# Communicating Preferences to Get Better Recommendations

Amir Habibi*

January 24, 2022

**Abstract**

I study a cheap talk model between two players: a buyer/decision maker and a seller/expert. There is two-sided private information and I allow for sequential, two-way communication. The buyer faces a decision about purchasing one of a number of goods, and the buyer's valuation depends on: her own private preferences, and the seller's private information about quality. In the first stage of communication, the buyer can communicate about her private preferences to the seller. In the second stage of communication, the seller can communicate about the quality of the goods to the buyer. When the good has multiple attributes, I show that the buyer can strictly benefit from communicating about her preferences, whereas when the good only has a single attribute this is no longer the case.

Keywords: cheap talk, communication, product recommendations
JEL Classification: D82, L15

## 1 Introduction

*Economic models of communication have little to say about real conversations – dynamic exchanges in which people take turns.* ∼ Joel Sobel.[1]

Should a consumer reveal anything about her preferences to a salesperson before getting a recommendation? Consider the following situation. A consumer is considering buying a new phone and faces a choice between the new model of phone for two different brands. If she does not purchase either of the phones, she can continue to use her

*Humboldt University of Berlin. Email: amir.habibi@hu-berlin.de. I am grateful to Valeria Burdea, Antonio Cabrales, Martin Cripps, Nathan Hancart, Sebastian Schweighofer-Kodritsch, Roland Strausz and especially Dilip Ravindran for helpful comments. I am also grateful for financial support from the German Research Foundation (DFG) through CRC TRR 190 (Project number 280092119).

[1]Sobel (2013).

current phone which she has a (private) value for. A salesperson is incentivised to make a sale—she gets a fixed payment if the consumer buys either of the two phones and no payment if the consumer does not make a purchase. The salesperson privately knows the quality of each of the new model phones. The consumer's valuation for each phone is a combination of the quality and her (privately known) preferences. The salesperson can make a recommendation to influence the consumer's beliefs about the qualities. For example, he could say that brand X's new phone is better than brand Y's. However, he is not able to provide hard evidence—so communication is by cheap talk.

Before getting a recommendation, should the consumer communicate—again by cheap talk—her preferences to the salesperson? Revealing preferences may hurt the consumer because the salesperson is no longer able to make credible recommendations once he knows the consumer's preferences. For example, if the consumer reveals that she is only interested in one of the two brands on offer, the salesperson cannot credibly communicate any information about the quality of that particular phone. However, communicating preferences may also be beneficial since it allows the salesperson to make a recommendation that is more useful for the consumer. For example, if the unknown quality of the phone is made up of two attributes—battery life and camera quality—if the consumer communicates which of these attributes she is most interested in, the salesperson can make a recommendation across goods for the attribute that is most relevant.

I analyse a stylised model of the interaction described above. The main results formalise the intuition given. When the consumer has private information about her relative preference over the two goods on offer, she does not communicate these in equilibrium. On the other hand, when the consumer has private information about her relative preference over two different attributes, she communicates these in equilibrium and this is strictly beneficial for both parties compared to a setting in which the buyer could not communicate. I make use of techniques in Lipnowski and Ravid (2020) who study an abstract cheap talk game where the sender (seller) has state-independent preferences (as in my setting). In particular, their results allow me to find the seller's maximum payoff given a belief he holds about the buyer communicates before the seller. This intermediate step is necessary to solve my model in which the buyer communicates about her preferences before the seller communicates.[2]

The contribution of the paper is twofold. First, as suggested by the opening quotation, there has been little focus on cheap talk models with two-way communication. I analyse a model with two sided (independently drawn) private information and sequential, two-way communication. To the best of my knowledge, this has not been studied thus far.

[2]Their paper also analysed the buyer-seller set-up I do as an example to illustrate their results. They only consider one way communication (from seller to buyer) and they characterise the (seller's best) equilibrium for any symmetric distribution over good quality, for any number of goods, but in the specific case where the buyer (ex ante) values all goods equally. The buyer-seller set-up with cheap talk recommendations was originally proposed in Chakraborty and Harbaugh (2010).

The majority of the cheap talk literature focuses on a single round of cheap talk from an informed sender to a receiver who takes a payoff relevant action (as in the seminal model of Crawford and Sobel (1982)). However, there is also a significant literature with multiple rounds of cheap talk but these primarily focus on one-way communication or simultaneous two-way communication. I discuss the theoretical literature that is most closely related to my model in Section 4.

Second, my theoretical model provides insights into consumer privacy contributing to the current debate about regulating firms' access to consumer data. In a survey, Acquisti et al. (2016) point out that privacy of consumer data can both benefit and harm consumers. The literature is primarily focuses on settings with price discrimination. In my model I consider a different setting: sellers communicate with buyers through cheap talk—or in effect, through recommendations based on soft information. My model highlights when a consumer benefits from revealing her preferences to the seller. Furthermore, the model also provides insights into the marketing strategy of a firm: when should it try and gather consumer preferences, and when it would be damaging to do so because it's recommendations become less credible. In a related paper, Hoffmann et al. (2020) consider a model where firms provide information to consumers through hard information and analyse the implication of allowing firms to access consumer data before choosing what information to disclose. As in my model with multiple attributes, they find that consumers are hurt by not being able to disclose their private information.[3]

The remainder of the paper is organized as follows. Section 2 introduces the model. Section 3 provides the analysis. Section 4 discusses the related literature. Section 5 provides some discussion and concludes.

# 2 Model set-up

***Players.*** There is a buyer (she) and seller (he).[4]

***Information.*** There are two goods, each with two attributes. The quality of attribute $j$ of good $i$ is given by $\theta_{ij} \in \{0, 1\}$. Quality is negatively correlated across goods for each attribute: so $\theta_{1j} = 1 - \theta_{2j}$ with $\Pr[\theta_{1j} = 1] = \frac{1}{2}$ for each attribute $j$. Quality is drawn independently across attributes—this means there are four 'states of the world'. The buyer also has a preference parameter given by $\beta = (\beta_a, \beta_g) \in [0, 1]^2$ drawn from some distribution $F$. $\beta_a$ and $\beta_g$ are drawn independently and the respective marginal

[3]Another reason consumers may want to reveal preferences is to save on search costs (Varian (2002)). A very simple example is that if a consumer wants to buy a new phone and enters an electronics store, she benefits from revealing that she wants to buy a phone, meaning that she is directed to that section of the store.

[4]In most of the cheap talk literature, the seller would be the 'sender' or the 'expert' and the buyer would be 'receiver' or the 'decision maker'.

distributions are given by $F_a$ and $F_g$. The players share a common prior over $\theta$ and $\beta$, denote the buyer's prior over $\theta$ by $\mu_0$.

***Actions and timing.*** The timing of the game is as follows:

1. The buyer privately learns the realisation of $\beta$, and the seller privately learns the realisation of $\theta$.
2. The buyer sends a message $m^b \in \mathcal{M}^b$ to the seller;
3. The seller sends a message $m^s \in \mathcal{M}^s$ to the buyer;
4. The buyer learns the value of her outside option $u \sim U[0, 1]$;
5. The buyer takes an action, $a \in \{a_0, a_1, a_2\}$: her outside option ($a_0$) or one of the two goods ($a_1$) and ($a_2$);
6. The players get their payoffs and the game ends.

***Payoffs.*** The buyer's payoff depends on her preference parameter $\beta$, the quality of the attributes of the goods $\theta$, and her outside option $u$:

$$U = \begin{cases} \beta_g \; (\beta_a \theta_{11} + (1 - \beta_a)\theta_{12}) & \text{if } a = a_1, \\ (1 - \beta_g)\,(\beta_a \theta_{21} + (1 - \beta_a)\theta_{22}) & \text{if } a = a_2, \\ u & \text{if } a = a_0. \end{cases}$$

Here $\beta_a$ represents the relative preference *across attributes*, and $\beta_g$ represents the preference *across goods*. The seller's payoff is *state independent*—it simply depends on whether or not the buyer buys one of the goods:

$$V = \begin{cases} 1 & \text{if } a = a_1, \\ 1 & \text{if } a = a_2, \\ 0 & \text{if } a = a_0. \end{cases}$$

***Strategies.*** The buyer's strategy is to choose 1) a messaging strategy that maps her preference to a message $m^b : [0, 1]^2 \to \mathcal{M}^b$ and 2) an action strategy that maps her preferences, her message, and the message of the seller to a choice over goods: $a : [0, 1]^2 \times \mathcal{M}^b \times \mathcal{M}^s \to \{a_1, a_2, a_0\}$.[5] The seller's strategy is to choose a messaging strategy that maps the state $\theta$ and the buyer's message to a message: $m^b : \{0, 1\}^2 \times \mathcal{M}^b \to \mathcal{M}^s$.

[5] I restrict attention to pure strategies to easy notation, but this is restriction does not affect the analysis in any substantive way.

I refer to the seller's strategy as an (information) policy.

***Beliefs.*** The seller updates his belief over $\beta$ to $\hat{F}(m^b) \in \Delta[0,1]^2$ following the message of the buyer $m^b \in \mathcal{M}^b$. Following the message of the buyer $m^s \in \mathcal{M}^s$ and her own message $m^b \in \mathcal{M}^b$, the buyer updates her belief over $\theta$ to $\mu(m^b, m^s) = \big(\mu_1(m^b, m^s), \mu_2(m^b, m^s)\big)$ where $\mu_1(m^b, m^s) \equiv \Pr[\theta_{11} = 1 | m^b, m^s]$ and $\mu_2(m^b, m^s) \equiv \Pr[\theta_{12} = 1 | m^b, m^s]$.[6]

***Equilibrium.*** The solution concept is perfect Bayesian equilibrium. I allow for sufficiently rich spaces of messages $\mathcal{M}^b$ and $\mathcal{M}^s$. I rule out equilibria in which different messages have the same meaning. Formally, this means that in every subgame where there is communication, there cannot be two messages played with positive probability that result in the same posterior belief.[7]

# 3 Analysis

As in all cheap talk games, there will typically be multiple equilibria.[8] To select an equilibrium, consistent with much of the literature, I use the seller preferred equilibrium. Throughout the paper, this is what I refer to by 'an equilibrium'.

**Definition 1.** ***Seller preferred equilibrium***: *An equilibrium which maximises the seller's expected utility among the set of possible equilibrium payoffs.*

Note that from an ex ante point of view (meaning before the realisation of $\beta$), the seller preferred equilibrium will also be the equilibrium that maximises the buyer's utility.[9]

The key economic question of interest of the paper is whether in equilibrium, there can be benefits from the buyer communicating information about her preferences, $\beta$. In order to formalise this I introduce two further definitions:

**Definition 2.** ***Buyer influential equilibrium***: *An equilibrium in which there are two messages from the buyer played with positive probability that result in distinct beliefs for the seller.*

**Definition 3.** ***Beneficial conversation equilibrium***: *A buyer influential equilibrium in which the seller gets a strictly higher payoff compared a (seller preferred) equilibrium where the message space of the buyer is restricted to a single message:* $|\mathcal{M}^b| = 1$.

[6] The beliefs can be formulated in this way since it is assumed that $\theta_{1j} = 1 - \theta_{2j}$ for $j = 1, 2$.

[7] Note that this is a standard assumption and equilibria that are ruled out are payoff equivalent to an equilibrium that is not ruled out. See Section 4 of Sobel (2013) for a discussion.

[8] Existence is never a problem in cheap talk games since there always exists a 'babbling equilibrium' in which all messages are played by all types with equal probability and no information is transmitted.

[9] This may not be the case from an interim perspective—once the buyer knows her type. See the example in Result 2 below.

I am interested in whether the equilibrium is a beneficial conversation equilibrium—this is an equilibrium in which the seller gets a strictly higher payoff compared to the game where the buyer is not able to communicate.[10] In such an equilibrium the buyer also strictly benefits (ex ante) from her ability to communicate.

I consider two cases. First, where the buyer is only interested in a single attribute and her preferences are about which good she prefers. In this case the buyer cannot benefit from communicating her preferences. Second, where the buyer is interested in two attributes and her preferences are about which attribute she prefers. In this case the buyer can quite generally benefit from communication.

## 3.1 Preliminaries

Before analysing the specific cases described above, I start by discussing how to find the seller's optimal policy for a given belief he holds about the buyer's preferences, $\hat{F}$. To do this, I will introduce some additional notation and discuss how to characterise the maximum value that a sender (seller) can obtain in a cheap talk game where his preferences are state-independent. This methodology follows from Lipnowski and Ravid (2020) (henceforth, LR).[11]

Define $v(\mu, \hat{F})$ as the seller's expected payoff for a given buyer posterior belief $\mu \in \Delta \{0,1\}^2$ and belief that the buyer has preferences $\beta \sim \hat{F}$. Let $p$ be an information policy, and $s$ to be some possible seller payoff. A policy $p$ *secures* $s$ if $\{v \geq s\} = \left\{\mu : v(\mu, \hat{F}) \geq s\right\}$, and that $s$ is *securable* if an information policy exists that secures $s$. Informally, a payoff $s$ is securable if there is some information policy for which the worst payoff in its support is $s$.

**Theorem 1** (Lipnowski and Ravid (2020))**.** *Suppose $s \geq v(\mu_0, \hat{F})$. Then, an equilibrium inducing a seller payoff $s$ exists if and only if $s$ is securable.*

Note that the policy $p$ that secures $s$ need not be an equilibrium policy. The theorem does not provide any information about what the seller's optimal policy(ies) is (are). In order to find a policy in a seller preferred equilibrium, I make use of this theorem by using it to find an upper bound on the set of securable payoffs. If an (equilibrium) policy that achieves the highest securable payoff is found then this is clearly in the set of seller preferred policies.

## 3.2 Case 1: Buyer is only interested in a single attribute

**Assumption 1.** *$F_a$ satisfies the following:* $\Pr[\beta_a = 1] = 1$.

[10]Or equivalently, to an alternative equilibrium where in the first round of communication the buyer chooses an uninformative message (a babbling equilibrium).

[11]As they note, their model and results extend to games where the receiver (buyer) has private information that is not correlated with the sender's private information.

This means that the buyer is only interested in attribute 1. However, there is still potential uncertainty on how much the buyer is interested in each of the two goods based on the value of attribute 1. I maintain this assumption throughout this subsection.

I start by building intuition with a series of simple examples, and then provide the main result of this section (Proposition 1). This shows quite generally that the equilibrium is not a beneficial conversation equilibrium—meaning that the players never strictly benefit from the buyer communicating about her preferences.

**Assumption 2.** *$F_g$ satisfies the following:* $\Pr[\beta_g = \frac{1}{2}] = 1$.

This means there is no uncertainty in the preferences of the buyer, and (ex-ante) the buyer values each good equally. Clearly, given that the buyer does not have any private information, there is no buyer influential equilibrium. However, this will be a useful benchmark to analyse and the preferences are a simplified version of those in Chakraborty and Harbaugh (2010). With no communication, the buyer will value both products equally obtaining a utility of 1/4 for both $a = a_1$ and $a = a_2$—this means she buy a good with probability 1/4, which is the seller's payoff. With only one way communication from the seller to buyer, a recommendation of the best product is an equilibrium which influences the buyer's beliefs. The recommendation can be made by communicating the value of $\theta_{11}$:

$$m^s = \begin{cases} m_1^s & \text{if } \theta_{11} = 1, \\ m_2^s & \text{if } \theta_{11} = 0. \end{cases}$$

The message, $m^s = m_i^s$ can be interpreted as a recommendation for good $i$. When good $i$ is recommended the buyer chooses to buy good $i$ with probability 1/2. This is an equilibrium policy for the seller since either recommendation—good 1 or good 2—leads to the same probability of sale (and same expected payoff for the seller). Informally, it is clear that this is the best possible equilibrium for the seller, since it is not possible to induce a higher valuation—and hence probability of sale. Furthermore, this equilibrium is best for the buyer as well since she is able to make the most informed decision.

Now I consider more general distributions and I will return to formally verify the result above under Assumption 2. Denote the seller's belief that the buyer's preferences over goods ($\beta_g$) is distributed by $\hat{F}_g$ and recall the buyer's belief over $(\theta_{11}, \theta_{21})$ is given by $\mu_1$ where $\mu_1 \equiv \Pr[\theta_{11} = 1 | m^b, m^s]$. For this subsection, since $(\theta_{12}, \theta_{22})$ does not effect the buyer's utility, I will slightly abuse notation and write the seller's value function as a function of just $\mu_1$ and $\hat{F}_g$ and not the full state space $\mu$ and full set of seller beliefs over $\beta = (\beta_a, \beta_g)$. The seller's value function is:

$$v(\mu_1, \hat{F}_g) = \int_{\beta_g} \max\{\beta_g \mu_1, (1-\beta_g)(1-\mu_1)\} \, d\hat{F}_g(\beta_g).$$

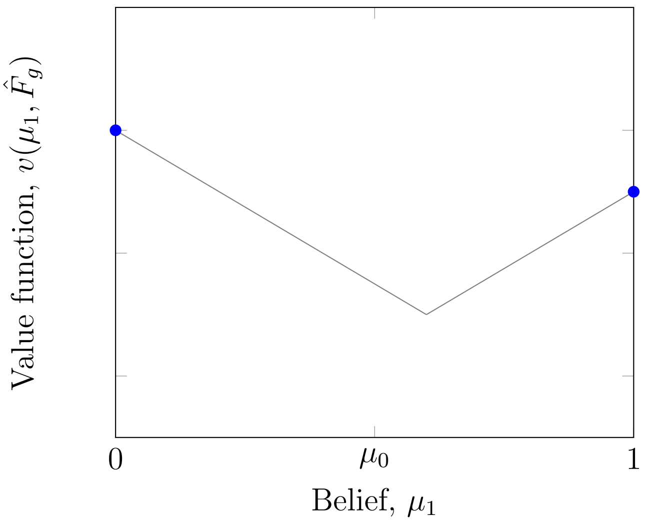


Figure 1: A possible value function $v(\mu_1, \hat{F}_g)$ in gray; the blue dots are the two possible maximum points.

$\max\{\beta_g\mu_1, (1-\beta_g)(1-\mu_1)\}$ is a convex function. Since the sum of convex functions is also a convex function, the seller's value function is convex. This means it attains a maximum at one of the end points $\mu_1 = 0$ or $\mu_1 = 1$—depicted by the blue dots in Figure 1.

If

$$\min\left\{v(0, \hat{F}), v(1, \hat{F})\right\} \geq v(\mu_0, \hat{F}),$$

the policy of fully revealing the state secures the seller a payoff of

$$\min\left\{v(0, \hat{F}), v(1, \hat{F})\right\}.$$

Since the value function convex it is clear that it is not possible for any policy to secure a strictly higher payoff. If

$$\min\left\{v(0, \hat{F}), v(1, \hat{F})\right\} < v(\mu_0, \hat{F}),$$

the policy of not revealing any information secures the seller a payoff of

$$v(\mu_0, \hat{F}).$$

Again, since the value function convex it is clear that it is not possible for any policy to secure a strictly higher payoff.

Define $\hat{\beta}_g \equiv \mathbb{E}_{\hat{F}_g}[\beta_g]$. Given the analysis above, the seller's value is summarised in the following lemma:

**Lemma 1.** *Under Assumption 1, if the seller has a belief that $\beta_g$ has distribution $\hat{F}_g$ and*

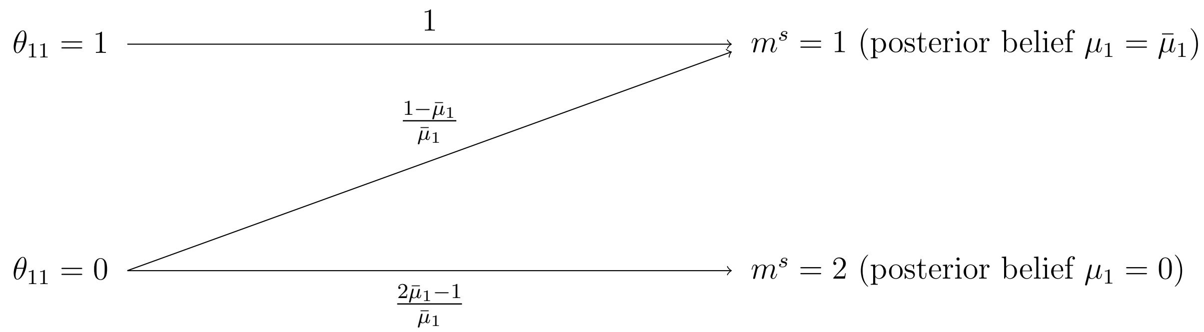


Figure 2: An example of an information policy for the seller.

*the buyer has a belief $\mu_0$ over $\theta_{i1}$, then the seller's expected payoff in equilibrium is*

$$\hat{v}(\mu_0, \hat{F}_g) = \begin{cases} v(\mu_0, \hat{F}_g) & \textit{if } v(\mu_0, \hat{F}_g) \geq \min\left\{\hat{\beta}_g, 1-\hat{\beta}_g\right\}, \\ \min\left\{\hat{\beta}_g, 1-\hat{\beta}_g\right\} & \textit{otherwise.} \end{cases}$$

As a corollary, returning to Assumption 2, where $\hat{F}_g = F_g$ is degenerate at $\beta_g = \frac{1}{2}$ and the prior is $\Pr[\theta_{11} = 1] = \frac{1}{2}$, it is confirmed that the seller obtains a value of 1/2 which can be uniquely achieved by the policy described above.

Now I construct an equilibrium seller policy for a general distribution $\hat{F}_g$. When $v(\mu_0, \hat{F}_g) \geq \min\left\{\hat{\beta}_g, 1-\hat{\beta}_g\right\}$, the seller does not provide any information; whereas when $v(\mu_0, \hat{F}_g) < \min\left\{\hat{\beta}_g, 1-\hat{\beta}_g\right\}$, the seller uses the following information policy when $\hat{\beta}_g \geq 1/2$.[12,13] The message space is $\mathcal{M}^s = \{m_1^s, m_2^s\}$ and the probability of sending a message is $\Pr[m^s = m_1^s|\theta_{11} = 1] = 1$ and $\Pr[m^s = m_1^s|\theta_{11} = 0] = \frac{1-\bar{\mu}_1}{\bar{\mu}_1}$, where $\bar{\mu}_1 \in [1/2, 1]$ is chosen to ensure indifference (see below). An example is depicted in Figure 2.

This induces posterior probabilities $\Pr[\theta_{11} = 1|m^s = m_1^s] = \bar{\mu}_1$ and $\Pr[\theta_{11} = 1|m^s = m_2^s] = 0$. Following $m^s = m_i^s$, the buyer will either buy good $i$ or take her outside option. To ensure that the seller is indifferent between sending each message when $\theta_{11} = 0$, $\bar{\mu}_1$ satisfies the following equation:

$$v(0, \hat{F}_g) = v(\bar{\mu}_1, \hat{F}_g).$$

This can be rewritten as

$$1 - \hat{\beta}_g = \int_{1-\bar{\mu}_1}^{1} \beta_g \bar{\mu}_1 d\hat{F}_g(\beta_g) + \int_0^{1-\bar{\mu}_1} (1-\beta_g)(1-\bar{\mu}_1) d\hat{F}_g(\beta_g). \tag{3.1}$$

This cannot be solved for a general distribution $\hat{F}_g$, but as an illustrative example consider

[12] When $\hat{\beta}_g < 1/2$ there is an analogous policy with the messages switched around.
[13] This is the unique policy that achieves his optimal payoff.

a distribution with binary support $\beta_g \in \left\{\frac{1}{4}, \frac{9}{10}\right\}$ where $\Pr[\beta_g = \frac{1}{4}] = 1/2$. Note that $\hat{\beta}_g = \frac{23}{40} > 1/2$. Solving equation 3.1 gives a unique solution: $\bar{\mu}_1 = 2/3$. The information policy of the seller is to send the message $m^s = 1$ when $\theta_{11} = 1$ and mix with equal probability between $m^s = m_1^s$ and $m^s = m_2^s$ when $\theta_{11} = 0$. Furthermore, since $1 - \hat{\beta}_g = \frac{17}{40} > v(\mu_0, \hat{F}_g) = \frac{33}{80}$, this policy is optimal (compared with a policy of no information).

Now consider the possibility of buyer communication. First, I define

$$I(x) \equiv x + \frac{1}{2}(1-x)^2 = \frac{1}{2}(1+x^2)$$

as the buyer's expected payoff (before learning her outside option) when the valuation of the more valuable good is $x$. Note $I(\cdot)$ is increasing and convex on $x \in [0,1]$. For a buyer with preference $\beta_g \in [0,1]$ and for a policy $\bar{\mu}_1 \in [1/2, 1]$ the buyer's expected payoff is given by

$$u(\bar{\mu}_1, \beta_g) = \begin{cases} \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1} I\left(1-\beta_g\right) + \frac{1}{2\bar{\mu}_1} I\left(\bar{\mu}_1 \beta_g\right) & \text{if } \hat{\beta}_g \geq 1/2, \\ \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1} I\left(\beta_g\right) + \frac{1}{2\bar{\mu}_1} I\left(\bar{\mu}_1 (1-\beta_g)\right) & \text{if } \hat{\beta}_g < 1/2. \end{cases}$$

This can be simplified to

$$u(\bar{\mu}_1, \beta_g) = \begin{cases} \frac{1}{2} + \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1}(1-\beta_g)^2 + \frac{1}{2}\bar{\mu}_1 \beta_g^2 & \text{if } \hat{\beta}_g \geq 1/2, \\ \frac{1}{2} + \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1}\beta_g^2 + \frac{1}{2}\bar{\mu}_1 (1-\beta_g)^2 & \text{if } \hat{\beta}_g < 1/2. \end{cases} \tag{3.2}$$

It is clear that the buyer's utility is strictly increasing in $\bar{\mu}_1$ since this gives her better information—and the best policy is always $\bar{\mu}_1 = 1$ meaning that the seller is fully revealing the state. So regardless of her preferences, the buyer would like to induce a belief that her preference is $\beta_g = \frac{1}{2}$ and get the most information from the seller's recommendation (but note that the preferences are not symmetric around $\beta_g = \frac{1}{2}$). Intuitively the 'single-peaked' nature of the preferences suggest that the buyer can never strictly benefit from communication. The reason is that if there are different messages that induce different seller beliefs, it must be that one of the messages induces 'better' beliefs for the buyer meaning she has an incentive to deviate and always induce these 'better' beliefs. This intuition is confirmed in the next result.

**Proposition 1.** *With a single attribute (Assumption 1), the (unique) equilibrium is never a beneficial conversation equilibrium.*

Formal proofs are all in the Appendix. The main steps of the proof are summarised as follows: First, I show that there cannot be more than one message sent in equilibrium which lead to expected beliefs either all above or all below 1/2. If this was the case clearly they would need to induce the same informational policy (summarised by $\bar{\mu}$), and I show that if this is the case, the messages can be replaced by a single message. Second, I show that there cannot be two messages sent where one message leads to an expected belief above 1/2 and the other to an expected belief below 1/2. The reason is that an equilibrium constructed by combining these two messages into a single message leads to an expected belief closer to 1/2 and on average result in more information being communicated by the seller.

### 3.3 Case 2: Buyer is potentially interested in both attributes

I now consider the possibility that the buyer is potentially interested in both attributes. In contrast to when the buyer is only interested in a single attribute, it will now be the case that the equilibrium is a beneficial conversation equilibrium. As in the previous subsection, I begin with some simple examples to build intuition and then present a more general result (Proposition 2). This result shows that quite generally, the equilibrium is a beneficial conversation equilibrium, and that it always takes a very simple form.

**Assumption 3.** *F satisfies the following:* $\Pr[\beta_g = \frac{1}{2}] = 1$.

To make things simple, I assume that the buyer values both goods equally and focus on preferences over attributes. I maintain this assumption throughout this subsection. The potential uncertainty will now be on how much the buyer values each attribute.

**Assumption 4.** *F satisfies the following:* $\Pr[\beta_a = \frac{1}{2}] = 1$.

Now the buyer values both attributes equally, so all that matters is the sum of the attributes for each good: $\theta_{i1} + \theta_{i2}$ for each good $i$. This is in effect the same situation as under Assumption 1 and 2 (although note that because $\theta_{i1}$ and $\theta_{i2}$ are independently drawn $\frac{1}{2}(\theta_{i1} + \theta_{i2})$ has a different distribution to $\theta_{i1}$). To find the optimal policy, consider the value function of the seller depicted in Figure 3. This is plotted in the two dimensional space below with the two axis being $\mu_1 \equiv \Pr[\theta_{11} = 1]$ and $\mu_2 \equiv \Pr[\theta_{12} = 1]$.

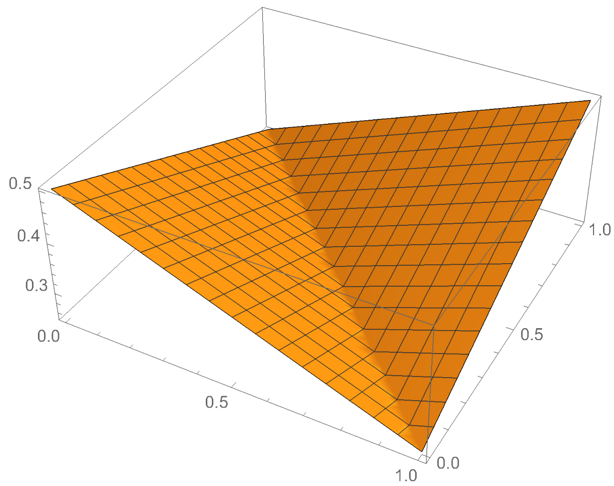


Figure 3: Value function $\left(v\left((\mu_1, \mu_2), \hat{F}\right)\right)$ under Assumption 4

The following policy secures a payoff of $v\left((1,\frac{1}{2}),\hat{F}_a\right) = 3/8$:[14]

$$m^s = \begin{cases} m_1^s & \text{if } \theta_{11} = 1, \\ m_2^s & \text{if } \theta_{11} = 0. \end{cases}$$

Effectively, this recommends the best good for attribute 1, and provides no information for attribute 2.[15] To verify that there is no policy that secures a higher payoff, I prove a more general lemma that can be applied to this specific example and will also be used for the more general results below. Recall that any payoff that the seller can secure (as defined in Theorem 1) is a payoff that the seller can achieve with some equilibrium policy. Thus the maximum value that he can secure, is his payoff in the seller preferred equilibrium.

**Lemma 2.** *For any posterior belief over $\beta_a$, $\hat{F}_a$, the maximum payoff the seller can secure is*

$$\hat{v}(\mu_0,\hat{F}_a) = \max\left\{v\left((1,\frac{1}{2}),\hat{F}_a\right), v\left((\frac{1}{2},1),\hat{F}_a\right), v\left((1,0),\hat{F}_a\right)\right\}.$$

The key intuition is that Bayes plausibility prevents the seller from securing a higher payoff. In Figure 3, the regions where the seller achieves a strictly higher payoff are in the right and left corners. However, there is no policy for which the posteriors of all messages lie in these two regions.

I illustrate the policies that secure these payoffs for the seller. When

$$\max\left\{v\left((1,\frac{1}{2}),\hat{F}_a\right), v\left((\frac{1}{2},1),\hat{F}_a\right), v\left((1,0),\hat{F}_a\right)\right\} = v\left((1,\frac{1}{2}),\hat{F}_a\right),$$

the policy depicted in Figure 4 secures this payoff and is also an equilibrium: it recommends the best good for attribute 1.[16] When

$$\max\left\{v\left((1,\frac{1}{2}),\hat{F}_a\right), v\left((\frac{1}{2},1),\hat{F}_a\right), v\left((1,0),\hat{F}_a\right)\right\} = v\left((1,0),\hat{F}_a\right),$$

the policy depicted in Figure 5 secures this payoff, however apart from when $\beta_a$ only takes extreme values (as in Assumption 5), this is not an equilibrium. The policy completely reveals the state. A typical example of an equilibrium policy is depicted in Figure 6. This policy recommends the best policy for attribute 1. For attribute 2 it makes a garbled recommendation biased towards recommending the good that was not recommended for attribute 1.

[14]Throughout this subsection, I will slightly abuse notation and summarise the seller's belief over $\beta$ by $\hat{F}_a$ since $\beta_g$ is known.

[15]Note that this is not the unique equilibrium policy that achieves this payoff. An alternative optimal policy is to recommend the best product for a single attribute, and be completely uninformative about the other attribute.

[16]The case where $\max\left\{v\left((1,\frac{1}{2}),\hat{F}_a\right), v\left((\frac{1}{2},1),\hat{F}_a\right), v\left((1,0),\hat{F}_a\right)\right\} = v\left((\frac{1}{2},1),\hat{F}_a\right)$ is similar—the policy recommends the best good for attribute 2, and reveals nothing about attribute 1.

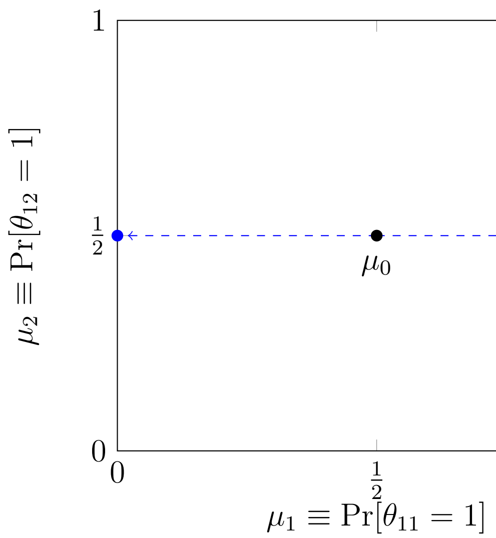


Figure 4: The blue dots represent the posteriors from the policy that secures $v\left((1, \frac{1}{2}), \hat{F}\right)$.

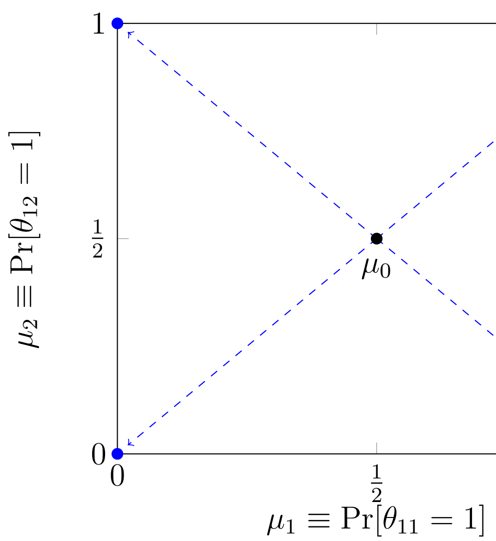


Figure 5: The blue dots represent the posteriors from the policy that secures $v\left((1, 0), \hat{F}\right)$.

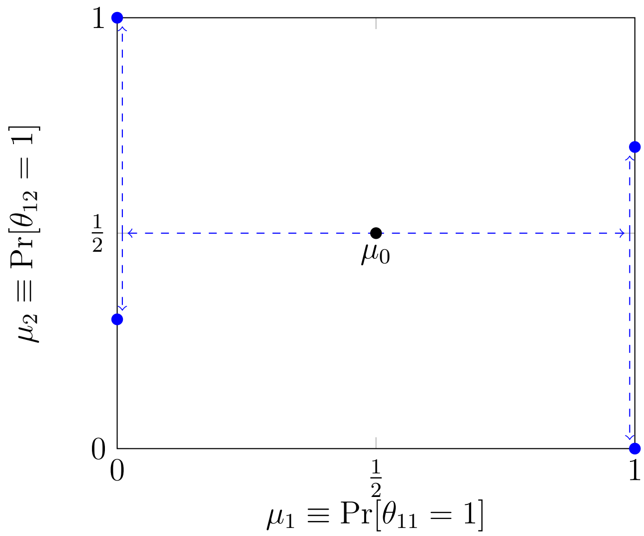


Figure 6: The blue dots represent the posteriors from the equilibrium policy that secures $v\left((1,0),\hat{F}\right)$. This completely reveals attribute 1, and partially reveals attribute 2.

**Assumption 5.** *$F$ satisfies the following: $\beta_a \in \{0,1\}$, and $\Pr[\beta_a = 0] = p \in (0,1)$.*

There is now uncertainty on the buyer's preferences over attributes. In particular, the buyer now only values one of the two attributes.

**Result 1.** *Under Assumptions 3 and 5, the equilibrium is not a beneficial conversation equilibrium.*

The seller's optimal policy is to recommend best good for each attribute—meaning he fully reveals the state. It is straightforward to show that this is an equilibrium, and clearly given that the state is fully revealed and the buyer's probability of buying is maximised, it is the optimal policy. To verify that this is optimal using Lemma 2, note that the result implies that the seller's value is $v\left((1,0),\hat{F}_a\right) = 1/2$. This is the payoff achieved by the policy of fully revealing the state.

Under this assumption the buyer cannot benefit from communicating her preferences before receiving the recommendation from the seller—she is already learning everything about the state. However, it turns out that this is a special case since the buyer's extreme preferences do not prevent the seller from communicating fully about both attributes. To see this, I now introduce a preference-type ($\beta_a = \frac{1}{2}$), who values both goods.

**Assumption 6.** *$F$ satisfies the following: $\beta_a \in \left\{0,\frac{1}{2},1\right\}$, and $\Pr[\beta_a = 0] = \Pr[\beta_a = 1] = p \in (0,\frac{1}{2})$.*

Under this assumption, the equilibrium will be a beneficial buyer informative equilibrium.

**Result 2.** *Under Assumptions 3 and 6, the equilibrium is a beneficial conversation equilibrium in which the buyer truthfully reveals her preferences. Compared to a game in which the buyer cannot communicate, the buyer's communication strictly benefits both players.*

I do not provide a proof of this result as it is a special case of Proposition 2 below, however, it is helpful to go through the intuition for this specific distribution over $\beta_a$. The key idea is that, compared to Assumption 5, there is now a friction in the seller communicating the state to the buyer and that the buyer can alleviate this friction by communicating her preferences. Consider when $p$ is close to $\frac{1}{2}$. It is likely that the buyer has a preference for just one attribute, and it is unlikely he just wants to buy the best good overall—so it is 'close' to Assumption 5. Consider what happens if the seller tries to use the same policy as before—recommending the best good for each attribute. Suppose for attribute 1, he recommends good 1. Then when making a recommendation for attribute 2 he is no longer indifferent between recommending good 1 and good 2—he has a strict preference to recommend good 1. The reason is that it is possible the buyer has a preference for the best good overall ($\beta_a = \frac{1}{2}$), and so if both attributes are better for one of the two goods—in this case good 1—then this increases the probability of a sale. So by revealing her preferences, the buyer allows the seller to make a recommendation that is better for her. This is also better for the seller, since the buyer is now more likely to buy one of the two goods.

Note that from an interim perspective the buyer with preference-type $\beta_a = \frac{1}{2}$, would prefer an equilibrium in which the buyer does not communicate. This is because the existence of preference-types $\beta_a = 0$ and $\beta_a = 1$ means that the seller's policy will be more informative than without them. However, from an ex ante perspective, before $\beta_a$ is learned, the buyer strictly prefers the equilibrium in which she reveals $\beta_a$.

Now I consider a general distribution of $F_a$. I fully characterise the buyer's communication in the seller optimal equilibrium. Furthermore, I show that under some mild assumptions, the equilibrium is always a beneficial conversation equilibrium.

**Assumption 7.** *The support of $F_a$ has positive mass in each of the intervals $\left(0, \frac{1}{2}\right)$ and $\left(\frac{1}{2}, 1\right)$.*

The assumption means that with positive probability each of the two attributes is potentially more important for the buyer. It also rules out the extreme case of Assumption 5, where the buyer has extreme preferences and she is interested only in one of attributes.

**Proposition 2.** *With two attributes and no bias towards either good (Assumption 3), there is an equilibrium that takes the following form:*

- *the buyer sends the message $m_1^b$ if $\beta_a \geq \frac{1}{2}$ and $m_2^b$ if $\beta_a < \frac{1}{2}$;*

- *following the message $m^b_j$, the seller sends the message $m^s_1$ if $\theta_{1j} = 1$ and $m^s_2$ if $\theta_{1j} = 0$.*

*If the distribution $F$ satisfies Assumption 7, the equilibrium is a beneficial conversation equilibrium. Furthermore, the equilibrium above is unique iff* $\Pr[\beta_a = \frac{1}{2}] = 0$.

In words, the equilibrium takes the following form. The buyer reveals which attribute she is more interested in, but not by how much more she is interested in that attribute. The message $m^b_j$ can be interpreted as saying: 'I am more interested in attribute $j$, tell me which good is better for this attribute.' Then the seller's policy fully reveals the best good for that attribute, and nothing about the other attribute.[17] This can be interpreted as the buyer saying: 'For the attribute you are most interested in, this is the best good.'

The formal proof is again in the Appendix, here I will discuss the intuition. If the seller has a belief that the buyer's preference is definitely towards one of the two attributes—so the updated belief $\hat{F}_a$ has support either above or below $\beta_a = \frac{1}{2}$—then the seller's optimal policy is just to fully reveal that attribute. Of the two attributes, the seller clearly benefits more from revealing information about the more favoured attribute. And once he has fully revealed about that attribute he is completely biased on the other attribute—he wants to recommend the same product as for the favoured attribute. This means he cannot reveal any information about this attribute. In order to see why the buyer's communication is to just reveal which attribute she prefers, it is straightforward that given the choice, the buyer wants to learn about the attribute she is most interested in. What is more subtle is why in equilibrium there is not a group of 'moderate' types close to $\beta_a = \frac{1}{2}$ who do not pool and learn about both attributes from the seller. In fact, this is the case under Assumption 6 for the type $\beta_a = \frac{1}{2}$, however it will never be the case for any other type. The reason is that buyers (other than type $\beta_a = \frac{1}{2}$) learn more from just learning about their preferred attribute, rather than from the seller's optimal policy when types above and below $\frac{1}{2}$ pool. In the latter case, the buyer learns about both attributes, but not everything about the attribute she is most interested in.

Under Assumption 7 in the equilibrium described, the buyer will send the message indicating a preference towards attribute 1 and 2 both with positive probability. Furthermore, since this equilibrium is the seller preferred equilibrium and gives the seller a strictly higher payoff than when the buyer does not communicate, the equilibrium is a beneficial conversation equilibrium. To see why Assumption 7 is necessary for the equilibrium to be a beneficial conversation equilibrium, consider the cases that it rules out. First, there is the case as under Assumption 5 where the buyer only has extreme preferences and there is no friction in communication about two attributes. Second, there is the case where the support of $F_a$ is either contained in $[0, \frac{1}{2}]$ or $[\frac{1}{2}, 1]$. In this case the

[17]Note that in Result 2, the buyer does not use this strategy for her messages. However, note that the equilibrium described is payoff equivalent (for both players) to one in which when $\beta_a = \frac{1}{2}$, the buyer randomises between reporting $\beta_a = 0$ and $\beta_a = 1$.

buyer is always interested in the same attribute and so the equilibrium is not a buyer influential equilibrium—she always sends the same message.

# 4 Related literature

The baseline model with one sided private information—a single attribute and the buyer having known and equal preferences over the two goods—was first analysed in Chakraborty and Harbaugh (2010).[18] LR use the tools they develop for more general state-independent cheap talk games to find the sender (seller) optimal equilibrium. As discussed above, I make use of these tools in the setting I study. Chakraborty and Harbaugh (2014) build on their example in their earlier paper to analyse a model in which a seller has a single good with multiple attributes. They focus on the potential value of 'puffery'—promoting one attribute over another. Their model does not consider a seller with multiple goods like I do, and in their model the buyer always has a strict preference for privacy (so there is never a buyer influential equilibrium in which strictly benefits the buyer).

Another paper that considers whether consumers benefit from having less private information is Gardete and Bart (2018). They study a model in which a seller (sender) tries to persuade a buyer (receiver) to purchase a good. The buyer and seller have partially aligned preferences—the seller always wants to make a sale, but more so when the match value is higher. The seller may have some information about the buyer's preferences. The question the paper considers is how much information is best? An intermediate level is optimal for the seller. Too much leads to recommendations not being credible. However, for the buyer, no information is optimal. A number of recent papers have considered whether a consumer (buyer) would want to communicate with a seller. For example, see Ali et al. (2020) and Hidir and Vellodi (2021). However, both of these papers consider a seller who is uninformed and can price discriminate. My model considers this question from a different perspective, when prices are fixed, but the seller has information that helps the buyer make the best decision. Closer to my model, a related applied theory paper is Levit and Tsoy (2022). They propose a theory to explain why 'one-size fits all' recommendations that are commonly made in cheap talk. The model has a sender (expert) who has to make recommendations to two different receivers (decision makers). The sender privately knows what is the best decision for each receiver, but for exactly one of them has a bias. This means making different recommendations will reveal the bias, and so recommendations are no longer followed. This is the basic intuition for why one-size fits all recommendations are optimal. This is related to the intuition in my paper where the sender cannot credibly say one attribute is better for one good and vis-versa for the other attribute.

[18] They consider a richer state space where $(\theta_{11}, \theta_{21}) \sim F[0,1]^2$.

As discussed in the introduction, there are very few papers where there are multiple rounds of cheap talk in a 'back-and-forth' manner between two privately informed players. Much of the literature on two way communication has either one-sided private information and/or simultaneous communication (Forges (1990), Krishna and Morgan (2004), Golosov et al. (2014)). One paper that has two way sequential communication is Chen (2009). However, this paper studies a model in which there is a one dimensional state of the world (as in Crawford and Sobel (1982)), and both players get a (private) informative signal about this—meaning that the private information is correlated.[19] A recent paper that has two way and sequential communication is Antic et al. (2020), however, this has a different focus since the two players have aligned interests and want to minimise what a third player, an outside observer, learns from their communication.

Finally, the analogue of my model in a setting with full commitment is studied in Kolotilin et al. (2017). They consider a model of Bayesian persuasion with a privately informed receiver. They show there is no benefit to the sender if he conditions the message (information structure) on a report made by the receiver. Their result relies on a binary decision based on linear preferences for the receiver and the sender having state-independent preferences.

# 5 Discussion and concluding remarks

In this paper I consider a cheap talk model with sequential two way communication and two sided private information. In order to make progress with this communication protocol, I have considered a specific setting where the sender/seller has state independent preferences. I have also assumed the simplest possible form for the state space of $\theta$. These assumptions allow me to provide clear conditions under which both players benefit from the decision maker/buyer communicating before the sender/seller. A natural question is to ask to what extent my results would generalise? With a richer state space, I conjecture that a similar intuition would hold. The benchmark model from Chakraborty and Harbaugh (2010) with a single attribute and no uncertainty on preferences is analysed with $\theta_i$ drawn i.i.d. from an arbitrary, full-support distribution on $[0, 1]$. Here LR show that the seller's best policy is to recommend the best product. However, even introducing a known bias towards one of the two goods (so no uncertainty for the seller) poses a technical challenge and it is unclear what the seller's optimal policy is.[20] Another natural extension is to consider uncertainty on both $\beta_a$ and $\beta_g$. Here I conjecture that in equilibrium, if there are multiple attributes, the buyer will still only reveal information

[19] In a recent theoretical and experimental paper Burdea and Woon (2021) study two way communication but with only one sided information. Their results rely on some sender's being 'truthful' types, who do not choose messages 'strategically'.

[20] In particular, for a distribution on the unit square, the 'corner' poses a problem. However, even if the distribution was on a quarter circle with centre $(0, 0)$, it is unclear what the optimal seller policy is.

about her preferences about attributes ($\beta_a$).

A further question of interest is what payoffs could be achieved in my setting if instead of the specified protocol, any possible communication protocol was possible. This could include simultaneous rounds of communication that allow for randomisations through 'jointly controlled lotteries' (as in Forges (1990) and Krishna and Morgan (2004)). Furthermore, one could consider communication through a mediator (as in Myerson (1986)). These possibilities clearly can only increase the set of payoffs (and increase the seller's maximum payoff). I believe that the communication protocol that I have studied is both novel and quite natural for the application to a buyer and seller. However, in future work it would be interesting to understand to what extent payoffs can be increased with more general protocols, and what form such a more complex 'conversation' takes with two sided private information.[21]

# References


Acquisti, A., Taylor, C., and Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature*, 54(2):442–492.

Ali, S. N., Lewis, G., and Vasserman, S. (2020). Voluntary Disclosure and Personalized Pricing. *working paper.*

Antic, N., Chakraborty, A., and Harbaugh, R. (2020). Subversive Conversations. *working paper.*

Burdea, V. and Woon, J. (2021). Questioning. *working paper.*

Chakraborty, A. and Harbaugh, R. (2010). Persuasion by Cheap Talk. *The American Economic Review*, 100(5):2361–2382.

Chakraborty, A. and Harbaugh, R. (2014). Persuasive Puffery. *Marketing Science*, 33(3):382–400.

Chen, Y. (2009). Communication with Two-sided Asymmetric Information. *working paper.*

Crawford, V. P. and Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, 50(6):1431–1451.

Forges, F. (1990). Equilibria with Communication in a Job Market Example*. *The Quarterly Journal of Economics*, 105(2):375–398.


[21]Forges (1990) asks such a question in a specific cheap talk game with one sided private information.

Gardete, P. M. and Bart, Y. (2018). Tailored Cheap Talk: The Effects of Privacy Policy on Ad Content and Market Outcomes. *Marketing Science*, 37(5):733–752.

Golosov, M., Skreta, V., Tsyvinski, A., and Wilson, A. (2014). Dynamic strategic information transmission. *Journal of Economic Theory*, 151:304–341.

Hidir, S. and Vellodi, N. (2021). Privacy, Personalization, and Price Discrimination. *Journal of the European Economic Association*, 19(2):1342–1363.

Hoffmann, F., Inderst, R., and Ottaviani, M. (2020). Persuasion Through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation. *Management Science*, 66(11):4958–4979.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017). Persuasion of a Privately Informed Receiver. *Econometrica*, 85(6):1949–1964.

Krishna, V. and Morgan, J. (2004). The art of conversation: eliciting information from experts through multi-stage communication. *Journal of Economic Theory*, 117(2):147–179.

Levit, D. and Tsoy, A. (2022). A Theory of One-Size-Fits-All Recommendations. *American Economic Journal: Microeconomics*.

Lipnowski, E. and Ravid, D. (2020). Cheap Talk With Transparent Motives. *Econometrica*, 88(4):1631–1660.

Myerson, R. B. (1986). Multistage Games with Communication. *Econometrica*, 54(2):323–358.

Sobel, J. (2013). Advances in Economics and Econometrics: Giving and Receiving Advice. Cambridge University Press.

Varian, H. R. (2002). Economic Aspects of Personal Privacy. In Lehr, W. H. and Pupillo, L. M., editors, *Cyber Policy and Economics in an Internet Age*, Topics in Regulatory Economics and Policy Series, pages 127–137. Springer US, Boston, MA.

# Appendix A Proofs

## A.1 Proof of Proposition 1

*Proof.* Assume throughout that $\mathbb{E}_F[\beta_g] \geq 1/2$. If $\mathbb{E}_F[\beta_g] < 1/2$, a very similar argument can be made. Given a preference-type $\beta_g$, the buyer chooses a message $m^b \in \mathcal{M}^b$ to maximise her utility. The seller then correctly updates his beliefs, and the message $m^b$ results in an information policy fully characterised by $\bar{\mu}_1$ (as described above).

I will prove the result by contradiction. Suppose there are two distinct messages played in equilibrium: $m$ and $m'$ and the distribution of types playing each message is given by $\hat{F}_g$ and $\hat{F}'_g$. Assume that both are played by types such that the expected value of the seller's posterior—given by $\hat{\beta}_g$ and $\hat{\beta}'_g$ respectively—are greater than 1/2. It is straightforward that the two messages must result in information policies that are equally informative, i.e. that $\bar{\mu}_1 = \bar{\mu}'_1$.[22] If this were not the case, then no type would choose the message with the less informative information policy (i.e. with $\min\{\bar{\mu}_1, \bar{\mu}'_1\}$). Now, I show that $m$ and $m'$ can be replaced by a single message $m''$ played by all types previously playing $m$ and $m'$ and that results in an information $\bar{\mu}''_1 = \bar{\mu}_1 = \bar{\mu}'_1$. So the equilibrium with $m''$ is outcome equivalent to the one with $m$ and $m'$. To see why this is the case, the information policies $\bar{\mu}_1$ and $\bar{\mu}'_1$ are given by the two equations

$$
\begin{aligned}
1 - \hat{\beta}_g &= \int_{1-\bar{\mu}_1}^{1} \beta_g \bar{\mu}_1 d\hat{F}_g(\beta_g) + \int_{0}^{1-\bar{\mu}_1} (1-\beta_g)(1-\bar{\mu}_1) d\hat{F}_g(\beta_g),\\
1 - \hat{\beta}'_g &= \int_{1-\bar{\mu}_1}^{1} \beta_g \bar{\mu}_1 d\hat{F}'_g(\beta_g) + \int_{0}^{1-\bar{\mu}_1} (1-\beta_g)(1-\bar{\mu}_1) d\hat{F}'_g(\beta_g).
\end{aligned}
$$

Note that in each equation $\bar{\mu}_1$ is the same. Let $p$ and $p'$ be the probability of the respective message being played and let $\hat{F}''_g$ be the distribution of types playing the new combined message. Multiplying the first equation by $\frac{p}{p+p'}$ and the second equation by $\frac{p'}{p+p'}$ and summing the two equations gives

$$
1 - \hat{\beta}''_g = \int_{1-\bar{\mu}_1}^{1} \beta_g \bar{\mu}_1 d\hat{F}''_g(\beta_g) + \int_{0}^{1-\bar{\mu}_1} (1-\beta_g)(1-\bar{\mu}_1) d\hat{F}''_g(\beta_g).
$$

Since $\bar{\mu}_1$ solves this equation, the information policy of the new message $m''$ is also $\bar{\mu}_1$. So, in equilibrium there must be at most one message played with $\hat{\beta}_g \geq 1/2$. A similar argument means that there must be at most one message played with $\hat{\beta}_g < 1/2$.

However, I have not ruled out that there may be one message played with $\hat{\beta}_g \geq 1/2$ and one with $\hat{\beta}_g < 1/2$. I now show that this is not possible. Suppose there are two distinct messages played in equilibrium $m$ and $m'$ such that the expected value of the seller's posterior are $\hat{\beta}_g \geq 1/2$ and $\hat{\beta}'_g < 1/2$ with respective distributions $\hat{F}_g$ and $\hat{F}'_g$. Following $m$, the seller's payoff from the optimal policy is $v(0, \hat{F}_g) = 1 - \hat{\beta}_g$. Similarly, for $m'$ the seller's payoff is $v(1, \hat{F}'_g) = \hat{\beta}'_g$. Now consider a babbling equilibrium, where the buyer sends a single message $m''$ for all types $\beta_g$. Denote the probability that in the original equilibrium, $m$ is played by $p$ and $m'$ by $1-p$. Since, by assumption,

[22] Note that this does not mean that it must be that $\hat{\beta}_g = \hat{\beta}'_g$. For example, returning to the earlier example, recall there was a distribution $\beta_g \in \{\frac{1}{4}, \frac{9}{10}\}$, with expectation $\hat{\beta}_g = \frac{23}{40}$, that resulted in a policy with $\bar{\mu} = \frac{2}{3}$. The degenerate distribution with $\beta_g = \frac{3}{5}$ has a different expectation, but results in the same policy $\bar{\mu}' = \frac{2}{3}$.

$p\hat{\beta}_g + (1-p)\hat{\beta}'_g = \mathbb{E}_F[\beta_g] \geq 1/2$, in the new equilibrium the seller's payoff from the policy with message $m''$ is given by

$$v(0, \hat{F}''_g) = 1 - \left(p\hat{\beta}_g + (1-p)\hat{\beta}'_g\right). \tag{A.1}$$

In contrast, the expected payoff in the original equilibrium is

$$pv(0, \hat{F}_g) + (1-p)v(0, \hat{F}'_g) = p(1-\hat{\beta}_g) + (1-p)\hat{\beta}_g. \tag{A.2}$$

By subtracting A.2 from A.1, it is straightforward that the seller's payoff is always higher under the babbling equilibrium with message $m''$ always being sent.

Finally, note that this equilibrium is unique. This is because there is a unique policy pinned down by $\bar{\mu}_1$. To verify this, differentiating the RHS of 3.1 gives

$$\frac{\partial}{\partial \bar{\mu}_1} \int_{1-\bar{\mu}_1}^{1} \beta_g \bar{\mu}_1 d\hat{F}_g(\beta_g) + \int_0^{1-\bar{\mu}_1} (1-\beta_g)(1-\bar{\mu}_1) d\hat{F}_g(\beta_g) = \int_{1-\bar{\mu}_1}^{1} \beta_g d\hat{F}_g(\beta_g) - \int_0^{1-\bar{\mu}_1} (1-\beta_g) d\hat{F}_g(\beta_g)$$
$$> 0.$$

The inequality follows from the fact that $\bar{\mu}_1 > 1/2$ and that $\hat{\beta}_1 > 1/2$. This means that the RHS of 3.1 is strictly increasing in $\bar{\mu}_1$ and so by the Intermediate Value Theorem, equation 3.1 has a unique solution.

□

## A.2 Proof of Lemma 2

*Proof.* Assume throughout that $\hat{\beta}_a \equiv \mathbb{E}_{\hat{F}_a}[\beta_a] \geq 1/2$. This means that $v\left((1, \frac{1}{2}), \hat{F}_a\right) \geq v\left((\frac{1}{2}, 1), \hat{F}_a\right)$. When $\hat{\beta}_a \leq 1/2$, have that $v\left((1, \frac{1}{2}), \hat{F}_a\right) \geq v\left((\frac{1}{2}, 1), \hat{F}_a\right)$, and a very similar argument can be made.

Throughout, I describe a policy which has binary support (which is not necessarily an equilibrium policy) as two sets of lotteries over the possible states of the world. Denote by $\pi^k_{ij} \in [0,1]$ the probability that message $m^s_k \in \{m^s_1, m^s_2\}$ is sent in the state $\theta_{11} = i \in \{0,1\}$ and in the state $\theta_{12} = j \in \{0,1\}$. Bayes plausibility requires that $\pi^1_{ij} + \pi^2_{ij} = 1$ for all $i, j$. Furthermore, the total probability of message $m^b_k$ being sent is $\pi^k = \pi^k_{11} + \pi^k_{10} + \pi^k_{01} + \pi^k_{00}$ for $k \in \{1, 2\}$.

***Case 1.*** $v\left((1, \frac{1}{2}), \hat{F}_a\right) \geq v\left((1, 0), \hat{F}_a\right)$

The following policy secures a payoff of $v\left((1, \frac{1}{2}), \hat{F}_a\right)$: $\pi^1_{11} = \pi^1_{12} = 1$ and $\pi^1_{21} = \pi^1_{22} = 0$. In words, this policy completely reveals the value of attribute 1, and says nothing about the value of attribute 2.

Now, I show that there is no policy that secures a higher payoff than $v\left((1, \frac{1}{2}), \hat{F}_a\right)$.

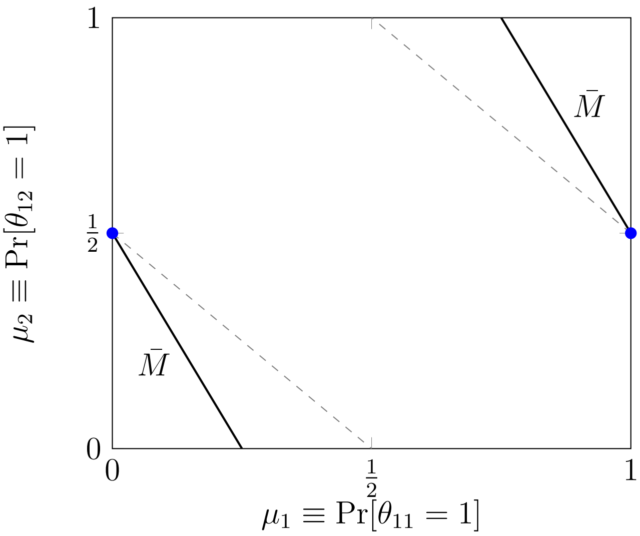


Figure 7: Set $\bar{M}$ in $(\mu_1, \mu_2)$ belief space. The gray dotted lines delineate the regions that any $\bar{M}$ are contained in. The blue dots represent the posteriors from the policy that secures $v\left(1, \frac{1}{2}), \hat{F}\right)$.

First, I restrict attention to binary policies, then I show that this extends to the set of all possible policies.

Denote by $\bar{M}$, the set of buyer posterior beliefs $(\mu_1, \mu_2)$ where the seller obtains a strictly higher payoff than the secured payoff $v\left((1, \frac{1}{2}), \hat{F}_a\right)$:

$$\bar{M} \equiv \left\{(\mu_1, \mu_2) : v\left((\mu_1, \mu_2), \hat{F}_a\right) > v\left((1, \frac{1}{2}), \hat{F}_a\right)\right\}.$$

I now introduce a lemma that restricts the possible beliefs in the set $\bar{M}$.

**Lemma 3.** *Assume that $\hat{\beta}_a \geq 1/2$.*

$$\bar{M} \subseteq \{(\mu_1, \mu_2) : \mu_1 + \mu_2 > 3/2\} \cup \{(\mu_1, \mu_2) : \mu_1 + \mu_2 < 1/2\}.$$

I depict an example of the set $\bar{M}$ in Figure 7.

*Proof.* I consider only the region where $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times [0, 1]$—the symmetry of the problem means an almost identical argument can be made for $(\mu_1, \mu_2) \in \left[0, \frac{1}{2}\right] \times [0, 1]$.

I proceed in two steps. First, I show that $v\left((1, \frac{1}{2}), \hat{F}\right) \geq v\left((\mu_1, \mu_2), \hat{F}\right)$ for any $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times \left[0, \frac{1}{2}\right]$ (Step 1). Second, I show that for any $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times \left[\frac{1}{2}, 1\right]$, $v\left((1, \frac{1}{2}), \hat{F}\right) < v\left((\mu_1, \mu_2), \hat{F}\right)$ only if $\mu_1 + \mu_2 > 3/2$ (Step 2).

The seller's value function is given by:

$$v(\mu, \hat{F}_a) = \frac{1}{2} \int_{\beta_a} \max\left\{\beta_a \mu_1 + (1 - \beta_a)\mu_2, \beta_a(1 - \mu_1) + (1 - \beta_a)(1 - \mu_2)\right\} d\hat{F}_a(\beta_a).$$

***Step 1.*** Consider lines where $\mu_2$ is fixed for some $\mu_2 \in \left[0, \frac{1}{2}\right]$ and $\mu_1$ takes values from $\frac{1}{2}$ to 1. For $\mu_1 \in \left[\frac{1}{2}, 1\right]$, $v$ can be written as

$$\begin{aligned} v(\mu, \hat{F}_a) &= \frac{1}{2}\int_{\beta_a} \max\left\{\beta_a\mu_1 + (1-\beta_a)\mu_2, \beta_a(1-\mu_1) + (1-\beta_a)(1-\mu_2)\right\} d\hat{F}_a(\beta_a) \\ &= \frac{1}{2}\int_{\beta_a \geq \frac{1-2\mu_2}{2(\mu_1-\mu_2)}} \beta_a\mu_1 + (1-\beta_a)\mu_2 d\hat{F}_a(\beta_a) \\ &+ \frac{1}{2}\int_{\beta_a < \frac{1-2\mu_2}{2(\mu_1-\mu_2)}} \beta_a(1-\mu_1) + (1-\beta_a)(1-\mu_2) d\hat{F}_a(\beta_a). \end{aligned}$$

Differentiating with respect to $\mu_1$ gives

$$\frac{dv}{d\mu_1} = \frac{1}{2}\int_{\frac{1}{2}}^{\frac{1-2\mu_2}{2(\mu_1-\mu_2)}} \beta_a d\hat{F}_a(\beta_a) - \frac{1}{2}\int_{\frac{1-2\mu_2}{2(\mu_1-\mu_2)}}^{1} \beta_a d\hat{F}_a(\beta_a).$$

It is clear from this that for every $\mu_2$, along $\mu_1 \in \left[\frac{1}{2}, 1\right]$ takes a 'V-shape' with the maximum in this range at either $\mu_1 = \frac{1}{2}$ or $\mu_1 = 1$. Furthermore, since $v\left((1, \frac{1}{2}), \hat{F}_a\right) > v\left((1, 0), \hat{F}_a\right)$, for $\mu_2 \in \left[0, \frac{1}{2}\right]$, $v((1, \mu_2), \hat{F}_a)$ is decreasing in $\mu_2$; and for $\mu_2 \in \left[0, \frac{1}{2}\right]$, $v\left((\frac{1}{2}, \mu_2), \hat{F}_a\right)$ is increasing in $\mu_2$. Combining these we have that $v\left((1, \frac{1}{2}), \hat{F}_a)\right) \geq v\left((\mu_1, \mu_2), \hat{F}_a\right)$ for all $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times \left[0, \frac{1}{2}\right]$.

***Step 2.*** To show that for any $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times \left[\frac{1}{2}, 1\right]$, $v\left((1, \frac{1}{2}), \hat{F}\right) < v\left((\mu_1, \mu_2), \hat{F}\right)$ only if $\mu_1 + \mu_2 > 3/2$, first observe that for any $(\mu_1, \mu_2) \in \left[\frac{1}{2}, 1\right] \times \left[\frac{1}{2}, 1\right]$, the buyer will choose good 1 regardless of her preference type $\beta_a$. The seller's payoff is:

$$\begin{aligned} v(\mu, \hat{F}_a) &= \frac{1}{2}\int_{\beta_a} \beta_a\mu_1 + (1-\beta_a)\mu_2 d\hat{F}_a(\beta_a), \\ &= \frac{1}{2}\left(\hat{\beta}_a\mu_1 + (1-\hat{\beta}_a)\mu_2\right). \end{aligned}$$

This is strictly greater than $v\left((1, \frac{1}{2}), \hat{F}\right)$ if

$$\frac{1}{2}\left(\hat{\beta}_a\mu_1 + (1-\hat{\beta}_a)\mu_2\right) > \frac{1}{2}\left(\hat{\beta}_a + (1-\hat{\beta}_a)\frac{1}{2}\right),$$

which simplifies to

$$\mu_2 > \frac{1}{2} + \frac{1-\hat{\beta}_a}{\hat{\beta}_a}(1-\mu_1).$$

It is straightforward that $\mu_1 + \mu_2 > 3/2$ is a sufficient condition for this to be satisfied. □

Now using this lemma, I return to show that the seller cannot secure a strictly higher payoff. For a binary policy to secure a strictly higher payoff than $v\left((1,\frac{1}{2}),\hat{F}_a\right)$, it must be that following both messages $m_1^b$ and $m_2^b$, the buyer's posterior belief is in $\bar{M}$: $\mu(m_1^b),\mu(m_2^b) \subseteq M$ . Clearly it cannot be the case that both posteriors either have $\mu_1+\mu_2 > 3/2$ or $\mu_1+\mu_2 < 1/2$. So, have $\mu(m_1^b)$ such that $\mu_1+\mu_2 > 3/2$; and $\mu(m_2^b)$ such that $\mu_1+\mu_2 < 1/2$. Calculating the posterior beliefs in terms of $\pi_{ij}^1$:

$$\begin{aligned}
\mu(m_1^b) &= \left(\frac{\pi_{11}^1+\pi_{10}^1}{\pi^1},\frac{\pi_{11}^1+\pi_{01}^1}{\pi^1}\right),\\
\mu(m_2^b) &= \left(\frac{(1-\pi_{11}^1)+(1-\pi_{10}^1)}{(4-\pi^1)},\frac{(1-\pi_{11}^1)+(1-\pi_{01}^1)}{(4-\pi^1)}\right).
\end{aligned}$$

To have $\mu(m_1^b),\mu(m_2^b) \subseteq M$, these must satisfy

$$\begin{aligned}
&\frac{\pi_{11}^1+\pi_{10}^1}{\pi^1}+\frac{\pi_{11}^1+\pi_{01}^1}{\pi^1} > 3/2,\\
&\frac{(1-\pi_{11}^1)+(1-\pi_{10}^1)}{(4-\pi^1)}+\frac{(1-\pi_{11}^1)+(1-\pi_{01}^1)}{(4-\pi^1)} < 1/2.
\end{aligned}$$

Rewriting these inequalities

$$\begin{aligned}
&\pi_{11}^1-\pi_{01}^1-\pi_{10}^1-3\pi_{00}^1 > 0,\\
&3\pi_{11}^1+\pi_{01}^1+\pi_{10}^1-\pi_{00}^1 > 4.
\end{aligned}$$

Since $\pi_{11}^1 \leq 1$ and $\pi_{00}^1 \geq 0$, this implies that

$$\begin{aligned}
&\pi_{01}^1+\pi_{10}^1 < 0,\\
&\pi_{01}^1+\pi_{10}^1 > 0,
\end{aligned}$$

which is a contradiction.

Now consider the possibility that there are more than two messages in the seller's policy. As before, there must be at least one message that leads to a posterior in either of the two sets $\mu_1+\mu_2 > 3/2$ or $\mu_1+\mu_2 < 1/2$. Note that both these sets are convex. Suppose that there was a policy with more than two messages where all posteriors were in these two regions. Combining all messages within each of the two sets would lead to posteriors that were still within the two sets. This would mean that there was a policy with two messages that secured a strictly higher payoff than $v\left((1,\frac{1}{2}),\hat{F}_a\right)$. However, as shown above this is not possible.

***Case 2.*** $v\left((1,0),\hat{F}_a\right) \geq v\left((1,\frac{1}{2}),\hat{F}_a\right)$

This case is very similar. The policy that secures a payoff of $v\left((1,0),\hat{F}_a\right)$ requires four messages:

$$\begin{aligned}&m_1^s \text{ if } \theta_{11}=\theta_{12}=1,\\&m_2^s \text{ if } \theta_{11}=1,\theta_{12}=0,\\&m_3^s \text{ if } \theta_{11}=0,\theta_{12}=1,\\&m_4^s \text{ if } \theta_{11}=\theta_{12}=0.\end{aligned}$$

In words, this policy completely reveals the value of both attributes.

To show that it is not possible to improve on this policy, again, it is the case that the set of buyer posterior beliefs that lead to a strictly higher payoff for the seller is

$$M \subseteq \{(\mu_1,\mu_2) : \mu_1+\mu_2 > 3/2\} \cup \{(\mu_1,\mu_2) : \mu_1+\mu_2 < 1/2\}.$$

Using the same argument as before, there is no policy that secures a strictly higher payoff than $v\left((1,0),\hat{F}_a\right)$.

□

## A.3 Proof of Proposition 2

*Proof.* I begin by showing that the strategies described form an equilibrium. Then, I show that this equilibrium is a seller preferred equilibrium. Next, I show that Assumption 7 is necessary for the equilibrium to be a beneficial conversation equilibrium. Finally, I show that the strategies are the unique seller preferred equilibrium if $\Pr[\beta_a=\frac{1}{2}]=0$.

To verify that the seller's policy is optimal given the buyer's strategy, consider the seller's problem following $m_1^b$. The seller's belief over $\beta_a$ is $\hat{F}_a(m_1^b)$ and has support $[\frac{1}{2},1]$. By Lemma 2, since $\hat{\beta}_a \geq \frac{1}{2}$ and $v\left((1,\frac{1}{2}),\hat{F}_a(m_1^b)\right) > v\left((1,0),\hat{F}_a(m_1^b)\right)$, the maximum payoff the seller can secure is $v\left((1,\frac{1}{2}),\hat{F}_a(m_1^b)\right)$. This is achieved by the policy of revealing only attribute 1 as in the statement of the proposition.

Next, given this choice of policy by the seller, the buyer's communication strategy described in the proposition is optimal. To see this, consider a buyer who has preference-type $\beta_a \geq \frac{1}{2}$ (there is a similar argument for $\beta_a < \frac{1}{2}$). Her payoff from choosing $m_1^b$ (and learning from the seller's optimal policy) is $I\left(\frac{1+\beta_a}{2}\right)$, while her payoff from choosing $m_2^b$ is $I\left(\frac{2-\beta_a}{2}\right)$. Since $I(\cdot)$ is an increasing function, it is clear that the buyer's communication strategy is optimal.

Now, I show that there cannot be another equilibrium that strictly improves the seller's payoff. If there are two messages that are played by a different distribution of types where all type $\beta_a$ are either above or below $\beta_a=\frac{1}{2}$, the seller's optimal policy following both messages will be the same. This means that an equilibrium in which these two messages are replaced by a single message is payoff equivalent. So, it is left to consider the possibility that there is a message played by types both above and below $\beta_a = \frac{1}{2}$.

Consider an equilibrium with a message $\bar{m}^b$ that is sent by at least two buyer types: $\beta_a \geq \frac{1}{2}$ and $\beta'_a < \frac{1}{2}$. Denote the set of types playing this message by $\bar{M}^b$. Define $\hat{\beta}_a^+ \equiv \mathbb{E}[\beta_a | \beta_a \geq \frac{1}{2}, \beta_a \in \bar{M}^b]$ and $\hat{\beta}_a^- \equiv \mathbb{E}[\beta_a | \beta_a < \frac{1}{2}, \beta_a \in \bar{M}^b]$, these are the conditional expectation of the types playing the new message given they are above and below $\frac{1}{2}$. Also define $\bar{p}_+ \equiv \Pr[\beta_a \geq \frac{1}{2} | \beta_a \in \bar{M}^b]$ and $\bar{p}_- \equiv \Pr[\beta_a < \frac{1}{2} | \beta_a \in \bar{M}^b]$ as the respective probabilities of these. Now, I show that an equilibrium in which these types play $m_1^b$ and $m_2^b$ respectively (as in the proposition) and the seller chooses the optimal policy (again, as in the proposition) is strictly better for the seller. The seller's value from all buyer types playing $\bar{m}^b$ can be derived from Lemma 2 as before, and is

$$(\bar{p}_+ + \bar{p}_-)\, v(\mu, \hat{F}_a(\bar{m}^b)) = (\bar{p}_+ + \bar{p}_-) \max \left\{ v\left((1,0), \hat{F}_a(\bar{m}^b)\right), v\left((1, \frac{1}{2}), \hat{F}_a(\bar{m}^b)\right) \right\}. \tag{A.3}$$

The different parts of the RHS of the expression above can be calculated as:

$$\begin{aligned} v\left((1,0), \hat{F}_a(\bar{m}^b)\right) &= \frac{1}{2}\bar{p}_+ \hat{\beta}_a^+ + \frac{1}{2}\bar{p}_- \left(1 - \hat{\beta}_a^-\right), \\ v\left((1, \frac{1}{2}), \hat{F}_a(\bar{m}^b)\right) &= \frac{1}{4}\bar{p}_+ \left(1 + \hat{\beta}_a^+\right) + \frac{1}{4}\bar{p}_- \left(1 + \hat{\beta}_a^-\right). \end{aligned}$$

In the first expression ($\mu = (1,0)$), when $\beta_a \geq \frac{1}{2}$ the payoffs are calculated using the buyer's valuation of the first good, and when $\beta_a < \frac{1}{2}$ the payoffs are calculated using the buyer's valuation of the second good. In contrast, in the second expression ($\mu = (1, \frac{1}{2})$), the payoffs are calculated using the buyer's value of the first good.

In the original equilibrium from the proposition, the payoff for the seller from the buyer types playing $\bar{m}^b$ is

$$(\bar{p}_+ + \bar{p}_-)\, \bar{p}_+ \frac{1}{4}\left(1 + \hat{\beta}_a^+\right) + \bar{p}_- \frac{1}{4}\left(2 - \hat{\beta}_a^-\right). \tag{A.4}$$

By comparing A.3 to A.4, it follows that the payoff in the original equilibrium is strictly greater than the payoff under the new equilibrium when they play $\bar{m}$.

Now I show that if the distribution $F$ satisfies Assumption 7, the equilibrium is a beneficial conversation equilibrium. To do this I compare the seller's payoff when the buyer is not able to communicate and the payoff in the equilibrium above and show that the latter is always greater.

Again, I assume that $\hat{\beta}_a \geq \frac{1}{2}$ (and again, a similar argument can be made when

$\hat{\beta}_a < \frac{1}{2}$). The seller's payoff when the buyer cannot communicate is

$$\max\left\{v\left((1,\frac{1}{2}),\hat{F}_a\right),v\left((1,0),\hat{F}_a\right)\right\} = \max\left\{\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a(\beta_a), \int \max\left\{\beta_a, 1-\beta_a\right\}d\hat{F}_a(\beta_a)\right\}$$

The seller's payoff in the equilibrium above when the buyer can communicate is

$$p^+\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^+ + (\beta_a) + p^-\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^-(\beta_a),$$

where $p^+ \equiv \mathbb{E}[\beta_a|\beta_a \geq \frac{1}{2}]$, $p^- \equiv \mathbb{E}[\beta_a|\beta_a < \frac{1}{2}]$ and $\hat{F}_a^+$, $\hat{F}_a^-$ are the conditional distributions of $\hat{F}_a$ above and below 1/2.

If

$$v\left((1,\frac{1}{2}),\hat{F}_a\right) \geq v\left((1,0),\hat{F}_a\right),$$

then the difference between the payoff in the equilibrium when the buyer can communicate and the equilibrium when he cannot is

$$\begin{aligned}
&\left(p^+\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^+ + (\beta_a) + p^-\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^-(\beta_a)\right) - \left(\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a(\beta_a)\right)\\
&= p^-\int_0^{\frac{1}{2}}\left(1-\frac{1}{2}\beta_a\right) - \left(\frac{1}{2}+\frac{1}{2}\beta_a\right)d\hat{F}_a^-(\beta_a)\\
&= p^-\int_0^{\frac{1}{2}}\left(\frac{1}{2}-\beta_a\right)d\hat{F}_a^-(\beta_a)\\
&> 0,
\end{aligned}$$

where the final inequality follows from the fact that $\beta_a \leq \frac{1}{2}$ for all $\beta_a$ and there is a positive mass of $\beta_a$ for which this holds with a strict inequality.

If

$$v\left((1,\frac{1}{2}),\hat{F}_a\right) < v\left((1,0),\hat{F}_a\right),$$

then the difference between the payoff in the equilibrium when the buyer can communicate and the equilibrium when he cannot is

$$\begin{aligned}
&\left(p^+\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^+ + (\beta_a) + p^-\int \beta_a + \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^-(\beta_a)\right) - \left(\int \max\left\{\beta_a, 1-\beta_a\right\}d\hat{F}_a(\beta_a)\right)\\
&= p^+\int \frac{1}{2}\left(1-\beta_a\right)d\hat{F}_a^-(\beta_a) + p^-\int \frac{1}{2}\beta_a d\hat{F}_a^-(\beta_a)\\
&> 0.
\end{aligned}$$

Therefore, the equilibrium is a beneficial conversation equilibrium.

When $\Pr[\beta_a = \frac{1}{2}] > 0$, there is a seller preferred equilibrium in which there are 3 messages from the buyer to seller:

- the buyer sends the message $m_1^b$ if $\beta_a > \frac{1}{2}$, $m_2^b$ if $\beta_a < \frac{1}{2}$, and $m_{\frac{1}{2}}^b$ if $\beta_a = \frac{1}{2}$;
- following the message $m_j^b$, $j = 1, 2$, the seller sends the message $m_1^s$ if $\theta_{1j} = 1$ and $m_2^s$ if $\theta_{1j} = 0$; and following the message $m_{\frac{1}{2}}^b$ with probability half the seller sends the message $m_{11}^s$ if $\theta_{11} = 1$ and $m_{12}^s$ if $\theta_{11} = 0$, and with probability half the seller sends the message $m_{21}^s$ if $\theta_{21} = 1$ and $m_{22}^s$ if $\theta_{21} = 0$.

Following the reasoning above, it is straightforward to verify that this is an equilibrium, and that the seller's payoff is the same as the the equilibrium above meaning that it is a seller preferred equilibrium.

When $\Pr[\beta_a = \frac{1}{2}] = 0$, the equilibrium above can be replaced with the equilibrium in the proposition. Since A.3 is strictly lower than A.4 the (seller preferred) equilibrium must take the form in the proposition. Furthermore, all types $\beta_a \neq \frac{1}{2}$ have a strict incentive to choose their specified strategy. Thus the equilibrium is unique. $\square$