

Contributions

Maxim Ivanov*

Dynamic Information Revelation in Cheap Talk

Abstract: This paper studies a multi-stage version of Crawford and Sobel's communication game. In every period the receiver determines a test about the unknown state whose result is privately observed by the sender. After the sender sends a costless message about an outcome of the test, the receiver selects a test in the next period. After a finite number of periods of interaction, the receiver makes a decision. The paper offers a sequence of tests that refine sender's information step-by-step and preserve truthtelling in every period. This sequence allows the receiver to learn the state in a subinterval of the state space with an arbitrary precision and has appealing theoretical properties. It consists of simple binary tests which reveal whether the state is above a certain cutoff, where the cutoffs are monotonic across periods and independent from results of the previous tests. Finally, we show that the relative payoff efficiency of multi-stage interaction compared to a single-stage game increases without a bound as the bias in preferences tends to zero.

Keywords: communication, information, cheap talk

JEL Classification: C72, D81, D82, D83

DOI 10.1515/bejte-2014-0017

1 Introduction

This paper focuses on the problem of strategic communication between a privately informed expert (the sender) and an uninformed decision maker (the receiver) à la Crawford and Sobel (1982), hereafter CS. In general, communication between the involved parties is characterized by two features. The first one is a conflict of interest. Different objectives create an incentive for the sender to misrepresent information in her favor, which results in the lower quality of disclosed information. The second feature is the imperfect primary information of the sender, since

*Corresponding author: Maxim Ivanov, Department of Economics, McMaster University, Kenneth Taylor Hall, 1280 Main Street West, Hamilton, ON, Canada L8S 4M4, E-mail: mivanov@mcmaster.ca

even the most knowledgeable experts may have imperfect information. However, while the difference in players' preferences is usually exogenous, the quality of the sender's information can sometimes be endogenized by the receiver. That is, though the sender's information is still private, the precision of this information is determined by the receiver.¹ Moreover, if players interact in multiple periods, the receiver may affect the precision of the sender's information and request a report from her in each period. In other words, the receiver selects a dynamic informational control policy – a sequence of tests – and makes a final decision, whereas the sender conducts each test and reports its result to the receiver.

Our major contribution is that we show how the receiver can use these instruments – controlling the quality of the sender's private information (hereafter, informational control) and dynamic interaction – to extract the sender's private information about an unknown state of nature. In particular, the paper offers a sequence of tests that allows the receiver to learn information with any precision over an interval of states which converges to the state space as the preference bias tends to zero. An implication of this finding is the Pareto improvement over single-stage informational control and, hence, communication with the perfectly informed sender. This result becomes even stronger given that single-stage informational control provides a higher ex-ante payoff to the receiver than other incentive tools such as delegation or mediation with the perfectly informed sender (Ivanov 2010).²

The sender's learning process and players' interaction are organized as follows. The sequence of tests consists of binary tests such that each test progressively reveals whether the state is above a certain cutoff. The cutoffs are monotonic across periods and independent of the results of previous tests. The sender privately observes the result of each test and sends a report to the receiver. Testing and communication continue until either the sender reports that the state is above the current cutoff or the limit on the number of tests is achieved. Then the receiver makes a decision on the basis of reported information.

There is the similarity between the receiver's process of learning the sender's information in our setup and a sequential sampling plan performed by a single decision maker (see Feldman and Fox 1991). This plan tests a hypothesis on the basis of a sample of a variable size. In particular, if the result of the test based

1 An answer to a natural question about the possibility of acquiring information directly by the receiver is often determined by different information acquisition, processing, and opportunity costs of the players. The difference in costs forces the receiver to delegate the information acquisition problem to the sender.

2 In delegation, the receiver delegates decision making to the sender (see, e.g. Dessein 2002). In mediation, players communicate through a disinterested mediator (see Goltsman et al. 2009).

on the current sample is negative, then the hypothesis is rejected and the testing procedure stops. Otherwise, if the test result is positive, then the sample size increases by one and an additional test is conducted (until the limit on the sample size is achieved). The key difference between the receiver's learning in our setup and the sequential sampling plan is the conflicting preferences of players in our setting. Consider, for example, a positively biased sender whose optimal decision exceeds the optimal decision of the receiver given any information about the state. Because the sender tends to exaggerate her information in an attempt to induce higher decisions, her strategic motives can potentially harm communication and, hence, the efficiency of decision making.

We show that the constructed sequence of tests preserves truthtelling, which stems from the sender's trade-off between the set of feasible actions in each period and future informational benefits. In particular, as the sender accumulates more precise information over time, the set of feasible decisions shrinks. In order to illustrate this trade-off in the case of a positively biased sender, consider a sequence of tests with decreasing cutoffs. In this case, the sender observes either a very precise signal about the state (when it is above the cutoff) which will not be updated in future periods or a very imprecise signal (when it is below the cutoff) which will be updated in future periods. Second, the receiver's best response to the sender's truthtelling has the "trigger" character. In particular, because the sender's signal "the state is above the current cutoff" will not be updated in future periods, her report about such a signal induces the receiver's decision (the current trigger decision), unconditionally on future reports. In contrast, reporting that the state is below the cutoff allows the sender to induce trigger decisions, which are strictly below the current trigger decision, in future periods. If the positively biased sender learns that the state is above the cutoff, her best feasible decision is the highest one, i.e. the current trigger decision. The sender induces it by reporting her information truthfully. Otherwise, if the sender learns that the state is below the cutoff, her current information is sufficiently imprecise. In this case, she prefers to learn more information and induce one of the feasible decisions in future periods by reporting her information truthfully.³

³ In short, sender's learning of information is counter to her bias. That is, the incentive of the risk-averse sender to obtain a more precise signal about the state implies that the receiver will rationally choose lower decisions. Because the sender is positively biased, she cannot gain by delaying the information that the state is high, while learning that the state is low and reporting this information would result in more precise information and the possibility of inducing lower actions in future periods. The key condition for sustaining the sender's truthtelling upon learning that the state is below the cutoff is that the precision of the sender's information is sufficiently low *and* the highest feasible action is sufficiently large. This implies that there is a cutoff such that precise extraction is possible if the state is above this cutoff (given a sufficiently large number of periods).

In comparing the equilibrium in our setup to the CS equilibria for the leading uniform-quadratic specification (with the uniform distribution of states and quadratic payoff functions of players), it is worth noting that the sender reveals less information about high states in the CS equilibria. This is reflected in the fact that the lengths of subintervals of states which induce the same decision in any CS equilibrium are increasing in the state. Intuitively, because the sender has both the incentive and the possibility of exaggerating her information, the receiver's credibility to the sender's messages about high states is low. In contrast, the sender reveals more information about high states in our setup. As described above, the sender eventually learns precise information about the state only if it is high enough. At that moment, however, the sender lacks the opportunity to exaggerate her information since her best feasible decision is the receiver's best response to the sender's current information.

As an application of our results, consider the communication problem in defense procurement. While the military is an expert in evaluating characteristics of weapons, the budget for them is determined by the Congress (the receiver). Moreover, the parties' interests are not aligned – it is argued that the military tends to be biased toward weapons with excessive costs (Rogerson 1990). Also, the Department of Defense has received many accusations of manipulating test results to yield the most favorable interpretation (U.S. General Accounting Office 1988, 1992). In this light, suppose that the issue of communication is the efficiency of a weapon, which monotonically depends on the unknown maximum operating temperature (the state), and the decision of the Congress is the budget spent on the weapon production. Assume that the optimal budget of the Congress monotonically depends on the state and is below the optimal budget desired by the military. Thus, the military is interested in exaggerating the actual state. In order to mitigate the resulting communication problem, our paper suggests that the receiver must determine a sequence of field tests of the weapon such that each test is performed by the military at a specific weather temperature. This specific temperature plays the role of the cutoff, since testing the weapon at this temperature allows the military to see whether the maximum operating temperature exceeds it or not. Also, the cutoffs must be decreasing over time, and the military must report about the result of each test.

Our work is related to the literature that investigates the role of the sender's imperfect information and multi-stage strategic communication on the efficiency of decision making. First, Green and Stokey (2007) show that the receiver's ex-ante payoff is not always monotone in the quality of the sender's information. Austen-Smith (1994) and Fischer and Stocken (2001) demonstrate this result in the CS setup. Ivanov (2010) extends this result by showing that communication with an imperfectly informed sender can be more efficient to the receiver than

delegating authority to the perfectly informed sender. This paper extends the works by Fischer and Stocken (2001) and Ivanov (2010) by allowing the sender to learn information over time. The work closest to this paper is Ivanov (2013), which shows that by determining the precision of the sender's information and communicating in multiple rounds, the receiver can elicit perfect information from the sender. However, there are two crucial differences between the two papers. First, the procedure of acquiring information in Ivanov (2013) is predicated upon stronger assumptions. In particular, the precision of the sender's future information is highly sensitive to her previous report(s). In other words, the type of the test which can provide additional information to the sender depends on her previous reports. In contrast, the sequence of tests in this paper is fixed from the beginning. Thus, though the receiver can stop the sequence at any moment or, equivalently, ignore the sender's future messages, he cannot modify the initial sequence of tests. Because of this, the information acquisition procedure in our paper is simpler and more applicable. Second, our procedure is more robust to the sender's prior information. In particular, if the sender is privately informed about whether the state is above or below some level, this information does not affect her incentives to communicate in our model. In contrast, this information can collapse informative communication in the setup of Ivanov (2013).

Our work also complements the literature on communication through multiple periods with a perfectly informed sender(s). Aumann and Hart (2003) consider two-person games with discrete types and two-sided asymmetric information. In their setup, one side is better informed than the other, and the players can communicate without time constraints. They demonstrate that the set of equilibrium outcomes in the dynamic setup can be significantly expanded compared to the static one. In the uniform-quadratic CS setup, Krishna and Morgan (2004) investigate multi-stage communication such that the sender and the receiver communicate simultaneously in every period. Golosov et al. (2013) consider a dynamic game consisting of repeating the CS model multiple times (i.e. in each period the sender sends a message, and the receiver takes an action). These papers demonstrate that two factors – multiple rounds of cheap talk conversation and active participation of the receiver in the communication process – can improve information transmission.⁴ However, these papers consider the case of the perfectly informed sender and thus utilize multi-stage

⁴ Also, Krishna and Morgan (2004) note that simply extending the CS model to the multi-stage case does not improve communication. Because the sender knows the state before communication starts, she sends the sequence of messages that induces the most preferable action. As a result, the receiver infers the same information about the state as in the single-stage game.

interaction for communication only. Such an interaction, however, is not effective enough for achieving the first-best outcome of the receiver. In contrast, this paper shows that using the dynamic setup for both information updating and communication allows the receiver to achieve the first-best outcome with any precision in a subinterval of states.⁵

The paper proceeds as follows. Section 2 presents the model. The analysis and motivating examples are provided in Section 3. Section 4 compares the receiver's ex-ante payoffs in static and dynamic models of informational control. Section 5 concludes the paper.

2 The model

We consider a model of multi-stage communication in which two players, the sender and the receiver, interact during $T + 1 < \infty$ periods. Players communicate about *state* θ which is constant through the game and distributed on $\Theta = [0, 1]$ according to a distribution function $F(\theta)$ with a positive and bounded density $f(\theta)$. The sender has access to imperfect information about θ , and the receiver makes a *decision* (or an *action*) $a \in \mathcal{R}$ that affects the payoffs of both players. The receiver's and sender's payoff functions are

$$U(a, \theta) = -(a - \theta)^2, \text{ and } V(a, \theta) = -(a - \theta - b)^2,$$

respectively, where the bias parameter $b > 0$ reflects the divergence in the players' interests.

2.1 Actions

At the beginning of each period $t = 1, \dots, T$, the receiver determines the sender's *information structure*, or *test* $I_t = \{F_t(s_t|\theta), \theta \in \Theta\}$, which consists of a family of conditional distributions of signals.⁶ Then, the sender privately observes a *signal* s_t from the signal space $\mathcal{S} \supset \Theta$, drawn from an associated distribution $F_t(s_t|\theta)$. At the end of period t , the sender sends a (possibly random) *message* m_t from the message space $\mathcal{M} \supset \mathcal{S}$ to the receiver. Finally, in the last period $T + 1$, the

⁵ Battaglini (2004) considers a model with multi-dimensional signals and multiple imperfectly informed senders. He shows that for a large number of senders with distinct preferences and limited commitment of the receiver there is an equilibrium in which the receiver learns information with an arbitrarily precision.

⁶ That is, given the state θ , signals across periods are independent.

receiver takes an action a . We assume that information structures are independent across periods conditional on θ , that is, I_t does not depend on signal realizations $\{s_\tau\}_{\tau=1}^{t-1}$, messages $\{m_\tau\}_{\tau=1}^{t-1}$, or information structures $\{I_\tau\}_{\tau=1}^{t-1}$ in previous periods. This restriction significantly decreases the receiver's possibilities of extracting information from the sender.⁷

2.2 Strategies

Let $I^{t-1} = \{I_1, \dots, I_{t-1}\}$ and $m^{t-1} = \{m_1, \dots, m_{t-1}\}$ be the history of the receiver's and the sender's actions, respectively, at the beginning of period t . Thus, $h_t^0 = \{m^{t-1}, I^{t-1}\}$ determines the public history (i.e. the history available to the receiver) before selecting I_t . Also, let $s^t = \{s_1, \dots, s_t\}$ be the sender's private information before sending a message m_t . Denote by $h_t = \{m^{t-1}, I^t, s^t\} = \{h_t^0, I_t, s^t\}$ the sender's (private) history. A behavioral strategy of the receiver in communication periods (*communication schedule*) $I_t : \mathcal{H}_t^0 \rightarrow \mathcal{I}, t = 1, \dots, T$, is a sequence of functions which maps the set of public histories $\mathcal{H}_t^0$ into the space of information structures $\mathcal{I}$. In the last period, the receiver's strategy (hereafter, *decision rule*) $a : \mathcal{H}_{T+1}^0 \rightarrow \mathcal{R}$ maps $\mathcal{H}_{T+1}^0$ into the decision space $\mathcal{R}$.⁸ The behavioral strategy of the sender $\sigma_t^E : \mathcal{H}_t \rightarrow \Delta\mathcal{M}, t = 1, \dots, T$ is the sequence of functions, which maps the set of sender's histories into the space of probability distributions on the message set $\mathcal{M}$.

Thus, for any θ , a strategy profile $\sigma = \left(\{\sigma_t(h_t)\}_{t=1}^T, \{I_t(h_t^0)\}_{t=1}^T, a(h_{T+1}^0) \right)$ determines a stochastic process $Q_\theta(\sigma)$ for a triple $q_t = (I_t, s_t, m_t)$, which takes values in $\mathcal{I} \times \mathcal{S} \times \mathcal{M}$, and a random variable $a_\theta \in \mathcal{R}$, which represents the decision. This completely describes what happens if the players follow σ . The continuation of the process $Q_\theta(\sigma)$ following a history h_t , denoted $Q_\theta(\sigma|h_t)$, is also a stochastic process which describes the future behavior of the triple following h_t . Similarly, a random variable $a_\theta(\sigma|h_t)$ describes the behavior of the decision following h_t .

2.3 Beliefs

The receiver's belief system $\mu_t : \mathcal{H}_t^0 \rightarrow \Delta\Theta, t = 1, \dots, T+1$, is the sequence of mappings which determines the probability distribution over θ for each public history h_t^0 . We call the belief system $\mu = \{\mu_t(h_t^0)\}_{t=1}^{T+1}$ *consistent* if it is derived from the player's strategies on the basis of Bayes' rule where applicable.

⁷ Ivanov (2013) shows that if I_t can depend on past messages $\{m_\tau\}_{\tau=1}^{t-1}$, then the receiver can perfectly elicit θ in two-period communication only.

⁸ Because $U(a, \theta)$ is strictly concave in a , the receiver never mixes over actions.

2.4 Equilibrium

Formally, a communication schedule $\{I_t(h_t^0)\}_{t=1}^T$ is a part of the receiver's strategy and hence is determined endogenously. However, we investigate the benefits of a particular class of $\{I_t(h_t^0)\}_{t=1}^T$, such that each $I_t(h_t^0)$ is a partition of Θ into a finite number of intervals. Because the receiver in our model does not optimize over all communication schedules, we provide a definition of an equilibrium for an arbitrary $\{I_t(h_t^0)\}_{t=1}^T$. In particular, given a communication schedule $\{I_t(h_t^0)\}_{t=1}^T$, a perfect Bayesian equilibrium (hereafter, an equilibrium) is a pair of a strategy profile σ^* and the belief system μ^* such that: (1) μ^* is consistent with σ^* ; (2) given $\{I_t(h_t^0)\}_{t=1}^T$, $a(h_{T+1}^0)$ and h_t , $\sigma_t^*(h_t)$ maximizes the sender's interim payoff in period $t = 1, \dots, T$:

$$V(\sigma|h_t) = E_{\theta}[E_{Q_{\theta}(\sigma|h_t)}V(a_{\theta}(\sigma|h_t), \theta)],$$

and (3) given $\mu_{T+1}^*(h_{T+1}^0)$, $\{\sigma_t^*(h_t)\}_{t=1}^T$, and any history h_t^0 , $a^*(h_{T+1}^0)$ maximizes the receiver's payoff $E_{\theta}[U(a, \theta)|\mu_{T+1}^*(h_{T+1}^0)]$ in period $T + 1$.

3 Dynamic informational control

Fischer and Stocken (2001) show that the receiver can increase his ex-ante payoff by restricting the quality of the sender's information in the CS setup, in particular, by partitioning Θ into intervals $\{\Theta_k\}_{k=1}^N = \{[\theta_k, \theta_{k+1}]\}_{k=1}^N$, such that the sender observes $\Theta_k \ni \theta$. Intuitively, the preferences of the less informed sender are more closely aligned with those of the receiver.⁹ As a result, the less informed sender communicates truthfully, which increases the receiver's ex-ante payoff. In this section, we investigate the benefits of multi-stage updating of the sender's information. Before focusing on the main setup, we consider a modified model in which the receiver commits to a communication schedule.

⁹ Any equilibrium CS partition $\{\Theta_k^{CS}\}_{k=1}^{N^{CS}} = \{[\theta_{k-1}^{CS}, \theta_k^{CS}]\}_{k=1}^{N^{CS}}$ is determined by cutoff types θ_k^{CS} , $k = 1, \dots, N - 1$ who are indifferent between adjacent decisions a_k^{CS} and a_{k+1}^{CS} . However, if the sender observes only $\Theta_k^{CS} \ni \theta$, she strictly prefer the associated decision a_k^{CS} , since all $\theta' \in \Theta_k^{CS}$ strictly prefer a_k^{CS} . That is, the partitional information structure replaces θ_k^{CS} by mean types $\omega_k = E[\theta'|\theta' \in \Theta_k^{CS}]$. Because of that, the receiver can specify a finer partition than the equilibrium CS partition without violating the sender's incentives to communicate truthfully.

3.1 Commitment to the communication schedule

Any equilibrium in the original game can be replicated in the modified game in which the receiver commits to a communication schedule $\{I_t(h_t^0)\}_{t=1}^T$ from the beginning but does not commit to an action. The following result shows that the converse is also true. All proofs are collected in the Appendix.

Lemma 1 *Any equilibrium in the game with commitment to a communication schedule $\{I_t(h_t^0)\}_{t=1}^T$ is outcome equivalent to some equilibrium in the game without commitment.*

The intuition for this result is most easily gained from the fact that in any cheap-talk game (i.e. unconditionally on T and $\{I_t(h_t^0)\}_{t=1}^T$) there is the *babbling* equilibrium in which the sender does not reveal any useful information and the receiver ignores the sender's messages. In the multi-period game, the sender can punish the receiver with a babbling equilibrium in the continuation subgame. If the receiver follows the equilibrium $\{I_t^*(h_t^0)\}_{t=1}^T$, the sender continues to update the receiver's beliefs about the state. Otherwise, no additional information will be provided. Because information cannot hurt the receiver, his equilibrium behavior is optimal.¹⁰

Because the lack of commitment to the communication schedule does not affect the outcome of communication, without loss of generality, we can restrict attention to the setup with the receiver's commitment to the information structures. This preserves all of the results about the efficiency of communication but simplifies the analysis, since we do not need to be concerned about the sender's reaction to out-of-equilibrium information structures.

3.2 Truthtelling equilibria

Our main focus is *truthtelling* equilibria in which the sender truthfully reports her information in each period, i.e. $m_t(h^t) = s_t, \forall h^t, t = 1, \dots, T$. A communication schedule $\{I_t(h_t^0)\}_{t=1}^T$ is *incentive-compatible* if there is a truthtelling equilibrium with such a schedule.

10 The sender's behavior is reminiscent of the grim-punishment strategy in the classical prisoners' dilemma game. That is, the sender "cooperates" by sending informative messages as long as the receiver updates her information according to some scenario. In contrast, any deviation by the receiver results in permanent punishment in terms of the quality of received information.

First, consider a game with single-stage communication, that is $T = 1$. Define a function

$$\Delta_w(z, x) = \omega_z^x - \omega_0^z, 0 \leq z \leq x \leq 1,$$

where $\omega_{\underline{\theta}}^{\bar{\theta}} = E[\theta | \theta \in [\underline{\theta}, \bar{\theta}]]$ is a posterior mean of θ conditional on $\theta \in [\underline{\theta}, \bar{\theta}]$. (For $\underline{\theta} = \bar{\theta}$, we put $\omega_{\underline{\theta}}^{\bar{\theta}} = \underline{\theta}$.) Since $f(\theta)$ is bounded, $\Delta_w(z, x)$ is continuous in (z, x) . Note that ω_q^s is the receiver's decision if the sender truthfully reports that $\theta \in [q, s]$. Thus, $\Delta_w(z, x)$ determines the distance between posterior means obtained by splitting an interval $[0, x]$ into two subintervals, $[0, z]$ and $[z, x]$. Equivalently, it is the distance between the receiver's decisions in the truth-telling equilibrium of the single-stage communication game with the prior distribution $F|\theta| \theta \in [0, x]$ and the sender's information structure given by the partition $\{[0, z], [z, x]\}$.

Suppose that b satisfies

$$\Delta_w(1, 1) = 1 - E[\theta] > 2b. \quad [1]$$

Intuitively, condition (1) guarantees that there is informative communication in the single-stage game. Consider a (degenerate) information structure that discloses only whether θ is 1 or not, and suppose that the receiver treats the sender's messages as truthful. In this case, the sender has a choice between decisions $\omega_1^1 = 1$ and $\omega_0^1 = E[\theta]$. Condition (1) states that the sender prefers to induce ω_0^1 by truthful reporting upon learning that $\theta \in [0, 1)$, since this decision is closer to the sender's optimal interim decision $\omega_0^1 + b$. Also, the sender prefers to induce ω_1^1 after learning $\theta = 1$. By the continuity of $\Delta(z, 1)$ in z , there is an informative equilibrium with the partitional information structure $\{[0, z], [z, 1]\}$, where z is close to 1.

Consider $\theta_c \in (0, 1)$ defined as

$$\theta_c = \max\{x | \Delta_w(x, x) = 2b\} = \max\{x | x - E[\theta | \theta \leq x] = 2b\}, \quad [2]$$

which exists because function $\Delta_w(x, x) = x - E[\theta | \theta \leq x]$ is continuous in x and $\Delta_w(0, 0) = 0 < 2b < 1 - E[\theta] = \Delta_w(1, 1)$.¹¹ The interpretation of θ_c is as follows. (For simplicity, let θ_c be a unique solution to $\Delta_w(x, x) = 2b$.) Suppose that both players believe that $\theta \in [0, x]$, where $0 < x < 1$. Then, θ_c is the smallest interval $[0, x]$ such that there exists informative single-stage communication with the partitional information structure $\{[0, x], \{x\}\}$. In other words, θ_c is a measure of

¹¹ There may exist multiple solutions to equation $\Delta_w(x, x) = 2b$. Consider $b = 1/50$ and $f(\theta) = 1/100, \theta \in [0, 1/10] \cup [1/5, 1]$, and $f(\theta) = 10, \theta \in (1/10, 1/5)$. Then, the solutions to $\Delta_w(x, x) = 2b$ are $x_1 = 0.08, x_2 = 0.1$, and $x_3 = 0.18$.

the lowest sender's uncertainty about θ relative to her bias, such that the sender is willing to disclose whether $\theta = x$ or $\theta < x$.

3.2.1 Approximately full information revelation in $[\theta_c, 1]$

We restrict the analysis to a particular class of communication schedules such that the sender's information structure in each period $t = 1, \dots, T$ is the two-interval partition $\{\Theta_0^t, \Theta_1^t\} = \{[0, \theta_t], [\theta_t, 1]\}$ for any history h_t^0 . That is, the sender's signal $s_t \in \{\Theta_0^t, \Theta_1^t\}$ is such that $\theta \in s_t$, $t = 1, \dots, T$. Equivalently, a communication schedule in this class can be determined by the sequence of cutoffs $\{\theta_t\}_{t=0}^{T+1}$, where we put $\theta_0 = 1$ and $\theta_{T+1} = 0$.

Consider a *decreasing* communication schedule, i.e. a schedule with a strictly decreasing $\{\theta_t\}_{t=0}^{T+1}$. Such a communication schedule is characterized by two important properties which can be clearly seen if T is large and cutoffs θ_t are distributed uniformly over Θ , i.e. the distance between any adjacent cutoffs $\theta_{t-1} - \theta_t > 0$ is small. The first property reflects the quality of the sender's current information – in each period her information is either very precise or very imprecise. As long as the sender observes the lower interval $[0, \theta_t]$, the quality of her information is low because shrinking the set of possible states from $\Theta_0^{t-1} = [0, \theta_{t-1}]$ to $\Theta_0^{t-1} \cap \Theta_0^t = [0, \theta_t]$ updates the sender's information insignificantly. However, if the sender observes the higher interval $[\theta_t, 1]$, she infers that θ is in the set $\Theta_0^{t-1} \cap \Theta_1^t = [\theta_t, \theta_{t-1}]$. This substantially updates her previous information. The second property determines the sender's future informational benefits. If the sender's current information is precise, it will not be updated in the future. In contrast, if her information is still vague, then it will surely be improved in the future periods. An example below illustrates how a combination of these factors allows the receiver to extract all available information from the risk-averse sender.

Example 1. Suppose that θ is distributed uniformly on $[0, 1]$ and let $b = 3/14$. Consider the two-period communication schedule $\{\Theta_0^1, \Theta_1^1\} = \{[0, 6/7], [6/7, 1]\}$, $\{\Theta_0^2, \Theta_1^2\} = \{[0, 3/7], [3/7, 1]\}$ and the decision rule $a(m_0, m_0) = a_{00} = 3/14$, $a(m_0, m_1) = a_{01} = 9/14$, $a(m_1, \cdot) = a_1 = 13/14$ depicted in Figure 1.

Suppose that the sender observes the higher subinterval Θ_1^1 in the first period. Hence, her information is not updated in the second period. Then, the truthful message m_1 induces the decision a_1 unconditionally on the sender's message in the second round. In contrast, reporting m_0 allows the sender to induce two decisions, a_{00} and a_{01} . Because these decisions are strictly inferior to a_1




Figure 1: Two-stage communication

given Θ_1^1 , the sender reveals her information truthfully.¹² If the sender observes Θ_0^1 , her information will be updated in the second period. Then, distorting information by sending message m_1 induces the decision a_1 unconditionally on the future messages. This deprives the sender of all future informational benefits because new information does not affect the receiver's decision. As a result, the sender's interim payoff $V(a_1|\Theta_0^1) = -1/7$ is lower than that of $-3/49$ in the case of truthful reporting in both periods.¹³

Now, suppose that the sender observes the subinterval Θ_1^2 in the second period after observing Θ_0^1 in the first period. She infers that $\theta \in \Theta_0^1 \cap \Theta_1^2 = [3/7, 6/7]$. In this period, however, only decisions a_{01} and a_{00} are feasible, where a_{01} is strictly preferable to a_{00} .¹⁴ That is, the sender induces a_{01} by truthful communication. If the sender observes Θ_0^2 in the second round, she deduces that $\theta \in [0, 3/7]$. Given this information, a_{00} and a_{01} results in the same interim payoff to the sender, so that the sender (weakly) prefers to communicate truthfully. Also, the induced decisions are the receiver's best response to the sender's truthtelling strategy. Finally, note that the receiver's ex-ante payoff is approximately $-1/75$, which exceeds that of $-1/48$ in the most informative equilibrium in the single-stage informational control.¹⁵

The following theorem generalizes the above example and is the main result of the paper.

Theorem 1 *For a decreasing communication schedule $\{\theta_t\}_{t=0}^{T+1}$, such that $\theta_T \geq \theta_c$, there is a truthtelling equilibrium in the game with $\{\theta_t\}_{t=0}^{T+1}$.*

¹² This is because $V(a|\Theta_1^1)$ is increasing in a for $a < a^*(\Theta_1^1)$, and $\max\{a_{00}, a_{01}\} < a_1 < a^*(\Theta_1^1)$, where $a^*(\Theta_1^1) = \arg\max_a V(a|\Theta_1^1) = 8/7$.

¹³ $E[V(a(m_0, \cdot), \theta)|\Theta_0^1] = \frac{1}{\theta_1} \int_0^{\theta_2} V(a_{00}, \theta) d\theta + \frac{1}{\theta_1} \int_{\theta_2}^{\theta_1} V(a_{01}, \theta) d\theta = -\frac{3}{49} > -\frac{1}{7}$

¹⁴ Since $a_{00} < a_{01} < a''$, where $a'' = 6/7$ is the sender's optimal decision given updated information.

¹⁵ Note that the receiver cannot replicate the equilibrium outcome in the single-stage game via specifying the partition with cutoffs $\{3/7, 6/7\}$ and reacting to sender's messages as if they were truthful. The reason is that if $\theta \in [3/7, 6/7]$, the sender can profitably deviate from inducing a_1 .

An implication of this result is that if there are no exogenous restrictions on the communication horizon T , the receiver can learn $\theta \in [\theta_c, 1]$ with an arbitrary precision.

Corollary 1 *For any $\varepsilon > 0$, there is $T < \infty$ and a decreasing $\{\theta_t\}_{t=0}^{T+1}$ such that: $\theta_T = \theta_c$, $\max_{t=1, \dots, T} (\theta_{t-1} - \theta_t) < \varepsilon$, and there is a truthtelling equilibrium in the game with $\{\theta_t\}_{t=0}^{T+1}$.*

It is worthy to note that dynamic informational control demonstrates some similarity to the model of communication investigated by Krishna and Morgan (2004). In both cases, the main incentives for the risk-averse sender to provide more information stem from her uncertainty about the future receiver's actions or the informational benefits, which are affected by the sender's current message. However, the mechanisms of information extraction in the two models are different. In Krishna and Morgan's model, the uncertainty stems from the random outcome of simultaneous communication in the first round which is generated by the jointly-controlled lottery. In particular, if the outcome of communication is "success", the sender can update her information in the next period. Thus, even though the sender might not be allowed to update a report in the second round, the uncertainty about future interaction affects the sender's incentives in the first period. In addition, in the case of "success", the sender reveals more information in the second round for sufficiently high states. Together, these factors provide an overall improvement over CS communication. In the informational control case, uncertainty arises directly from updating the sender's information over periods. The important feature of informational control is a possibility of generating a smaller number of sender's types than that in the models with the perfectly informed sender, say, two versus a continuum. This substantially simplifies the problem of aligning the sender's incentives with the receiver's ones. As a result, although the number of sender's types is smaller, each type fully reveals herself in each round.

3.2.2 Partial information extraction in $[0, \theta_c]$

Condition (2) states that if the sender learns that $\theta < \theta_c$ at $t = T$, so she can induce two actions, $\omega_{T+1} = E[\theta | \theta \leq \theta_c]$ and $\omega_T = E[\theta | \theta \in [\theta_c, \theta_{T-1}]]$, then she prefers to induce ω_{T+1} by truthfully reporting the lower interval. However, if θ_{T-1} is close to θ_c and so is ω_T , splitting the interval $[0, \theta_c]$ into two subintervals, $[0, z]$ and $[z, \theta_c]$, by specifying $\theta_{T-1} = \theta_c$ and $\theta_T = z$ can violate the sender's

incentives to report truthfully about the lower interval. This is because the action $\omega_c = E[\theta|\theta \in [z, \theta_c]]$ is feasible in the last communication period and is below θ_c . Hence, condition (2) seems to be too restrictive for informative communication about $\theta \in [0, \theta_c]$. Nevertheless, this logic is imperfect, and the receiver can extract partial information about θ even if it is below θ_c . In fact, if $\theta < z$, then partitioning the interval $[0, \theta_c]$ updates the sender's information in the last communication period, which decreases her posterior valuation of θ from ω_{T+1} to $\omega_z = E[\theta|\theta \leq z] < \omega_{T+1}$. Hence, the sender's optimal action decreases from $\omega_{T+1} + b$ to $\omega_z + b$. This suppresses the sender's incentive to overstate information. We use this observation in order to show the following result.

Lemma 2 *For a decreasing $\{\theta_t\}_{t=0}^{T+1}$ such that $\theta_{T-1} \geq \theta_c$ and $\Delta_w(\theta_T, \theta_{T-1}) \geq 2b$, there is a truthtelling equilibrium in the game with $\{\theta_t\}_{t=1}^{T+1}$.*

According to Lemma 2, given a decreasing $(T-1)$ – period communication schedule $\{\theta_t\}_{t=0}^T$ such that $\theta_{T-1} \geq \theta_c$ and $\Delta_w(z, \theta_{T-1}) \geq 2b$ for some $z \in (0, \theta_{T-1})$, the receiver can extract partial information in $[0, \theta_{T-1}]$ by specifying the decreasing T – period communication schedule $\{\hat{\theta}_t\}_{t=0}^{T+1}$ which replicates the original one in periods $t = 1, \dots, T-1$ and is such that $\hat{\theta}_T = z$. Also, the condition $\Delta_w(\theta_T, \theta_{T-1}) \geq 2b$ in the Lemma holds, for example, if

$$E[\theta|\theta \leq \theta_{T-1}] > 2b,$$

and θ_T is sufficiently small. Intuitively, suppose that the sender in the last communication period knows that θ is either 0 or is in $(0, \theta_{T-1}]$ and the receiver believes that the sender is truthful. If the sender observes $\theta = 0$, she has a choice between inducing the lowest rationalizable decision $a = 0$ and $a' = E[\theta|\theta \leq \theta_{T-1}]$. If $E[\theta|\theta \leq \theta_{T-1}] > 2b$, the sender prefers to report truthfully that $\theta = 0$ as the decision a' is too far from her optimal decision b . Therefore, partitioning $[0, \theta_{T-1}]$ into $[0, \theta_T]$ and $[\theta_T, \theta_{T-1}]$, such that $\theta_T \downarrow 0$ preserves the sender's incentive to reveal her information truthfully.

4 Welfare analysis: the uniform-quadratic case

In this section, we evaluate the receiver's ex-ante benefits in dynamic informational control versus those in one-stage communication and other organizational forms. For that purpose, we restrict attention to partitional information structures and the leading uniform-quadratic setup.¹⁶ First, note that in any

¹⁶ The uniform-quadratic setup is widely used in the literature on strategic communication. See, for example, Blume, Board, and Kawamura (2007), Gilligan and Krehbiel (1989, 1987),

equilibrium of a sender–receiver communication game (i.e. the game without receiver’s commitment to actions) with quadratic preferences of players, the receiver’s ex-ante payoff is given by¹⁷

$$EU = \text{Var}[a] - \text{Var}[\theta], \quad [3]$$

where $\text{Var}[y]$ is a variance of a random variable y .

Now, we evaluate the upper limit on the receiver’s ex-ante payoff in equilibria with decreasing communication schedules when $T \rightarrow \infty$. If $\lim_{T \rightarrow \infty} \max_{t=1, \dots, T-1} |\theta_t - \theta_{t-1}| = 0$, then the receiver learns $\theta \in [\theta_{T-1}, 1]$ perfectly in the limit as $T \rightarrow \infty$. Thus, the upper limit on the receiver’s ex-ante payoff depends only on θ_{T-1} and θ_T :

$$\begin{aligned} \overline{EU}(\theta_{T-1}, \theta_T) &= \lim_{\substack{T \rightarrow \infty, \\ \max_{t=1, \dots, T-1} |\theta_t - \theta_{t-1}| \rightarrow 0,}} EU^T = - \sum_{\tau=T-1}^T \int_{\theta_{\tau+1}}^{\theta_{\tau}} \left(\frac{\theta_{\tau+1} + \theta_{\tau}}{2} - \theta \right)^2 d\theta \\ &= -\frac{1}{12} \theta_T^3 - \frac{1}{12} (\theta_{T-1} - \theta_T)^3, \end{aligned} \quad [4]$$

where

$$EU^T = - \sum_{\tau=1}^T \int_{\theta_{\tau+1}}^{\theta_{\tau}} \left(\frac{\theta_{\tau+1} + \theta_{\tau}}{2} - \theta \right)^2 d\theta.$$

Given the constraint $\theta_{T-1} = \theta_c = 4b$, $\overline{EU}(\theta_{T-1}, \theta_T)$ is maximized at $\theta_T = 2b$, which results in the ex-ante payoff

$$\overline{EU} = -\frac{4}{3} b^3. \quad [5]$$

We can now compare the ex-ante benefits of the single-stage and multi-stage communication games. In the single-stage communication game, the incentive-compatibility constraints are

$$\Delta\theta_{k+1} + \Delta\theta_k \geq 4b, \forall k, \quad [6]$$

where $\Delta\theta_k = \theta_{k+1} - \theta_k$ is the length of an interval $[\theta_k, \theta_{k+1}]$ in the partitioned information structure.¹⁸ By (6) and $\theta_{T-1} = 4b$, there exist informative equilibria

Goltsman et al. (2009), Krishna and Morgan (2008, 2004, 2001), Melumad and Shibano (1991), Ottaviani and Squintani (2006).

¹⁷ See, for example, Goltsman et al. (2009).

¹⁸ Given the message that $\theta \in \Theta_k$, the receiver’s best response to the sender’s truth-telling strategy is $a_k = \omega_k = E[\theta | \theta \in \Theta_k] = \frac{\theta_k + \theta_{k+1}}{2}$. Upon observing Θ_k , the sender prefers a_k to

in both one-stage and multi-stage versions of the game if $b \leq 1/4$. In multi-stage communication the inequality $\Delta\theta_T + \Delta\theta_{T-1} \geq 4b$ only must hold, while the distance $\theta_{t-1} - \theta_t, t < T$ can be arbitrarily small. Also, since the players' ex-ante payoffs differ by b^2 , all equilibria are Pareto ranked.¹⁹ This results in the following theorem.

Theorem 2 *For $b < 1/4$ and any equilibrium ex-ante payoff EU^1 in the single-stage communication game, there is an equilibrium in the multi-stage communication game such that $EU^T > EU^1$. Also, $\frac{EU^T}{EU^1} \rightarrow 0$ as $b \rightarrow 0$.*

According to this theorem, if informative CS communication is feasible, i.e. $b < 1/4$, then dynamic informational control is ex-ante payoff superior to such organizational forms as optimal delegation (with the perfectly informed sender) and communication through a disinterested mediator. This is because these mechanisms are ex-ante inferior to a single-stage informational control.²⁰ Moreover, as $b \rightarrow 0$, the ex-ante benefits of multi-stage communication relative to one-stage communication rise without a bound. This is because the average length of intervals in the optimal partition of the single-stage communication game has an order of $2b$. Thus, the receiver's ex-ante payoff has an order of $-\frac{b^2}{3}$. In contrast, the receiver's ex-ante payoff in truthtelling equilibria given the optimal decreasing communication schedule is characterized by the residual variance of θ conditional on $\theta \in [0, 4b]$ multiplied by the probability of that event, which has an order of $-\frac{4}{3}b^3$.

5 Conclusion and discussion

This paper demonstrates that through communication with an imperfectly informed sender in multiple rounds in which the receiver controls the precision of the sender's information, the receiver can derive almost all information for a subinterval of the state space. This results in an ex-ante Pareto improvement compared to one-stage communication. Moreover, as the bias in players'

$a \geq a_{k+1} > a_k$ if $\frac{a_{k+1}+a_k}{2} \geq \omega_k + b = a_k + b$. This implies $a_{k+1} - a_k = \frac{\theta_{k+1}+\theta_{k+2}}{2} - \frac{\theta_k+\theta_{k+1}}{2} \geq 2b$ and $\theta_{k+2} - \theta_k = \Delta\theta_{k+1} + \Delta\theta_k \geq 4b$. Also, because $V(a|\theta_k)$ is decreasing in a for all $a < \omega_k + b$, the sender never prefers $a < a_k$ to a_k .

19 For quadratic preferences, CS show that $EV = EU - b^2$, where EV and EU are the sender's and the receiver's ex-ante payoffs, respectively. This relationship holds in any equilibrium of our model as well.

20 See Ivanov (2010) and Goltsman et al. (2009) show that the optimal delegation is an optimal arbitration mechanism, i.e. the communication protocol in which players communicate via the neutral arbiter who can enforce his recommendations. Because mediation is a special case of arbitration, it is ex-ante payoff inferior to the optimal delegation.

preferences decreases, the relative ex-ante performance of multi-period interaction versus single-stage game rises without a bound.

It is important to highlight several factors that influence our results. First, the performance of dynamic informational control depends on the shape of the distribution of states. At the first glance, it seems that if the distribution is concentrated near 0, for example, if the density of the distribution is decreasing, then dynamic informational control is not highly effective as the receiver only learns precise information about high states, which is unlikely. This is not true in general. For example, for $b = 3/14$ and a decreasing communication schedule $\{\theta_t\}_{t=0}^{T+1}$, such that $\theta_T = \theta_c$, the receiver's ex-ante payoff $\overline{EU}$ in the limit as $T \rightarrow \infty$ is approximately equal to $-1/56$ for the triangle distribution skewed to the left, and to the uninformative payoff $-1/12$ for the triangle distribution skewed to the right. This is because the value of θ_c depends on the shape of the distribution also. By construction, θ_c is the smallest interval $[0, \theta_{T-1}]$ which sustains informative communication in the subgame at the last communication period, which is a single-stage communication game with the partitional information structure $\{[0, \theta_{T-1}), [\theta_{T-1}, 1]\}$. Then, if the density of the distribution is increasing and θ_{T-1} is low, the sender strictly benefits from distorting information and inducing decision $a_{T-1} = \theta_{T-1}$ upon learning that $\theta \in [0, \theta_{T-1})$ since this decision is closer to the sender's optimal decision $E[\theta | \theta < \theta_{T-1}] + b$ than decision $a_T = E[\theta | \theta < \theta_{T-1}]$ induced by truthful communication. As a result, θ_c is higher for distributions in which high states are more likely.

Second, the efficiency of dynamic informational control is affected by the communication horizon T . The perfect learning of $\theta \in [\theta_c, 1]$ requires an infinite communication horizon T . Given the receiver's utility EU^T in the game with a finite T , the relative difference between EU^T and $\overline{EU}$ can serve as the measure of imperfection of extracted information $\varepsilon = \left| \frac{\overline{EU} - EU^T}{\overline{EU}} \right|$. Consider, for example, the uniform-quadratic setup. For a fixed T and decreasing communication schedules $\{\theta_t\}_{t=0}^{T+1}$ which satisfy the conditions in Lemma 2, the receiver's ex-ante payoff is maximized for $\{\theta_t\}_{t=0}^{T+1}$ such that $\theta_T = 2b$ and $\theta_t = 1 - \frac{1-4b}{T-1}t$, $t = 1, \dots, T-1$. This results in the ex-ante payoff $EU^T = \overline{EU} - \frac{1}{12} \frac{(1-4b)^3}{(T-1)^2}$. It follows then that the communication horizon T_ε , which guarantees that the loss in the ex-ante payoff does not exceed ε , increases as $\varepsilon^{-1/2}$. Referring to the above example of $b = \frac{3}{14}$, the receiver's ex-ante payoff $\overline{EU}$ is approximately $-\frac{1}{76}$. However, only two periods of communication provide the ex-ante payoff $-\frac{1}{75}$, so that $\varepsilon \simeq 2\%$.²¹

²¹ Also, T_ε is quite sensitive to the value of b and the order of $b^{-3/2}$. For example, for $\varepsilon = \frac{1}{10}$, it follows that $T_\varepsilon = 11$ if $b = \frac{1}{10}$ and $T_\varepsilon = 51$ if $b = \frac{1}{20}$.

In the analysis above we assumed that the sender's learning of information is costless and the players do not discount the time taken by learning the state. Evaluating the efficiency of communication with costly experiments or with a discount factor seems to be an interesting avenue for future research. At the same time, the characterization of effective information schedules becomes a complicated question. For example, if the cost of a single test is positive but does not depend on the structure of the test, then multi-stage informational control can be suboptimal. This is because such a schedule updates the receiver's information gradually, so the accumulated costs of learning may exceed the benefits of multi-stage learning and communication. In the case of a decreasing communication schedule, if the receiver in stage $t + 1$ knows that θ is in the interval $[0, \theta_{t+1}]$ instead of $[0, \theta_t]$, this updates his information insignificantly. Therefore, if the cost of a single test is sufficiently high, it might be beneficial to specify a single-period information structure with two cutoffs, θ_{t+1} and θ_t , instead of a two-period communication schedule with single cutoffs.

Acknowledgements: I am grateful to Dirk Bergemann, Oliver Board, Kalyan Chatterjee, Paul Fischer, Vijay Krishna, Bart Lipman, Robert Marshall, Marek Pycia, Neil Wallace, and three anonymous referees for their valuable suggestions and discussions. Misty Ann Stone provided invaluable help with copy editing the manuscript. All mistakes are mine.

Appendix

In this section we provide proofs of the results.

Proof of Lemma 1. Denote $\tilde{\sigma} = \left(\{\tilde{\sigma}_t^E(h_t)\}_{t=1}^T, \tilde{a}(h_{T+1}^0) \right)$ an equilibrium profile of the game with commitment to the information structures $\{\tilde{I}_t(h_t^0)\}_{t=1}^T$. Note that for any history h_{T+1}^0 , the optimal decision rule is $\tilde{a}(h_{T+1}^0) = E[\theta|h_{T+1}^0]$. In the game without commitment, define the following “trigger” function

$$\tau_d(h_t^0, I^t) = \begin{cases} T + 1 & \text{if } I^t = \{\tilde{I}_i(h_t^0)\}_{i=1}^t, \\ \min\{i : I_i \neq \{\tilde{I}_i(h_t^0)\}_{i=1}^t\} & \text{otherwise,} \end{cases}$$

which indicates the first period in which the receiver deviates from $\{\tilde{I}_t(h_t^0)\}_{t=1}^T$. Consider the sender's strategy

$$\bar{\sigma}_t^E(h_t) = \begin{cases} \tilde{\sigma}_t^E(h_t) & \text{if } \tau_d(h_t^0, I^t) = T + 1 \text{ and} \\ m_t^0 & \text{otherwise,} \end{cases} \quad [7]$$

which implies that the sender follows the strategy $\{\tilde{\sigma}_t^E(h_t)\}_{t=1}^T$ as long as the receiver follows $\{\tilde{I}_t(h_t^0)\}_{t=1}^T$. In contrast, any deviation by the receiver results in completely uninformative messages m_t^0 in the current and all future periods. The receiver's best response to this strategy is

$$\bar{a}(h_{T+1}^0) = \tilde{a}\left(h_{\tau_d(h_t^0, I^T)}^0\right) = E\left[\theta|h_{\tau_d(h_t^0, I^T)}^0\right].$$

That is, the receiver makes a decision on the basis of information up to the period of deviation from $\{\tilde{I}_t(h_t^0)\}_{t=1}^T$ only and ignores all other messages.

Consider the strategy profile $\bar{\sigma} = \left(\{\bar{\sigma}_t^E(h_t)\}_{t=1}^T, \{\tilde{I}_t(h_t^0)\}_{t=1}^T, \bar{a}(h_{T+1}^0)\right)$, and suppose that the receiver deviates from it by choosing $\sigma_P = \left(\{I_t(h_t^0)\}_{t=1}^T, a(h_{T+1}^0)\right)$, where $I_t \neq \tilde{I}_t(h_t^0)$ for some h_t^0 . It follows that $\tau_d(h_t^0, I^t) = \tau < T + 1, t = \tau, \dots, T$, so that $\bar{\sigma}_t^E(h_t) = m_t^0, t = \tau, \dots, T$. Since the receiver does not receive useful information in periods $t = \tau, \dots, T$, his interim payoff is bounded from above by

$$U^P(\sigma^P, \bar{\sigma}^E|h_t^0) \leq E_\theta[U^P(\tilde{a}(h_t^0), \theta)|h_t^0].$$

However, $U^P(\bar{\sigma}^P, \bar{\sigma}^E|h_t^0) \geq E_\theta[U^P(\tilde{a}(h_t^0), \theta)|h_t^0]$, because following $\{\tilde{I}_t(h_t^0)\}_{t=\tau}^T$ results in a history h_{T+1}^0 , which includes h_t^0 . Therefore, the receiver may choose $a(h_{T+1}^0) = \tilde{a}(h_t^0)$ and replicate the interim payoff $E_\theta[U^P(\tilde{a}(h_t^0), \theta)|h_t^0]$.

Also, the sender's strategy (7) is the best response to the receiver's strategy. First, for all h_t such that the receiver follows $\tilde{I}_t(h_t^0)$, the sender cannot beneficially deviate from $\tilde{\sigma}_t^E(h_t)$, since it is optimal in the game with commitment to the communication schedule. Finally, if h_t involves deviation from $\tilde{I}_t(h_t^0)$, then sending uninformative messages $\{m_t^0\}_{t=\tau}^T$ is also optimal, since $\bar{a}(h_{T+1}^0)$ does not depends on them. ■

Proof of Theorem 1. Define $\Delta\Theta_t = [\theta_t; \theta_{t-1}], t = 1, \dots, T + 1$. Given the sender's truthtelling strategy $\{m_t(h^t)\}_{t=1}^T = \{s_t\}_{t=1}^T$, the receiver's consistent posterior beliefs are such that θ is distributed on $\Delta\Theta_{\tau_1}$ according to distribution $F[\theta|\theta \in \Delta\Theta_{\tau_1}]$,²² where

$$\tau_1 = \begin{cases} \min\{\tau : m_\tau = \Theta_1^s\} & \text{if } \{m_t\}_{t=1}^T \neq \{\Theta_0^t\}_{t=1}^T, \text{ and} \\ T + 1 & \text{if } \{m_t\}_{t=1}^T = \{\Theta_0^t\}_{t=1}^T. \end{cases}$$

For such beliefs, the receiver's best response is

$$a(h_{T+1}^0) = \omega_{\tau_1} = E[\theta|\theta \in \Delta\Theta_{\tau_1}]. \quad [8]$$

²² The receiver's posterior beliefs are consistent for both equilibrium and off-equilibrium messages of the sender. If the sender sends, say, a sequence $\{m_1, m_0\}$, which implies $\theta \in [\theta_1, 1] \cap [0, \theta_0] = \phi$, the receiver interprets these messages as $\theta \in [\theta_1, 1]$.

We show now that the sender cannot strictly benefit by distorting information at any stage conditional on any history.

A.1 Equilibrium path

Given the state θ and the receiver's best response (8), consider the sender's incentives in any period conditional on truthful reporting in previous periods. Let $t = \min\{\tau : s_\tau = \Theta_1^t\}$ be first stage in which the sender observes the higher interval Θ_1^t . For a decreasing communication schedule, we have $\Theta_1^t \subset \Theta_1^{t+1}, \forall \tau$. Thus, $s_\tau = \Theta_1^t, \tau \geq t$, and the sender's information is not updated in periods $\tau > t$. Also, we have $\{s_\tau\}_{\tau=1}^{t-1} = \{\Theta_0^t\}_{\tau=1}^{t-1}$. Hence, the sender infers that $\theta \in \Delta\Theta_t$, so that her optimal interim decision is $\omega_t + b$. Truthtelling in the previous periods implies $\{m_\tau\}_{\tau=1}^{t-1} = \{\Theta_0^t\}_{\tau=1}^{t-1}$. Then, message $m_t = \Theta_1^t$ induces the decision $a_t = \omega_t$ unconditionally on messages $\{m_\tau\}_{\tau=t+1}^T$. In contrast, if $m_t = \Theta_0^t$, then the set of feasible decisions in the future periods is $\{a_\tau\}_{\tau=t+1}^{T+1} = \{\omega_\tau\}_{\tau=t+1}^{T+1}$. Because $\omega_\tau < a_t < \omega_\tau + b, \tau > t$, this means $V(a_\tau | \Delta\Theta_t) < V(a_t | \Delta\Theta_t), \tau > t$. Hence, the sender is strictly worse off by sending $m_t = \Theta_0^t$.

If $s_t = \Theta_0^t$, the sender infers that $\theta \in \Theta_0^t$. Given $\{m_\tau\}_{\tau=1}^{t-1} = \{\Theta_0^t\}_{\tau=1}^{t-1}$, message $m_t = \Theta_1^t$ induces action $a_t = \omega_t$ unconditionally on $\{m_\tau\}_{\tau=t+1}^T$. This results in the interim payoff to the sender:

$$\begin{aligned} V(a_t | \Theta_0^t) &= E[V(\omega_t, \theta) | \Theta_0^t] = -E[(\omega_t - \theta - b)^2 | \Theta_0^t] \\ &= -E[(\omega_t - \theta)^2 | \Theta_0^t] + 2bE[\omega_t - \theta | \Theta_0^t] - b^2 \\ &= -E[(\omega_t - v_t + v_t - \theta)^2 | \Theta_0^t] + 2bE[\omega_t - \theta | \Theta_0^t] - b^2 \\ &= -(\omega_t - v_t)^2 - 2(\omega_t - v_t)E[v_t - \theta | \Theta_0^t] \\ &\quad - E[(\theta - v_t)^2 | \Theta_0^t] + 2bE[\omega_t - \theta | \Theta_0^t] - b^2 \\ &= -(\omega_t - v_t)^2 - E[(\theta - v_t)^2 | \Theta_0^t] + 2b(\omega_t - v_t) - b^2 \\ &= -(\omega_t - v_t)(\omega_t - v_t - 2b) - E[(\theta - v_t)^2 | \Theta_0^t] - b^2, \end{aligned}$$

where $v_t = E[\theta | \Theta_0^t] = E[\theta | \theta \leq \theta_t] < \theta_t < \omega_t$. By $\omega_t > \theta_t \geq \theta_c$ and (2), we have $\omega_t - v_t - 2b > \theta_t - E[\theta | \theta \leq \theta_t] - 2b \geq 0$. This leads to $(\omega_t - v_t)(\omega_t - v_t - 2b) > 0$ and

$$V(a_t|\Theta_0^t) \leq -E[(\theta - v_t)^2|\Theta_0^t] - b^2.$$

Now, consider the sender's interim payoff from sending $\{m_\tau\}_{\tau=t}^T = \{s_\tau\}_{\tau=t}^T$:

$$\begin{aligned} E[V(a(\{s_\tau\}_{\tau=t}^T), \theta)|\Theta_0^t] &= \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} V(\omega_{\tau+1}, \theta) dF(\theta|\Theta_0^t) \\ &= - \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\omega_{\tau+1} - \theta - b)^2 dF(\theta|\Theta_0^t) \\ &= - \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\theta - \omega_{\tau+1})^2 dF(\theta|\Theta_0^t) + 2b \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\omega_{\tau+1} - \theta) dF(\theta|\Theta_0^t) - b^2 \\ &= - \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\theta - \omega_{\tau+1})^2 dF(\theta|\Theta_0^t) - b^2 \geq \int_0^{\theta_t} (\theta - v_t)^2 dF(\theta|\Theta_0^t) - b^2 \\ &= -E[(\theta - v_t)^2|\Theta_0^t] - b^2, \end{aligned}$$

where the inequality follows from

$$\begin{aligned} - \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\theta - \omega_{\tau+1})^2 dF(\theta|\Theta_0^t) &= - \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} ((\omega_{\tau+1} - v_t) - (\theta - v_t))^2 dF(\theta|\Theta_0^t) \\ &= - \int_0^{\theta_t} (\theta - v_t)^2 dF(\theta|\Theta_0^t) + 2 \sum_{\tau=t}^T \int_{\theta_{\tau+1}}^{\theta_t} (\omega_{\tau+1} - v_t)(\theta - v_t) dF(\theta|\Theta_0^t) \\ &\quad - \sum_{\tau=t}^T P(\Delta\Theta_{\tau+1}|\Theta_0^t)(\omega_{\tau+1} - v_t)^2 \\ &= - \int_0^{\theta_t} (\theta - v_t)^2 dF(\theta|\Theta_0^t) + 2 \sum_{\tau=t}^T P(\Delta\Theta_{\tau+1}|\Theta_0^t)(\omega_{\tau+1} - v_t)^2 \\ &\quad - \sum_{\tau=t}^T P(\Delta\Theta_{\tau+1}|\Theta_0^t)(\omega_{\tau+1} - v_t)^2 \\ &= - \int_0^{\theta_t} (\theta - v_t)^2 dF(\theta|\Theta_0^t) + \sum_{\tau=t}^T P(\Delta\Theta_{\tau+1}|\Theta_0^t)(\omega_{\tau+1} - v_t)^2 \\ &\geq \int_0^{\theta_t} (\theta - v_t)^2 dF(\theta|\Theta_0^t), \end{aligned}$$

where $P(\Delta\Theta_{\tau+1}|\Theta_0^t) = \frac{F(\theta_\tau) - F(\theta_{\tau+1})}{F(\theta_t)}$. Therefore,

$$E\left[V\left(a\left(\{s_\tau\}_{\tau=t}^T\right), \theta\right) | \Theta_0^t\right] = E\left[(\theta - v_t)^2 | \Theta_0^t\right] - b^2 \geq V(a_t | \Theta_0^t),$$

so the sender cannot beneficially deviate from reporting truthfully at any $t = 1, \dots, T$ conditional on truthtelling in all previous stages.

A.2 Off-equilibrium path

Consider the sender's incentives at stage $t = 2, \dots, T$ given that she was not truthful in previous stage(s), i.e. h_t is such that $\{m_\tau\}_{\tau=1}^{t-1} \neq \{s_\tau\}_{\tau=1}^{t-1}$. If $\{m_\tau\}_{\tau=1}^{t-1} \neq \{\Theta_0^\tau\}_{\tau=1}^{t-1}$, then the receiver takes an action ω_{τ_1} , where $\tau_1 = \min\{\tau : m_\tau = \Theta_1^\tau\} < t$. Because the receiver's decision does not depend on $\{m_\tau\}_{\tau=t}^T$, the sender cannot benefit by distorting information at t .

Now, suppose that $\{m_\tau\}_{\tau=1}^{t-1} = \{\Theta_0^\tau\}_{\tau=1}^{t-1}$. Since $\{s_\tau\}_{\tau=1}^{t-1} \neq \{m_\tau\}_{\tau=1}^{t-1}$, it follows that $s_i = \Theta_1^i$, where $i = \min\{\tau : s_\tau = \Theta_1^\tau\} < t$, and $\{s_\tau\}_{\tau=i}^T = \{\Theta_1^\tau\}_{\tau=i}^T$. Hence, the sender's information at t is given by $\Delta\Theta_i$ and is not updated afterward. Since $\{m_\tau\}_{\tau=1}^{t-1} = \{\Theta_0^\tau\}_{\tau=1}^{t-1}$, then message $m_t = \Theta_1^t$ induces action $a_t = \omega_t$ unconditionally on $\{m_\tau\}_{\tau=t+1}^T$. If $m_t = \Theta_0^t$, then the set of feasible actions is $\{a_\tau\}_{\tau=t+1}^{T+1} = \{\omega_\tau\}_{\tau=t+1}^{T+1}$. Because $a_\tau < \omega_t < \omega_t + b$, $\tau > t$, we have $V(a_\tau | \Delta\Theta_t) < V(a_t | \Delta\Theta_t)$, $\tau > t$. Hence, the sender is strictly worse off by sending $m_t = \Theta_0^t$. Thus, she cannot beneficially deviate from truthtelling at any t conditional on misreporting at previous stages. ■

Proof of Lemma 2. Given $\theta \in [0, \theta_T]$ and truthful reporting in periods $t = 1, \dots, T - 1$, consider the subgame at the beginning of the last communication period $t = T$. This sub-game is a game with single-period communication, the prior distribution $F[\theta | \theta \in [0, \theta_{T-1}]]$, and the sender's information structure given by the two-interval partition $\{[0, \theta_T], [\theta_T, \theta_{T-1}]\}$. Then, the receiver's best responses to the sender's truthful reporting of signals $s_t = \Theta_0^T$ and $s_t = \Theta_1^T$ are $\omega_{T+1} = \omega_0^{\theta_T} = E[\theta | \theta \leq \theta_T]$ and $\omega_T = \omega_{\theta_T}^{\theta_{T-1}} = E[\theta | \theta_T \leq \theta \leq \theta_{T-1}]$, respectively. Because the sender's incentive-compatibility constraint

$$\Delta_w(\theta_T, \theta_{T-1}) = \omega_{\theta_T}^{\theta_{T-1}} - \omega_0^{\theta_T} = \omega_T - \omega_{T+1} \geq 2b$$

holds by conditions of the Lemma, then Lemma 1 in Ivanov (2010) implies that there is a truthtelling equilibrium in this game. Because $\theta_{T-1} \geq \theta_c$, then following the same lines as those in Theorem 1 shows that (8) and truthful reporting in periods $t = 1, \dots, T - 1$ are the receiver's and sender's equilibrium strategies in the original game. ■

Proof of Theorem 2. Consider a single-stage communication game with the partitional information structure given by strictly increasing cutoffs $\{\theta_\tau\}_{\tau=0}^K$, where $\theta_0 = 0$ and $\theta_T = 1$, and focus on equilibria which maximize the receiver's ex-ante payoff. By Lemma 2 in Ivanov (2010), we can restrict attention to incentive-compatible information structures and truthtelling equilibria under these structures. In these equilibria, $a_\tau = E[\theta | \theta \in [\theta_{\tau-1}, \theta_\tau]] = \frac{\theta_{\tau-1} + \theta_\tau}{2}$, $\tau = 1, \dots, K$, and the sender's incentive-compatibility constraints are

$$\theta_{\tau+1} - \theta_{\tau-1} = \Delta\theta_{\tau+1} + \Delta\theta_\tau \geq 4b, \tau = 1, \dots, K-1, \quad [9]$$

where $\Delta\theta_\tau = \theta_\tau - \theta_{\tau-1}$, $\tau = 1, \dots, K$. If $b < 1/4$, Lemma 3 in Ivanov (2010) implies that informative communication is feasible, i.e. there is a partitional information structure with $K \geq 2$ intervals which satisfies (9). The ex-ante payoff to the receiver in the truthtelling equilibrium under this structure is

$$EU^1 = - \sum_{\tau=1}^K \int_{\theta_{\tau-1}}^{\theta_\tau} (a_\tau - \theta)^2 d\theta = - \sum_{\tau=1}^K \Delta\theta_\tau^3.$$

Now, consider the multi-period communication game. For the uniform distribution, we have $\theta_c = 4b$. For $K \geq 3$, consider $(K+1)$ - period communication with the decreasing communication schedule $\{\hat{\theta}_\tau\}_{\tau=0}^{K+1}$:

$$\hat{\theta}_\tau = \begin{cases} \hat{\theta}_\tau = \theta_{T-\tau}, \tau = 0, \dots, K-3, \\ \hat{\theta}_{T-2} = \frac{\theta_2 + \theta_3}{2}, \text{ and} \\ \hat{\theta}_\tau = \theta_{T+1-\tau}, \tau = K-1, K, K+1. \end{cases}$$

Since $\Delta\theta_1 + \Delta\theta_2 \geq 4b$ by (9), we have $\hat{\theta}_{K-1} = \theta_2 = \Delta\theta_1 + \Delta\theta_2 \geq 4b = \theta_c$, and $\Delta_w(\hat{\theta}_K, \hat{\theta}_{K-1}) \geq \frac{\hat{\theta}_K}{2} + \left(\frac{\hat{\theta}_{K-1}}{2} - \frac{\hat{\theta}_K}{2}\right) = \frac{\hat{\theta}_{K-1}}{2} \geq 2b$. Hence, by Lemma 2 there exists the truthtelling equilibrium in the game. The ex-ante payoff to the receiver in this equilibrium is $EU^{K+1} = - \sum_{\tau=1}^{K+1} \Delta\hat{\theta}_\tau^3$, where $\Delta\hat{\theta}_\tau = \hat{\theta}_{\tau-1} - \hat{\theta}_\tau$, $\tau = 1, \dots, K+1$. Then,

$$EU^{K+1} - EU^1 = -\Delta\hat{\theta}_{K-2}^3 - \Delta\hat{\theta}_{K-1}^3 + \Delta\theta_3^3 = -\Delta\hat{\theta}_{K-2}^3 - \Delta\hat{\theta}_{K-1}^3 + \left(\Delta\hat{\theta}_{K-2} + \Delta\hat{\theta}_{K-1}\right)^3 > 0,$$

since $\Delta\theta_\tau > 0$, $\Delta\hat{\theta}_\tau > 0$, and $\Delta\theta_3 = \Delta\hat{\theta}_{K-2} + \Delta\hat{\theta}_{K-1}$.

If $K = 2$, then $\Delta\theta_1 + \Delta\theta_2 = 1 > 4b$, and the highest ex-ante payoff to the receiver $\overline{EU}^1 = -\frac{1}{48}$ in the single-stage communication game is achieved in the truthtelling equilibrium for $\theta_1 = \frac{1}{2}$. Consider a 3 - period communication game with the decreasing communication schedule $\{\hat{\theta}_\tau\}_{\tau=0}^3$, such that

$\hat{\theta}_1 = \max\{4b, \frac{3}{4}\} < \hat{\theta}_0 = 1$ and $\hat{\theta}_2 = \frac{1}{2} < \hat{\theta}_1$. Because $\hat{\theta}_1 \geq 4b = \theta_c$ and $\Delta_w(\hat{\theta}_2, \hat{\theta}_1) = \frac{\hat{\theta}_1}{2} \geq 2b$, by Lemma 2 there is the truthtelling equilibrium which provides the ex-ante payoff to the receiver $EU^3 = -\sum_{\tau=1}^3 \Delta \hat{\theta}_\tau^3$. Thus,

$$EU^3 - EU^1 \geq EU^3 - \overline{EU}^1 = -\Delta \hat{\theta}_1^3 - \Delta \hat{\theta}_2^3 + \Delta \theta_2^3 = -\Delta \hat{\theta}_1^3 - \Delta \hat{\theta}_2^3 + \left(\Delta \hat{\theta}_1 + \Delta \hat{\theta}_2\right)^3 > 0,$$

since $\Delta \hat{\theta}_1 > 0$, $\Delta \hat{\theta}_2 > 0$, and $\Delta \theta_2 = \Delta \hat{\theta}_1 + \Delta \hat{\theta}_2$.

If $b \rightarrow 0$, then Lemma 3 in Ivanov (2010) implies

$$EU^1(b) \leq -\frac{1}{12\bar{K}^2(b)} \leq -\frac{1}{12(\frac{1}{2b} + 1)^3} = -\frac{b^2}{3} + O(b^3),$$

where $\bar{K}(b) \leq \frac{1}{2b} + 1$ is the maximum size of incentive-compatible partitions and $O(b)$ has the order of b . From (5), we have $\frac{\overline{EU}(b)}{EU^1(b)} = 4b + O(b^2)$, that is, $\lim_{b \rightarrow 0} \frac{\overline{EU}(b)}{EU^1(b)} = 0$. ■

References

- Aumann, R., and S. Hart. 2003. "Long Cheap Talk." *Econometrica* 71:1619–60.
- Austen-Smith, D. 1994. "Strategic Transmission of Costly Information." *Econometrica* 62:955–63.
- Battaglini, M. 2004. "Policy Advice with Imperfectly Informed Senders." *Advances in Theoretical Economics* 4(1), Article 1.
- 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: 1431–51.
- Dessein, W. 2002. "Authority and Communication in Organizations." *Review of Economic Studies* 69:811–38.
- Feldman, D., and M. Fox. 1991. *Probability: The Mathematics of Uncertainty*. New York, NY: Marcel Dekker Inc.
- Fischer, P., and P. Stocken. 2001. "Imperfect Information and Credible Communication." *Journal of Accounting Research* 39:119–34.
- General Accounting Office. 1988. *Weapons Testing: Quality of DoD Operational Testing and Reporting*. Washington, DC: U.S. Government Printing Office.
- General Accounting Office. 1992. *Electronic Warfare: Established Criteria Not Met for Airborne Self-Protection Jammer Production*. Washington, DC: U.S. Government Printing Office.
- 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.

- Gilligan, T., and K. Krehbiel. 1989. "Asymmetric Information and Legislative Rules with a Heterogeneous Committee." *American Journal of Political Science* 33:459–90.
- Golosov, M., V. Skreta, A. Tsyvinski, and A. Wilson. 2013. "Dynamic Strategic Information Transmission." Working paper.
- Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani. 2009. "Mediation, Arbitration and Negotiation." *Journal of Economic Theory* 144:1397–420.
- Green, J., and N. Stokey. 2007. "A Two-Person Game of Information Transmission." *Journal of Economic Theory* 135:90–104.
- Ivanov, M. 2010. "Informational Control and Organizational Design." *Journal of Economic Theory* 145:721–51.
- Ivanov, M. 2013. "Dynamic Learning and Strategic Communication." Working paper.
- Krishna, V., and J. Morgan. 2001. "Asymmetric Information and Legislative Rules: Some Amendments." *American Political Science Review* 95:435–52.
- Krishna, V., and J. Morgan. 2004. "The Art of Conversation: Eliciting Information from Senders Through Multi-Stage Communication." *Journal of Economic Theory* 117:147–79.
- Krishna, V., and J. Morgan. 2008. "Contracting for Information under Imperfect Commitment." *RAND Journal of Economics* 39:905–25.
- Melumad, N., and T. Shibano. 1991. "Communication in Settings with No Transfers." *RAND Journal of Economics* 22:173–98.
- Ottaviani, M., and F. Squintani. 2006. "Naive Audience and Communication Bias." *International Journal of Game Theory* 35:129–50.
- Rogerson, W. 1990. "Quality Vs. Quantity in Military Procurement." *American Economic Review* 80:83–92.