Cheap Talk with Costly Outside Information

Christian Baker *

November 20, 2020

Abstract

We study a model of cheap talk in which a receiver (agent), after having received a message from a sender (principal), may endogenously acquire additional information by paying an entropy cost. Using tools from the Bayesian persuasion and rational inattention literature, we examine the structure of the receiver's learning strategy and how this influences that ability of the sender to engage in persuasive communication. While low cost information will benefit the receiver, we find that for intermediate cost levels it can benefit the receiver to be able to commit ex ante to not engage in learning.

1 Introduction

In many settings of strategic communication between an informed principal and an uniformed agent, the agent has access to information beyond what is communicated by the principal. A customer may research a product online after consulting with a salesperson, or a policy maker may commission additional research after consulting with a think tank. Starting with the work of Crawford and Sobel (1982), the literature on strategic information transmission has highlighted the limitations of communication via cheap talk messages. We build on that literature by showing how the information availability affects the ability to persuasively communicate.

We consider a model of cheap talk with two players, a sender (the principal) and a receiver (the agent). The receiver must choose an action which affects the payoffs of both players. The sender has more information than the receiver and can send a message to

*Department of Economics, University of Chicago, cbaker8128@gmail.com

influence the action of the receiver, but the sender has no commitment mechanism and is able to lie. The receiver, after observing the message from the sender, may choose to acquire additional information by paying a cost which depends on the amount of additional information obtained. We assume the sender has state-independent preferences, only caring about the action taken by the receiver, while the receiver's payoff depends on both the action taken as well as the state of the world, which is unknown to the receiver.

As an example, consider a customer at an electronics store trying to decide which of two differently priced TVs to purchase. Unsure of their relative quality and not wanting to buy the inferior one, her default decision is to leave without making any purchase. If a salesman working on commission were to advise her that the more expensive TV is of higher quality, the customer may be suspicious of his motives and not believe him. The salesman, recognizing this, may decide to instead recommend the cheaper TV, since he prefers making any sale to no sale. But then the customer would still have no reason to believe the salesman's advice was actually truthful. Hence the salesman cannot just directly tell the customer which TV is superior. If the customer's sole source of information is the salesman, the best the salesman can do in equilibrium is to provide just enough information about the relative merits of the TVs to make the customer indifferent between the two.

But now consider the case in which after listening to the salesman the customer can perform her own research, for instance by reading online reviews or consulting professional consumer reports. This could lead to a breakdown in communication between the salesman and customer. If this information were were cheap and reliable, then nothing the salesman could say would matter, as the customer would just find out which product to buy by herself. But generally the quality of this information would be imperfect, and obtaining it would require time, psychological effort, and money. This could result in the customer being worse off in the end, since she would have to pay for the information instead of receiving it from the salesman, and would still not be guaranteed to make the correct decision. Moreover, the salesman could actually benefit from the customer having access to outside information, since it would increase the probability of a sale.

Related Literature Kamenica and Gentzkow (2011) characterize the sender's benefit under communication with commitment in terms of a value function, giving the highest value the sender can obtain from the receiver's optimal behavior for given posterior beliefs.

The concave envelope of this function gives the maximal equilibrium payoff the sender can receive. We use a similar approach in the present paper to solve the receiver's problem. Essentially, the receiver is sending a message to herself, but her value function must also incorporate the information costs.

Ravid and Lipnowski (2017) build on Kamenica and Gentzkow (2011) to examine cheap talk when the sender's preferences are state-independent. They show that in the absence of commitment, the sender's maximal value is given by the *quasi*concave envelope of the sender's value function. This is the approach we use to analyze the sender's problem.

While a major focus of the cheap talk literature has been biased experts (Krishna and Morgan, 2001; Deimen and Szalay, 2018), we focus on the case of a sender who has state-independent preferences. This allows us to build on Ravid and Lipnowski (2017), which uses the belief-based approach from Bayesian persuasion to study cheap talk under state-independent sender preferences. We also draw on Matějka and McKay (2015) and Kamenica and Gentzkow (2011) in finding and characterizing the solution to the receiver's problem.

Two closely related papers are (Matyskova, 2018) and Bloedel and Segal (2018) which study Bayesian persuasion when the receiver has information costs. In Matyskova (2018) the receiver can acquire costly information after receiving a message from the sender, but unlike our present model the sender has commitment power. The main result in Matyskova (2018) is that every equilibrium outcome can be achieved without the receiver learning: any information that the receiver would learn can just be transmitted by the sender to begin with. In contrast, we find that without commitment the receiver generally will learn in equilibrium, and this can expand or shrink the set of equilibrium outcomes. Bloedel and Segal (2018) studies a model of Bayesian persuasion where the receiver is rationally inattentive to the sender. Both papers use the same information-cost framework as in Matějka and McKay (2015), which we also employ.

Also related is the literature on communication with endogenous information acquisition on the part of the sender. In both Di Pei (2015) and Argenziano et al. (2016), the sender must pay a cost to acquire information, before transmitting a message. In contrast, it is our receiver who has the option of paying a cost to acquire supplemental information after receiving a signal. Additionally, we use an entropy-based cost function.

The structure we place on the receiver's information costs draw on the rational inattention literature, using Shannon entropy (Sims, 2003). Much of this literature focuses on single-agent rational inattention problems (Van Nieuwerburgh and Veldkamp, 2009; Matějka and McKay, 2015; Caplin and Dean, 2015; Mackowiak and Wiederholt, 2009). Our model is more closely related to the literature examining rational inattention in strategic situations. For instance, Ravid (2017) considers bargaining with a rationally inattentive buyer, Lipnowski et al. (2018) considers disclosure of information by a well-intentioned principal, and Bloedel and Segal (2018) studies information disclosure when the receiver's action set is binary. A common theme is that the sender will strategically manipulate the receiver's attention. Similarly, we find that the sender will change his messaging strategy in order to manipulate what the receiver decides to learn. Also related is Gentzkow and Kamenica (2014), who model Bayesian persuasion where the sender must pay an entropy-based cost when transmitting information.

The remainder of the paper is organized as follows. Section 2 sets up the basic model. Sections 3 and 4 characterize the receiver's problem and sender's problem, primarily by way of a leading example. Section 5 concludes.

2 Model

2.1 Setup

There are two players, a sender and a receiver. At the beginning of the game an unknown payoff relevant state $\theta \in \Theta$ is realized, which the sender perfectly observes. After observing the state, the sender sends the receiver a message. The receiver observes the message and, after updating her beliefs, may choose to acquire costly information about the state via an experiment. Finally, the receiver chooses an action $a \in A$ and payoffs are realized.

We assume that both Θ and A are finite sets with at least two elements. The state θ is drawn from a full-support prior distribution $\mu_0 \in \Delta\Theta$ known to both players.[1] The payoff function for the sender is $u_s : A \to \mathbb{R}$, and the (gross) payoff to the receiver is $u_r : A \times \Theta \to \mathbb{R}$. Hence while the receiver's payoff may depend on the state, the sender's payoff depends solely on the action of the receiver. This state-independence assumption plays a key role in our analysis and results.

Instead of directly modeling the signal structure, we model information transmission and

[1] $\Delta\Theta$ is the set of all probability distributions on Θ.

acquisition using information policies, which are distributions over posterior beliefs (Kamenica and Gentzkow, 2011; Ravid and Lipnowski, 2017). While this is making a strong assumption about the richness of information transmission, this abstraction greatly aids in tractability, and allows us to focus on outcomes of the game rather than the structure of the signals.

When either the sender or the receiver chooses an information policy, the distribution of posteriors must average back to the prior. Hence

$$\mathcal{I}(\mu) := \left\{ p \in \Delta\Delta\Theta : \int_{\Delta\Theta} \nu \, \mathrm{d}p(\nu) = \mu \right\}$$

is the set of feasible information policies for a prior μ.

Throughout, we will use μ_0 to refer to the receiver's prior beliefs, μ_1 to refer to the receiver's interim beliefs after receiving a message from the sender, and μ_2 to refer to the receiver's posterior beliefs after having engaged in learning.

The timing of the game is then

1. The sender chooses an information policy $p \in \mathcal{I}(\mu_0)$.

2. Given interim beliefs μ_1 drawn from p, the receiver chooses an information policy $q_{\mu_1} \in \mathcal{I}(\mu_1)$.

3. The receiver's posterior belief μ_2 is drawn from q_{μ_1}, and the receiver chooses an action.

We can analyze this game by working backwards. We will first look at the receiver's optimal action given her posterior beliefs, then her optimal information policy given her interim beliefs, and finally determine the sender's optimal information policy.

2.2 Receiver's Problem

In the last stage of the receiver's problem, given posterior beliefs $\mu_2 \in \Delta\Theta$ the receiver chooses an action with the highest expected payoff. Define the receiver's value function

$$v_r(\mu_2) := \max_{a \in A} \mathbb{E}_{\theta \sim \mu_2}[u_r(a, \theta)]. \tag{1}$$

This is the maximum gross payoff that the receiver can achieve given posterior beliefs μ_2. Let $\alpha^* : \Delta\Theta \to \Delta A$ be the receiver's strategy which optimizes (1).

In the first stage of the receiver's problem, given interim beliefs μ_1, the receiver chooses an information policy $q_{\mu_1} \in \mathcal{I}(\mu_1)$, a distribution over posteriors. Note that the optimal information policy will be a function of interim beliefs.

Following the rational inattention literature,[2] we will assume that the cost of an information policy is proportional to the change in entropy of the receiver's beliefs,

$$c(q; \mu_1) = \kappa \left(H(\mu_1) - \mathbb{E}_{\mu_2 \sim q}[H(\mu_2)] \right),$$

where $\kappa \geq 0$ is a cost parameter and $H : \Delta\Theta \to \mathbb{R}$ is Shannon entropy. For discrete Θ, the entropy of a distribution $\mu \in \Delta\Theta$ is[3]

$$H(\mu) = -\sum_{\theta \in \Theta} \mu(\theta) \ln(\mu(\theta)).$$

Entropy captures how much uncertainty about the state θ is expected to be reduced by the information policy, in the sense of Blackwell (1953). For instance, the entropy of a discrete distribution with N equally probable events is $\ln(N)$, while the entropy of a distribution which places all weight on a single state is 0. With this assumption on the structure of information costs, the receiver will only choose to acquire more information, in the Blackwell ordering sense, if it increases her expected payoff.

Given interim beliefs μ_1, the receiver's learning problem is

$$\begin{aligned} \max_{q_{\mu_1} \in \Delta\Delta\Theta} \quad & \mathbb{E}_{\mu_2 \sim q_{\mu_1}}[v_r(\mu_2)] - c(q_{\mu_1}; \mu_1) \\ \text{s.t.} \quad & q_{\mu_1} \in \mathcal{I}(\mu_1). \end{aligned} \tag{2}$$

Matějka and McKay (2015) prove the existence of a solution to (2). We will denote an information policy which solves (2) as $q^*_{\mu_1}$. Let $Q^* : \Delta\Theta \to \Delta\Delta\Theta$ be a function which maps interim beliefs to optimal information policies, so that for all $\mu_1 \in \Delta\Theta$ we have $Q^*(\mu_1) = q^*_{\mu_1}$.

[2] e.g. Sims (2003); Yang (2017); Caplin and Dean (2013); Matějka and McKay (2015).

[3] By convention, we let $0 \ln 0 = 0$.

2.3 Sender

Given a prior μ_0 and the receiver's strategy (α^*, Q^*), the sender chooses an information policy $p \in \Delta\Delta\Theta$ to maximize his expected payoff, subject to the constraint that $p \in \mathcal{I}(\mu_0)$. Let

$$v_s(\mu_1) := \mathbb{E}_{\mu_2 \sim q^*_{\mu_1}}[u_s(\alpha^*(\mu_2)] \tag{3}$$

be the sender's expected utility when the receiver has interim beliefs μ_1.

Since the sender lacks commitment power, the sender has the additional constraint that all interim beliefs that can be induced by the information policy must yield the sender the same expected payoff.

The sender solves the following problem

$$\max_{p \in \Delta\Delta\Theta} \quad \mathbb{E}_p[v_s(\mu_2)] \tag{4}$$

$$\text{s.t.} \quad p \in \mathcal{I}(\mu_0) \tag{5}$$

$$v_s(\mu_2) = v_s(\mu_2'), \forall\ \mu_2, \mu_2' \in \text{supp}(p)$$

3 Characterization of Receiver's Problem

We now present examples demonstrating the receiver's learning strategy. The next section will look at the sender's problem.

Example 1 We will first consider the leading example from Ravid and Lipnowski (2017). Let $\Theta = \{1, 2\}$, $\mu_0 = \Pr[\theta = 2] = 1/2$, $A = \{0, 1, 2\}$, $u_s(a) = a$, and

$$u_r(a, \theta) = \begin{cases} 0, & \text{if } a = 0 \\ 1, & \text{if } a = \theta \\ -3, & \text{otherwise} \end{cases}$$

An interpretation of this example is that the receiver is a policy maker, and the sender is a political think tank. There are two possible policies, 1 and 2, one of which is good and the other bad. The think tank knows which policy is best, but the policy maker thinks either

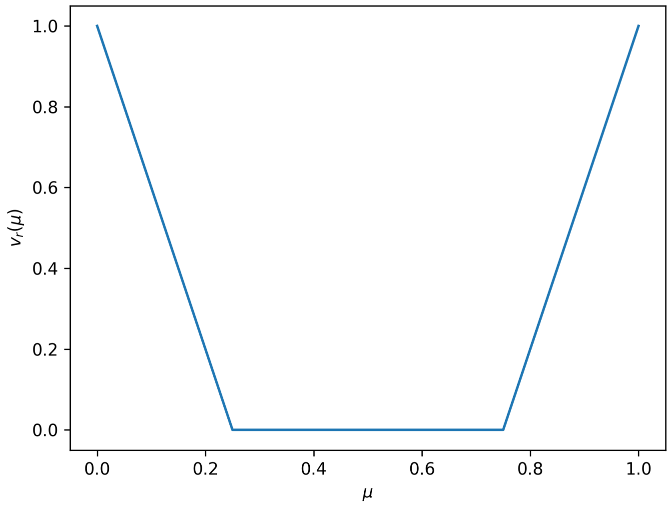

Figure 1: Receiver's gross value function

policy is equally likely to be the best. The policy maker can choose to implement policy 1, policy 2, or no policy ($a = 0$). If the policy maker chooses the wrong policy she gets a payoff of -3, the correct policy carries a payoff of 1, and maintaining the status quo a payoff of 0. The think tanks political leanings are transparent, favoring the second policy over the first, and favoring the first policy over the status quo.

If the policy maker has posterior beliefs $\mu_2 = \Pr[\theta = 2]$, then $a = 2$ is a optimal if and only if $\mu_2 \geq 3/4$, $a = 1$ is optimal if and only if $a \leq 1/4$, and $a = 0$ is optimal for $\mu_2 \in [1/4, 3/4]$. Figure 1 shows the receivers gross value function $v_r(\mu_1)$ from equation (1). For instance, if the receiver's posterior is $\mu_2 = 1/2$, then the receiver chooses $a = 0$ and receives a payoff of 0. If the receiver's posterior is $\mu_2 = 1/8$, then he chooses $a = 1$ and has an expected payoff of $-3(1/8) + 1(7/8) = 1/2$.

Now suppose that the receiver has interim beliefs μ_1, and may acquire additional information at a cost. The receiver's optimization problem (2) is now

$$\begin{aligned} \max_{q \in \Delta[0,1]} \quad & \mathbb{E}_{\mu_2 \sim q}[v_r(\mu_2)] - \kappa(H(\mu_1) - \mathbb{E}_{\mu_2 \sim q}[H(\mu_2)]) \\ \text{s.t.} \quad & \mathbb{E}_{\mu_2 \sim q}[\mu_2] = \mu_1. \end{aligned}$$

Letting $\hat{v}_r(\mu_2) = v_r(\mu_2) + \kappa H(\mu_2)$ and noting that $\kappa H(\mu_1)$ is a constant, the receiver's

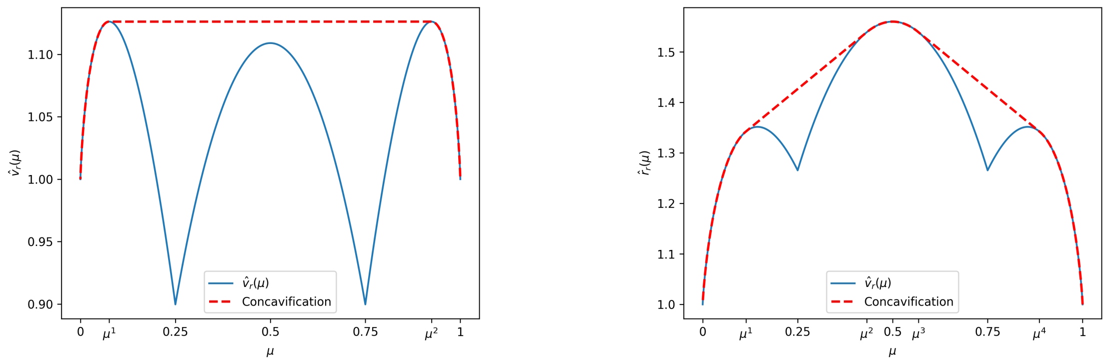

(a) $\kappa = 1.6$

(b) $\kappa = 2.25$

Figure 2: Receiver learning value functions

problem is equivalent to

$$\max_{q \in \Delta[0,1]} \quad \mathbb{E}_{\mu_2 \sim q}[\hat{v}_r(\mu_2)]$$
$$\text{s.t.} \quad \mathbb{E}_{\mu_2 \sim q}[\mu_2] = \mu_1.$$

The value function $\hat{v}_r$ captures both the payoff to the receiver from having a particular posterior belief, as well as the learning cost that must be paid in order to move to that belief. The constant $\kappa H(\mu_1)$ only depends on the receiver's interim beliefs, and while it will affect the receiver's final payoff it will not affect the optimal learning strategy. Hence the interim beliefs induced by the sender only affect the receiver's learning choice through the constraint $q \in \mathcal{I}(\mu_1)$, not by changing the shape of the value function.

The value function $\hat{v}_r$ is plotted in blue in Figure 2 for two different values of κ. Plotted in red is the concave envelope of $\hat{v}_r$, which we will call $\hat{V}_r$, the point-wise lowest concave function that majorizes $\hat{v}_r$.

The support of the optimal information policy q^* will be the posterior beliefs that support the concavification at the interim belief μ_1. So when $\hat{V}_r(\mu_1) = \hat{v}_r(\mu_1)$ the receiver does not learn. When $\hat{V}_r(\mu_1) > \hat{v}_r(\mu_1)$, the receiver will learn. In Figure 1a with $\kappa = 1.6$, the receiver learns whenever $\mu_1 \in (\mu^1, \mu^2)$, and the support of the optimal policy will be $\text{supp}(q^*) = \{\mu^1, \mu^2\}$. The weights on these two posteriors will be such that they average back to the interim belief. Similarly, in Figure 1b with $\kappa = 2.25$, the receiver will learn whenever $\mu_1 \in (\mu^1, \mu^2)$ or $\mu_1 \in (\mu^3, \mu^4)$.

Posterior beliefs μ_2 close to $1/2$ have a low entropy cost because they are not very informative, while more informative beliefs incur a higher cost. This leads to the central hump in Figure 1. But while beliefs close to $1/2$ are cheap, they are not as valuable as the more informative beliefs. This leads to the two humps on the left and right in Figure 1.

Note how the receiver's learning strategy changes based on the cost of learning. When learning is relatively cheap with $\kappa = 1.6$, as in Figure 1a, the receiver learns for all intermediate values of μ_1, and only refrains from learning when she is relatively certain. When learning does take place, the receiver chooses a learning strategy which will help her decide between taking $a = 1$ or $a = 2$. Hence for $\kappa = 1.6$, the receiver will never choose $a = 0$ regardless of her interim beliefs.

However, when learning is more costly with $\kappa = 2.25$ the receiver refrains from learning when she is relatively uncertain of the true state, when interim beliefs are between μ^2 and μ^3 in Figure 1b. For beliefs in this range, learning would be too costly. Learning only occurs when the receiver is close to being indifferent. Additionally, when learning does occur, the receiver chooses a learning strategy which will help her decide between $a = 1$ and $a = 0$ (when $\mu_1 \in (\mu^1, \mu^2)$) or $a = 2$ and $a = 0$ (when $\mu_1 \in (\mu^3, \mu^4)$).

Example 2 Next we will consider an example of the receiver's problem with three states. Let $\Theta = \{0, 1, 2\}$ and $A = \{a_1, a_2, \ldots, a_n\}$, where $a_1 = 0$, $a_n = 2$, and $a_{i+1} = a_i + 2/(n-1)$ for $i = 1 \ldots, n-1$ (i.e. n evenly spaced actions from 0 to 2). The receiver's (gross) preferences are $u_r(a, \theta) = -(a - \theta)^2$. Hence the receiver wishes to choose the action which is closest to her posterior expectation of the state.

Figure 3 shows two simplices. For instance, the lower left corner represents beliefs which place all weight on state $\theta = 0$, the upper corner represents beliefs which place all weight on state $\theta = 1$, and the lower right corner represents beliefs which place all weight on state $\theta = 2$. Beliefs which place equal weight on all three states would be in the center, and beliefs which place zero weight on $\theta = 1$ would lie along the bottom edge.

For the case of binary actions with $n = 2$, the receiver will choose a_1 whenever $\mathbb{E}_{\theta \sim \mu_2}[\theta] < 1$, and will choose a_2 whenever $\mathbb{E}_{\theta \sim \mu_2}[\theta] > 1$. When the receiver can learn, we have the same basic results as in the previous example: the receiver will want to learn only if her interim beliefs are relatively uncertain. Figure 3 depicts the range of interim beliefs for which the receiver will not learn in grey, and the range of interim beliefs for which the receiver will

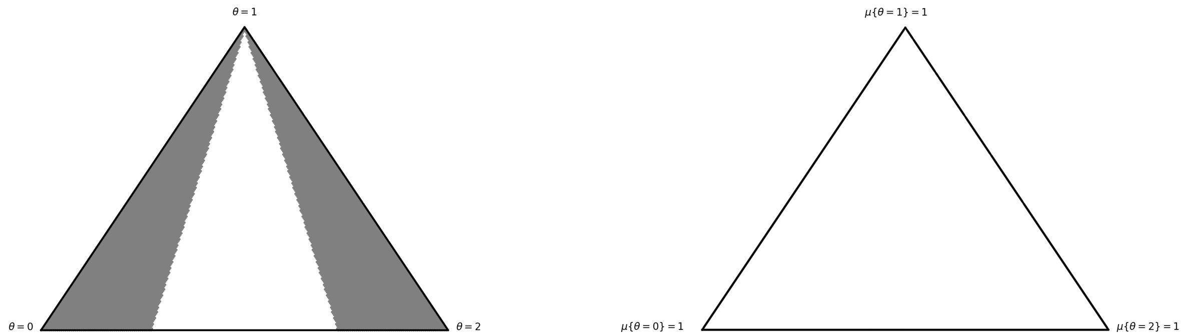

(a) $n = 2$ (b) $n = 200$

Figure 3: Simplices with non-learning (grey) and learning (white) regions.

learn in white.

The receiver will focus her learning on distinguishing between the extreme states of $\theta = 0$ and $\theta = 2$. Also note that while the receiver's action at any given posterior μ_2 depends only on her expectation of the state, The receiver's decision to learn depends on more than just the interim expectation of the state. Distinguishing between states $\theta = 0$ and $\theta = 2$ is just as costly as distinguishing between $\theta = 0$ and $\theta = 1$, but it is more valuable to the receiver to distinguish between $\theta = 0$ and $\theta = 2$. We can see this by looking at Figure 3b, where the receiver learns only when she is relatively sure that the state is not $\theta = 1$, as additional information is then more valuable.

4 Characterization of Sender's Problem

We now revisit the first example from the previous section to see how receiver learning impacts the sender's ability to transmit information.

Example 1 Recall that the sender's preferences are given by $u_s(a) = a$ and the prior is $\mu_0 = 1/2$. We will first consider the sender-optimal equilibrium if the receiver does not have the option to learn. If the sender could commit to perfectly reveal the state to the receiver, then the sender would receive an ex post payoff of 1 when the state is $\theta = 1$ and an ex post payoff of 2 when the state is 2. If the sender could reveal no information, then the sender's ex post payoff would be 0 regardless of the state.

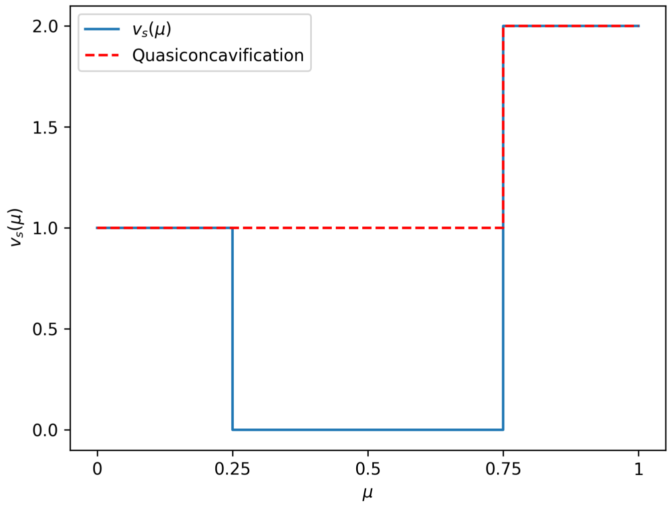

Figure 4: Sender's value function with no receiver learning.

Figure 4 shows the sender's value function $v_s(\mu_1)$ in blue. If the receiver has beliefs $\mu < 1/4$ when choosing an action, she will choose $a = 1$. If she has beliefs $\mu > 3/4$ when choosing an action, she will choose $a = 2$. Otherwise she chooses $a = 0$.

Now suppose that the sender were to send messages recommending the receiver to choose $a = 2$ using the strategy

$$\Pr\{\text{recommend } a = 1|\theta = 1\} = 2/3 \qquad \Pr\{\text{recommend } a = 1|\theta = 2\} = 0$$
$$\Pr\{\text{recommend } a = 2|\theta = 1\} = 1/3 \qquad \Pr\{\text{recommend } a = 2|\theta = 2\} = 1$$

The receiver's unique best response to the sender recommending $a = 1$ is $a = 1$. But when the sender recommends $a = 2$, the receiver will be indifferent between $a = 2$ and $a = 0$, since his posterior beliefs will be $\mu_2 = 3/4$.

In order for this to be a cheap talk equilibrium, upon receiving the recommendation to choose $a = 2$ the receiver must randomize equally between $a = 0$ and $a = 2$. Then the sender will receive an ex post payoff of 1 regardless of the message that is sent, and will have no incentive to deviate from the equilibrium. If the receiver were to instead always choose $a = 2$ when the sender recommends $a = 2$, then since the sender has no commitment power he would have an incentive to lie when the true state is $\theta = 1$. The mixed strategy of the receiver helps to provide credibility to the sender's messages.

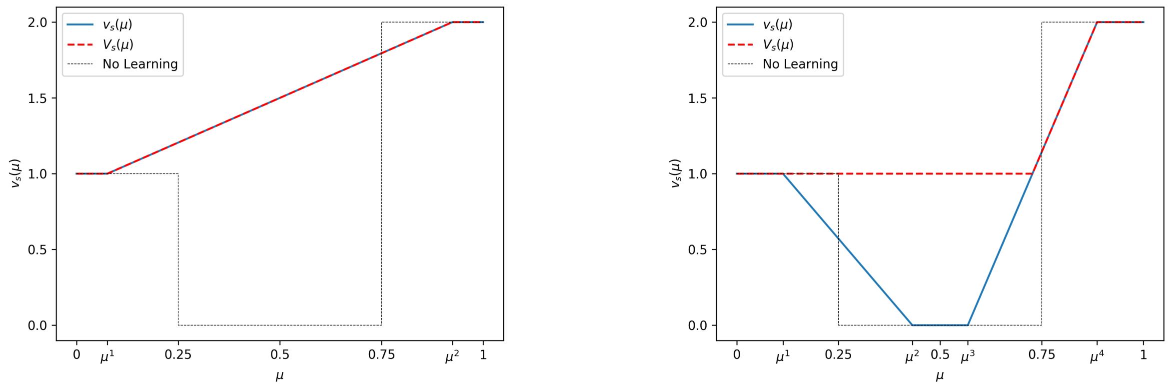

(a) $\kappa = 1.6$

(b) $\kappa = 2.25$

Figure 5: Sender value functions with receiver learning. Cutoff values correspond to Figure 2.

The dashed red line depicts the quasiconcave envelope of the sender's value function, which we will call V_s. This is the point-wise lowest quasiconcave and upper semicontinuous function that majorizes v_s. As shown in Ravid and Lipnowski (2017), the function V_s gives the sender's ex ante expected payoff in the sender-optimal equilibrium as a function of the prior μ_0. The quasiconcavity is a result of the sender's lack of commitment power: every interim belief in the support of p must yield the sender the same payoff.

Now we consider the case where $\kappa < \inf$ and the receiver can learn. Figure 5 shows the sender's value functions for the cases of $\kappa = 1.6$ and $\kappa = 2.25$. The values of μ^1, μ^2, μ^3 and μ^4 come from the cutoff values in Figure 2. As discussed in the previous section, when $\kappa = 1.6$ the receiver will always choose either $a = 1$ or $a = 2$. In Figure 5b this results in $v_s(\mu) = V_s(\mu)$ for all $\mu \in [0, 1]$, meaning that the sender can no longer persuade the receiver.

However, when learning is more costly with $\kappa = 2.25$, there is still room for persuasion. In Figure 5b we see how the sender's value function has changed. In the region (μ^1, μ^2) the receiver will be learning, with a positive probability of choosing $a = 1$ or $a = 0$. Similarly, for $\mu_1 \in (\mu^3, \mu^4)$ the receiver will learn, with a positive probability of choosing $a = 2$ or $a = 0$. This has a similar result to the receiver choosing a mixed strategy in the example when there was no learning option; the randomization between $a = 0$ and $a = 2$ due to learning can give the sender credibility, allowing for persuasive communication even in the absence of mixed strategies.

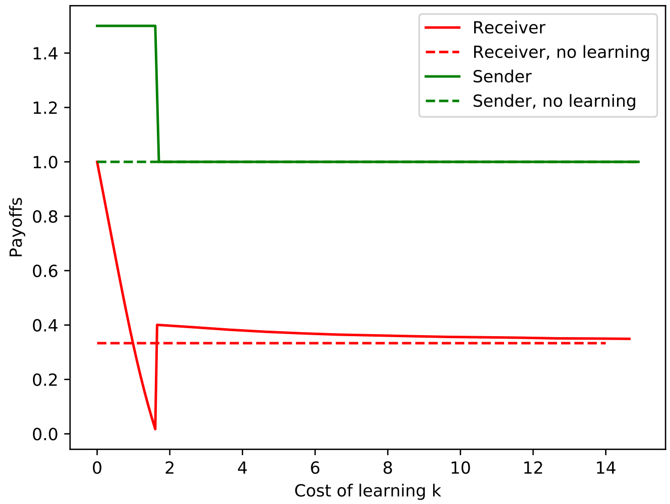

Figure 6: Sender and Receiver payoffs in sender-optimal equilibrium as a function of the cost of learning κ.

Comparative Statics Figure 6 shows the payoffs to the sender and receiver in the sender-optimal equilibrium when $\mu_0 = 1/2$. Note that there is a cutoff at $\bar{\kappa} = 1.64$. Below this level the receiver learns for all intermediate beliefs and always choose either $a = 1$ or $a = 2$, and the sender cannot effectively persuade. Above this level the receiver will not learn for intermediate values of μ_1, and when learning does there is a positive probability of choosing $a = 0$. This opens up the possibility for the sender to engage in persuasive communication.

Note the receiver's payoffs are non-monotonic in the cost of learning. The receiver is best off when $\kappa = 0$, because then she can costlessly learn the state and always choose $a = \theta$. For values of κ just below the cutoff (as in Figures 2a and 5a), the sender is unable to credibly communicate any information. The receiver will invest in some learning, and choose either $a = 1$ or $a = 2$. But the receiver can be made better off by having κ just above the cutoff, because then the receiver will have to learn less, and the sender will be able to supplement the information. This highlights the way in which higher costs of learning can actually be beneficial to the receiver.

Note that for different priors the sender could be made worse off by receiver learning. For instance, if $\kappa = 1.6$ but $\mu_0 \in (0.75, \mu^2)$, then in if the receiver could not learn she would choose $a = 2$, but with learning she has a positive probability of choosing $a = 1$.

5 Conclusion

We have analyzed strategic communication when the receiver has the option of acquiring additional information after receiving a signal from the sender. Drawing on methods developed in the literature on Bayesian persuasion and cheap talk, we have characterized the problems of the sender and receiver and analyzed them through an example. This example has highlighted how the option of the receiver to learn may either help or harm the sender and receiver. Receiver learning can help to sustain an equilibrium by providing the sender with more credibility, but it may also mean that the sender cannot communicate any information if the receiver will just go out and learn the true state anyway.

References

Argenziano, R., S. Severinov, and F. Squintani (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics 8*(3), 119–55.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, 265–272.

Bloedel, A. W. and I. R. Segal (2018). Persuasion with rational inattention.

Caplin, A. and M. Dean (2013). Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research.

Caplin, A. and M. Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review 105*(7), 2183–2203.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Deimen, I. and D. Szalay (2018). Delegated expertise, authority, and communication.

Di Pei, H. (2015). Communication with endogenous information acquisition. *Journal of Economic Theory 160*, 132–149.

Gentzkow, M. and E. Kamenica (2014). Costly persuasion. *American Economic Review 104*(5), 457–62.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017). Persuasion of a privately informed receiver. *Econometrica 85*(6), 1949–1964.

Krishna, V. and J. Morgan (2001). A model of expertise. *The Quarterly Journal of Economics 116*(2), 747–775.

Lipnowski, E., L. Mathevet, and D. Wei (2018). Attention management.

Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review 99*(3), 769–803.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–98.

Matyskova, L. (2018). Bayesian persuasion with costly information acquisition.

Ravid, D. (2017). Bargaining with rational inattention.

Ravid, D. and E. Lipnowski (2017). Cheap talk with transparent motives.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics 50*(3), 665–690.

Van Nieuwerburgh, S. and L. Veldkamp (2009). Information immobility and the home bias puzzle. *The Journal of Finance 64*(3), 1187–1215.

Yang, M. (2017). Optimality of debt under flexible information acquisition.