-\delta\left(b\right) \leq a_p\left(s+\frac{d}{2}\right) \leq y\left(s,b\right), \forall s,$$

which gives

$$V\left(a_p\left(s+\frac{d}{2}\right),s,b\right) > V\left(y_1^*\left(s,s+\frac{d}{2},b\right),s,b\right), \forall s,$$

where the inequality holds because $V_a'\left(a,s,b\right) > 0, a < y\left(s,b\right)$. As a result,

$$V\left(a_p\left(s\right),s,b\right)+V\left(a_p\left(s+\frac{d}{2}\right),s+\frac{d}{2},b\right) \tag{18}$$
$$> V\left(y_1^*\left(s,s+\frac{d}{2},b\right),s,b\right)+V\left(y_1^*\left(s,s+\frac{d}{2},b\right),s+\frac{d}{2},b\right).$$

Note that (17) does not depend on p_s. Also, the optimal decision in the case of deviation from truthtelling in the first stage,

$$y_1\left(s,s+\frac{d}{2},b\right) = \arg\max_{a\in\mathcal{A}} p_s\left(V\left(y,s,b\right)+\frac{1-p_s}{p_s}V\left(y,s+\frac{d}{2},b\right)\right),$$

is continuous in $p_s > 0$ and $y_1\left(s,s+\frac{d}{2},b\right) = y_1^*\left(s,s+\frac{d}{2},b\right)$ if $p_s = \frac{1}{2}$. Because (18) holds strictly, there exists an ε–neighborhood of 1 for each (s,b), such that if it contains $\frac{1-p_s}{p_s} = \frac{f\left(s+\frac{d}{2}\right)}{f(s)}$, then

$$E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] = p_s\left(V\left(a_p\left(s\right),s,b\right)+\frac{f\left(s+\frac{d}{2}\right)}{f\left(s\right)}V\left(a_p\left(s+\frac{d}{2}\right),s+\frac{d}{2},b\right)\right)$$
$$> p_s\left(V\left(y_1\left(s,s+\frac{d}{2},b\right),s,b\right)+\frac{f\left(s+\frac{d}{2}\right)}{f\left(s\right)}V\left(y_1\left(s,s+\frac{d}{2},b\right),s+\frac{d}{2},b\right)\right)$$
$$= E\left[V\left(y_1\left(s,s+\frac{d}{2},b\right),\theta,b\right)|s\right].$$

■

Proof of Lemma 4. First, construct the communication schedule as that in the proof of Theorem 2 with the distance $\frac{d}{2}$ between the joint states, where the value of $d$is determined below. Given the signal $s \in [kd, kd+\frac{d}{2}), k \in \mathbb{N}_0$ in the first stage, truthtelling in both periods provides the payoff to the expert:

$$E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] = -p_s\beta^2\left(b,s\right)-p_s^c\beta^2\left(b,s+\frac{d}{2}\right) \geq -\delta^2.$$

Inducing the optimal decision

$$y_1\left(s, s+\frac{d}{2}, b\right) = p_s\left(s+\beta\left(b,s\right)\right) + p_s^c\left(s+\frac{d}{2}+\beta\left(b, s+\frac{d}{2}\right)\right),$$

provides the payoff:

$$E\left[V\left(y_1\left(s, s+\frac{d}{2}, b\right), \theta, b\right)|s\right] = -p_s p_s^c\left(\frac{d}{2}+\beta\left(b, s+\frac{d}{2}\right)-\beta\left(b,s\right)\right)^2.$$

Because $a_p\left(s+\frac{d}{2}\right) - a_p\left(s\right) = s+\frac{d}{2}-s = \frac{d}{2}$ and $V\left(a,s,b\right)$ is symmetric in a around $y\left(s,b\right) = s+\beta\left(s,b\right)$, the second-stage incentive compatibility constraints are given by:

$$a_p\left(s+\frac{d}{2}\right) - a_p\left(s\right) = \frac{d}{2} \geq 2\beta\left(b,s\right). \tag{19}$$

Since $\beta\left(b,s\right) \leq \delta$, (19) are satisfied if $d \geq 4\delta$. Also, (19) implies $\frac{d}{2} - \beta\left(b,s\right) \geq \beta\left(b,s\right) > 0$ and

$$E\left[V\left(y_1\left(s, s+\frac{d}{2}, b\right), \theta, b\right)|s\right] < -p_s p_s^c\left(\frac{d}{2}-\delta\right)^2 = -\frac{f\left(s\right) f\left(s+\frac{d}{2}\right)}{\left(f\left(s\right)+f\left(s+\frac{d}{2}\right)\right)^2}\left(\frac{d}{2}-\delta\right)^2.$$

Suppose $f\left(s\right) \geq f\left(s+\frac{d}{2}\right)$. (The case of $f\left(s\right) < f\left(s+\frac{d}{2}\right)$ is symmetric). If $f\left(\theta\right)$ satisfies condition (A4) for $\alpha > 0$ and $d > 0$, then

$$\frac{f\left(s\right) f\left(s+\frac{d}{2}\right)}{\left(f\left(s\right)+f\left(s+\frac{d}{2}\right)\right)^2} = \frac{f\left(s+\frac{d}{2}\right)}{f\left(s\right)} \frac{1}{\left(1+\frac{f\left(s+\frac{d}{2}\right)}{f(s)}\right)^2} \geq \frac{e^{-\alpha\frac{d}{2}}}{\left(1+\frac{f\left(s+\frac{d}{2}\right)}{f(s)}\right)^2} \geq \frac{e^{-\alpha\frac{d}{2}}}{4}.$$

This implies

$$E\left[V\left(y_1\left(s, s+\frac{d}{2}, b\right), \theta, b\right)|s\right] < -\frac{e^{-\alpha\frac{d}{2}}}{4}\left(\frac{d}{2}-\delta\right)^2 = \phi\left(d,\alpha\right),$$

where $\phi\left(d,\alpha\right)$ attains the minimum $\phi\left(d^*\left(\alpha\right),\alpha\right) = -\frac{e^{-2-\alpha\delta}}{\alpha^2}$ at $d^*\left(\alpha\right) = \frac{4}{\alpha}+2\delta$ for each $\alpha > 0$. Equivalently, each $d > 2\delta$ is the minimizer of $\phi\left(x,\alpha\right)$ over x for $\alpha^*\left(d\right) = \frac{4}{d-2\delta} > 0$.

Because (A4) holds for $\alpha = \frac{z^*}{\delta}$ and $d = d^*\left(\alpha\right) = 2\delta\left(\frac{2}{z^*}+1\right)$, where $z^* \simeq 0.314$ is a unique solution to the equation $z = e^{-1-\frac{z}{2}}$, we have

$$\alpha\delta = z^* = e^{-1-\frac{z^*}{2}} = e^{-1-\frac{\alpha\delta}{2}},$$

and $e^{-2-\alpha\delta} = \alpha^2\delta^2$. This results in

$$E\left[V\left(a_p\left(\theta\right),\theta,b\right)|s\right] \geq -\delta^2 = -\frac{e^{-2-\alpha\delta}}{\alpha^2} = \phi\left(d^*\left(\alpha\right),\alpha\right) = \phi\left(d,\alpha\right) > E\left[V\left(y_1\left(s, s+\frac{d}{2}, b\right), \theta, b\right)|s\right].$$

Finally, (19) holds, because $z^* < 1$ implies $d = 2\delta\left(\frac{2}{z^*}+1\right) > 4\delta$. ■

Proof of Theorem 3. Because $f\left(\theta\right) > 0, \theta \in \Theta$, the probabilities of the joint states conditional on observing $s \in Q_k$ for some $k = 1, ..., N$ in the first period are

$$p_s = \Pr[\theta = s|s] = \frac{f(s)}{f(s) + f(\varphi_k(s))} > 0, \text{ and } \Pr[\theta = \varphi_k(s)|s] = 1 - p_s > 0.$$

The differentiability of $V(a, \theta, b)$ and $y(\theta, b)$ and the continuity of $\varphi_k(\theta)$ in (a, θ, b) imply that $y_1(s, b), V(a_p(s), s, b) - V(a_p(\varphi_k(s)), s, b)$,$E[V(y_1(s, b), \theta, b)|s]$, and $E[V(a_P(\theta), \theta, b)|s]$ are continuous in (s, b).

First, consider the expert's second-stage incentive compatibility constraints (7) given truthtelling in $t = 1$. That is, upon reporting $s \in Q_k$ for some $k = 1, ..., N$, the expert can induce $a_p(s)$ or $a_p(\varphi_k(s))$ in the second stage. Suppose that the expert observes $\theta = s \in Q_k$ for some $k = 1, ..., N$ in the second period. If $b = 0$, then $a_p(s) = y(s, 0), s \in \Theta$. Because $a_p(\theta) \neq a_p(\theta'), \theta' \neq \theta$, then the strict concavity of $V(a, \theta, b)$ in a results in

$$V(a_p(\theta), \theta, 0) - V(a_p(\varphi_k(\theta)), \theta, 0) = V(y(\theta, 0), \theta, 0) - V(y(\varphi_k(\theta), 0), \theta, 0) > 0, \theta \in Q_k. \quad (20)$$

Similarly, if the expert observes $\theta = \varphi_k(s) \in \Theta_k \backslash Q_k$ for some $s \in Q_k$, then

$$V(a_p(\theta), \theta, 0) - V(a_p(s), \theta, 0) = V(y(\theta, 0), \theta, 0) - V(y(s, 0), \theta, 0) > 0, \theta \in \Theta_k \backslash Q_k. \quad (21)$$

Due to the continuity of $V(a_p(\varphi_k(s)), s, b) - V(a_p(s), s, b)$ in (s, b), it follows that there exists $b_k^c > 0$, such that (20) and (21) hold for all $b \leq b_k^c$ and $\theta \in \Theta_k$.

In the first period, the expert has two options upon observing $s \in Q_k$. First, she can induce the optimal decision $y_1(s, b) \in \mathcal{A}$ given her current information. Second, the expert can learn θ perfectly in the second stage by reporting s truthfully. In the latter case, the expert can induce only $a \in \{a_p(s), a_p(\varphi_k(s))\}$ in the second period.

In the case of inducing $y_1(s, b)$, suppose first that $b = 0$. Then, $a_p(\theta) \neq a_p(\theta'), \theta' \neq \theta$ implies $y(s, 0) = a_p(s) \neq a_p(\varphi_k(s)) = y(\varphi_k(s), 0)$. Thus, $y_1(s, 0) \neq y(s, 0)$ and/or $y_1(s, 0) \neq y(\varphi_k(s), 0)$, which gives

$$\begin{aligned}
E[V(y_1(s, 0), \theta, 0)|s] &= p_s V(y_1(s, 0), s, 0) + (1 - p_k) V(y_1(s, 0), s, 0) \\
&< p_s V(y(s, 0), s, 0) + (1 - p_k) V(y(\phi_k(s), 0), \phi_k(s), 0) \\
&= p_s V(a_p(s), s, 0) + (1 - p_k) V(a_p(\phi_k(s)), \phi_k(s), 0) \\
&= E[V(a_p(\theta), \theta, 0)|s], s \in Q_k.
\end{aligned}$$

where $E[V(a_p(\theta), \theta, b)|s]$ is the expert's payoff in the case of truthful reporting in both stages. Because $E[V(y_1(s, b), \theta, b)|s]$ and $E[V(a_p(\theta), \theta, b)|s]$ are continuous in (s, b), there exists $b_k^y > 0$, such that

$$E[V(a_p(\theta), \theta, b)|s] \geq E[V(y_1(s, b), \theta, b)|s], s \in Q_k, b \leq b_k^y.$$

Thus, taking $b \leq \bar{b} = \min_{k=1,...,N} \{b_k^c, b_k^y\}$ sustains truthtelling of the expert in both stages.

Finally, the expert has no incentives to deviate from reporting $\theta = \{1\}^N$ in both stages, since $V(a_p(\theta), \theta, b) \geq V(a, \theta, b), \theta = \{1\}^N, a \in \mathcal{A}$. ∎

Appendix B

This section highlights an example, which demonstrates how the principal can extract information from the expert with an arbitrary precision for a range of biases if the density function is zero for some states.

Suppose that $U(a,\theta) = -(a-\theta)^2, V(a,\theta,b) = -(a-\theta-b)^2$, and $f(\theta) = 2\theta, \theta \in [0,1]$. In this case, there is no communication schedule (2) and (3), which sustains the fully informative equilibrium for any $b > 0$. This is because the second-stage incentive-compatibility constraints require $\varphi(s) - s \geq 2b$, which implies $p_s = \frac{s}{s+\varphi(s)} \leq \frac{s}{2(s+b)}$. As $s \to 0$, it follows that $p_s \to 0$ and

$$E[V(y_1(s,b),\theta,b)|s] = -p_s(1-p_s)(\varphi(s)-s)^2 > -b^2 = E[V(a_p(\theta),\theta,b)|s], b > 0.$$

It is, however, possible to elicit information from the expert with an arbitrary precision. For that purpose, fix $\bar{d} \in (0,1/2)$ and split the interval $[0,\bar{d})$ into $N > 0$ equal subintervals of length $d = \bar{d}/N$. For each subinterval $\Theta_i = [(i-1)d, id), i = 1,...,N$, assign the complement subinterval $\mathcal{Z}_i = [1/2 + z_{i-1}, 1/2 + z_i)$, where $z_0 = 0$ and z_i is determined recursively:

$$z_i = \frac{\left(4z_{i-1}^2 + 4z_{i-1} + 8d^2 i - 4d^2 + 1\right)^{1/2} - 1}{2}.$$

The complement intervals are chosen to preserve the expert's uncertainty over states $\theta \in \Theta_i \cup \mathcal{Z}_i$:

$$\Pr[\theta \in \Theta_i] = \int\limits_{\Theta_i} dF(\theta) = \int\limits_{\mathcal{Z}_i} dF(\theta) = \Pr[\theta \in \mathcal{Z}_i], i = 1,...,N. \tag{22}$$

Define the expert's information structure as follows. Suppose first that $\theta \in \bar{\Theta} = [0,\bar{d}) \cup [1/2, 1/2 + z_N)$. Note that $[0,\bar{d}) \cap [1/2, 1/2 + z_N) = \varnothing$ and $\bar{\Theta} \subset [0,1]$, since $\bar{d} < 1/2$ and $z_N \to \bar{d}^2 < 1/2$ as $N \to \infty$. If $\theta \in \Theta_i \cup \mathcal{Z}_i$ for some $i = 1,...,N$, the expert observes the signal $s_i = \theta_i$ in the first stage. From (22), the posterior probabilities of the intervals Θ_j and $\mathcal{Z}_j$ are equal:

$$\Pr[\theta \in \Theta_i | s_i] = \Pr[\theta \in \mathcal{Z}_i | s_i] = 1/2, i = 1,...,N.$$

For each $\theta \in [\bar{d}, 1/2)$, assign the complement state $\varphi(\theta) \in [1/2 + z_N, 1)$, such that

$$\varphi(\theta) = \frac{1 - 2z_N}{1 - 2\bar{d}}\left(\theta - \bar{d}\right) + z_N + \frac{1}{2}.$$

In the second stage, the expert is able to distinguish between the intervals Θ_j and $\mathcal{Z}_j$ only upon sending the message $s_j \in \{\theta_i\}_{i=1}^N$ and between states s and $\varphi(s)$ only upon sending the message $s \in [\bar{d}, 1/2)$. By construction, $\bar{\Theta} \cup \{\theta, \varphi(\theta)\}_{\theta \in [\bar{d},1/2)} \cup \{1\} = [0,1]$. If $N \to \infty$, then $z_i \to \theta_i^2$, and the lengths of intervals Θ_i and $\mathcal{Z}_i$, $\theta_i - \theta_{i-1} = \bar{d}/N \to 0$ and $z_i - z_{i-1} \to \theta_i^2 - \theta_{i-1}^2 \to 0$, respectively. That is, the optimal decision given $s_i = \theta_i$:

$$y_1(s_i, b) = \frac{1}{2}E[\theta | \theta \in \Theta_i] + \frac{1}{2}E[\theta | \theta \in \mathcal{Z}_i] + b,$$

converges to $y_1^i = \frac{1}{2}\theta_i + \frac{1}{2}\left(\frac{1}{2} + \theta_i^2\right) + b$. Also, the principal's best responses upon learning Θ_i and $\mathcal{Z}_i$ are $a_p(\Theta_i) = E[\theta|\Theta_i] \to \theta_i$ and $a_p(\mathcal{Z}_i) = E[\theta|\mathcal{Z}_i] \to 1/2 + z_i \to 1/2 + \theta_i^2$, respectively. Hence, the expert's payoff in the case of inducing $y_1(s_i, b)$:

$$E[V(y_1(s_i,b),\theta,b)|s_i] = \frac{1}{2}E[V(y_1(s_i,b),\theta,b)|\Theta_i] + \frac{1}{2}E[V(y_1(s_i,b),\theta,b)|\mathcal{Z}_i]$$
$$\underset{N\to\infty}{\to} -\frac{1}{2}\left(y_1^i-\theta_i-b\right)^2 - \frac{1}{2}\left(y_1^i-\frac{1}{2}-\theta_i^2-b\right)^2 = -\frac{1}{16}\left(1-2\theta_i+2\theta_i^2\right)^2 < -\frac{1}{16}\left(1-2\bar{d}\right)^2.$$

Similarly, the expert's payoff in the case of truthful reporting in both stages,

$$E[V(a_p(\theta),\theta,b)|s_i] = \frac{1}{2}E[V(a_p(\Theta_i),\theta,b)|\Theta_i] + \frac{1}{2}E[V(a_p(\mathcal{Z}_i),\theta,b)|\mathcal{Z}_i]$$
$$\underset{N\to\infty}{\to} -\frac{1}{2}(\theta_i-\theta_i-b)^2 - \frac{1}{2}\left(\frac{1}{2}+\theta_i^2-\frac{1}{2}-\theta_i^2-b\right)^2 = -b^2.$$

The second-stage incentive-compatibility constraints hold if $a_p(\mathcal{Z}_i) - a_p(\Theta_i) \geq 2b$. Note that $a_p(\mathcal{Z}_i) - a_p(\Theta_i) \to 1/2 + \theta_i^2 - \theta_i > 1/2 - \bar{d}$ as $N \to \infty$. It follows then that $E[V(a_p(\theta),\theta,b)|s_i] > E\left[V\left(y_1^i,\theta,b\right)|s_i\right]$ and $a_p(\mathcal{Z}_i) - a_p(\Theta_i) \geq 2b$ if $b \leq \hat{b} = \frac{1}{4} - \frac{\bar{d}}{2}$ and N is sufficiently large.

Suppose now that $\theta \in [\bar{d}, 1/2) \cup [1/2 + z_N, 1)$. Because $f(\theta) \geq 2\bar{d} > 0$ if $\theta \geq \bar{d}$, and

$$\varphi(s) - s \geq \varphi\left(\bar{d}\right) - \bar{d} = \frac{1}{2} + z_N - \bar{d} > \frac{1}{2} - \bar{d} > 0, s \in [\bar{d}, 1/2),$$

by Theorem 1 there exists $\bar{b}$, such that

$$E[V(y_1(s,b),\theta,b)|s] \geq E[V(a_p(\theta),\theta,b)|s],$$

and $\varphi(s) - s \geq 2b$ for all $s \in [\bar{d}, 1/2)$ and $b \leq \bar{b}$. Therefore, if $b \leq \min\{\bar{b}, \hat{b}\}$, then there is an equilibrium, in which the expert fully reveals $\theta \in [\bar{d}, 1/2) \cup [1/2 + z_N, 1]$ and approximately fully reveals $\theta \in [0, \bar{d}) \cup [1/2, 1/2 + z_N)$ as $N \to \infty$.

References

Alonso, R. and N. Matouschek (2008), Optimal delegation, Review of Economic Studies, 259–293

Aumann, R. and S. Hart (2003), Long cheap talk, Econometrica 71, 1619–1660

Austen-Smith, D. (1994), Strategic transmission of costly information, Econometrica 62, 955–963

Battaglini, M. (2002), Multiple referrals and multidimensional cheap talk, Econometrica 70, 1379–1401

Blume, A., O. Board, and K. Kawamura (2007), Noisy talk, Theoretical Economics 2, 395–440

Crawford, V. and J. Sobel (1982), Strategic information transmission, Econometrica 50, 1431–1451

Dessein, W. (2002), Authority and communication in organizations, Review of Economic Studies 69, 811–838

Eso, P. and Y. Fong (2008), Wait and see, Working paper

Fischer, P. and P. Stocken (2001), Imperfect information and credible communication, Journal of Accounting Research 39, 119–134

General Accounting Office (1992), Electronic warfare: established criteria not met for airborne self-protection jammer production, U.S. Government Printing Office, Washington, D.C.

General Accounting Office (1988), Weapons testing: quality of DoD operational testing and reporting, U.S. Government Printing Office, Washington, D.C.

Gilligan, T. and K. Krehbiel (1989), Asymmetric information and legislative rules with a heterogeneous committee, American Journal of Political Science 33, 459–490

Gilligan, T. and K. Krehbiel (1987), Collective decision-making and standing committees: an informational rationale for restrictive amendment procedures, Journal of Law, Economics, and Organization 3, 287–335

Golosov, M., V. Skreta, A. Tsyvinski, and A. Wilson (2010), Dynamic strategic information transmission, Working paper

Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani (2009), Mediation, arbitration and negotiation, Journal of Economic Theory 144, 1397–1420

Green, J. and N. Stockey (2007), A two-person game of information transmission, Journal of Economic Theory 135, 90–104

Ivanov, M. (2010a), Informational control and organizational design, Journal of Economic Theory 145, 721-751

Ivanov, M. (2010b), Communication via a strategic mediator, Journal of Economic Theory 145, 869–884

Kartik, N., M. Ottaviani, and F. Squintani (2007), Credulity, lies, and costly talk, Journal of Economic Theory 134, 93–116

Kovác, E. and T. Mylovanov (2009), Stochastic mechanisms in settings without monetary transfers: the regular case, Journal of Economic Theory 144, 1373–1395

Krishna, V. and J. Morgan (2008), Contracting for information under imperfect commitment, RAND Journal of Economics 39, 905–925

Krishna, V. and J. Morgan (2004), The art of conversation: eliciting information from experts through multi-stage communication, Journal of Economic Theory 117, 147–179

Krishna, V. and J. Morgan (2001a), A model of expertise, Quarterly Journal of Economics 116, 747–775

Krishna, V. and J. Morgan (2001b), Asymmetric information and legislative rules: some amendments, American Political Science Review 95, 435-452

Mathis J. (2008), Full revelation of information in sender–receiver games of persuasion, Journal of Economic Theory 143, 571–584

Melumad, N. and T. Shibano (1991), Communication in settings with no transfers, RAND Journal of Economics 22, 173–198

Myerson, R. (1991), Game theory: analysis of conflict, Harvard University Press, Cambridge, MA

Ottaviani, M. and F. Squintani (2006), Naive audience and communication bias, International Journal of Game Theory 35, 129–150

Radner R. (1993), The organization of decentralized information processing, Econometrica 62, 1109–46

Rogerson, W. (1990), Quality vs. quantity in military procurement, American Economic Review 80, 83–92

Seidmann, D. and E. Winter (1997), Strategic information transmission with verifiable messages, Econometrica 65, 163–169

Shaked, M. and J. G. Shanthikumar (2007), Stochastic orders, Springer