



On Strategic Transmission of Gradually Arriving Information

Alexander Frug

January 2020

Barcelona GSE Working Paper Series

Working Paper n° 1146

On Strategic Transmission of Gradually Arriving Information

Alexander Frug*

December 24, 2019

Abstract

The main insight of the literature on strategic information transmission is that even a small conflict of interest between a fully informed sender (e.g., a financial adviser) and an uninformed receiver (an investor) often poses considerable difficulties for effective communication. However, in many real-life situations, the sender is not fully informed at the outset but gradually studies the case before offering advice. The gradual arrival of information to the sender weakens the strategic barriers between the players and significantly improves communication.

JEL classification: D82, D83

*Universitat Pompeu Fabra and Barcelona Graduate School of Economics. E-mail: alexander.frug@upf.edu.

1 Introduction

Decision-makers often rely on experts' advice. If the interests of the parties involved are not perfectly aligned, the expert may have an incentive to conceal or misrepresent information in an attempt to manipulate the decision-maker's actions. In Crawford and Sobel's (1982) model of strategic costless communication, a fully informed sender (expert) reports to an uninformed receiver (decision-maker), who then takes an action that affects the payoffs of both players.¹ A fundamental insight of the model is that even a small conflict of interest between the players may significantly restrict the amount of information that can be transmitted in equilibrium.

The assumption that the sender is fully informed at the beginning of the interaction may be a reasonable approximation in some cases. However, many real-world advising interactions involve an important stage of learning on the expert's part. For example, a therapist needs to inspect the patient before advising an appropriate treatment; a financial adviser has to study the client's financial situation prior to offering certain plans and products; an appraiser needs to inspect many different characteristics of an asset before suggesting a value or a sale price to the buyer. Sometimes, while acquiring information, the expert may communicate intermediate and inconclusive impressions, even though it is well understood that the receiver will make a decision only upon the expert's final word.

This paper is concerned with advising interactions in which the expert's information *gradually* improves over time. I argue that, in many situations, this natural friction together with a simple "update-as-you-go" reporting protocol significantly weakens the adverse effect that slightly conflicting interests are believed to have on the overall quality of communication.

Recently, several models of cheap-talk communication in which the sender's learning process is endogenously controlled by one of the players have been

¹Numerous variants and extensions of the model were developed in recent decades in an attempt to improve our understanding of various situations of strategic communication. Sobel (2013) provides a comprehensive review of the communication literature.

studied in the literature. Ivanov (2015, 2016) and Frug (2016) assume that the dynamic arrival of information can be prespecified by the receiver, and Frug (2018) assumes that the sender covertly selects which experiment to perform in each period. The ability to determine endogenously what information will arrive in each period turns out to be an exceptionally powerful tool; careful design of the sender’s information structures in different periods may fully relax the sender’s incentive constraints and lead to extremely informative communication.²

Models in which the players determine what information arrives in each period offer insights into *optimal dynamic information control*. In many cases, however, the players have a limited effect (if any) on the process of information arrival. Understanding and quantifying the effect of gradualness of information arrival in such cases is challenging: the effect varies from one learning process to another, and it is typically not as stark as that of endogenously designed learning. Nonetheless, identifying the substantial and rather general benefits inherent in such gradualness is essential since it may put many real-life interactions of communication in a different perspective. For instance, a common component that is present in many advising interactions—the expert’s need to study the case before offering advice—may restore high-quality communication in spite of a small conflict of interest.

In this paper, I consider the canonical uniform-quadratic constant-bias specification³ of the model by Crawford and Sobel (1982), and examine the implications of a rich family of learning processes on the overall quality of communication. Specifically, I consider all *binary interval learning processes*. Under any such learning process, in every period, the sender learns whether the state is above or below some threshold. While this class of learning processes is special in some aspects (e.g., the information set in each period is

²See, in particular, Ivanov (2016) and Frug (2016) where a complete separation of the state space is consistent with equilibrium.

³In this I follow many other theoretical papers, such as Blume, Board, and Kawamura (2007), Goltsman, Hörner, Pavlov, and Squintani (2009), Ivanov (2010), and Krishna and Morgan (2001, 2004). Under this specification, the state is uniformly distributed on the unit interval and the players have quadratic loss functions. The conflict of interest between the players is reflected in the sender’s bias $b > 0$.

an interval), it is also very rich: these learning processes differ from one another in the thresholds at which the sender's information is refined over time. Various specifications of the thresholds may induce very different learning dynamics in terms of the general speed of learning and the relative quality of information on different regions of the state space.

The main result of the paper identifies a simple (and tight) condition on the sender's learning, that guarantees the existence of an equilibrium where every interval of the receiver's information partition has the following property: the communication game à la Crawford and Sobel in which the state is uniformly distributed on that interval does not have informative equilibria. In brief, the condition requires that the learning process not contain signals that generate significantly superior information on low states, relative to the information they provide on high states of the world.

To illustrate the result, suppose that the sender's bias is $b = \frac{1}{100}$. Under the most informative Crawford–Sobel equilibrium, the receiver's information partition consists of 7 intervals and the length of the rightmost interval is greater than $\frac{1}{4}$. Now suppose that the sender is initially uninformed, but every period he learns whether the state belongs to the upper or lower half of his previous information set. The general construction developed in this paper will show that, in this case, there exists an equilibrium where the receiver's information partition consists of 32 intervals of equal length. What if this simple learning process on the part of the sender is replaced by an arbitrary binary interval learning process? Provided that it satisfies the aforementioned condition, there exists an equilibrium under which the receiver's information partition consists of more than 25 intervals, all of which are shorter than $\frac{1}{25}$, *irrespectively of the details of the learning process*.

The paper proceeds as follows. Section 2 presents the model and the class of admissible learning processes. Section 3 presents the benchmark and the terminology of short and long intervals. The main result is given in Section 4. Concluding remarks are offered in Section 5. The proof of Proposition 2 appears in Section 6.

2 Model

There are two players, sender and receiver. When the receiver chooses $a \in \mathbb{R}$ at state $\theta \in [0, 1]$, his payoff is

$$-(\theta - a)^2,$$

and the sender's payoff is

$$-(\theta + b - a)^2,$$

where the constant $b > 0$ is the sender's bias. At the start of the game, both players are uninformed and share the common prior belief that θ is drawn from a uniform distribution over the unit interval. In each period, conditional on reaching that period, the sender observes a signal realization and then reports to the receiver (the signals and reports are described below). At the end of each period, the receiver decides whether to *postpone action*, in which case the interaction proceeds to the next period, or to *choose* $a \in \mathbb{R}$, in which case the interaction ends and the players receive payoffs.

Binary interval learning process. In each period $t \in \mathbb{N}$, the sender observes a signal realization $s_t \in \{\text{left}, \text{right}\}$ that reveals whether θ is below ($s_t = \text{left}$) or weakly above ($s_t = \text{right}$) a given threshold in the interior of his information set prior to observing s_t . A particular binary interval learning process determines the thresholds of the signals. Rather than specifying the values of the thresholds directly, it is convenient to describe a learning process by a function that assigns to each finite sequence of signal realizations σ the probability to observe the signal realization “left” immediately after σ ,

$$q(\sigma) = \text{Prob}[(\sigma, \text{left})|\sigma].$$

The set of all binary interval learning processes is fully characterized by the set of all functions $\Sigma \rightarrow (0, 1)$, where Σ denotes the set of all finite binary sequences of $\{\text{left}, \text{right}\}$.

Example 1. The function $q(\cdot) \equiv \frac{1}{2}$ corresponds to the “completely balanced”

learning process in which, every period, the sender's information set is partitioned into two *equal* intervals. From the ex-ante perspective, the sender's information structure at the end of period t is given by 2^t intervals of length 2^{-t} each.

Example 2. Let $q(\sigma_t) = \frac{1}{3^t}$, where σ_t is a sequence of signal realizations before period t . A special feature of this example is that, as time goes by, the sender's information partition does not approach full information. For instance, if $\theta \geq \frac{1}{2}$, the sender's information set is an interval of length greater than $\frac{1}{2}$, regardless of the number of observed signal realizations.⁴

The focus of this paper is on the effect of *gradualness* in information arrival rather than on the effect of the sender's permanent lack of high-quality information. Therefore, only learning processes in which the sender eventually learns the state with an arbitrary accuracy will be considered. The following set of binary interval learning processes is considered admissible:

$$Q = \{q : \Sigma \rightarrow (0, 1) \mid \exists \varepsilon > 0, \forall \sigma \in \Sigma, q(\sigma) \in (\varepsilon, 1 - \varepsilon)\}.$$

It is easy to see that any binary interval learning process in Q has the property that, for any $\theta \in [0, 1]$ and $\delta > 0$, there exists $t \in \mathbb{N}$ such that if the state is θ , then the sender's information set in period t is an interval of length at most δ . From now on, elements of the set Q will be referred to as binary interval learning processes.

Reports. In each period t , after observing the realization of s_t , the sender reports $m_t \in \{\text{left}, \text{right}\}$. The particular reporting protocol is inessential; the result of this paper holds as long as there are at least two available messages for the sender in each period.

⁴Let $[x_T, 1]$ denote the sender's information set at the beginning of period $T \in \mathbb{N}$ after he observes $s_t = \text{right}$, for all $t < T$. Note that $x_T \leq \sum_{t=1}^T \frac{1}{3^t} < \sum_{t=1}^{\infty} \frac{1}{3^t} = \frac{1}{2}$ for all $T \in \mathbb{N}$.

3 Fully Informed Sender Benchmark

The benchmark of the analysis in this paper is the case where a fully informed sender submits a single report to the receiver.⁵ (For our present purposes, I do not restrict the report to be binary but allow it, for instance, to be an element of the unit interval.) By an argument akin to the revelation principle, it can be shown that allowing the sender to transmit multiple reports would not expand the set of equilibrium outcomes, as long as the reporting opportunities are deterministic. Krishna and Morgan (2004) have shown that allowing for multiple reports and introducing randomness into future reporting opportunities may lead to more informative communication. The authors also show how such randomness can be induced endogenously if *both* players are active in the communication phase.⁶

In the gradual learning model of the present paper, multi-stage communication is not used as a means to introduce randomness into future communication opportunities. The value of organizing information transmission in multiple reports will come directly from the gradual arrival of information *to the sender*. To focus on this aspect of multi-stage communication, I assume that the sender's reporting opportunities are not random and that the receiver does not actively participate in a "conversation."

The uninformative equilibrium always exists.⁷ A well-known property of the benchmark is that informative equilibria exist if and only if $b \leq \frac{1}{4}$. An equivalent statement that will be useful in this paper is the following.

Observation 1 *Fix $b > 0$ and an interval J . Consider the information transmission game á la Crawford and Sobel where, instead of the unit interval, the state is uniformly distributed on J . The game has informative equilibria if and only if the length of J is at least $4b$.*

⁵For details on the properties mentioned in this discussion, see Crawford and Sobel (1982), Section 4.

⁶For a more general treatment of a similar idea see Aumann and Hart (2003).

⁷If the sender's report affects the receiver's action in equilibrium, then the equilibrium is *informative*; otherwise the equilibrium is *uninformative*.

Short and Long Intervals

Following Observation 1, for a given value of $b > 0$, an interval of length at least $4b$ will be referred to as a *long* interval and an interval of length below $4b$ will be referred to as a *short* interval. This classification of intervals will play a main role in the analysis.

In the most informative Crawford–Sobel equilibrium, only the leftmost interval in the receiver’s information partition is short, and the quality of the receiver’s information (weakly) decreases with the state (the intervals become longer as the state increases). In the next section, I show that both of these properties no longer hold when the sender’s information arrives gradually over time.

4 Main Result

I now return to the model presented in Section 2. For the next result, recall the definition of the *golden ratio conjugate*,

$$\Phi = \frac{1}{\varphi},$$

where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. It is useful to note the approximate values, $\Phi \approx 0.618$ and $\Phi^2 \approx 0.382$.

Proposition 1 *Let $q(\cdot) \in Q$ be a binary interval learning process such that $q(\sigma) \geq \Phi^2$ for all $\sigma \in \Sigma$. Then, there exists an equilibrium in which every element of the receiver’s information partition is a short interval.*

The Proposition will follow as an immediate corollary of Proposition 2. To interpret the condition in Proposition 1, recall that if $q(\sigma) = \frac{1}{2}$, the signal that follows σ partitions the sender’s information set into two equal intervals. Thus, it can be thought of as a *balanced* signal that provides information of

the same quality on low and high states of the world (given σ). If, for example, $q(\sigma) < \frac{1}{2}$, then the sender’s information set after (σ, right) is a longer interval—and thus reflects a lower quality of information—than that after (σ, left) . In this case, the signal that follows σ is *more informative at the bottom*. By contrast, a signal that is *more informative at the top* is one for which the q -value is above $\frac{1}{2}$.

Proposition 1 states that whenever the sender’s learning process does not contain signals that are significantly more informative at the bottom, the receiver can obtain, in equilibrium, “high-quality information” in all states, in the sense that the length of his information set (an interval) would fall below $4b$, regardless of the specifics of the learning process. To see the effect of such gradual learning on the players’ payoffs relative to the benchmark, note that the receiver’s expected payoff from any (finite) information partition \mathcal{P} is the negative of the expected variance,

$$-\mathbb{E}[\text{var}(\theta)|\mathcal{P}] = -\sum_{J \in \mathcal{P}} \left(\int_{\theta \in J} (\theta - \mathbb{E}[\theta|\theta \in J])^2 d\theta \right).$$

In addition, in any equilibrium, from the ex-ante perspective, the sender’s and receiver’s expected payoffs are identical up to a constant. Thus, the expected variance of the receiver’s equilibrium information partition is a natural measure of the loss of information in that equilibrium.

Proposition 1 offers an easy (and tight) upper bound on the expected variance of the receiver’s information partition: since all intervals in this partition are short, the expected variance is bounded from above by the variance of the smallest long interval (an interval of length $4b$), which equals $\frac{4}{3}b^2$. Figure 1 plots the parabola $\frac{4}{3}b^2$ and the expected variance under the most informative Crawford–Sobel equilibrium (the non-smooth function)⁸ as a function of b .

⁸For details, see Crawford and Sobel (1982), Section 4.




Figure 1: Expected variance as a function of $b \in [0, \frac{1}{4}]$.

For example, when $b = \frac{1}{100}$, any learning process for which the condition in Proposition 1 holds allows for an equilibrium where all intervals in the receiver’s information partition are shorter than $\frac{1}{25}$. The upper bound on the expected variance of the receiver’s information partition in this case is, approximately, 25 times lower than the expected variance induced by the most informative Crawford–Sobel equilibrium (which partitions the unit interval into only 7 intervals).

4.1 Equilibrium of Short Intervals

An important aspect of the result in Proposition 1 is that it can be attained in an equilibrium where the sender uses a simple and intuitive strategy by which he reports truthfully and without delay all the gradually arriving information, until a short interval is reported. I now provide some definitions and then turn to a formal description of the players’ strategies.

Definitions

Let \hat{S} be the sender’s reporting policy under which, in each period t , he reports truthfully the realization of s_t if and only if his information set *prior*

to observing s_t is a long interval; otherwise, when s_t refines an already short interval, the sender reveals no information (e.g., chooses m_t at random in a manner that is uncorrelated with s_t).

Any information partition that consists of finitely many intervals will sometimes be written as $\langle \theta_1, \dots, \theta_k \rangle$, where $0 < \theta_i < \theta_{i+1} < 1$ for all $i < k$, and any consecutive pair (together with the bounds 0 and 1) represents the endpoints of an interval in the information partition. The receiver's information partition induced by \hat{S} is denoted by \mathcal{I}^R .

For each finite sequence of sender's reports h , let $I^R(h)$ denote the *receiver's information set induced by h and \hat{S}* .

Fix $b > 0$. The sequence of reports \bar{h} is called *a terminal sequence of reports* if $I^R(\bar{h})$ is a short interval and any proper prefix of \bar{h} corresponds to a long interval.⁹ Denote the set of all terminal sequences of reports by

$$\bar{H} = \{\bar{h} : |I^R(\bar{h})| < 4b \text{ and } (\forall h \sqsubseteq \bar{h}; h \neq \bar{h} \Rightarrow |I^R(h)| \geq 4b)\}.$$

For $\bar{h} \in \bar{H}$, denote by $a(\bar{h})$ the receiver's optimal action under the belief that the state is uniformly distributed on $I^R(\bar{h})$, and let $\bar{A} = \{a(\bar{h}) : \bar{h} \in \bar{H}\}$ be the set of all actions that are optimal for the receiver under beliefs induced by terminal sequences of reports.

Let $H = \{h : \exists \bar{h} \in \bar{H} \text{ s.t. } h \sqsubseteq \bar{h}\}$ be the set of all prefixes of elements of \bar{H} , and let $\hat{H} = H \setminus \bar{H}$ be the set of all proper prefixes of elements of¹⁰ \bar{H} .

Strategies

Let R^* be the receiver's strategy by which he *postpones action* at any $h \in \hat{H}$, and *chooses* $a(\bar{h})$ at any h for which $h \sqsupseteq \bar{h} \in \bar{H}$ (i.e., h contains the terminal

⁹Given a sequence $y = \{y_j\}_{j=1}^n$, $n \in \mathbb{N}$, a sequence y' is a *prefix of y* , $y' \sqsubseteq y$, if $y' = \{y_j\}_{j=1}^{n'}$ for some $n' \leq n$. If $n' < n$, then y' is a *proper prefix of y* .

¹⁰The sets \bar{H} , H , and \hat{H} depend on b ; to ease the exposition, this is not reflected explicitly in the notation.

sequence of reports \bar{h} as a prefix).

Let S^* be the sender's strategy where at histories consistent with \hat{S} the sender reports as specified by \hat{S} , and at histories inconsistent with¹¹ \hat{S} the reports are determined by backward induction given the receiver's strategy¹² R^* .

Proposition 2 *Fix $b > 0$ and consider a binary interval learning process $q(\cdot) \in \mathcal{Q}$. If $q(\sigma) \geq \Phi^2$ for all $\sigma \in \hat{H}$, then the pair of strategies (S^*, R^*) constitutes an equilibrium.*

The proof of Proposition 2 appears in Section 6. When (S^*, R^*) constitutes an equilibrium, it is referred to as an *equilibrium of short intervals*. To develop an intuition for the result, it is useful to start with an observation regarding the following single-period communication game.

Observation 2 *Fix $b > 0$, an interval J , and $x \in (\inf J, \sup J)$. Consider the information transmission game in which it is commonly known that the state is uniformly distributed on J and that the sender knows whether the state is below or above x . Truth-telling is incentive compatible if and only if J is a long interval.*

To see this, suppose that θ is uniformly distributed over $[0, z]$ and fix $x \in (0, z)$. Suppose that the sender knows that $\theta \in [0, x]$. In this case, the sender's "bliss-point" equals $\mathbb{E}[\theta | \theta \in [0, x]] + b = \frac{x}{2} + b$. Given the symmetry of the quadratic loss function, the sender will choose a report that induces the receiver to take the action nearest to $\frac{x}{2} + b$. Truth-telling induces the receiver's action $\mathbb{E}[\theta | \theta \in [0, x]] = \frac{x}{2}$, while a dishonest report that $\theta \in [x, z]$

¹¹Histories that are inconsistent with \hat{S} from the perspective of the sender are histories of the form $\langle s_1, m_1 \rangle, \dots, \langle s_T, m_T \rangle$ along which there exists $t \leq T$ such that $s_t \neq m_t$ and the sender's information set before observing s_t is a long interval.

¹²Reports at histories inconsistent with \hat{S} can be calculated backwards since attention is restricted to learning processes where after finitely many reports, the receiver considers all subsequent reports uninformative. In case of multiple optimal reports at a given history, any selection would be equally good since different selections correspond to the same sender's expected payoff in this as well as all previous information sets.

induces the action $\frac{x+z}{2}$. Thus, truth-telling is incentive compatible¹³ if and only if

$$\left(\frac{x}{2} + b\right) - \frac{x}{2} \leq \frac{x+z}{2} - \left(\frac{x}{2} + b\right),$$

which is equivalent to $z \geq 4b$.

Now consider the original interaction where information arrives gradually. Observation 2 shows that, given the receiver's strategy R^* , when the game reaches a report that partitions a long interval into two short intervals, truth-telling is incentive compatible. Interestingly, Observation 2 is completely independent of x . This suggests that a myopic sender (who considers any information arrival as definitive) is willing to reveal all of his gradually arriving information (regardless of the thresholds), until the receiver's information set becomes a short interval. The requirement in Proposition 2 that the q -values of the relevant signals be at least Φ^2 is only needed to provide inter-temporal incentives for truth-telling to the forward-looking sender. The asymmetry between signals that are more informative at the top and those more informative at the bottom is illustrated in the following example.

Example 3. Consider the learning process $q^{low}(\cdot) \equiv \frac{1}{3}$. For $b = \frac{1}{6}$, there are three terminal sequences of reports

$$\bar{H} = \{(m_1 = \text{left}), (m_1 = \text{right}, m_2 = \text{left}), (m_1 = \text{right}, m_2 = \text{right})\},$$

and the receiver's information partition is $\mathcal{I}^R = \langle \frac{1}{3}, \frac{5}{9} \rangle$. In this case, S^* cannot be part of an equilibrium: if $s_1 = \text{left}$, the sender prefers to deviate to $m_1 = \text{right}$ with the intent to report $m_2 = \text{left}$, regardless of s_2 . Consistent with the logic of Observation 2, the union of the sender's information set given $s_1 = \text{left}$, $[0, \frac{1}{3})$, and the right-adjacent interval in \mathcal{I}^R , $[\frac{1}{3}, \frac{5}{9})$, forms a short interval since $\frac{5}{9} < \frac{2}{3} = 4b$.

Now consider the learning process $q^{high}(\cdot) \equiv \frac{2}{3}$. Note that, in any period, $q^{low}(\cdot)$ and $q^{high}(\cdot)$ generate information partitions that reflect one another

¹³An upwardly biased sender never finds it optimal to lie if he learns that $\theta \in [x, z]$.

around $\frac{1}{2}$. Therefore, for $b = \frac{1}{6}$, $\mathcal{I}^R = \langle \frac{4}{9}, \frac{2}{3} \rangle$ and

$$\bar{H} = \{(m_1 = \text{left}, m_2 = \text{left}), (m_1 = \text{left}, m_2 = \text{right}), (m_1 = \text{right})\}.$$

Conditional on $s_1 = m_1 = \text{left}$, since $I^R(\text{left})$ is a long interval, truth-telling is optimal in period 2. Truth-telling in period 1 is optimal as well: since the sender's information set given $s_1 = \text{left}$ is a long interval, misreporting upwards would induce an action that is too high, even from the perspective of the upwardly biased sender (also, since $b > 0$, he would never find it beneficial to lie if $s_1 = \text{right}$). In this case, the profile of strategies (S^*, R^*) constitutes an equilibrium.

Note that the union of the sender's information set given $(s_1, s_2) = (\text{left}, \text{right})$, $[\frac{4}{9}, \frac{2}{3})$, and the right-adjacent interval in \mathcal{I}^R , $[\frac{2}{3}, 1]$, forms a short interval (as in the case of $q^{low}(\cdot)$). The difference between the two learning processes is that, by the end of period 2 (when $[\frac{4}{9}, \frac{2}{3})$ becomes an element of the sender's information partition), the sender cannot deviate to induce the right-adjacent interval in \mathcal{I}^R since the belief that $\theta \in [\frac{2}{3}, 1]$ can only be induced in period 1.

The intuition of the proof of Proposition 2 is as follows. Conditional on being truthful before period t , it would never be attractive for the sender to distort his report downwards (say $m_t = \text{left}$ while $s_t = \text{right}$), since he is upwardly biased. Therefore, suppose that $s_t = \text{left}$. The forward-looking sender compares, given his information, the expected payoff under S^* with the payoff he can obtain by deviating to $m_t = \text{right}$ and choosing an optimal post-deviation plan of reports from $t + 1$ onwards. The first part of the proof shows that an optimal post-deviation plan of reports has a very simple structure: regardless of future signal realizations, the sender will induce the receiver to take the lowest possible action. The second part of the proof shows that truth-telling is optimal in period t since misreporting would induce the receiver to take an action that is too high, even from the perspective of the upwardly biased sender. Both parts of the proof rely on the fact that signals cannot be too informative at the bottom.

Tightness

The following example illustrates how modifying a signal along the learning process affects the long-run considerations of the sender, and shows that Φ^2 is a tight lower bound on the q -values of the relevant signals under which Proposition 2 holds.

Example 4. Consider a learning process where $q(\cdot) \equiv \Phi^2$ and suppose that $4b = \Phi$. In this case, $\bar{H} = \{(\text{left}), (\text{right}, \text{left}), (\text{right}, \text{right})\}$, and¹⁴ $\mathcal{I}^R = \langle \Phi^2, \Phi \rangle$. By Proposition 2, S^* is consistent with an equilibrium. However, since $|I^R(\text{left}) \cup I^R(\text{right}, \text{left})| = 4b$, upon observing $s_1 = \text{left}$, the sender is indifferent between truth-telling and reporting $m_1 = \text{right}$, with the intent to report $m_2 = \text{left}$ in period 2. Replacing either s_1 or s_2 (which follows $s_1 = \text{right}$) with a signal whose q -value is $\Phi^2 - \varepsilon$ for a small $\varepsilon > 0$ shifts the original threshold Φ to the left, and so $I^R(\text{left}) \cup I^R(\text{right}, \text{left})$ becomes a short interval (Figure 2 illustrates the effect of modifying only s_1 where both thresholds are affected). S^* is no longer consistent with an equilibrium as the sender would deviate in period 1.




Figure 2: Receiver's information partition under S^* .

It is worth noting that the value of the bias $b = \frac{\Phi}{4}$ in the above illustration does not play any special role. For any $b \in (0, \frac{\Phi}{4})$, one can construct a similar example where, after several periods of learning and reporting truthfully, the players arrive at an information set of length $\frac{4b}{\Phi}$. Modifying the continuation of learning on that interval as before would have similar effects to those illustrated above for the interval $[0, 1]$.

¹⁴By the definition of Φ , we have $\Phi^2 + (1 - \Phi^2)\Phi^2 = \Phi$.

5 Remarks

Gradual Learning, Partial Information, and Intermediate Reports. Recall the completely balanced learning $q(\cdot) \equiv \frac{1}{2}$ from Example 1. When $b = \frac{1}{100}$, in the equilibrium of short intervals, the sender reports truthfully in the first 5 periods and the receiver's information partition \mathcal{I}^R consists of 32 intervals of equal length. If the periodic reports are replaced by only one communication opportunity at the end of period 5, it is incentive compatible for the sender to reveal all of his information. Moreover, the same equilibrium outcome can be achieved if the sender is endowed with a static information structure \mathcal{I}^R prior to communication. The value of intermediate reports is thus unclear, and the overall effect of gradualness in information arrival boils down to the observation that the amount of information that can be transmitted in equilibrium is not monotone in the quality of the sender's information (e.g., Fischer and Stocken 2001, Ivanov 2010).

These points are special to completely balanced learning. To see the dynamic aspect of learning consider $q^{high}(\cdot) \equiv \frac{2}{3}$ from Example 3 and recall that when $b = \frac{1}{6}$, $\mathcal{I}^R = \langle \frac{4}{9}, \frac{2}{3} \rangle$. If the sender's gradual learning is replaced by the static information structure $\langle \frac{4}{9}, \frac{2}{3} \rangle$, truth-telling is no longer incentive compatible: reporting falsely that $\theta \geq \frac{2}{3}$ when the information set is $[\frac{4}{9}, \frac{2}{3})$ is strictly profitable. This does not happen when learning is gradual since the sender decides to avoid the report that $\theta \geq \frac{2}{3}$ at an early stage of learning while his information set is still $[0, \frac{2}{3})$.

In the above example, the reports in both periods are terminal for at least some signal realizations. To illustrate the value of intermediate reports, suppose that the sender's information partitions in periods 1 and 2 are, respectively, $\langle \frac{1}{2} \rangle$ and $\langle \frac{1}{3}, \frac{1}{2}, \frac{3}{4} \rangle$. The equilibrium of short intervals for $b = \frac{1}{8}$ yields $\mathcal{I}^R = \langle \frac{1}{3}, \frac{1}{2}, \frac{3}{4} \rangle$. In this equilibrium, the sender reports whether the state is above or below $\frac{1}{2}$, at the end of period 1. This report is essential. Since the union of the middle elements of \mathcal{I}^R is a short interval, \mathcal{I}^R is inconsistent with equilibrium without the report in period 1. Such intermediate

reports facilitate information transmission not only because they are transmitted while the sender's incentive constraints are weaker but also because they provide a context for future reports.

Strategic Information Withholding. In the equilibrium of short intervals, the sender never withholds “transmittable” information: whenever a new piece of information arrives, he reports it immediately if and only if revelation of this information is incentive compatible in the static game in which no additional information is expected. While this equilibrium provides a general lower bound on the effect of gradualness in information arrival, it need not be optimal. Consider again $q^{high}(\cdot) \equiv \frac{2}{3}$ and suppose, for simplicity, that $b = \frac{1}{4}$. In the equilibrium of short intervals, only the *first* signal realization is revealed, so $\mathcal{I}^R = \langle \frac{2}{3} \rangle$. If no information is transmitted in period 1, it becomes incentive compatible to reveal whether the state is above or below $\frac{4}{9}$ in period 2. Ex ante, both players strictly prefer the receiver's information partition $\langle \frac{4}{9} \rangle$ over $\langle \frac{2}{3} \rangle$ since the former is “more balanced.” Whether and when information withholding is desired depends on the specifics of the learning process. Characterizing the circumstances in which complicated forms of communication outperform the simple update-as-you-go dynamics of the equilibrium of short intervals is left as an open question.

6 Proof of Proposition 2

The following additional notation will be used. For every sequence of signal realizations σ , let $I^S(\sigma)$ be the *sender's information set induced by σ* , and for every sequence of reports $h \in H$, let $\bar{A}(h) = \{a(\bar{h}) : \bar{h} \in \bar{H} \text{ and } h \sqsubseteq \bar{h}\}$ be the set of all receiver's actions that are optimal under some terminal sequence of reports that is consistent with h .

By the definition of R^* , it is immediate that R^* is a best response to S^* . It will be shown that, if the condition specified in the proposition is satisfied, S^* is a best response to R^* .

Throughout the proof, let $\sigma = (s_1, \dots, s_t) \in H$ and $h = (m_1, \dots, m_t) \in H$ denote, respectively, the sequences of signal realizations and reports such that $m_{t'} = s_{t'}$ for all $t' < t$, and $s_t \neq m_t$.

Step 1 *Suppose that the sender observed σ and reported h . Then, an optimal continuation of reports for the sender does not depend on future signal realizations.*

Misreporting downwards: Let $s_t = \text{right}$ and $m_t = \text{left}$. From the assumption that $h \in H$, it follows that, for all $a \in \bar{A}(h)$, $a < \inf(I^S(\sigma))$. Since $b > 0$, for every value of $\theta \in I^S(\sigma)$, the sender strictly prefers the receiver's maximal action among those that can be induced given h . Therefore, given σ and h , it is uniquely optimal for the sender to choose future reports that induce the receiver's action $\max \bar{A}(h)$, regardless of future signal realizations.

Misreporting upwards: Let $s_t = \text{left}$ and $m_t = \text{right}$. By supermodularity of the sender's preferences, if the sender prefers the action $\min \bar{A}(h)$ over all other actions in $\bar{A}(h)$ when $\theta = \sup(I^S(\sigma))$, then the same is true for all $\theta \in I^S(\sigma)$. To prove this for $\theta = \sup(I^S(\sigma))$, it is sufficient to show that¹⁵

$$|\bar{A}(h) \cap (\sup(I^S(\sigma)), \sup(I^S(\sigma)) + 2b)| \leq 1.$$

If $h \in \bar{H}$, then $|\bar{A}(h)| = 1$. Suppose that $h \in \hat{H}$, and let $a_1 = \min \bar{A}(h)$ and $a_2 = \min (\bar{A}(h) - \{a_1\})$. Denote by h_{left} the terminal sequence of reports that continues h and induces a_1 . Observe that h_{left} ends with $m_\tau = \text{left}$ for some $\tau > t$. Consider $h_{\text{right}} = (m_1, \dots, m_\tau)$ that differs from h_{left} only in m_τ , which equals "right" under h_{right} . Note that $|I^R(h_{\text{left}}) \cup I^R(h_{\text{right}})| \geq 4b$, and that h_{right} need not be a terminal sequence of reports.

Case 1.1: $h_{\text{right}} \in \bar{H}$. In this case, $a_2 = \mathbb{E}[I^R(h_{\text{right}})]$. Since $I^R(h_{\text{left}})$ and $I^R(h_{\text{right}})$ are disjoint, a_1 and a_2 are separated by at least $2b$.

¹⁵At state $\theta = \sup(I^S(\sigma))$, for receiver's actions above $\sup(I^S(\sigma)) + b$, the sender's payoff is monotonically decreasing in the receiver's action. Moreover, the sender's payoff from any receiver's action $a \geq \sup(I^S(\sigma)) + 2b$ is below his payoff from any action in the interval $(\sup(I^S(\sigma)), \sup(I^S(\sigma)) + 2b)$.

Case 1.2: $h_{\text{right}} \in \mathring{H}$. In this case,

$$4b \leq |I^R(h_{\text{right}})| \leq (1 - \Phi^2) \cdot |I^R(h_{\text{left}}) \cup I^R(h_{\text{right}})|,$$

where the first inequality follows from $h_{\text{right}} \in \mathring{H}$, and the second from the assumption $q(\cdot) \geq \Phi^2$. Dividing by $(1 - \Phi^2)$ we get

$$|I^R(h_{\text{left}}) \cup I^R(h_{\text{right}})| \geq \frac{4b}{1 - \Phi^2}.$$

The following lower bound on the length of $I^R(h_{\text{left}})$ can be obtained

$$|I^R(h_{\text{left}})| \geq \Phi^2 \frac{4b}{1 - \Phi^2} = \Phi \cdot 4b \geq 2b;$$

the first inequality follows from $q(\cdot) \geq \Phi^2$, the equality is equivalent to the definition of Φ , and the second inequality is obvious as $\Phi > \frac{1}{2}$. In particular, it follows that $a_2 > \sup(I^R(\sigma)) + 2b$, which completes the proof of step 1.

Step 2 Truth-telling in period t is optimal, provided that all previous reports were truthful.

Suppose that the sender observes and reports σ truthfully. Before observing additional signal realizations, the players have the same information set and rank all receiver's information partitions over this set identically.¹⁶ This observation gives a simple but useful lower bound on the sender's expected payoff from continuing with the strategy S^* , namely, his expected payoff in the event that communication terminates immediately,

$$\mathbb{E} [u^S(a = \mathbb{E}[I^S(\sigma)], \theta) | \theta \in I^S(\sigma)],$$

where $u^S(a, \theta)$ denotes the sender's payoff at state θ when the receiver's action is a . It will be shown that, if there was truth-telling in all previous periods, the sender's payoff from misreporting the last signal in σ , i.e., s_t , falls below that lower bound.

¹⁶Provided that, given any information, the receiver chooses an action that is equal to the expected state, the players' expected payoffs differ by a constant b^2 .

Misreporting downwards is not profitable: Suppose that $s_t = \text{right}$ and $m_t = \text{left}$. As shown in step 1, it is optimal for the sender to induce the receiver to take a single action $a = \max \bar{A}(h)$ (where h is as defined in the beginning of the proof). Since $\max \bar{A}(h) < \mathbb{E}[I^S(\sigma)]$ and $b > 0$,

$$\mathbb{E}[u^S(a = \mathbb{E}[I^S(\sigma)], \theta) | \theta \in I^S(\sigma)] > \mathbb{E}[u^S(a = \max \bar{A}(h), \theta) | \theta \in I^S(\sigma)].$$

Misreporting upwards is not profitable: Suppose that $s_t = \text{left}$ and $m_t = \text{right}$. First consider the case where $\sigma \in \overset{\circ}{H}$. In this case, $|I^S(\sigma)| \geq 4b$, and hence

$$\mathbb{E}[u^S(a = \mathbb{E}[I^S(\sigma)], \theta) | \theta \in I^S(\sigma)] \geq \mathbb{E}[u^S(a = \sup(I^S(\sigma)), \theta) | \theta \in I^S(\sigma)]$$

with equality if $|I^S(\sigma)| = 4b$. Since $\sup(I^S(\sigma)) \geq \mathbb{E}[I^S(\sigma)] + 2b$, any *single* action $a' > \sup(I^S(\sigma))$ yields a lower expected payoff for the sender. Thus, from the result in step 1, it follows that misreporting the realization of s_t is not profitable.

Now, suppose that $\sigma \in \bar{H}$. The interval $I^S(\sigma)$ is an element of \mathcal{I}^R . Let I^+ denote the interval right-adjacent to $I^S(\sigma)$ in this partition (i.e., $\sup(I^S(\sigma)) = \inf(I^+)$). The argument in step 1 showed that, following a dishonest report $m_t = \text{right}$, it is optimal for the sender to induce the receiver's belief that $\theta \in I^+$. To complete the proof, it is sufficient to show that $|I^S(\sigma) \cup I^+| \geq 4b$.

If the receiver's belief $\theta \in I^+$ is induced by a sequence of reports of length t , then, since S^* prescribes an informative report of s_t after the $(t-1)$ -prefix of σ , it follows that $|I^S(\sigma) \cup I^+| \geq 4b$ and the proof is complete. Assume now that the receiver's belief $\theta \in I^+$ is induced by a sequence of $\tau > t$ informative reports. This has two implications. First, since $\sup(I^S(\sigma)) = \inf(I^+)$, it must be the case that $m_\tau = \text{left}$. The assumption that $q(\cdot) \geq \Phi^2$ implies that

$$|I^+| \geq \Phi^2 \cdot 4b.$$

Second, $h \in \overset{\circ}{H}$ and hence $|I^S(h)| \geq 4b$ (recall that h is the sequence of reports of length t that differs from σ only in the last component). By an

argument similar to the one given in Case 1.2, the assumption that $q(\cdot) \geq \Phi^2$ implies that

$$|I^S(\sigma)| \geq \Phi^2 \frac{4b}{1 - \Phi^2}.$$

Since $I^S(\sigma)$ and I^+ are disjoint intervals, by the definition of Φ ,

$$|I^S(\sigma) \cup I^+| \geq \left(\frac{\Phi^2}{1 - \Phi^2} + \Phi^2 \right) \cdot 4b = 4b.$$

This completes the proof of step 2 and the proposition follows. \square

Reference

Aumann, Robert J., and Sergiu Hart. "Long cheap talk." *Econometrica* 71.6 (2003): 1619–1660.

Blume, Andreas, Oliver J. Board, and Kohei Kawamura. "Noisy talk." *Theoretical Economics* 2.4 (2007): 395–440.

Crawford, Vincent P., and Joel Sobel. "Strategic information transmission." *Econometrica* 50.6 (1982): 1431–1451.

Fischer, Paul E., and Phillip C. Stocken. "Imperfect information and credible communication." *Journal of Accounting Research* 39.1 (2001): 119–134.

Frug, Alexander. "Dynamic cheap talk with static informational control." *Economics Letters* 143 (2016): 118–120.

Frug, Alexander. "Strategic gradual learning and information transmission." *Journal of Economic Theory* 177 (2018): 594–615.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. "Mediation, arbitration and negotiation." *Journal of Economic Theory* 144 (2009): 1397–1420.

Ivanov, Maxim. "Informational control and organizational design." *Journal of Economic Theory* 145.2 (2010): 721–751.

Ivanov, Maxim. "Dynamic learning and strategic communication." *International Journal of Game Theory* 45.3 (2016): 627–653.

Ivanov, Maxim. "Dynamic information revelation in cheap talk." *The BE Journal of Theoretical Economics* 15.2 (2015): 251–275.

Krishna, Vijay, and John Morgan. "A model of expertise." *The Quarterly Journal of Economics* 116.2 (2001): 747–775.

Krishna, Vijay, and John Morgan. "The art of conversation: Eliciting information from experts through multi-stage communication." *Journal of Economic Theory* 117.2 (2004): 147–179.

Sobel, Joel. "Giving and receiving advice." *Advances in Economics and Econometrics* 1 (2013): 305–341.