

In Communication Games, When Does *Ex Ante* Fact-Finding Benefit the Receiver?

Mark Whitmeyer*

March 5, 2020

Abstract

Always, if the number of states is equal to two; or if the number of receiver actions is equal to two, and

- i. The number of states is three or fewer, or
- ii. The game is cheap talk, or
- iii. There are just two available messages for the sender.

We provide a counterexample for each failure of these conditions.

Keywords: Cheap Talk, Costly Signaling, Information

JEL Classifications: C72; D82; D83

*Department of Economics, University of Texas at Austin

Email: mark.whitmeyer@utexas.edu. Thanks to Rosemary Hopcroft, Vasudha Jain, Meg Meyer, Alex Teytelboym, Thomas Wiseman, and Joseph Whitmeyer for their comments and suggestions. This paper contains material that had previously been in the job market paper, “Bayesian Elicitation”.

1 Introduction

A little learning is a dangerous thing.

Alexander Pope

An Essay on Criticism

This is a paper on the value of information. The setting is a two player sender-receiver, signaling, or communication, game. There is an unknown state of the world, about which the receiver is uninformed. The receiver is faced with a decision problem but has no direct access to information about the state. Instead, there is an informed sender, who, after learning the state, chooses a (possibly costly) action (message)¹, which the receiver observes.

We make a minimal number of assumptions. Each player has a von Neumann-Morgenstern utility function that may depend on the message chosen by the sender, the state of the world, and the action chosen by the receiver. Thus, these games include cheap talk games as in Crawford and Sobel (1982) [8], and signaling games as in Spence (1978) [27] or Cho and Kreps (1986) [7]. Throughout we assume that the number of states, messages, and actions are finite.

In this setting, we pose a simple question. Is the receiver's maximal equilibrium payoff convex in the prior? That is, restricting attention to the equilibrium that maximizes the receiver's expected payoff, does *ex ante* learning always benefit the receiver? If not, then are there conditions that guarantee this convexity?

We show that the answer to the first question is no: the receiver's maximal equilibrium payoff is not generally convex in the prior. However, there are broad conditions that guarantee convexity. If the game is **simple**—the sender's message has only instrumental value to the receiver—then the receiver's payoff is convex in the prior provided either

1. There are at most two states; or
2. The receiver has at most two actions, and

¹Throughout, in order to distinguish the sender's action from the receiver's action, the sender's action is termed a message. In some settings, like cheap talk games, this moniker is literal. In some settings, like e.g. the classic Spence scenario, in which the sender chooses a level of education, message is less fitting as a label. Hence, the reader should keep in mind that the message is simply the receiver's action.

- i. The game is cheap talk; or
- ii. There are at most two messages; or
- iii. There are at most three states.

If there are three or more messages, four or more states, and the game is not cheap talk, then even if the game is simple and the receiver has just two actions, the receiver's payoff may fail to be convex in the prior. Moreover, if there are three or more states and the receiver has three or more actions, then the receiver's payoff may fail to be convex in the prior, even if the game is simple and cheap talk with transparent motives (a cheap talk game in which each sender has identical preferences over the action chosen by the receiver). Furthermore, if the game is non-simple then the receiver's payoff may fail to be convex in the prior, even if there are just two states and two actions.

Why is the receiver's payoff convex in those scenarios described above? Why may the payoff fail to be convex otherwise? There is a crucial trade-off that belongs to *ex ante* information acquisition: there is an initial gain in information that, all else equal, benefits the receiver. However, all else may not be equal: the initial learning may result in a belief at which the receiver-optimal equilibrium may be quite bad for the receiver. Hence, the two effects may have opposite effects on the receiver's welfare, in which case the magnitude of each effect determines whether learning is beneficial.

The conditions described above guarantee that the first effect dominates—even if the resulting beliefs after learning lead to worse equilibria for the receiver, her welfare loss is guaranteed to be less than the welfare gain from the information acquisition itself. If the conditions do not hold, then the first effect may not dominate. Even though the receiver gains information initially, the resulting equilibria may be so bad that the receiver may strictly prefer not to learn.

Thanks to the ubiquity of communication games, there are numerous interpretations of *ex ante* information acquisition. In the Spence (1978) [27] setting, this paper's question becomes, "when does any test (prior to the sender's education choice) benefit the hiring firm(s)?" A seminal paper in finance is Leland and Pyle (1977) [18], who explore an entrepreneur signaling through his equity retainment decision. There, "when does any background information

or access to the entrepreneur’s history benefit a prospective investor?” In a political economy setting in which an incumbent signals through his policy choice (see e.g. Angeletos, Hellwig, and Pavan 2006, and Caselli, Cunningham, Morelli, and de Barreda 2014) [1, 6], we ask, “when does any initial news article benefit a representative member of the populace?”

More applications of *ex ante* information acquisition include reports about the state of the economy, in the case of a central bank signaling through its monetary policy (Melosi 2016) [20]; product reviews, in the case of a firm signaling through advertising (Nelson 1974, and Milgrom and Roberts 1986) [23, 22], or through its warranty offer (Gal-Or 1989) [9]; and financial reports or audits, in the case of a firm signaling through dividend provision (Bhattacharyya 1980) [3].

The remainder of Section 1 discusses related work, and Section 2 describes the formal model. Sections 3 and 4 contain the main results of the paper, Theorems 3.1 and 4.1, which provide sufficient conditions for convexity and show that the receiver’s payoff may not be convex should those conditions not hold, respectively. Section 5 concludes.

1.1 Related Work

One way to rephrase this paper’s research question is, “if information is free prior to a communication game, then does it benefit the receiver in expectation to acquire it?” Ramsey (1990) [26] asks this question in the context of a decision problem and answers in the affirmative, and this result also follows from Blackwell (1951, 1953) [4, 5] among many others.

There are a number of papers that investigate the value of information in strategic interactions (games). Neyman (1991) [24] shows that information can only help a player in a game if other players are unaware that she has it. Kamien, Tauman, and Zamir (1990) [13] explore an environment in which an outside agent, “the Maven”, possesses information relevant to an n -player game in which he is not a participant. There they look at the outcomes that the maven can induce in the game and how (and for how much) the maven should sell the information. Bassan, Gossner, Scarsini, and Zamir (2003) [2] establish necessary and sufficient conditions for the value of information to be socially positive in a class of games with incomplete information.

In two-player (simultaneous-move) Bayesian games, Lehrer, Rosenberg, and Shmaya (2013) [17] forward a notion of equivalence of information structures as those that induce the same distributions over outcomes. They characterize this equivalence for several solution concepts, including Nash equilibrium. In a companion paper, they (Lehrer, Rosenberg, and Shmaya 2010) [16] look at the same set of solution concepts in (two-player) common interest games and characterize which information structures lead to higher (maximal) equilibrium payoffs. Gossner (2000) [10] compares information structures through their ability to induce correlated equilibrium distributions, and Gossner (2010) [11] introduces a relationship between “ability” and knowledge: not only does more information imply a broader strategy set, but a converse result holds as well.

Ui and Yoshizawa (2015) [28] explore the value of information in (symmetric) linear-quadratic-Gaussian games and provide necessary and sufficient conditions for (public or private) information to increase welfare. Kloosterman (2015) [14] explores (dynamic) Markov games and provides sufficient conditions for the set of strongly symmetric subgame perfect equilibrium payoffs of a Markov game to decrease in size (for any discount factor) as the informativeness of a public signal about the next period’s game increases. Gossner and Mertens (2001) [12], Lehrer and Rosenberg (2006) [15], Pełski (2008) [25], and De Meyer, Lehrer, and Rosenberg (2010) [21] all study the value of information in zero-sum games.

In a sense, this paper explores the decision problem faced by the receiver in which the information she obtains is endogenously generated by equilibrium play by the sender. That is, the receiver’s problem is one in which *ex ante* information acquisition results in a (possibly) different information generation process at the resulting posterior belief. Outside of that there are no strategic concerns; and the sender is perfectly informed, so there is no learning on his part. Consequently, this paper is more similar in spirit to the original question asked by Ramsey, and we need not concern ourselves with the possible complexity of information structures for multiplayer games of incomplete information.

Furthermore, this paper investigates the value of information in communication games, which are by definition games of information transmission. In contrast to the broad class of games of incomplete information, in communication games the transfer of information between sender and receiver is of paramount importance. The main results of this paper

pertain to a restriction of that class of games—**simple** games—in which the sender’s message affects the receiver’s payoff only through the information that it contains.

The paper closest to this one is its companion paper, Whitmeyer (2019) [29], which investigates how a receiver can design an information structure in order to optimally elicit information from a sender in a communication game. Namely, in a two player communication game, we allow the receiver to commit *ex ante* to a signal $\pi : M \rightarrow \Delta(X)$, where X is a (compact) set of signal realizations. Instead of observing the sender’s message, the receiver observes a signal realization correlated with the message. In one of the main results of that paper, we discover that in simple two-action games, this ability guarantees that the value of information is always positive. Contrast this to the negative result that we find in this paper—that in simple two-action games the value of information is not generally positive—the other paper turns this on its head and shows that information design guarantees a positive value of information.

2 The Model

There are two players: an informed sender, S ; and a receiver, R , who share a common prior over the state of the world, $\mu_0 \in \Delta(\Theta)$, where $\mu_0(\theta) = \Pr(\Theta = \theta)$. There are two stages to the scenario—first, there is a learning stage.

Stage 1 (Learning Stage): There is some finite (or at least compact) set of signal realizations Y and a signal or Blackwell experiment, mapping $\zeta : \Theta \rightarrow \Delta(Y)$ whose realization is public. This experiment leads to a distribution over posteriors, where the posterior following signal realization y is μ_y . Call ζ the **Initial Experiment**. Each signal realization begets (via Bayes’ law) a posterior distribution, μ_y . Thus, experiment ζ leads to a distribution over posterior distributions, $P \in \Delta\Delta(\Theta)$, whose average is the prior distribution:

$$\mathbb{E}_P[\mu] \equiv \int_{\Delta(\Theta)} \mu dP(\mu) = \mu_0$$

Each posterior is the prior for the ensuing communication game. That is, following each realization of the experiment, the sender and receiver then take part in a second stage, the communication game.

Stage 2 (Communication Game): In this stage, S and R share the common prior μ_y . The sender has private information, his type (or the state of the world), $\theta \in \Theta$: he observes his type before choosing a message, m , from a set of messages M . The receiver observes m , but not θ , updates her belief about the receiver's type and message using Bayes' law, then chooses a mixture over actions, A . We assume that these sets, M , A and Θ , are finite.

Each player, S and R , has preferences over the message sent, the action taken, and the type of the sender. These are represented by the continuous utility functions u_i , $i \in \{S, R\}$: $u_i : M \times A \times \Theta \rightarrow \mathfrak{R}$.

Let us revisit the timing. First, there is an initial experiment which begets a distribution over (common) posterior beliefs, which are each respectively (common) prior beliefs in the ensuing communication game. Second, S observes his private type $\theta \in \Theta$, and chooses a message $m \in M$ to send to R . R observes m , updates his belief, and chooses action $a \in A$.

We extend the utility functions for the players to behavioral strategies. A behavioral strategy for S , $\sigma_\theta(m)$ is a probability distribution over M ; it is the probability that a type θ sender sends message m . Similarly, a behavioral strategy for R , $\rho(a|m)$ is a probability distribution over A ; it is the probability that the receiver chooses action a following message m .

We focus on receiver-optimal Perfect Bayesian Equilibrium (PBE), which we define in the standard manner. Henceforth by equilibrium or PBE, we refer to those particular equilibria, and by receiver's payoff we mean the receiver's payoff in the receiver-optimal PBE.

Throughout, we consider various sub-classes of communication games. These sub-classes are defined as follows

Definition 2.1. We say a communication game is **Simple** if the receiver has preferences over the action taken, A , and the type of the sender, Θ , but not over the message chosen by the sender, M . Equivalently, a game is simple provided the receiver's preferences are represented by the continuous utility function $u_R : A \times \Theta \rightarrow \mathfrak{R}$.

On occasion, we derive results that hold for two other classes of communication games; cheap talk, and cheap talk with transparent motives. We remind the reader of their definitions:

Definition 2.2. We say a communication game is **Cheap Talk** if the sender has preferences over the action taken by the receiver, A , and the type of the sender, Θ , but not over the message he chooses, M . Namely, for each type, each message is equally costless. Equivalently, a game is cheap talk provided the sender's preferences are represented by the continuous utility function $u_S : A \times \Theta \rightarrow \mathfrak{R}$.

A subclass of the class of cheap talk games are those with transparent motives, which term was introduced in Lipnowski and Ravid (2017) [19]:

Definition 2.3. We say a communication game is **Cheap Talk with Transparent Motives** if the game is a cheap talk game in which the sender's preferences over the action taken by the receiver, A , are independent of his type, Θ . Equivalently, a game is cheap talk with transparent motives provided the sender's preferences are represented by the continuous utility function $u_S : A \rightarrow \mathfrak{R}$.

3 When the Value of Information is Always Positive

This section is devoted to establishing the following theorem, which provides sufficient conditions for the value of information to always be positive in communication games.

Theorem 3.1. *In simple communication games, the value of information is always positive for the receiver provided*

1. *There are two states (or fewer); or*
2. *The receiver has two actions (or fewer); and*
 - i. *There are three states (or fewer); or*
 - ii. *There are two messages (or fewer); or*
 - iii. *The game is cheap talk.*

To begin, we show that if there are two states of the world (or two types of sender), the receiver's payoff is convex in the prior. Observe that if there is no initial experiment, and

the sender and receiver participate in the signaling game with common prior μ_0 , then there exists a signal or experiment $\eta : \Theta \rightarrow \Delta(M)$ that is induced by the optimal equilibrium. This experiment leads to a distribution over posteriors, where the posterior following message m is μ_m . Call this experiment the **Null-Optimal Experiment**.

Lemma 3.2. *In any simple communication game with two states and n actions, the receiver's payoff is convex in the prior.*

Proof. We sketch the proof here and leave the details to Appendix A.1. The first step is to establish Claim A.1, which allows us to restrict the number of messages in the game to two without loss of generality. As a result, there are just three cases that we need to consider: first, where the sender types pool in the receiver-optimal equilibrium at belief μ_0 ; second, where one sender type mixes and the other chooses a pure strategy (in the receiver-optimal equilibrium at belief μ_0); and third, where both sender types mix. Note that this lemma holds trivially if there exists a separating equilibrium, so we need not consider that case.

Next, following any realization of the initial experiment, y , there exists a receiver-optimal equilibrium. Equivalently, there exists a signal or experiment $\gamma_y : \Theta \rightarrow \Delta(M)$ that is induced by the optimal equilibrium. This experiment leads to a distribution over posteriors, where the posterior following message m is μ_m . Call this experiment the **y-equilibrium experiment**.


Then, we define ξ as the experiment that corresponds to the information ultimately acquired by the receiver following the initial learning and the resulting equilibrium play in the signaling game. All that remains is to show in each of the three cases that the null-optimal experiment, η , is less Blackwell informative than ξ and so the receiver prefers ξ —the receiver prefers any learning.

Note that it is possible to “prove this result without words”, which proof is depicted in Figure 1. In each case, the red point corresponds to the prior, the blue arrows and points to the initial experiment and posteriors, the yellow arrows and points to the null-optimal experiment, the green arrows to the y-equilibrium experiments, and the purple arrows and points to experiment ξ .


■

Next, we explore convexity when the receiver has only two actions. First, we establish

Case 1:



Case 2:



Case 3:




Figure 1: Lemma 3.2 Proof

that it is without loss of generality to restrict attention to equilibria in which no type mixes over messages at which the receiver strictly prefers different actions.

Lemma 3.3. *In simple games, there exists a receiver-optimal equilibrium in which no type of sender mixes over messages that induce beliefs at which the receiver strictly prefers different actions.*

Proof. The proof is left to Appendix [A.2](#) ■

Second, we discover that if there is a receiver optimal equilibrium at belief μ_0 in which at most two messages are used, then any information benefits the receiver. Formally,

Lemma 3.4. *Consider any simple communication game. If there is a receiver-optimal equilibrium at belief μ_0 in which at most two messages are used, then any initial experiment benefits the receiver.*

Proof. The full proof is left to Appendix [A.3](#). ■

Because of the costless nature of messages in cheap talk games, in conjunction with Lemma [3.3](#), it is clear that there must be a receiver-optimal equilibrium at belief μ_0 in which at most two messages are used. Accordingly, Lemma [3.4](#) implies

Corollary 3.5. *In any n state, two action, simple cheap talk game, the receiver's payoff is convex in the prior.*

From Lemma [3.2](#) we know that in two state, two action simple communication games, the value of information is always positive for the receiver. Perhaps surprisingly, the value of information is also always positive for the receiver in three state, two action simple communication games. *Viz,*

Lemma 3.6. *In simple communication games, for three states and two actions, the receiver's payoff is convex in the prior.*

Proof. Again, we leave the detailed proof to Appendix [A.4](#) but provide a sketch here. From Lemma [3.3](#), we may conclude that at most three messages will be used in the receiver-optimal

equilibrium at μ_0 . If two messages or fewer are used, then from Lemma 3.4, we have convexity. Thus, it remains to consider the case in which three messages are used. Fortunately, we show that there is just one such equilibrium that we need to consider.

Like Lemma 3.2, it is also possible to prove Lemma 3.6 without words, which proof is depicted in Figure 2. The red point corresponds to the prior, the blue arrows and points to the initial experiment and posteriors, the yellow arrows and points to the null-optimal experiment, the green arrows to the y -equilibrium experiments (or rather experiments that are payoff-equivalent to the y -equilibrium experiments), and the purple arrows and points to experiment ξ .

■

4 When the Value of Information is not Always Positive

This section tempers the optimism inspired by the Section 3. Namely, we establish Theorem 4.1, which states that if none of the sufficient conditions from Theorem 3.1 hold in some communication game, then there may be information that hurts the receiver

Theorem 4.1. *In the following communication games, ex ante information may hurt the receiver:*

1. *Simple games with four or more states, three or more messages, and two actions;*
2. *Simple games with three or more states and actions, and two or more messages;*
3. *Non-simple games with two or more states, actions, and messages.*

We begin by proving Lemma 4.2, the first result listed in Theorem 4.1.

Lemma 4.2. *In simple communication games, for four or more states, three or more messages, and two actions, the receiver's payoff is not generally convex in the prior.*

Proof. Proof is via counter-example. There are four states, $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4\}$, and a belief is a quadruple $(\mu_1, \mu_2, \mu_3, \mu_4)$, where $\mu_i := \Pr(\Theta = \theta_i)$ for all $i = 1, 2, 3, 4$ and $\mu_1 + \mu_2 + \mu_3 + \mu_4 = 1$.




Figure 3: Lemma 4.2 Game

The belief can be fully described with just three variables; hence, depicting the receiver’s payoff as a function of the belief requires four dimensions. This is (rather) difficult to do, so instead we will restrict attention to a family of experiments that involve learning on just one dimension. That is, we fix $\mu_1 = 1/3$ and $\mu_3 = 1/8$, and consider only the receiver’s payoff as a function of her (prior) belief about states θ_2 and θ_4 . Learning is on just one dimension, and so (abusing notation) we rewrite the receiver’s belief μ_2 as μ and μ_4 as $13/24 - \mu$, where $\mu \in [0, 13/24]$.

In states θ_1 and θ_2 , action a_2 is the correct action for the receiver; and in states θ_3 and θ_4 , action a_1 is correct:

Action	θ_1	θ_2	θ_3	θ_4
a_1	0	0	1	2
a_2	1	1	0	0

Likewise, the sender’s state (type)-dependent payoffs from message, action pairs are given as follows:

type	θ_1			θ_2			θ_3			θ_4		
message	m_1	m_2	m_3	m_1	m_2	m_3	m_1	m_2	m_3	m_1	m_2	m_3
a_1	1	0	0	0	1	0	0	1	3	-1	2	-2
a_2	1	0	0	0	1	0	2	4	0	5/4	0	-1

Note that types θ_1 and θ_2 have messages that are *strictly dominant* (m_1 and m_2 , respectively), and that θ_4 has a message that is *strictly dominated* (m_3).

Figure 4 depicts the receiver’s equilibrium payoff as a function of μ , V^T .² Explicitly, that function is

$$V^T(\mu) = \begin{cases} \frac{37}{24} - 2\mu & \mu \leq \frac{13}{36} \\ \frac{1}{3} + \mu & \frac{13}{36} < \mu \leq \frac{13}{24} \end{cases}$$

²The super-script T refers to “transparency.” This notation is due to the fact that in Whitmeyer (2019) [29] we explore the value of information in the case when the receiver can choose the information structure in the ensuing game. There, we contrast the value of information in that “optimal transparency” setting to the setting with full transparency, the focus of this paper.




Figure 4: Receiver Payoffs (Lemma 4.2 Proof)

and its derivation is left to Appendix A.5. The receiver's payoff is no longer convex in the belief—in fact, it is no longer lower semicontinuous. If $\mu > 13/36$ —the receiver becomes too sure that the sender is not type θ_3 or θ_4 —the only equilibria beget the pooling payoff, which correspond to no (or at least no useful) information transmission. The dotted line is a secant line that corresponds to a binary initial experiment that strictly hurts the receiver. ■

With three or more states and actions, information may harm the receiver:

Lemma 4.3. *If there are at least three states, three actions and two messages then the receiver's payoff is not generally convex in the prior.*

Proof. Proof is via counterexample. Consider the game depicted in Figure 5. There are three types θ_L , θ_M , and θ_H . Write a belief as a triple (μ_L, μ_M, μ_H) . Note that the game is cheap talk with transparent motives: each type gets utility 1 if the receiver chooses l or s , and 0 if the receiver chooses x . The receiver's preferences are given as follows:

Action	θ_L	θ_M	θ_H
l	0	1	2
s	13/24	13/24	1
x	1	0	1

Consider the following three beliefs

$$\mu_0 := \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\right), \quad \mu_1 := \left(\frac{1}{12}, \frac{1}{4}, \frac{2}{3}\right), \quad \text{and} \quad \mu_2 := \left(\frac{5}{12}, \frac{1}{4}, \frac{1}{3}\right)$$

and note that μ_0 is a convex combination of μ_1 and μ_2 , each with weight $1/2$. That is, for some prior μ_0 , μ_1 and μ_2 are the realizations of a binary initial experiment, ζ .

We depict the three prior distributions in the (x, y) -coordinate plane, where the x -axis corresponds to μ_H , and the y -axis corresponds to μ_M . There exist three convex regions of beliefs, l , s , and x , in which actions l , s , and x , respectively, are optimal. These regions and the three beliefs, μ_0 , μ_1 , and μ_2 , are illustrated in Figure 6.

After some effort (relegated to Appendix A.6), we conclude that at beliefs μ_0 and μ_1 the receiver-optimal equilibrium is one in which θ_H and θ_L choose different messages, say g and




Figure 5: Lemma 4.3 Game




Figure 6: Lemma 4.3 Action Regions

b , respectively; and θ_M mixes between those messages (g and b). The receiver's payoffs at these beliefs are $67/52$ and $83/52$, respectively.

In contrast, at belief μ_2 , such an equilibrium does not exist. Instead, the receiver-optimal equilibrium sees θ_H and θ_M choose different messages, say g and m , respectively; and θ_L mix between those messages (g and m). The receiver's payoff is $127/132$.

The posteriors corresponding to the y -equilibrium experiment for each of the three beliefs are depicted in Figure 7, where each x denotes a posterior distribution.

Finally, we can directly calculate and compare the receiver's expected payoff from this information acquisition to her payoff without obtaining the information:

$$\frac{67}{52} > \frac{2195}{1716} = \frac{1}{2} \cdot \frac{83}{52} + \frac{1}{2} \cdot \frac{127}{132}$$

whence we conclude that the receiver's optimal equilibrium payoff is not convex in the prior. Note that this game is a cheap talk game with transparent motives—even these restrictions are not enough to guarantee convexity. ■

An analog to Figure 1 is depicted in Figure 8. There, the red point corresponds to the prior μ_0 , the blue arrows and points to the initial experiment and posteriors, the yellow arrows and points to the null-optimal experiment, the green arrows to the y -equilibrium experiments, and the purple arrows and points to experiment ξ .

What goes wrong when there are three or more states and actions? Recall the two state case. In such a setting, because there are only two states, the actors' beliefs are one dimensional. As a result, any additional information can only shift the belief to the left or right on the one-dimensional simplex of beliefs. Moreover, because of the lack of diversity of sender types, the set of possible equilibrium vectors of strategies is quite small (qualitatively)—they either pool, separate, both mix, or only one mixes. Consequently, a change in the prior cannot have too great of an effect on the resulting equilibrium distribution of posteriors: as long as the change is in the correct direction (and remember, there are only two possible directions) and/or is sufficiently small, the receiver-optimal equilibrium from the original prior remains feasible, and hence the same vector of posteriors can be generated at equilibrium (albeit with different probabilities).




Figure 7: Lemma 4.3 Optimal Posteriors




Figure 8: Lemma 4.3 Induced Experiments

Furthermore, any change in the prior that eliminates the receiver-optimal equilibrium from the original prior must be large, so large that the resulting belief is more extreme than any posterior generated by the original equilibrium. Thus, the beliefs that correspond to ξ must be the same, or more extreme than the beliefs from η , and hence learning must be beneficial. Put another way, there is a trade-off to initial learning—the gain in information from the initial experiment versus the (possibly decreased) gain in information from the receiver-optimal equilibria at the new priors. With two states, the first effect dominates, and makes up for the fact that the gain in information from the equilibria may be diminished.

With more than two states, this is no longer true. As in the two state case, initial learning can result in priors for the communication game for which the receiver-optimal equilibrium under the original prior is no longer feasible. However, due to the fact that the belief space is now multi-dimensional, the beliefs generated by the receiver-optimal equilibrium at these new priors, while more extreme in *some* direction, do not correspond in general to a more informative experiment. Hence, ξ and η may not be Blackwell comparable, in which case comparisons of the receiver welfare between the no learning and learning scenarios must rely on the specific details of the initial experiment and payoffs of the game. As in the two state case, there is the same trade-off to initial learning, but now the initial gain may not dominate.

In the counterexample constructed in Lemma 4.3, the proposed initial experiment is harmful, since it involves too much learning about whether the state is θ_L . In particular, for belief μ_2 , the receiver is too confident that the state is θ_L , which precludes the existence of an equilibrium in which the receiver can distinguish between the high type and the low type. Instead, the receiver-optimal equilibrium is one in which she can distinguish between the high type and the medium type, which is much less helpful for the receiver.

As discussed in Whitmeyer (2019) [29], the value of information may not be positive in this game even when the receiver can choose the optimal information structure in the signaling game. As we discover in Whitmeyer (2019), the equilibria described above at beliefs μ_0, μ_1, μ_2 , yield the maximum payoffs to the receiver of any equilibrium under any information structure.

Finally, if there are only two states and actions, the receiver's payoff may fail to be convex




Figure 9: Lemma 4.4 Game

in the prior if the game is not simple. To wit,

Lemma 4.4. *If the communication game is not simple, then the receiver's payoff is not generally convex in the prior.*

Proof. Proof is via counter example. Consider the modified Beer-Quiche game (*cf.* Cho and Kreps (1987) [7]) depicted in Figure 9, in which the receiver now obtains an additional payoff of 1 if the sender chooses Quiche and the receiver chooses the "correct" action (i.e. F if the sender is θ_W and NF if the sender is θ_S).

Observe that if $\mu \geq 1/2$, the receiver optimal equilibrium is one in which the senders pool on Q . The receiver's best response is NF and his expected payoff is 2μ . If $\mu < 1/2$, the receiver optimal equilibrium is one in which θ_W mixes, choosing B with probability $\sigma = \mu/(1 - \mu)$, and θ_S chooses B . The receiver's expected payoff for $\mu < 1/2$ is thus $\Pr(B)/2 + 2\Pr(Q) = 2 - 3\mu$. Hence,

$$V^T(\mu) = \begin{cases} 2 - 3\mu, & \mu < \frac{1}{2} \\ 2\mu, & \mu \geq \frac{1}{2} \end{cases}$$

Figure 10 depicts V^T . The receiver's payoff is neither convex nor lower semicontinuous. The dotted line in the figure is a secant line that corresponds to a binary initial experiment that strictly hurts the receiver. ■




Figure 10: Receiver Payoffs (Lemma 4.4 Proof)

5 Conclusion

This paper comprehensive answers the question of when information always benefits the receiver in two player communication games. As Theorem 4.1 illustrates, Theorem 3.1 is as strong as possible—should none of its conditions hold, the value of information may be strictly negative.

Naturally, there is room for more work on related questions. What can we say, for instance, about the value of information for the sender? Answering such a question would pose a challenge since the proof techniques used in this paper would no longer work. Here we were able to bypass the details of the sender’s incentives and work with distributions of beliefs. This allowed us to tackle the problem as a decision problem for the receiver, in which we applied Blackwell’s theorem. Such an approach would not work when exploring sender welfare because he is not a decision maker.

References

- [1] George-Marios Angeletos, Christian Hellwig, and Alessandro Pavan. Signaling in a global game: Coordination and policy traps. *Journal of Political Economy*, 114(3):452–484, 2006.
- [2] Bruno Basset, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, Dec 2003.
- [3] Sudipto Bhattacharya. Nondissipative signaling structures and dividend policy. *The Quarterly Journal of Economics*, 95(1):1–24, 08 1980.
- [4] David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.
- [5] David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

- [6] Francesco Caselli, Tom Cunningham, Massimo Morelli, and Inés Moreno de Barreda. The incumbency effects of signalling. *Economica*, 81(323):397–418, 2014.
- [7] In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221, 1987.
- [8] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.
- [9] Esther Gal-Or. Warranties as a signal of quality. *The Canadian Journal of Economics*, 22(1):50–61, 1989.
- [10] Olivier Gossner. Comparison of information structures. *Games and Economic Behavior*, 30(1):44 – 63, 2000.
- [11] Olivier Gossner. Ability and knowledge. *Games and Economic Behavior*, 69(1):95 – 106, 2010. Special Issue In Honor of Robert Aumann.
- [12] Olivier Gossner and Mertens Jean-François. The value of information in zero-sum games. *Mimeo*, July 2001.
- [13] Morton I Kamien, Yair Tauman, and Shmuel Zamir. On the value of information in a strategic conflict. *Games and Economic Behavior*, 2(2):129 – 153, 1990.
- [14] Andrew Kloosterman. Public information in markov games. *Journal of Economic Theory*, 157:28 – 48, 2015.
- [15] Ehud Lehrer and Dinah Rosenberg. What restrictions do bayesian games impose on the value of information? *Journal of Mathematical Economics*, 42(3):343 – 357, 2006.
- [16] Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670 – 682, 2010.
- [17] Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Garbling of signals and outcome equivalence. *Games and Economic Behavior*, 81:179 – 191, 2013.

- [18] Hayne E. Leland and David H. Pyle. Informational asymmetries, financial structure, and financial intermediation. *The Journal of Finance*, 32(2):371–387, 1977.
- [19] Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Mimeo*, March 2017.
- [20] Leonardo Melosi. Signalling Effects of Monetary Policy. *The Review of Economic Studies*, 84(2):853–884, 09 2016.
- [21] Bernard De Meyer, Ehud Lehrer, and Dinah Rosenberg. Evaluating information in zero-sum games with incomplete information on both sides. *Mathematics of Operations Research*, 35(4):851–863, 2010.
- [22] Paul Milgrom and John Roberts. Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821, 1986.
- [23] Phillip Nelson. Advertising as information. *Journal of Political Economy*, 82(4):729–754, 1974.
- [24] Abraham Neyman. The positive value of information. *Games and Economic Behavior*, 3(3):350 – 355, 1991.
- [25] Marcin Peški. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62(2):732 – 735, 2008.
- [26] F. P. Ramsey. Weight or the value of knowledge. *The British Journal for the Philosophy of Science*, 41(1):1–4, 1990.
- [27] Michael Spence. Job market signaling. In Peter Diamond and Michael Rothschild, editors, *Uncertainty in Economics*, pages 281 – 306. Academic Press, 1978.
- [28] Takashi Ui and Yasunori Yoshizawa. Characterizing social value of information. *Journal of Economic Theory*, 158:507 – 535, 2015. Symposium on Information, Coordination, and Market Frictions.
- [29] Mark Whitmeyer. Bayesian elicitation. *Mimeo*, December 2018.

A Proofs

A.1 Lemma 3.2 Proof

Proof. If there exists a separating equilibrium in the game then the result is trivial; henceforth we consider only simple signaling games that do not admit separating equilibria. Since there are just two states of the world, $\Theta := \{\theta_L, \theta_H\}$, any belief (probability distribution over states) can be completely characterized by the parameter $\mu := \Pr(\Theta = \theta_H)$. The interval of beliefs can be partitioned into finitely many partitions with boundaries $0 < \mu_1 < \mu_2 < \dots < \mu_k < 1$. If an action is optimal for a receiver given some belief, μ , then it is either optimal only at that belief or it is optimal for a closed interval of beliefs $[\mu_i, \mu_j]$ of which μ is a member.

We first show that it is without loss of generality to restrict the sender to two messages.

Claim A.1. *For any equilibrium that yields the receiver a payoff of v in which $m > 2$ messages are used, there exists an equilibrium in which at most 2 messages are used that yields the receiver a payoff that is weakly higher than v .*

Proof. Suppose that $l > 2$ messages are used. Since there exist no separating equilibria, there are only two feasible l -message equilibria: i. Both types choose a mixed strategy with full support, or ii. One type (say θ_H) chooses a mixed strategy with full support, and the other type chooses a mixed strategy with support on all but one message.

In both cases, there will be l resulting equilibrium beliefs $\mu'_1 < \mu'_2 < \dots < \mu'_l$, where in case ii. $\mu'_l = 1$. However since each type is mixing, they must be indifferent over each message in the support of their mixed strategy. Hence, there must also be an equilibrium in both cases in which only two messages are used, which induce beliefs μ'_1 and μ'_l . Indeed such an equilibrium can be constructed by taking each on-path message m_i with the associated induced belief μ'_i with $i \neq 1, l$ and moving weight from each player's mixed strategy on m_i to message m_l at the ratio

$$\frac{\Delta(\sigma_H(m_i))}{\Delta(\sigma_L(m_i))} = \frac{\Delta(\sigma_H(m_l))}{\Delta(\sigma_L(m_l))} = \frac{(1 - \mu_0)\mu'_l}{(1 - \mu'_1)\mu_0}$$

Such a process decreases μ'_i and by construction maintains μ'_l . This can be done until $\mu'_i = \mu'_1$ for each i .

The Blackwell experiment that corresponds to this new, binary, distribution of posteriors is more informative than in the original situation, where l messages were used. Hence, the receiver's payoff must be weakly higher in the two-message equilibrium. ■

Thus, suppose that there are just two messages in the game. There are three cases to consider.

Case 1: For prior μ_0 , the receiver-optimal equilibrium is one in which both types pool. Observe that in this case, the null-optimal experiment is a completely uninformative experiment.

Consider any initial experiment ζ with $k \geq 2$ realizations. Following this experiment, there are k posteriors, μ_y , $y \in \{1, \dots, k\}$; or equivalently there are k priors in the resulting signaling game.

For each realization of the initial experiment, y , the receiver's equilibrium payoff in the resulting signaling game is clearly bounded below by the payoff from a pooling equilibrium, since that corresponds to the least informative (in the Blackwell sense) y -equilibrium experiment. Hence, suppose that in each of the k signaling games, there is a pooling equilibrium, and that is the receiver-optimal equilibrium.

Finally, define ξ as the experiment that corresponds to the information ultimately acquired by the receiver following the initial learning and the resulting equilibrium play in the signaling game. The null-optimal experiment, η , is less Blackwell informative than ξ and so the receiver prefers ξ —the receiver prefers any learning.

Case 2: For prior μ_0 , the receiver optimal equilibrium is one in which one type mixes and the other type chooses a pure strategy. Observe that in this case, the null-optimal experiment begets two posteriors: one that is in the interior on $[0, 1]$ and the other that is either 0 or 1.

Without loss of generality (the other cases follow analogously) suppose that θ_H mixes and chooses message m_1 with probability σ and θ_L chooses message m_1 . Following an observation of message m_2 , the receiver's belief is 1 and following message m_1 it is $\mu_j < \mu_0$. Moreover, using Bayes' law we obtain

$$\sigma = \frac{(1 - \mu_0)\mu_j}{(1 - \mu_j)\mu_0}$$

Consider any initial experiment, ζ , with $k \geq 2$ realizations. Observe that for any realiza-

tion that yields a belief $\mu_i \geq \mu_j$, an equilibrium in which θ_H mixes and θ_L does not must also exist, and hence the receiver's equilibrium payoffs for each of these beliefs must be bounded below by the payoff for that equilibrium. As in case 1, suppose that in each case that this equilibrium is optimal (and hence that the receiver's payoffs are at their lower bounds).

For each y such that $\mu_y \geq \mu_j$, the y -equilibrium experiment is one that sends the posteriors to μ_j and 1. Consequently, it is without loss of generality to suppose that ζ just has a single experiment realization that yields a belief above μ_j . Moreover, as in the first case, any realization of experiment ζ that yields a posterior $\mu_y < \mu_j$ must beget an equilibrium payoff bounded below by the pooling payoff. Hence, we suppose that for each such realization y , the optimal equilibrium is the pooling equilibrium. Moreover, the resulting payoff from this distribution over pooling payoffs itself is bounded below by the payoff were the initial experiment to have merely a single signal y that begets a belief below μ_j . Hence we suppose that is the case.

To summarize, ζ has just two signal realizations, y_1 and y_2 , corresponding to beliefs $\mu_1 < \mu_j$ and $\mu_2 > \mu_j$, respectively. γ_1 has just one signal realization, corresponding to belief μ_1 . γ_2 has two signal realizations, corresponding to beliefs μ_j and 1. Hence, ξ has three signal realizations, corresponding to beliefs μ_1, μ_j and 1. The null-optimal experiment η has two signal realizations, corresponding to beliefs μ_j and 1.

The resulting distribution over posteriors induced by η has support on 1 and μ_j . Likewise, the resulting distribution over posteriors induced by ξ has support on 1, μ_j and μ_1 . Since $\mu_1 < \mu_j$, ξ is more Blackwell informative than η and so the receiver prefers ξ —the receiver prefers learning.

Case 3: For prior μ_0 , the receiver optimal equilibrium is one in which both types mix. Observe that in this case, the null-optimal experiment begets two posteriors, both of which are in the interior of $[0, 1]$.

Let the high type choose a mixed strategy σ_H and let the low type choose a mixed strategy σ_L . The receiver will have two posteriors, $\mu_j > \mu_0 > \mu_l$ and using Bayes' law, we have

$$\sigma_H = \frac{\mu_j (\mu_0 - \mu_l)}{\mu_0 (\mu_j - \mu_l)}, \quad \text{and} \quad \sigma_L = \frac{(1 - \mu_j) (\mu_0 - \mu_l)}{(1 - \mu_0) (\mu_j - \mu_l)}$$

The remainder proceeds in the same way as in the first two cases, any experiment realization that yields a belief in the interval $[\mu_l, \mu_j]$ leads to an equilibrium payoff bounded below by the optimal equilibrium payoff at belief μ_0 , and any experiment realization that yields a belief outside that interval leads to an equilibrium payoff bounded below by the pooling payoff.

Ultimately, the null-optimal experiment, η , is less Blackwell informative than ξ (which corresponds to an information structure that serves as a *lower-bound* for the receiver's payoff from learning), so learning must always be beneficial.

We have gone through each case, and the result is shown. ■

A.2 Lemma 3.3 Proof

Proof. Let each action be strictly optimal in at least one state (or else the result is trivial). We may partition the set of types $\Theta = \Theta_1 \sqcup \Theta_2$, where Θ_1 is the set of types for whom the receiver strictly prefers to choose action a_1 , and Θ_2 is the set of types for whom the receiver strictly prefers to choose action a_2 . It is without loss of generality to suppose that there are no types for whom the receiver is indifferent between her two actions.

Equivalently, $v_i := u_R(a_1, \theta_i) > u_R(a_2, \theta_i) =: w_i$ for all $\theta_i \in \Theta_1$, and $w_i > v_i$ for all $\theta_i \in \Theta_2$.

Denote

$$\Theta_1 = \{\theta_1, \dots, \theta_t\}$$

$$\Theta_2 = \{\theta_{t+1}, \dots, \theta_n\}$$

Consider an equilibrium in which at least one type, θ_k , mixes. Without loss of generality, let $\theta_k \in \Theta_1$, i.e. he is a type for whom the receiver would strictly prefer to choose action a_1 . Next, suppose that θ_k mixes over a subset of the set of messages, M_k , where

$$M_k = \{m_1, \dots, m_l\}$$

with a generic element $m_k \in M_k$. Moreover, M_k is partitioned by three sets, M_k^0 , M_k^1 and M_k^2 , where M_k^0 is the set of messages after which the receiver is indifferent between her two actions, M_k^1 is the set of messages after which the receiver strictly prefers a_1 , and M_k^2 is the set of messages after which the receiver strictly prefers a_2 . By assumption, neither M_k^1 nor

M_k^2 is the empty set, in which case we can pick two messages, $m_1 \in M_k^1$ and $m_2 \in M_k^2$. The receiver's expected payoff at this equilibrium can be written as

$$V(\mu) = \mu(\theta_k)(\sigma_k(m_1)v_k + \sigma_k(m_2)w_k) + \gamma$$

where γ is the remainder of the receiver's payoff that—crucially for the sake of this proof—does not depend on $\sigma_k(m_1)$ or $\sigma_k(m_2)$. However, the receiver's payoff strictly increases if instead θ_k modified his mixed strategy so that $\hat{\sigma}_k(m_1) = \sigma_k(m_1) + \sigma_k(m_2)$ and $\hat{\sigma}_k(m_2) = 0$ since $v_k > w_k$ (recall that we stipulated that $\theta_k \in \Theta_1$). Moreover, it is easy to see that this is also an equilibrium: θ_k is indifferent over any pure strategy in the support of his mixed strategy; and following messages m_1 and m_2 under the new mixture, the receiver still finds it optimal to choose a_1 and a_2 , respectively (and the receiver's beliefs and payoffs following any other message are unchanged).

Since $m_1 \in M_k^1$ and $m_2 \in M_k^2$ were two arbitrary messages, and θ_k was an arbitrary type, the result follows. ■

A.3 Lemma 3.4 Proof

Proof. Again, let each action be uniquely optimal in at least one state, and let two messages be used in the receiver optimal equilibrium at belief μ_0 (if only one message is used, the receiver obtains the pooling payoff at μ_0 , and hence any initial experiment must be to her profit).

In addition, we may, without loss of generality, impose that at μ_0 there is an equilibrium such that, following each message, m_1 and m_2 , different actions, a_1 and a_2 , respectively, are strictly optimal. Otherwise, this would just yield the pooling payoff and the result would be trivial. By Lemma 3.3, this imposition ensures that each type is choosing a pure strategy.

Next, partition set Θ_1 into two sets, G_1^e and G_1^d , which correspond to the types who choose m_1 and m_2 , respectively. Likewise, partition set Θ_2 into two sets G_2^e and G_2^d which correspond to the types who choose m_2 and m_1 , respectively.

These sets are, explicitly,

$$\begin{aligned}
G_1^e &:= \{\theta_1, \dots, \theta_m\}, & G_1^d &:= \{\theta_{m+1}, \dots, \theta_t\} \\
G_2^e &:= \{\theta_{t+1}, \dots, \theta_r\}, & G_2^d &:= \{\theta_{r+1}, \dots, \theta_n\}
\end{aligned}$$

Moreover, note that it is possible that some are the empty set; although of course if G_2^d is nonempty then G_1^e cannot be empty, and similarly for G_1^d and G_2^e .

Next, since we have imposed that an equilibrium of the above form exists at μ_0 , such an equilibrium must exist at any belief μ such that the following condition holds:

Condition A.2.

$$\sum_{i=1}^m \mu(\theta_i) v_i + \sum_{j=r+1}^n \mu(\theta_j) v_j \geq \sum_{i=1}^m \mu(\theta_i) w_i + \sum_{j=r+1}^n \mu(\theta_j) w_j \quad (A1)$$

and

$$\sum_{i=t+1}^r \mu(\theta_i) w_i + \sum_{j=m+1}^t \mu(\theta_j) w_j \geq \sum_{i=t+1}^r \mu(\theta_i) v_i + \sum_{j=m+1}^t \mu(\theta_j) v_j \quad (A2)$$

Accordingly, at μ_0 , the receiver's payoff is

$$V^T(\mu_0) = \sum_{i=1}^m \mu_0(\theta_i) v_i + \sum_{j=r+1}^n \mu_0(\theta_j) v_j + \sum_{i=t+1}^r \mu_0(\theta_i) w_i + \sum_{j=m+1}^t \mu_0(\theta_j) w_j \quad (A3)$$

Without loss of generality, we may suppose that there are just three signal realizations: one, A , after which Condition A.2 holds; one, B , after which there is a pooling equilibrium for which a_1 is optimal; and one, C , after which there is a pooling equilibrium for which a_2 is optimal. We may make this assumption since the receiver is only aided by multiple signal realizations in each pooling region. Then the receiver's expected payoff from the initial signal is bounded below by

$$\begin{aligned}
& p \left[\sum_{i=1}^m \mu_A(\theta_i) v_i + \sum_{j=r+1}^n \mu_A(\theta_j) v_j + \sum_{i=t+1}^r \mu_A(\theta_i) w_i + \sum_{j=m+1}^t \mu_A(\theta_j) w_j \right] \\
& + q \left[\sum_{i=1}^m \mu_B(\theta_i) v_i + \sum_{j=r+1}^n \mu_B(\theta_j) v_j + \sum_{i=t+1}^r \mu_B(\theta_i) v_i + \sum_{j=m+1}^t \mu_B(\theta_j) v_j \right] \\
& + s \left[\sum_{i=1}^m \mu_C(\theta_i) w_i + \sum_{j=r+1}^n \mu_C(\theta_j) w_j + \sum_{i=t+1}^r \mu_C(\theta_i) w_i + \sum_{j=m+1}^t \mu_C(\theta_j) w_j \right] \quad (A4)
\end{aligned}$$

where

$$p := \Pr(A), \quad q := \Pr(B), \quad s := \Pr(C), \quad p + q + s = 1,$$

and $p\mu_A(\theta_i) + q\mu_B(\theta_i) + r\mu_C(\theta_i) = \mu_0(\theta_i)$ for all i . Expression A4 can be simplified to

$$V^T(\mu_0) + qY + s\Gamma$$

where

$$Y := \left[\sum_{i=t+1}^r \mu_B(\theta_i)(v_i - w_i) + \sum_{j=m+1}^t \mu_B(\theta_j)(v_j - w_j) \right] \quad (A5)$$

and

$$\Gamma := \left[\sum_{i=1}^m \mu_C(\theta_i)(w_i - v_i) + \sum_{j=r+1}^n \mu_C(\theta_j)(w_j - v_j) \right] \quad (A6)$$

Since a_1 is optimal in the pooling equilibrium following message B and a_2 is optimal in the pooling equilibrium following C we must have

$$\sum_{i=1}^m \mu_B(\theta_i)(v_i - w_i) + \sum_{j=r+1}^n \mu_B(\theta_j)(v_j - w_j) \geq \sum_{i=t+1}^r \mu_B(\theta_i)(w_i - v_i) + \sum_{j=m+1}^t \mu_B(\theta_j)(w_j - v_j) \quad (A7)$$

and

$$\sum_{i=1}^m \mu_C(\theta_i)(v_i - w_i) + \sum_{j=r+1}^n \mu_C(\theta_j)(v_j - w_j) \leq \sum_{i=t+1}^r \mu_C(\theta_i)(w_i - v_i) + \sum_{j=m+1}^t \mu_C(\theta_j)(w_j - v_j) \quad (A8)$$

Moreover, since Condition A.2 does not hold for belief μ_B , we must have either

$$\sum_{i=1}^m \mu_B(\theta_i)v_i + \sum_{j=r+1}^n \mu_B(\theta_j)v_j < \sum_{i=1}^m \mu_B(\theta_i)w_i + \sum_{j=r+1}^n \mu_B(\theta_j)w_j \quad (A9)$$

and/or

$$\sum_{i=t+1}^r \mu_B(\theta_i)w_i + \sum_{j=m+1}^t \mu_B(\theta_j)w_j < \sum_{i=t+1}^r \mu_B(\theta_i)v_i + \sum_{j=m+1}^t \mu_B(\theta_j)v_j \quad (A10)$$

Suppose that Inequality A9 holds. then we may substitute it into Inequality A7 and cancel:

$$\sum_{i=t+1}^r \mu_B(\theta_i)(v_i - w_i) + \sum_{j=m+1}^t \mu_B(\theta_j)(v_j - w_j) \geq 0$$

Hence, Y is positive. On the other hand, if Inequality A10 holds then we may substitute it directly into Y , which again must be positive.

A symmetric procedure works at belief μ_C to establish that Γ also must be positive. Since Γ and Y are both positive, Expression A4, the receiver's payoff from learning must be at least weakly greater than $V^T(\mu_0)$. We have exhausted every case, and so conclude that any initial experiment benefits the receiver. ■

A.4 Lemma 3.6 Proof

Proof. Denote the set of (three) states by $\Theta = \{\theta_L, \theta_M, \theta_H\}$. For convenience we use the shorthand $v_i := u_R(a_1, \theta_i)$ and $w_i := u_R(a_2, \theta_i)$, for all $i = L, M, H$. Without loss of generality, we may assume that action a_1 is strictly optimal in states θ_L and θ_H , and action a_2 is strictly optimal in state θ_M : $v_L > w_L$, $v_H > w_H$, and $w_M > v_M$.

Next, observe that we can picture any belief in the (x, y) -coordinate plane, where the x -axis corresponds to μ_M , and the y -axis corresponds to μ_H . Define region R_1 as the region in which a_1 is optimal and R_2 as the region in which a_2 is optimal. Each region, R_1 and R_2 , is compact and convex, and the two regions share a boundary that is a line segment. Define R to be the simplex of beliefs, $R = R_1 \cup R_2$.

From Lemma 3.3, for some prior μ_0 , there are just three possible arrangements of the posteriors that are induced by the null-optimal experiment:

Case 1: All posteriors lie in one region.

Case 2: All posteriors that follow messages chosen by θ_M fall in region R_2 , where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$; and all posteriors that follow messages chosen by θ_L and θ_H fall in region R_1 , where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$.

Case 3: All posteriors that follow messages chosen by θ_H (θ_L) fall in region R_1 , where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$; and all posteriors that follow messages chosen by θ_L (θ_H) and θ_M fall in region R_2 , where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$. By symmetry, we need focus only on the

case where the posteriors that follow θ_H 's messages fall in region R_1 .

Note that throughout this proof, by Lemma 3.3, each belief that is not on the line segment $R_1 \cap R_2$ must lie on the boundary of the triangle (2-simplex) of beliefs.

In the first case, the receiver clearly benefits from any initial experiment. The payoff under the prior is the pooling payoff, and so *ex ante* learning can only aid the receiver. The second case is trickier: there are two sub-cases that we need to examine:

Case 2a: The mixed strategy of each type has support on at least one message that induces a belief that is *not* on the boundary $R_1 \cap R_2$.

Case 2b: There is one type, say θ_L , that mixes only over messages that induce beliefs that are on the boundary $R_1 \cap R_2$.

In case 2a, it is easy to see that there must also be a receiver optimal equilibrium in which θ_H and θ_L each choose one message (possibly the same message) that induces a belief in $R \setminus R_2$, and θ_M chooses one message that induces a belief in $R \setminus R_1$. This is clearly an equilibrium, since each type already has support of its mixed strategy on its respective message; is optimal for the receiver, since this yields the receiver the maximum possible payoff (the separating payoff); and, moreover, does not depend on the prior. Hence, any initial experiment benefits the receiver.

In case 2b, there must exist a receiver-optimal equilibrium in which θ_H sends just one message, m_H , which induces a belief in $R \setminus R_2$; θ_L sends just one message, m_L , which induces a belief on the boundary $R_1 \cap R_2$; and θ_M mixes between m_L and m_M , the latter which induces a belief in $R \setminus R_1$. We will return to this distribution of posteriors shortly.

Case 3 also must be divided into two cases:

Case 3a: θ_L mixes only over messages that induce beliefs that are on the boundary $R_1 \cap R_2$.

Case 3b: θ_L mixes over at least one message that induces a belief in $R \setminus R_1$.

Case 3a is identical to case 2b. In case 3b, there must exist a receiver-optimal equilibrium in which θ_H sends just one message, m_H , that induces a belief in $R \setminus R_2$; and θ_L and θ_M pool on one message m_p , that induces a belief in $R \setminus R_1$. Here, only two messages are used and so by Lemma 3.4 any initial experiment benefits the receiver.

Consequently, it remains to consider the scenario that case 2b reduces to: for prior μ_0 just three messages are used as follows: θ_H separates and chooses message m_H , θ_L chooses message m_L and θ_M mixes between two messages, m_L and m_M , in such a way that the receiver is indifferent over her actions following m_L (note, that there is an equivalent scenario that is obtained by interchanging θ_H and θ_L).

The prior μ_0 must be such that

$$\mu_L^0 \leq \frac{w_M - v_M}{v_L - w_L} \mu_M^0$$

and the receiver's payoff is

$$v := V(\mu_0) = \mu_H^0 v_H + \mu_M^0 w_M + \mu_L^0 w_L$$

Call this equilibrium S^\dagger . It is clear that without loss of generality we may focus on an initial experiment that is binary, and which yields just two beliefs μ_1 and μ_2 , where μ_1 is a belief such S^\dagger is feasible, and μ_2 is a belief such that S^\dagger is infeasible. To see that this is without loss of generality, note that if there are multiple initial experiment realizations after which S^\dagger is feasible, the receiver achieves at least the payoff as in the case when there is just one such initial experiment realization. Likewise, if there are multiple initial experiment realizations after which S^\dagger is infeasible, since we need only assume the pooling payoff in this case, it is again clear that the receiver achieves at least the payoff as in the case where there is just one such initial experiment realization.

Thus, the initial experiment, ζ , yields $\mu_1 = (\mu_L^1, \mu_M^1, \mu_H^1)$ with probability p and $\mu_2 = (\mu_L^2, \mu_M^2, \mu_H^2)$ with probability $(1 - p)$, where $p\mu_1 + (1 - p)\mu_2 = \mu_0$. For belief μ_2 , the receiver's payoff is bounded below by the pooling equilibrium payoff, and so we assume that that is indeed the payoff. Note that since μ_2 is not a belief for which S^\dagger is feasible, we must have

$$\mu_L^2 > \frac{w_M - v_M}{v_L - w_L} \mu_M^2 \tag{A11}$$

Claim A.3. *For belief μ_2 , action a_1 is optimal.*

Proof. Suppose for the sake of contradiction that a_1 is not optimal. That is

$$\mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 w_H > \mu_L^2 v_L + \mu_M^2 v_M + \mu_H^2 v_H$$

But then

$$\begin{aligned}\mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 w_H &> \mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 v_H \\ \mu_H^2 w_H &> \mu_H^2 v_H\end{aligned}$$

where the second inequality follows from Inequality A11. This is a contradiction. ■

Thus, a_1 is optimal and so the receiver's expected payoff is

$$V = p [\mu_H^1 v_H + \mu_M^1 w_M + \mu_L^1 w_L] + (1 - p) [\mu_L^2 v_L + \mu_M^2 v_M + \mu_H^2 v_H]$$

which reduces to

$$V = v + (1 - p) [\mu_L^2 (v_L - w_L) - \mu_M^2 (w_M - v_M)]$$

which is greater than v by Inequality A11. ■

A.5 Lemma 4.2 Proof and Payoff Function Derivation

Proof. We derive the receiver's payoff as a function of the belief μ through a pair of claims. First,

Claim A.4. *For any belief $\mu > 13/36$, there exists no equilibrium in which a message is played that induces a belief such that action a_1 is strictly optimal.*

Proof. We can exhaustively proceed through each message:

1. Suppose a_1 is strictly optimal following m_1 . Then, θ_4 must have support of his mixed strategy on m_1 . That gives him a payoff of -1 , so he can deviate profitably to m_2 .
2. Suppose a_1 is strictly optimal following m_2 . Both θ_3 and θ_4 must have support of their mixed strategies on m_2 . Moreover, so much of their support must be on m_2 that the receiver must choose a_2 following m_1 (which will always be chosen by θ_1). Hence, θ_3 can deviate profitably to m_1 .

3. Suppose a_1 is strictly optimal following m_3 . Consequently θ_3 must have some support of his mixed strategy on m_3 (since θ_4 will never choose m_3). Moreover, θ_3 cannot be choosing a pure strategy since otherwise he would have a profitable deviation to m_2 . Hence, θ_3 must be mixing over m_3 and m_2 (since 3 is strictly larger than either of θ_3 's payoffs for m_1). The receiver must also mix following m_2 , so as to leave θ_3 willing to mix. In particular, the receiver must choose a_2 with probability $2/3$ following m_2 and a_1 with probability $1/3$. But note that θ_4 must also have support on m_2 , which message would thus yield it a payoff of $2/3$, which is less than $5/4$, the payoff he would get from deviating profitably to m_1 .

■

As a result $V^T = \frac{1}{3} + \mu$ for all $\mu > 13/36$. Second,

Claim A.5. *For any belief $\mu \leq 13/36$, the receiver optimal equilibrium begets a payoff of $37/24 - 2\mu$.*

Proof. First, an equilibrium that begets such a payoff exists. Type θ_1 chooses m_1 , types θ_2 and θ_4 choose m_2 , and type θ_3 chooses m_3 . Upon observing m_1 , the receiver chooses a_2 , and upon observing either m_2 or m_3 the receiver chooses a_1 .

It is immediately evident that neither θ_1 nor θ_2 have profitable deviations, since they are choosing strictly dominant strategies. On path, θ_4 obtains 2, whereas his payoff from deviating would be less than 2. Finally, θ_3 obtains 3 following m_3 and less than 3 following any other message. The receiver's equilibrium payoff is

$$\frac{1}{3} + \frac{1}{8} + 2 \left(\frac{13}{24} - \mu \right) = \frac{37}{24} - 2\mu$$

Second, it is easy to verify that this equilibrium is optimal for the receiver: θ_4 is unwilling to choose message m_3 and θ_1 and θ_2 always choose m_1 and m_2 , respectively. Thus, the receiver will always get her decision "wrong" with respect to some mixture of θ_1 , θ_2 or θ_4 . For $\mu \leq 1/3$, she prefers to get her decision wrong with respect to θ_2 , so our proposed equilibrium is obviously optimal. For $\mu \in (1/3, 13/36]$, she prefers to get her decision wrong with respect to

θ_1 but then θ_4 would always have a profitable deviation to m_2 . Thus, she gets her decision wrong with respect to θ_2 .

Alternatively, as noted in Whitmeyer (2019) [29], this is the receiver-optimal payoff under any information structure (or any degree of transparency). Therefore, since the receiver can obtain this payoff with full transparency, the corresponding equilibrium must be optimal. ■

■

A.6 Lemma 4.3 Proof

Proof. It is easy to see that in this game there are no separating equilibrium—type θ_l can deviate profitably by mimicking either of the other types. For each of the three possible priors (μ_0 , μ_1 , and μ_2) any pooling equilibrium results in the receiver choosing l . There exist pooling equilibria for each prior, and any off-path belief sustains such equilibria, since the senders receive their maximal payoff on path.

Compiling the remaining equilibria appears daunting (or at least unpleasant), but we can thankfully bypass this, since we are interested in the receiver optimal equilibria. Instead of finding each equilibrium, we take a belief-based approach and construct and solve the appropriate maximization problem for the receiver.

Observe that the receiver has at most three posterior distributions at equilibrium, which corresponds to the case in which all three messages are used on path, which messages beget different posterior distributions. Note that any distribution or belief in the context of this example corresponds to a point in Figure 6. It is a standard result that beliefs are a *martingale*—hence, an equilibrium pair of posterior beliefs is feasible only if there exists a line segment between the two beliefs that intersects the prior (for it to be an *equilibrium* pair of beliefs, each posterior must be generated by an equilibrium vector of strategies). Likewise, a triplet of posteriors is feasible only if the prior lies in the convex hull of the three points.

From this, we see immediately that for each (prior) belief, μ_0 , μ_1 , and μ_2 , it is impossible for none of the equilibrium posterior beliefs to lie in l . Even more, at least one of the posterior beliefs must lie in the *interior* of l . Hence, there are four cases: i. All of the posteriors lie in l , ii. Some of the posteriors lie in s and some in l , iii. Some of the posteriors lie in x and some

in l , or iv. There is one posterior in each of l , s and x .

However, cases iii. and iv. are impossible: each type strictly prefers l or s to x and thus some type must have a profitable deviation to a message that induces a posterior in l . Moreover, if all of the posteriors lie in l , then this yields the same payoff to the receiver as the case in which each type pools and so we may ignore case i.

Consequently, it remains to consider case ii. In addition, note that without loss we may focus on the situation in which there are just two posteriors since if the optimum consisted of three posteriors, two of them must lie in the same region, and we could just take their average, which would also lie in the same region and yield the receiver the same (total) payoff.

As a result, for each (prior) belief Q , $Q \in \{\mu_0, \mu_1, \mu_2\}$, the receiver solves

$$\max_{\lambda, \mu_1, \mu_2} \{ \lambda V(\mu_1) + (1 - \lambda)V(\mu_2) \}$$

subject to

$$\lambda\mu_1 + (1 - \lambda)\mu_2 = Q$$

and $\mu_1 \in B$, $\mu_2 \in A$, and $\lambda \in [0, 1]$. Substituting in the payoffs, and with the aid of Figure 6, we have

$$\max_{\lambda, \mu_H^1, \mu_M^1, \mu_H^2, \mu_M^2} \left\{ \lambda \left(\mu_H^1 + \frac{13}{24}(1 - \mu_H^1) \right) + (1 - \lambda)(2\mu_H^2 + \mu_M^2) \right\}$$

subject to

$$\lambda\mu_H^1 + (1 - \lambda)\mu_H^2 = \mu_H^0, \quad \lambda\mu_M^1 + (1 - \lambda)\mu_M^2 = \mu_M^0,$$

$$1 \geq \lambda \geq 0, \quad \mu_H^1 \geq 0,$$

$$\frac{13 - 37\mu_H^1}{24} \geq \mu_M^1 \geq \frac{11}{24}(1 - \mu_H^1),$$

$$\mu_M^2 \geq 0, \quad \mu_H^2 + \mu_M^2 \leq 1$$

For belief μ_0 , $\lambda^* = 6/13$, and the optimal pair of equilibrium posteriors is

$$\left(\frac{13}{24}, \frac{11}{24}, 0 \right), \quad \text{and} \quad \left(0, \frac{1}{14}, \frac{13}{14} \right) \quad (\text{A12})$$

which yields the receiver a payoff of $67/52$. This corresponds to an equilibrium in which θ_H and θ_L choose different messages, say g and b , respectively; and θ_M mixes between those messages (g and b).

For belief μ_1 , $\lambda_1^* = 2/13$, and the optimal pair of equilibrium posteriors is

$$\left(\frac{13}{24}, \frac{11}{24}, 0\right), \quad \text{and} \quad \left(0, \frac{7}{33}, \frac{26}{33}\right) \quad (\text{A12})$$

which yields the receiver a payoff of $83/52$. This corresponds to the same type of equilibrium as for μ_0 (the receiver can distinguish between θ_H and θ_L).

For belief μ_2 , $\lambda_2^* = 6/11$, and the optimal pair of equilibrium posteriors is

$$\left(\frac{13}{24}, \frac{11}{24}, 0\right), \quad \text{and} \quad \left(\frac{4}{15}, 0, \frac{11}{15}\right) \quad (\text{A13})$$

which yields the receiver a payoff of $127/132$. This corresponds to an equilibrium in which θ_H and θ_M choose different messages, say g and m , respectively; and θ_L mixes between those messages (g and m).

The pooling equilibrium payoffs for each prior are $5/4$, $19/12$, and $11/12$, for μ_0 , μ_1 , and μ_2 , respectively. Hence, the pooling equilibrium is not optimal for any of the priors, and so the equilibria that correspond to the posterior pairs given in Expressions [A12](#), [A12](#), and [A13](#) are optimal for their respective priors.

It remains to verify

$$\frac{67}{52} > \frac{2195}{1716} = \frac{1}{2} \cdot \frac{83}{52} + \frac{1}{2} \cdot \frac{127}{132}$$

and thus the receiver's optimal equilibrium payoff is not convex in the prior. Note that this game is a cheap talk game with transparent motives—even these restrictions are not enough to guarantee convexity. ■