

Habibi, Amir

Working Paper

Communicating Preferences to Improve Recommendations

Discussion Paper, No. 394

Provided in Cooperation with:

University of Munich (LMU) and Humboldt University Berlin, Collaborative Research Center Transregio 190: Rationality and Competition

Suggested Citation: Habibi, Amir (2023) : Communicating Preferences to Improve Recommendations, Discussion Paper, No. 394, Ludwig-Maximilians-Universität München und Humboldt-Universität zu Berlin, Collaborative Research Center Transregio 190 - Rationality and Competition, München und Berlin

This Version is available at:

<https://hdl.handle.net/10419/282086>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Communicating Preferences to Improve Recommendations

Amir Habibi (HU Berlin)

Discussion Paper No. 394

May 10, 2023

Communicating Preferences to Improve Recommendations

Amir Habibi*

April 11, 2023

Abstract

I study a cheap talk model between a buyer and a seller with two goods for sale. There is two-sided (independent) private information with sequential, two-way communication. In the first stage, the buyer communicates her private preferences to the seller. In the second stage, the seller communicates the quality of the goods to the buyer. When the buyer's preference is about which attribute common to both goods she prefers, the seller strictly benefits from the buyer communicating her preferences. Whereas when the buyer's preference is about which good she prefers, this is never the case.

Keywords: cheap talk, strategic communication, product recommendations

JEL Classification: D82, L15

1 Introduction

Recommendations widely used in e-commerce. Before making a recommendation, should a seller try and elicit a consumer's private information? Consider the following situation. A consumer is considering buying a new phone and is deciding between the latest model from two different brands on offer from a seller. If she does not purchase either of the phones, she can continue to use her current

*Humboldt University of Berlin. Email: amir.habibi@hu-berlin.de. I am especially grateful to Martin Cripps, Nathan Hancart, Dilip Ravindran and Roland Strausz for a number of helpful conversations. I also want to thank Valeria Burdea, Antonio Cabrales, Philippe Jehiel, Sebastian Schweighofer-Kodritsch, Peter Sørensen, Peter Spittal, Dezső Szalay and seminar audiences in Munich (CRC TRR 190 Retreat 2021), HU Berlin, UCL and Milan (ESEM 2022) for their helpful comments. I am also grateful for financial support from the German Research Foundation (DFG) through CRC TRR 190 (Project number 280092119).

phone for which she has a (private) value.¹ The seller is incentivised to make a sale—prices are fixed and he gets a fixed profit if the consumer buys either of the two phones and zero profit if the consumer does not make a purchase. The seller privately knows the quality of each of the new phones. The consumer’s valuation for each phone is a combination of the quality and her (privately known) preferences. The seller can make a recommendation to influence the consumer’s beliefs about quality. For example, he could say that brand X’s new phone is better than brand Y’s—such a comparative statement makes it more likely that the consumer purchases brand X and less likely that she purchases brand Y. However, he is not able to provide hard evidence—so communication is only by cheap talk.

Before getting a recommendation, is it helpful for the seller to allow the consumer to communicate her preferences (again by cheap talk)? This could be done, for example, by allowing the consumer to type into a search bar. So a search for ‘brand X phones’ would indicate she is interested only in brand X phones. Learning the consumer’s preferences may hurt the seller because he is no longer able to make credible recommendations. This happens if the consumer reveals that she is only interested in one of the two brands on offer. Now the seller cannot credibly communicate any information about the quality of that particular brand. However, communicating preferences may also be beneficial to the seller since it allows him to make a recommendation that is more useful for the consumer. This happens when the quality of the phones consist of two attributes—e.g., camera quality and battery life. Suppose the seller chooses an attribute and makes a comparative statement, for example saying that brand X has the better camera. Now if he wants to make a comparative statement about the battery he is biased towards saying brand X has the better battery. In this respect he is limited in what he can credibly recommend for the second attribute. This creates an opportunity for the seller to benefit from communication by the consumer. If the consumer says which attribute she is most interested in, the seller can make a recommendation for the phone that is best for that particular attribute. This recommendation is more helpful for the consumer (than a recommendation for a random attribute) and ultimately increases the chance of a sale—which clearly benefits the seller.

I analyse a stylised model of the interaction described above between a buyer and a seller with two goods for sale. The main results formalise the intuition above and characterize the seller optimal equilibrium. When the buyer’s private information is about which of the two goods she prefers, there is never any benefit for the seller if the buyer is able to communicate her preferences by cheap talk.

¹Throughout this article I use male pronouns for the seller and female pronouns for the buyer.

This is in contrast to the case where the buyer’s private information is about which attribute common to both goods she prefers. Here there is an equilibrium in which the buyer can communicate about her preferences and in this equilibrium the seller obtains a strictly higher payoff than the best equilibrium in a setting in which the buyer cannot communicate. Furthermore, in the best equilibrium for the seller, communication always takes a simple form with the buyer indicating her preferred attribute, and the seller recommending the best good for that attribute and revealing no information about the other attribute.

The intuition for the results are as follows. When the buyer’s private information is about which good she prefers, to maximize the likelihood of a sale, the seller wants to provide unbiased information by making a recommendation that reveals all his private information—the true ranking of the goods. In this sense the buyer and seller are aligned. However, the seller’s lack of commitment—since he must communicate by cheap talk—means he will want to pander given his belief about the buyer’s preferences. This means that typically he cannot credibly fully communicate his private information. Instead he makes a biased recommendation. So, to avoid the seller pandering, the buyer always wants to try and make the seller believe she values both goods equally. However, the fact that the buyer wants to do this regardless of her true preference means she, herself, cannot communicate credibly. On the other hand, when the buyer’s private information is about her relative preference across attributes, she can communicate which attribute she is most interested in, thus allowing the seller to provide an unbiased recommendation for that attribute. Note that the seller could still make a recommendation without learning anything about the buyer’s preferences. Such a recommendation could even credibly communicate (some) information about both attributes. However, this would be less effective and ultimately lead to a lower likelihood of a sale.

To solve the model, I make use of the ‘securization’ tools developed in Lipnowski and Ravid (2020) who study an abstract cheap talk game where the sender (seller) has state-independent preferences (as in my setting). Their results allow me to find the seller’s maximum payoff from communicating with the buyer given a belief he holds about the buyer’s preferences. This intermediate step is necessary to solve my model in which the buyer communicates about her preferences before the seller communicates.²

²Their paper also analyses the same buyer-seller set-up as an example to illustrate their results. However, they only consider one way communication (from seller to buyer). 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

The contribution of the paper is two-fold. First, my theoretical model provides insights about online markets and in particular recommender systems. Second, my model contributes to the theoretical cheap talk literature by considering a novel setting: two sided (independent) private information with sequential, two-way communication.

One reason a seller may want to elicit a consumer’s preferences is to overcome search costs (Varian (2002)). A simple example is that if a consumer wants to buy a new phone and visits a website selling electronics, it is beneficial for her to indicate that she is searching for a phone. Doing so means that the website can display phones and not other goods she is not interested in this reduces her search costs and increases the probability of a sale. However, assuming that there are no such search frictions to overcome—so the consumer can observe all the relevant goods on offer—it is unclear if it is beneficial for the seller to elicit further information.

To help users navigate a wide range of products, e-commerce employs recommender systems.³ Existing research on recommender systems, particularly outside the economics literature, typically do not consider credibility of such recommendations—that sellers may bias recommendations towards the most profitable products. Instead they focus on how the system uses its information to make the ‘best’ recommendation to fit the buyers preferences. When consumers are aware that sellers may bias their recommendation, this affects their strategic responses. They will take this into account both when interpreting recommendations and also what they communicate about their own preferences before receiving a recommendation. In the cheap talk literature this is described as ‘pandering’ (Che et al. (2013)). My model explicitly considers such pandering and provides insights into the strategy of a firm when providing recommendations. How does pandering affect the firm’s recommendations? When should it try and elicit consumer preferences? And when would it be damaging to do so because its recommendations become less credible?

Turning to theory, there has been little focus on cheap talk models with two-way sequential communication which is how many economic interactions happen. This is emphasised by the following quotation in a recent survey of the literature: ‘Economic models of communication have little to say about real conversations – dynamic exchanges in which people take turns.’—Sobel (2013). I analyse a model with two sided (independently drawn) private information and sequential,

Chakraborty and Harbaugh (2010).

³See Aggarwal (2016) for a recent textbook that details the various forms of algorithms used.

two-way communication. Within the environment I study, I identify when, in equilibrium, both players communicate information that influences the decision taken and when this improves the sender’s (seller’s) payoffs. To the best of my knowledge, this type of cheap talk model has not been studied thus far. 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)). Although there is also a growing literature with multiple rounds of cheap talk, these primarily focus on models with one sided private information and one-way communication or simultaneous two-way communication.⁴ My model also makes use of the securitization tools of Lipnowski and Ravid (2020), and to the best of my knowledge it is the first paper to make extensive use of them and apply them in the context of pandering. I discuss the related theoretical literature in Section 6.

The remainder of the paper is organized as follows. Section 2 introduces the model. Section 3 discusses equilibrium selection and how to apply the methodology of Lipnowski and Ravid (2020) to my model. Section 4 analyzes the model when the good only has a single attribute and buyer’s private information is about which of the two goods she prefers. Section 5 analyzes the model when the good has two attributes and buyer’s private information is about which attribute she prefers. Section 6 discusses the related literature. Section 7 provides some discussion and concludes.

2 Model

In this section I describe the model and then discuss the assumptions. The main parts of the paper focus on two cases which specialise the model. In Section 4, I analyse when the buyer’s private information is about which good she prefers; and in Section 5, I analyse when the buyer’s private information is about which attribute she prefers. To ease exposition and streamline analysis, for the main part of the analysis (Sections 4 and 5) I use the specific functional form for utility given below. In Section 7, I discuss how the results extend to more general set-ups.

⁴For example, in Aumann and Hart (2003), sequential cheap talk is referred to as ‘polite talk’ to contrast it with simultaneous two-way cheap talk which is the main focus on their paper. Krishna and Morgan (2004) demonstrates that two-way communication can improve outcomes in the standard Crawford and Sobel (1982) setting, but this relies on simultaneous communication, and given there is one-sided private information, information is only transmitted in one direction.

2.1 Model set-up

Players, information and payoffs. There is a buyer (she) and a seller (he).⁵ The buyer faces a choice over two goods and an outside option. Denote the buyer's action by $a \in A \equiv \{a_0, a_1, a_2\}$, where a_0 represents taking her outside option and a_1 and a_2 buying the respective good.

Each good has two attributes. Quality is negatively correlated across goods for each attribute. Formally, quality is determined by two random variables $(\theta_1, \theta_2) \in \Theta \equiv \{0, 1\}^2$. For each attribute $j = 1, 2$: $\theta_j = 1$ ($\theta_j = 0$) means that good 1 has high (low) quality for that attribute and good 2 has low (high) quality. The buyer has a preference parameter given by $\beta = (\beta_g, \beta_a) \in [0, 1]^2$. As will become clear below, these represent relative preferences across goods and across attributes respectively. The buyer's payoff depends on her preference parameter, the quality of the attributes of the goods, and her outside option $u_0 \in [0, 1]$:

$$U = \begin{cases} u_1(\theta, \beta) \equiv \beta_g (\beta_a \theta_1 + (1 - \beta_a) \theta_2) & \text{if } a = a_1, \\ u_2(\theta, \beta) \equiv (1 - \beta_g) (\beta_a (1 - \theta_1) + (1 - \beta_a) (1 - \theta_2)) & \text{if } a = a_2, \\ u_0 & \text{if } a = a_0. \end{cases}$$

Here β_g represents the relative preference *across goods*—higher values of β_g mean a stronger preference for good 1; and β_a represents the relative preference *across attributes*—higher values of β_a indicate a stronger preference for attribute 1.

Quality θ_1 and θ_2 are drawn independently and identically with $\Pr[\theta_j = 1] = \Pr[\theta_j = 0] = \frac{1}{2}$, for $j = 1, 2$. The buyer's preferences β are drawn from a distribution F , with β_g and β_a having marginal distributions given by F_g and F_a . In Section 4, I analyse the case where there is uncertainty on β_g and F_a is degenerate; and in Section 5, I analyse the case where there is uncertainty on β_a and F_g is degenerate. The buyer's outside option is drawn from $u_0 \sim U[0, 1]$.

The seller's payoff is *state independent*—it simply depends on whether or not the buyer buys one of the goods:

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

⁵In 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'.

Timing. The timing of the game is as follows:

1. The buyer privately learns the realisation of β , and the seller privately learns the realisation of θ ;
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_0 ;⁶
5. The buyer takes an action, $a \in A \equiv \{a_0, a_1, a_2\}$: 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.

Strategies. The buyer's strategy is to choose i) a messaging strategy that maps her preference to a message $m^b : [0, 1]^2 \rightarrow \Delta\mathcal{M}^b$ and ii) an action strategy that maps her preferences, her message, and the message of the seller to a choice over goods: $a : [0, 1]^2 \times \mathcal{M}^b \times \mathcal{M}^s \rightarrow A$.⁷ The seller's strategy is to choose a messaging strategy that maps the state θ and the buyer's message to a message: $m^s : \{0, 1\}^2 \times \mathcal{M}^b \rightarrow \Delta\mathcal{M}^s$. I refer to the seller's strategy as an (information) policy.

Beliefs. The seller updates his belief over β 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 θ to $\mu(m^b, m^s) = (\mu_1(m^b, m^s), \mu_2(m^b, m^s))$ where $\mu_1(m^b, m^s) \equiv \Pr[\theta_1 = 1 | m^b, m^s]$ and $\mu_2(m^b, m^s) \equiv \Pr[\theta_2 = 1 | m^b, m^s]$.

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.⁸

⁶In Section 7 I discuss how results are unchanged if this happens at step 1. The current timing helps ease exposition.

⁷I restrict the buyer's action to pure strategies to ease notation, but this restriction does not affect the analysis in any substantive way. This is due to the continuously distributed outside option u_0 .

⁸Note 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.

2.2 Discussion of model

Multi-product seller. I have assumed that the seller has two goods for sale, this is the simplest model that allows for credible recommendations. With only a single good there is no opportunity for influential cheap talk communication from seller to buyer—this insight follows from Chakraborty and Harbaugh (2010).

Buyer’s preferences. I have assumed a specific, but natural, functional form for the buyer’s preferences.⁹ The formulation I have employed imposes three negative correlations: i) perfect negative correlation in preferences across attributes, ii) perfect negative correlation in preferences across goods, and iii) perfect negative correlation in attribute quality across goods. This means that there are only two quality parameters and two preference parameters to keep track of allowing me to focus attention on the key economic forces. In Section 7, I consider a more general functional form for u_1 and u_2 and demonstrate that this does not affect the results when preferences are across goods (Section 4).¹⁰ The perfect negative correlation for horizontally differentiated goods is a simplifying assumption is used elsewhere in the cheap talk literature (Inderst and Ottaviani (2012) and Chung and Harbaugh (2019)) and in the information design literature (Armstrong and Zhou (2022)). The assumptions also mean that there is no aggregate uncertainty on the quality of the goods. In Section 7, I discuss how introducing aggregate uncertainty in the quality does not affect results when preferences are across goods (Section 4).

Hard versus soft information. I have assumed that the buyer’s information is ‘soft’, so she must communicate by cheap talk. In contrast, some information, for example whether to allow a website to use cookies is ‘hard’ information. In practice a consumer has both types of information. I do not explicitly model disclosing hard information, but my set-up can be thought of as having the hard information already being revealed and the buyer choosing whether to communicate the remaining soft information.

⁹An alternative formulation would be a ‘standard’ random utility model with

$$u_i = \beta_1 \theta_{1i} + \beta_2 \theta_{2i} + \epsilon_i.$$

For example, such a utility function is used in a related paper Chakraborty and Harbaugh (2014).

¹⁰Specifically, I consider a more general function $u_1(u_2)$ which is increasing (decreasing) in both β_g and θ , (weakly) convex in the first argument, and is such that the two arguments are compliments.

Prices. I have assumed that there are no pricing decisions for the seller. In many situations sellers cannot price discriminate, but can personalised recommendations given to an individual buyer. This happens both in e-commerce and with a salesperson in a ‘brick and mortar’ store.

3 Preliminaries

3.1 Equilibrium selection

As in all cheap talk games, there will typically be multiple equilibria.¹¹ To select an equilibrium, consistent with much of the literature, I use the seller preferred equilibrium. The equilibrium selection is discussed further in Section 7.

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

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, β . In order to formalise this I define a property of the equilibrium:

Definition 2. *Persuasive equilibrium with buyer communication:* *An equilibrium in which the seller gets a strictly higher payoff compared to a (seller preferred) equilibrium where the message space of the buyer is restricted to a single message: $|\mathcal{M}^b| = 1$.*

I am interested in whether the (seller preferred) equilibrium is a persuasive equilibrium with buyer communication—this is an equilibrium in which the seller gets a strictly higher payoff compared to a game where the buyer is not able to communicate.¹²

3.2 Solving the seller’s problem using the ‘Securability Theorem’

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

¹¹Existence 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.

¹²Or equivalently, to an alternative equilibrium where in the first round of communication the buyer chooses an uninformative message (a babbling equilibrium).

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).¹³

Define $v(\mu, \hat{F})$ as the seller's expected payoff for a given buyer posterior belief $\mu \in \Delta\Theta$ and belief that the buyer has preferences $\beta \sim \hat{F}$.¹⁴ Let p be an information policy, and s to be some possible seller payoff. A policy p *secures* s if $\Pr\{\mu : v(\mu, \hat{F}) \geq s \mid p\} = 1$, 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 at least 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.*

The policy p that secures s need not be an equilibrium policy. The intuition behind the result is that if such a (non-equilibrium) policy p secures s , then the seller strictly prefers some message over the message that obtains payoff s . However, the value of the preferred message can be lowered to s by adding states that are closer to the prior. This is demonstrated in Example 1 below.

Note that the theorem does not provide any information about what the seller's optimal policy is. 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 equilibrium payoffs. If an (equilibrium) policy that achieves the highest securable payoff is found then this is clearly in the set of seller preferred policies.

4 One attribute: Buyer's preferences are across goods

In this section the buyer has private information about her relative preference across goods. I maintain the following assumption throughout this section.

Assumption 1 (Buyer uncertainty is about preference over goods). F_a is degenerate such that $\Pr[\beta_a = 1] = 1$.

¹³As 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.

¹⁴The seller anticipates the buyer's best response given her belief. This is used to compute the seller's expect payoff.




Figure 1: The seller's value in Example 1 as a function of his belief in gray. The blue dots are the two posterior beliefs from each message and the corresponding value/prob of sale in the equilibrium policy.

This means that the uncertainty about the buyer's preferences can only be about β_g . The buyer's preferences simplify to:

$$\begin{aligned} u_1(\theta, \beta) &= \beta_g \theta_1, \\ u_2(\theta, \beta) &= (1 - \beta_g)(1 - \theta_1). \end{aligned}$$

Note that with this utility function, the assumption on the prior that $\Pr[\theta_1 = 1] = \frac{1}{2}$ is without loss. If this is relaxed, then β_g can be re-weighted accordingly.

4.1 Seller communication

I start by considering the communication from the seller to the buyer, taking the seller's beliefs about the buyer's preferences as given. First, I build intuition with a simple example where there is no uncertainty on how much the buyer values each good. Then, I provide a lemma that characterises the seller's value given any belief over the buyer's preferences. I use this to construct the seller's (unique) information policy that gives this payoff. Finally, I discuss the intuition for the seller's behavior that occurs in equilibrium and in particular how he panders towards one good.

Example 1 (No uncertainty on buyer preferences). F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$. This means there is no uncertainty on the preferences of the buyer. Clearly, given that the buyer does not have any private information, there

is no persuasive equilibrium with buyer communication. However, this will be a useful benchmark to analyse. With no communication, the buyer will value good 1 more than good 2 and obtains utility $u_1 = \frac{3}{10}$ if she chooses a_1 . This means she buys a good with probability $\frac{3}{10}$, which is the seller's payoff. The seller can improve his payoff by making a recommendation which influences the buyer's beliefs. Suppose she sent a message fully revealing the value of θ_1 , so

$$m^s = \begin{cases} m_1^s & \text{if } \theta_1 = 1, \\ m_2^s & \text{if } \theta_1 = 0. \end{cases}$$

If the buyer updated her beliefs given the information policy above, she buys good i when she received the message m_i^s . This means she buys a good with probability $\frac{1}{2} \times \frac{3}{5} + \frac{1}{2} \times \frac{2}{5} = \frac{1}{2}$. However, notice that this is not an equilibrium because regardless of the value of θ_1 , the seller strictly prefers to send the message m_1^s and pander towards the good that the buyer prefers. Although this is not an equilibrium, the policy above secures a payoff of $\frac{2}{5}$ which is strictly greater than when there is no communication. To obtain this payoff in equilibrium, the seller degrades the value of sending m_1^s , by sending this message when $\theta_1 = 0$ with some probability. He does this so that $\Pr[\theta_1 = 1 | m_1^s] = \frac{2}{3}$. This means that the buyer buys with a lower probability of $\frac{2}{5}$. In Figure 1 I depict the seller's value (or probability of sale) as a function of the posterior belief of the buyer. Below I also prove that $\frac{2}{5}$ is the seller's highest payoff and that this is uniquely achieved with the equilibrium policy above.

Now I consider general distributions F_g with uncertainty over β_g . Denote the seller's belief that the buyer's preferences over goods β_g are distributed by \hat{F}_g . Let the buyer's belief over θ_1 be given by μ_1 where $\mu_1 \equiv \Pr[\theta_1 = 1 | m^b, m^s]$. The seller's value function is:

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

Note that $\max \{ \beta_g \mu_1, (1 - \beta_g)(1 - \mu_1) \}$ 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 $\mu_1 = 1$ (as, for example, is the case in Figure 1).

If

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

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

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

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

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

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

$$v \left(\frac{1}{2}, \hat{F}_g \right).$$

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

Given the analysis above, the seller's value is summarised in the following lemma:

Lemma 1. *If the seller has a belief that β_g has distribution \hat{F}_g and the buyer has a belief μ_0 over θ_1 , then the seller's expected payoff in equilibrium is*

$$\hat{v}(\mu_0, \hat{F}_g) = \max \left\{ \begin{array}{l} v \left(\frac{1}{2}, \hat{F}_g \right), \\ \min \left\{ v \left(1, \hat{F}_g \right), v \left(0, \hat{F}_g \right) \right\} \end{array} \right\}.$$

As a corollary, returning to Example 1, where $\hat{F}_g = F_g$ is degenerate at $\beta_g = \frac{3}{5}$, it is confirmed that the seller obtains a value of $\frac{2}{5}$ which can be uniquely achieved by the policy described—uniqueness follows from the argument below.

Now I construct an equilibrium seller policy for a general distribution \hat{F}_g . In what follows I assume that $\hat{\beta}_g \equiv \mathbb{E}_{\hat{F}_g}[\beta_g] \geq \frac{1}{2}$.¹⁵ When

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

the seller does not provide any information. Whereas when

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

The policy will take a similar form to the one in Example 1. The message space

¹⁵When $\hat{\beta}_g < \frac{1}{2}$ there is an analogous policy with the messages switched around.

is $\mathcal{M}^s = \{m_1^s, m_2^s\}$ and the probability of sending each message given θ_1 is

$$\begin{aligned}\Pr[m^s = m_1^s | \theta_1 = 1] &= 1, \\ \Pr[m^s = m_1^s | \theta_1 = 0] &= \frac{1 - \bar{\mu}_1}{\bar{\mu}_1},\end{aligned}$$

where $\bar{\mu}_1$ is chosen to ensure the seller is indifferent between sending each message when $\theta_1 = 0$. This induces posterior probabilities

$$\begin{aligned}\Pr[\theta_1 = 1 | m^s = m_1^s] &= \bar{\mu}_1, \\ \Pr[\theta_1 = 1 | m^s = m_2^s] &= 0.\end{aligned}$$

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 $\theta_1 = 0$, $\bar{\mu}_1$ satisfies the following equation:

$$v(0, \hat{F}_g) = v(\bar{\mu}_1, \hat{F}_g). \quad (4.1)$$

The behavior of the seller in equilibrium can be summarised as follows. If the prior on the buyer's preferences are sufficiently skewed towards one good, then there is no communication. If this is not the case then the seller panders towards the good that the buyer is more likely to be interested in. In particular, when the seller learns that this good has high quality he *always* recommends this good. When he learns that this good has low quality, he still sometimes recommends this good. This means that when receiving a recommendation for the good that she knows she was more likely to be interested in, the buyer discounts this recommendation since she knows that the buyer was pandering towards this good. The concept of pandering in cheap talk is well known in the literature (Che et al. (2013)).¹⁶ The results in this section demonstrate how the securitization tools in Lipnowski and Ravid (2020) can be applied to study pandering in cheap talk games with state independent preferences. Next, I use these results to analyse the full game where the buyer can communicate before the seller makes a recommendation.

¹⁶Che et al. (2013) consider a model where the sender (seller) has state dependent preferences. Closer to my model, Chung and Harbaugh (2019) consider a model where pandering occurs with state independent preferences.

4.2 Buyer communication

Now I consider the first stage of communication—from the buyer to seller. First, I analyse the buyer’s preferences over different information policies. Then, I provide the main result of this section (Proposition 1). This shows quite generally that the equilibrium is not a persuasive equilibrium with buyer communication—meaning that the seller never strictly benefits from the buyer communicating about her preferences.

In order to analyse the buyer’s incentives, it is necessary to obtain the buyer’s payoff given a preference β_g and an information policy $\bar{\mu}_1$ from the seller. First, I define

$$I(x) \equiv x + \frac{1}{2}(1 - x)^2 = \frac{1}{2}(1 + x^2) \quad (4.2)$$

as the buyer’s expected payoff (before learning her outside option) when the valuation of the more valuable good is x .

For a buyer with preference $\beta_g \in [0, 1]$ and for a policy with $\bar{\mu}_1 \in [\frac{1}{2}, 1]$ (and continuing to assume that $\hat{\beta}_g \geq \frac{1}{2}$) the buyer’s expected payoff is given by

$$u(\bar{\mu}_1, \beta_g) = \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1} I(1 - \beta_g) + \frac{1}{2\bar{\mu}_1} I(\bar{\mu}_1 \beta_g).$$

This can be simplified to

$$u(\bar{\mu}_1, \beta_g) = \frac{1}{2} + \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1} (1 - \beta_g)^2 + \frac{1}{2} \bar{\mu}_1 \beta_g^2. \quad (4.3)$$

The buyer’s utility is strictly increasing in $\bar{\mu}_1$ for any β_g . This is intuitive, higher $\bar{\mu}_1$ gives her better information with the best policy always being $\bar{\mu}_1 = 1$ 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 all the information from the seller’s recommendation. Next I state the main result of this section.

Proposition 1. *Under Assumption 1 (so buyer uncertainty is about preference over goods), the (unique seller preferred) equilibrium is never a persuasive equilibrium with buyer communication.*

Formal proofs are all in the Appendix. The main steps of the proof and intuition are summarised as follows: First, I show that there cannot be more than one message sent in equilibrium which leads to expected beliefs either all

¹⁷A similar expression can be obtained for $\hat{\beta}_g < \frac{1}{2}$.

above or all below $\frac{1}{2}$. If this was the case clearly they would need to induce the same informational policy (summarised by $\bar{\mu}_1$)—if not the buyer would have a strict incentive to send only one of the messages. I show that when all such messages induce the same informational policy, these messages can be replaced by a single message and not affect the seller’s payoff. Second, I show that in the seller preferred equilibrium there cannot be two messages sent where one message leads to an expected belief above $\frac{1}{2}$ and the other to an expected belief below $\frac{1}{2}$. The reason is that an equilibrium constructed by combining these two messages into a single message leads to an expected belief closer to $\frac{1}{2}$ and on average results in more information being communicated by the seller—and thus a higher payoff for the seller.

5 Two attributes: Buyer’s preferences are across attributes

In this section the buyer has private information about her relative preference across two different attributes common to both goods. I maintain the following assumption throughout this section.

Assumption 2 (Buyer uncertainty is about preference over attributes).
 F_g is degenerate such that $\Pr[\beta_g = \frac{1}{2}] = 1$.

This means that the uncertainty about the buyer’s preferences can only be about β_a . The buyer’s preferences simplify to:

$$\begin{aligned} u_1(\theta, \beta) &= \frac{1}{2}(\beta_a \theta_1 + (1 - \beta_a) \theta_2), \\ u_2(\theta, \beta) &= \frac{1}{2}(\beta_a (1 - \theta_1) + (1 - \beta_a) (1 - \theta_2)). \end{aligned}$$

Note that the assumptions mean that ex ante the buyer has an equal valuation for the two goods. I discuss relaxing this assumption at the end of this section and provide a formal result in Appendix B.

As in Section 4, I begin with some simple examples to build intuition. First, I illustrate the link to the previous analysis by considering when there is no uncertainty on β_a . Then, I consider when there is uncertainty on β_a , but only with ‘extreme’ values $\beta_a \in \{0, 1\}$. Here there is no tradeoff between recommendations on the two attributes and the seller can still fully reveal the state. Next, I introduce an intermediate type $\beta_a = \frac{1}{2}$ that creates a friction in the seller’s ability to communicate and illustrates the value of the buyer’s communication. Finally, I

present the main result of this section (Proposition 2). This result shows that quite generally, the equilibrium is a persuasive equilibrium with buyer communication, and that it always takes a very simple form.

Example 2 (No uncertainty on the preferences of the buyer). F_a satisfies the following: $\Pr[\beta_a = \frac{1}{2}] = 1$. As in Example 1, this means there is no uncertainty on the preferences of the buyer. Furthermore, what matters is the total value across both attributes: $\theta_1 + \theta_2$. With no information, the buyer values both goods equally with utility $\frac{1}{4}$. To find the optimal policy, consider the value function of the seller depicted in Figure 2. This is plotted in the two dimensional space below with the two axis being $\mu_1 \equiv \Pr[\theta_1 = 1]$ and $\mu_2 \equiv \Pr[\theta_2 = 1]$.




Figure 2: Value function $(v((\mu_1, \mu_2), \hat{F}_a))$ for Example 2.

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

$$m^s = \begin{cases} m_1^s & \text{if } \theta_1 = 1, \\ m_2^s & \text{if } \theta_1 = 0. \end{cases}$$

The two posterior beliefs are indicated in Figure 2 by the blue dots. Effectively, this recommends the best good for attribute 1, and provides no information for attribute 2. I verify below that there is no policy that secures a higher payoff.

Next I provide a lemma—a generalisation of Lemma 1 in the one attribute case—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.




Figure 3: The blue dots represent the posteriors from the policy that secures $v\left((1, \frac{1}{2}), \hat{F}_a\right)$.

Lemma 2. *For any posterior belief over β_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 2, the regions where the seller achieves a strictly higher payoff are in the right and left corners—these correspond to one good being better for both attributes, $(\mu_1, \mu_2) = (0, 0)$ or $(1, 1)$. 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 3 secures this payoff and is also an equilibrium: it recommends the best good for attribute 1.¹⁸

When

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

¹⁸The 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 $v\left((1, 0), \hat{F}_a\right)$.

the policy depicted in Figure 4 secures this payoff, however apart from when β_a only takes extreme values (as will be the case in Example 3), this is not an equilibrium. The policy completely reveals the state.

A typical example of an equilibrium policy is depicted in Figure 5. This policy recommends the best policy for attribute 1. For attribute 2 it makes a recommendation biased towards recommending the good that was not recommended for attribute 1.

Returning to Example 2, the lemma verifies that the seller has a payoff $\frac{3}{8}$ in equilibrium. Next, I continue to build intuition by going through examples where there is uncertainty on β_a .

Example 3 (Buyer has only ‘extreme’ preferences). F_a 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. However, the (seller preferred) equilibrium is not a persuasive equilibrium with buyer communication. The seller’s optimal policy is to recommend best good for each attribute—meaning he fully reveals the state (as in Figure 4). 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 $v\left((1, 0), \hat{F}_a\right) = \frac{1}{2}$. This is the payoff achieved by the policy of fully revealing the state.




Figure 5: The blue dots represent the posteriors from the equilibrium policy that secures $v\left((1, 0), \hat{F}_a\right)$. This completely reveals attribute 1, and partially reveals attribute 2.

In Example 3 there is no benefit from the buyer communicating her preferences before receiving the recommendation—the recommendation already fully reveals the state of the world. 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 introduce a final example with a preference-type ($\beta_a = \frac{1}{2}$), who values both goods.

Example 4 (Buyer has a richer set of preferences). F_a satisfies the following: $\beta_a \in \{0, \frac{1}{2}, 1\}$, and $\Pr[\beta_a = 0] = \Pr[\beta_a = 1] = p \in (0, \frac{1}{2})$. Under this assumption, the equilibrium will be a persuasive equilibrium with buyer communication. The key idea is that, compared to Example 3, 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 Example 3 where there were only extreme preferences. 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 more effective recommendation that improves the seller's payoff.

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

Assumption 3. *The support of F_a is not contained in either $[0, \frac{1}{2}]$, $[\frac{1}{2}, 1]$ or $\{0, 1\}$.*

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 Example 3, where the buyer has extreme preferences and she is interested only in one of attributes.

Proposition 2. *If the distributional Assumption 3 is satisfied, then all (seller preferred) equilibria are persuasive equilibria with buyer communication. 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_j^b , the seller sends the message m_1^s if $\theta_j = 1$ and m_2^s if $\theta_j = 0$.*

Furthermore, the equilibrium above is unique if and only if $\Pr[\beta_a = \frac{1}{2}] = 0$.

If Assumption 3 is not satisfied then no equilibrium is a persuasive equilibrium with buyer communication.

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_j^b 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.¹⁹ 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

¹⁹Note that in Example 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 $\beta_a = 1$.

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.²⁰ Of the two attributes, the seller clearly benefits more from revealing information about the more favoured attribute. Once he has fully revealed about that attribute, he is completely biased on the other attribute—he wants to recommend the same good 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 seller benefits from the buyer learning 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 in Example 4 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. More formally, I make use of Lemma 2. I show that the maximum payoff the seller can achieve if he could choose any information structure (distribution over posterior beliefs) over β_a is precisely the one described above—simply whether β_a is above or below $\frac{1}{2}$. Clearly, if this information structure corresponds to an equilibrium—meaning it is incentive compatible for the buyer to report her type β_a truthfully—then this is the seller preferred equilibrium.

This can be seen graphically in Figure 6. Here I depict the seller’s value from the different beliefs in Lemma 2 for all possible values of β_a . Recall that the the beliefs in Lemma 2 give the maximum value the seller can achieve from communication. Notice how for $\beta_a \in [\frac{1}{2}, 1]$ the value function $v((1, \frac{1}{2}), \beta_a)$ gives the highest value for the seller; and for $\beta_a \in [0, \frac{1}{2}]$ this is reversed so $v((\frac{1}{2}, 1), \beta_a)$ gives the highest value for the seller.

Under Assumption 3 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 persuasive equilibrium with buyer communication. To see why Assumption 3 is necessary for the equilibrium to be a persuasive equilibrium with buyer communication, consider the cases that it rules out. First,

²⁰Note that given the buyer’s preferences, this is better for the seller than revealing the ‘best’ good overall.




Figure 6: Value functions: blue $v((1, \frac{1}{2}), \beta_a)$, orange $v((\frac{1}{2}, 1), \beta_a)$, green $v((1, 0), \beta_a)$.

there is the case as Example 3 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 buyer is always interested in the same attribute and so the equilibrium is not a persuasive equilibrium with buyer communication—she always sends the same message.

The assumption that the buyer values both goods equally has simplified the analysis in this section. It is natural to ask how the results extend when this assumption is relaxed. In Appendix B I consider what happens when the buyer values one good more (but maintain the assumption that there is only uncertainty on the buyer's preferences across attributes). I show that if F_a is such that the support is sufficiently 'close' to $\beta_a = \frac{1}{2}$ (i.e. ruling out extreme preferences) and that she does not have a strong preference towards either good, then the seller preferred equilibrium takes a similar form to Proposition 2. The buyer sends a message that indicates which attribute she is most interested in. Then the seller makes a recommendation based on this information. However, the recommendation is slightly more nuanced than before. The seller recommends the best good for the attribute the buyer communicated about with a bias away from the good that the buyer prefers. The seller's recommendation also communicates partial information about the buyer's less preferred attribute. The strategy is described more formally in Appendix B.

6 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). LR apply the securitization tools they develop for more general state-independent cheap talk games to find the sender (seller) optimal equilibrium in the buyer-seller game. As discussed above, I also make use of these tools in the setting I study. Chakraborty and Harbaugh (2014) build on the 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.

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.

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)). A 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.²¹ A recent paper that has two way

²¹Moreno de Barreda (2013) and Lai (2014) also have two sided private information where

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.²² A number of other earlier papers point out that with one-way, one-shot cheap talk communication, if the receiver (seller) has private information this may facilitate communication where otherwise it would not be possible—see Seidmann (1990) and Watson (1996).

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.

7 Discussion and concluding remarks

7.1 State space and preferences

In order to make progress with the novel communication protocol I am interested in, I have considered a specific setting where the sender/seller has state independent preferences allowing me to leverage the securitization tools of LR. I have also assumed the simplest possible form for the state space of θ . These assumptions allow me to provide clear conditions under which both players benefit from the decision maker/buyer communicating before the sender/seller.

In the one attribute case (Section 4), the results remain unchanged with aggregate uncertainty. Suppose that the model is changed such that $\theta_1 = (\theta_g, \theta_q) \in \{0, 1\} \times [0, 1]$, where the first component represents the good with the best quality as before, and the second component, drawn independently, is the quality of the best good. It is straightforward that the seller cannot communicate anything in the second dimension since she would always want to inflate this as much as possible. For distributions over more general state spaces Θ , the model becomes less tractable because it is challenging to characterise the optimal policy.²³

information is correlated.

²²In 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'.

²³For example, in the analysis of the buyer-seller example in LR, they focus on the case where the goods are symmetrically distributed. 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

Furthermore, in the one attribute case the main result (Proposition 1) remain unchanged if more general preferences are considered. These preferences and the accompanying analysis are in Appendix B. Intuitively these still ensure that all buyer types want the seller to believe that the buyer has no bias towards either good—such a belief induces the seller to fully reveal θ .

One further natural consideration is the possibility that the seller might have a preference over which good he is able to sell—for example, he gets paid a different commission for each good. If this preference is public, then as long as there is not a large difference in the preference towards one good, the results will remain qualitatively unchanged—it is similar to changing the preference of the buyer towards each good, β_g . However, in reality, it might be the case that the seller’s preferences are privately known—as a consumer, one might be unsure what commission a salesperson gets for selling a specific good. In Chakraborty and Harbaugh (2010), such a possibility is discussed in their Online Appendix. They show that their results—that there exists an informative equilibrium in the one way communication game—are robust under ‘almost certain motives’. Loosely speaking, this means that the prior the buyer has over the seller’s ‘type’ is almost degenerate. I do not model this formally within my framework, but expect that my results would remain qualitatively unchanged with a small amount of uncertainty in the seller’s payoffs.

7.2 Buyer preferred equilibrium

I have focused on the seller preferred equilibrium throughout. This facilitated analysis by allowing me to leverage the tools of LR. This also makes sense for the applications. In online interactions between a buyer and an e-commerce site, the seller typically determines the form of communication, for example by prompting search queries. Even in an off-line setting, a traditional salesperson elicits information from the buyer by approaching her in his store.

Despite this, it may be of interest to know when the seller preferred equilibrium is also the buyer preferred equilibrium. Although both the buyer and the seller want ‘more’ information to be transmitted from seller to buyer in the final stage of communication, it is not necessarily the case that seller preferred equilibrium is also the buyer preferred one. This is driven by the fact the seller has linear preferences over information whereas the buyer’s preferences are convex—and prefers ‘riskier’ prospects. In the one attribute case, given that the seller optimal policy is.

is going to pander, the buyer would prefer this is towards the good she is more interested in. This means that although the seller preferred equilibrium is pooling (i.e. all buyer types pool together), the buyer preferred equilibrium could be a separating equilibrium where the buyer reveals information about β_g .²⁴ In the model where the buyer's preferences are across attributes (Section 5), consider F_a such that $\beta_a \sim U\{\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon\}$ where $\epsilon \in (0, \frac{1}{2}]$. The seller preferred equilibrium always has the buyer communicate her β_a —as proved in Proposition 2. For small values of ϵ , this equilibrium is also the buyer preferred equilibrium. However, for $\epsilon \approx \frac{1}{2}$, the buyer prefers an equilibrium in which the buyer does not communicate (babbles) and the seller randomly chooses one of the two attributes and fully reveals the state for that one attribute and reveals almost all the information about the other attribute as well.²⁵ This creates risk which the buyer prefers to only revealing information about the attribute she is most interested in.

7.3 Alternative communication protocols and commitment

I make the assumption that u_0 is learned after communication. Suppose instead that this is privately learned by the buyer at the same time she learns β . It is straightforward that the seller cannot achieve a higher payoff in an equilibrium in which information about u_0 is revealed as part of the message m^b . Suppose that she could and that there were messages m^b that revealed information about u_0 (as well as potentially about β). Consider the seller's optimal policy. Following the logic of Lemma 1 and 2, with a different belief over u_0 , the value that the seller can secure may be different, but the policy that he secures will remain the same. Given that the seller's policy is unchanged (or equivalently depends only on his belief over β), his expected payoff across all messages is the same as if no information about u_0 was revealed by the buyer. Note that this argument would no longer hold if u_0 was not drawn independently of β .

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), Aumann and Hart (2003) and Krishna and Morgan (2004)). Furthermore, one

²⁴Such an equilibrium would have two messages with messages inducing a posterior $\hat{\beta}_g$ above and below $\frac{1}{2}$.

²⁵In the limit when $\epsilon = \frac{1}{2}$, then the seller reveals all the information as in Example 3 and the buyer is indifferent between this and the equilibrium where she communicates and only learns about her preferred attribute.

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. Furthermore, note that as shown in the proofs of Proposition 2, the seller cannot improve his payoff if he could commit to an information structure over the buyer’s private information. However, in future work it would be interesting to understand to what extent payoffs can be increased with more general protocols, and what form a more complex ‘conversation’ takes with two sided private information.²⁶

References

- Aggarwal, C. C. (2016). *Recommender systems*, volume 1. Springer.
- 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*.
- Armstrong, M. and Zhou, J. (2022). Consumer Information and the Limits to Competition. *American Economic Review*, 112(2):534–577.
- Aumann, R. J. and Hart, S. (2003). Long Cheap Talk. *Econometrica*, 71(6):1619–1660.
- 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.
- Che, Y.-K., Dessein, W., and Kartik, N. (2013). Pandering to Persuade. *American Economic Review*, 103(1):47–79.
- Chen, Y. (2009). Communication with Two-sided Asymmetric Information. *working paper*.

²⁶Forges (1990) asks such a question in a specific cheap talk game with one sided private information.

- Chung, W. and Harbaugh, R. (2019). Biased recommendations from biased and unbiased experts. *Journal of Economics & Management Strategy*, 28(3):520–540.
- 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.
- 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.
- Inderst, R. and Ottaviani, M. (2012). Competition through Commissions and Kickbacks. *American Economic Review*, 102(2):780–809.
- 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.
- Lai, E. K. (2014). Expert advice for amateurs. *Journal of Economic Behavior & Organization*, 103:1–16.
- Lipnowski, E. and Ravid, D. (2020). Cheap Talk With Transparent Motives. *Econometrica*, 88(4):1631–1660.
- Moreno de Barreda, I. (2013). Cheap Talk With Two-Sided Private Information. *working paper*.
- Myerson, R. B. (1986). Multistage Games with Communication. *Econometrica*, 54(2):323–358.
- Seidmann, D. J. (1990). Effective cheap talk with conflicting interests. *Journal of Economic Theory*, 50(2):445–458.

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.

Watson, J. (1996). Information Transmission When the Informed Party Is Confused. *Games and Economic Behavior*, 12(1):143–161.

Appendix A Proofs

A.1 Proof of Proposition 1

Proof. Assume throughout that $\mathbb{E}_F[\beta_g] \geq \frac{1}{2}$. If $\mathbb{E}_F[\beta_g] < \frac{1}{2}$, a very similar argument can be made. Given a preference-type β_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 characterised by $\bar{\mu}_1$ (in equation 4.1).

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 .

First, 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 $\bar{\mu}_1 = \bar{\mu}'_1$.²⁷ 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' . If $\bar{\mu}_1 = \bar{\mu}'_1 = \frac{1}{2}$ (i.e. the seller sends an uninformative message following both messages from the buyer) it is straightforward that this must be the case. If

²⁷Note that this does not mean that it must be that $\hat{\beta}_g = \hat{\beta}'_g$. Suppose F_g is uniform over $\beta_g \in \{\frac{1}{4}, \frac{9}{10}\}$, with expectation $\hat{\beta}_g = \frac{23}{40}$. This results in a policy with $\bar{\mu}_1 = \frac{2}{3}$. The degenerate distribution with $\beta_g = \frac{3}{5}$ has a different expectation, but results in the same policy $\bar{\mu}'_1 = \frac{2}{3}$.

$\bar{\mu}_1 = \bar{\mu}'_1 > \frac{1}{2}$, 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 $\hat{\beta}_g \geq \frac{1}{2}$. A similar argument means that there must be at most one message played with $\hat{\beta}_g < \frac{1}{2}$.

The analysis above does not rule out that there may be one message played with $\hat{\beta}_g \geq \frac{1}{2}$ and one with $\hat{\beta}_g < \frac{1}{2}$. 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 $\hat{\beta}_g \geq \frac{1}{2}$ and $\hat{\beta}'_g < \frac{1}{2}$ with respective distributions \hat{F}_g and \hat{F}'_g . Assume that in both cases this leads to the seller sending an informative message, i.e. $\bar{\mu}_1 > \frac{1}{2}$. 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 β_g . Denote the probability that in the original equilibrium, m is played by p and m' by $1 - p$. Since, by assumption, $p\hat{\beta}_g + (1 - p)\hat{\beta}'_g = \mathbb{E}_F[\beta_g] \geq \frac{1}{2}$, in the new equilibrium the seller's payoff from the policy with message m'' is given by

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

In contrast, the expected payoff in the original equilibrium is

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

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. Returning to rewrite equation 4.1 as

$$1 - \hat{\beta}_g = \int_{1-\bar{\mu}_1}^1 \beta_g \bar{\mu}_1 d\hat{F}_g(\beta_g) + \int_0^{1-\bar{\mu}_1} (1 - \beta_g)(1 - \bar{\mu}_1) d\hat{F}_g(\beta_g) \equiv R(\bar{\mu}_1). \quad (\text{A.3})$$

It must be that

$$R(\tfrac{1}{2}) < 1 - \hat{\beta}_g < R(1) = \hat{\beta}_g,$$

the first inequality from the fact $v(0, \hat{F}_g) > v(\tfrac{1}{2}, \hat{F}_g)$ and the second from the fact $\hat{\beta}_g > 1/2$. The intermediate value theorem guarantees existence, and the solution to A.3 is unique since R is convex.

From this it can be shown that there is a unique policy pinned down by $\bar{\mu}_1$. To verify this, differentiating the RHS of A.3 gives

$$\begin{aligned} & \frac{\partial}{\partial \bar{\mu}_1} \left[\int_{1-\bar{\mu}_1}^1 \beta_g \bar{\mu}_1 d\hat{F}_g(\beta_g) + \int_0^{1-\bar{\mu}_1} (1 - \beta_g)(1 - \bar{\mu}_1) d\hat{F}_g(\beta_g) \right] \\ &= \int_{1-\bar{\mu}_1}^1 \beta_g d\hat{F}_g(\beta_g) - \int_0^{1-\bar{\mu}_1} (1 - \beta_g) d\hat{F}_g(\beta_g) \\ &> 0. \end{aligned}$$

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 A.3 is strictly increasing in $\bar{\mu}_1$ and so by the Intermediate Value Theorem, equation A.3 has a unique solution. □

A.2 Proof of Lemma 2

Proof. Assume throughout that $\hat{\beta}_a \equiv \mathbb{E}_{\hat{F}_a}[\beta_a] \geq \tfrac{1}{2}$. This means that $v\left((1, \tfrac{1}{2}), \hat{F}_a\right) \geq v\left((\tfrac{1}{2}, 1), \hat{F}_a\right)$. When $\hat{\beta}_a \leq 1/2$, $v\left((1, \tfrac{1}{2}), \hat{F}_a\right) \geq v\left((\tfrac{1}{2}, 1), \hat{F}_a\right)$, 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 $\pi_{ij}^k \in [0, 1]$ the probability that message $m_k^s \in \{m_1^s, m_2^s\}$ is sent in the state $\theta_1 = i \in \{0, 1\}$ and in the state $\theta_2 = j \in \{0, 1\}$. Bayes plausibility requires that $\pi_{ij}^1 + \pi_{ij}^2 = 1$ for all i, j . Furthermore, the total probability of message m_k^b being sent is $\pi^k = \pi_{11}^k + \pi_{10}^k + \pi_{01}^k + \pi_{00}^k$ 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 with binary support secures a payoff of $v\left((1, \frac{1}{2}), \hat{F}_a\right)$: $\pi_{00}^1 = \pi_{01}^1 = 0$ and $\pi_{10}^1 = \pi_{11}^1 \in (0, 1]$. In words, this policy reveals nothing about attribute 2 and reveals information about attribute 1: if good 1 has high quality for attribute 1 this is learned perfectly, while if good 2 has high quality for attribute 1 this is learned imperfectly.

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

I assume that $v\left((\frac{1}{2}, 1), \hat{F}_a\right) \geq v\left((\frac{1}{2}, 0), \hat{F}_a\right)$, meaning that $v\left((1, \frac{1}{2}), \hat{F}_a\right) \geq v\left((\frac{1}{2}, 0), \hat{F}_a\right)$. If the assumption in the first inequality is reversed, the same argument as below can be made, but for the set with $\mu_2 \geq \frac{1}{2}$ rather than $\mu_1 \geq \frac{1}{2}$.

Denote by \bar{M} , the set of buyer posterior beliefs (μ_1, μ_2) where the seller obtains a (weakly) lower payoff than the secured payoff $v\left(1, \frac{1}{2}\right)$ and the posterior beliefs are such that $\mu_1 \geq \frac{1}{2}$:

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

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

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

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

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

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

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

The seller's value function is given by:

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

Step 1. Consider lines where μ_2 is fixed for some $\mu_2 \in [0, \frac{1}{2}]$ and μ_1 takes values from $\frac{1}{2}$ to 1. For $\mu_1 \in [\frac{1}{2}, 1]$, since v is the integral over a convex function, the maximum in this range is either $\mu_1 = \frac{1}{2}$ or $\mu_1 = 1$. Furthermore, since $v(1, \frac{1}{2}) > v(1, 0)$, for $\mu_2 \in [0, \frac{1}{2}]$, $v(1, \mu_2)$ is decreasing in μ_2 ; and for $\mu_2 \in [0, \frac{1}{2}]$,




Figure 7: Set \bar{M} in (μ_1, μ_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 $v\left((1, \frac{1}{2}), \hat{F}_a\right)$.

$v\left(\frac{1}{2}, \mu_2\right)$ is increasing in μ_2 . Combining these we have that $v\left(1, \frac{1}{2}\right) \geq v(\mu_1, \mu_2)$ for all $(\mu_1, \mu_2) \in [\frac{1}{2}, 1] \times [0, \frac{1}{2}]$.

Step 2. To show that for any $(\mu_1, \mu_2) \in [\frac{1}{2}, 1] \times [\frac{1}{2}, 1]$, $v\left((1, \frac{1}{2}), \hat{F}\right) < v\left((\mu_1, \mu_2), \hat{F}\right)$ only if $\mu_1 + \mu_2 > \frac{3}{2}$, first observe that for any $(\mu_1, \mu_2) \in [\frac{1}{2}, 1] \times [\frac{1}{2}, 1]$, the buyer will choose good 1 regardless of her preference type β_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 $v\left((1, \frac{1}{2}), \hat{F}_a\right)$ 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

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

It is straightforward that $\mu_1 + \mu_2 > \frac{3}{2}$ is a necessary condition for this to be satisfied. \square

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} : $\mu(m_1^b), \mu(m_2^b) \subseteq \bar{M}$. Clearly it cannot be the case that both posteriors either have $\mu_1 + \mu_2 > \frac{3}{2}$ or $\mu_1 + \mu_2 < \frac{1}{2}$. So, have $\mu(m_1^b)$ such that $\mu_1 + \mu_2 > \frac{3}{2}$; and $\mu(m_2^b)$ such that $\mu_1 + \mu_2 < \frac{1}{2}$. Calculating the posterior beliefs in terms of π_{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 $\mu(m_1^b), \mu(m_2^b) \subseteq \bar{M}$, these must satisfy

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

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 $\mu_1 + \mu_2 > \frac{3}{2}$ or $\mu_1 + \mu_2 < \frac{1}{2}$. 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:

$$\begin{aligned} m_1^s & \text{ if } \theta_1 = \theta_2 = 1, \\ m_2^s & \text{ if } \theta_1 = 1, \theta_2 = 0, \\ m_3^s & \text{ if } \theta_1 = 0, \theta_2 = 1, \\ m_4^s & \text{ if } \theta_1 = \theta_2 = 0. \end{aligned}$$

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

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

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 3 is necessary for the equilibrium to be a persuasive equilibrium with buyer communication. 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 β_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 $I\left(\frac{1+\beta_a}{2}\right)$,²⁸ 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 β_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 $1/2$. 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

$$(\bar{p}_+ + \bar{p}_-) v(\mu, \hat{F}_a(\bar{m}^b)) = (\bar{p}_+ + \bar{p}_-) \max \left\{ v\left((1, 0), \hat{F}_a(\bar{m}^b)\right), v\left((1, \frac{1}{2}), \hat{F}_a(\bar{m}^b)\right) \right\}. \quad (\text{A.4})$$

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

²⁸Recall $I(\cdot)$ is defined in 4.2.

the buyer types playing \bar{m}^b is

$$(\bar{p}_+ + \bar{p}_-) \bar{p}_+ \frac{1}{4} \left(1 + \hat{\beta}_a^+\right) + \bar{p}_- \frac{1}{4} \left(2 - \hat{\beta}_a^-\right). \quad (\text{A.5})$$

By comparing A.4 to A.5, 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 3, the equilibrium is a persuasive equilibrium with buyer communication. 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

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

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

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

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

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

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

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

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

If

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

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

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

Therefore, the equilibrium is a persuasive equilibrium with buyer communication.

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 $\theta_j = 1$ and m_2^s if $\theta_j = 0$; and following the message $m_{\frac{1}{2}}^b$ with probability half the seller sends the message m_{11}^s if $\theta_1 = 1$ and m_{12}^s if $\theta_1 = 0$, and with probability half the seller sends the message m_{21}^s if $\theta_2 = 1$ and m_{22}^s if $\theta_2 = 0$.

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.4 is strictly lower than A.5 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

Appendix B Extensions

B.1 One attribute model with more general buyer preferences

In this section I consider more general preference for the buyer within Section 4 and show that the main result of that section (Proposition 1) remains unchanged.

Specifically, Proposition 1 continues to hold if the preferences satisfy the following assumptions:²⁹

$$\begin{array}{llll} \frac{\partial u_1}{\partial \theta_1} > 0, & \frac{\partial^2 u_1}{\partial \theta_1^2} \geq 0, & \frac{\partial u_1}{\partial \beta_g} > 0, & \frac{\partial^2 u_1}{\partial \theta_1 \partial \beta_g} > 0, \\ \frac{\partial u_2}{\partial \theta_1} < 0, & \frac{\partial^2 u_2}{\partial \theta_1^2} \leq 0, & \frac{\partial u_2}{\partial \beta_g} < 0, & \frac{\partial^2 u_2}{\partial \theta_1 \partial \beta_g} > 0. \end{array}$$

I explain how the analysis in the main text and the proof of Proposition 1 can be adapted to these more general preferences.

First, note that

$$v(\mu, \hat{F}_g) = \mathbb{E}_\mu \int_{\beta_g} \max \{u_1(\theta_1, \beta_g), u_2(\theta_1, \beta_g)\} d\hat{F}_g(\beta_g).$$

is still convex.³⁰ This means that the equivalent of Lemma 1 can be obtained. It follows that given a buyer belief \hat{F}_g , the seller has a unique policy $\bar{\mu}_1$ as before.

Moving onto buyer preferences, for any β_g , the buyer's preferences are still strictly increasing in $\bar{\mu}_1$. This can be seen from the adapted expression in 4.2, which is now:

$$u(\bar{\mu}_1, \beta_g) = \frac{2\bar{\mu}_1 - 1}{2\bar{\mu}_1} I(u_2(0, \beta_g)) + \frac{1}{2\bar{\mu}_1} I(u_1(\bar{\mu}_1, \beta_g)).$$

The first part of the proof of Proposition 1—where different messages were played by types such that the expected value of the seller's posterior were greater than $\frac{1}{2}$ —follows very closely to the existing proof. Next consider the second part—where different messages were played by types such that the expected value of the seller's posterior was above and below $\frac{1}{2}$ for different messages. Here the RHS of equations A.1 and A.2 are exactly as before, and the result follows from the

²⁹A previous version of this paper incorrectly stated this result without requiring that u_1 and u_2 were (weakly) convex in their first arguments. I thank an anonymous referee for pointing out this error.

³⁰This follow from the assumption that u_1 is (weakly) convex in the first argument.

convexity of $v(0, F_g)$.

B.2 Two attribute model with a preference towards one good

In this section I extend the model in Section 5 allow a preference towards one good. More precisely the buyer's utility is now

$$\begin{aligned} u_1(\theta, \beta) &= \beta_g(\beta_a\theta_1 + (1 - \beta_a)\theta_2), \\ u_2(\theta, \beta) &= (1 - \beta_g)(\beta_a(1 - \theta_1) + (1 - \beta_a)(1 - \theta_2)). \end{aligned}$$

F is such that β_g is degenerate with $\beta_g \neq \frac{1}{2}$ and there is uncertainty on β_a . Without loss I assume that $\beta_g < \frac{1}{2}$. I also restrict the support of F such that both β_a and β_g are sufficiently close to $\frac{1}{2}$. It will become clear below why not allowing 'extreme' preferences facilitates the analysis.

As before, I start with the seller's problem given a belief \hat{F} following communication from the buyer. Lemma 2 needs to be adapted since it relied on $\beta_g = \frac{1}{2}$.

Lemma 4. *For any posterior belief over β_a , \hat{F}_a , the maximum payoff the seller can secure is*

$$\hat{v}(\mu_0, \hat{F}_a) = \max \left\{ v \left((1, \tfrac{1}{2} + \delta), \hat{F}_a \right), v \left((\tfrac{1}{2} + \delta', 1), \hat{F}_a \right), v \left((1, 0), \hat{F}_a \right), v \left((0, 1), \hat{F}_a \right) \right\},$$

for some δ and $\delta' \in (0, \frac{1}{2})$.

The difference between this result and Lemma 2 is that there are different values of (μ_1, μ_2) generating the maximum payoff the seller can secure—these differ from before by δ and δ' .

The proof of Lemma 2 can be adapted to show this result. In Figure 8, I depict the regions where the seller can get a higher payoff than the secured payoff in the lemma. Note that because of the restriction to 'non-extreme' values of β , in the upper right region the buyer will always buy good 1 and in the lower left good 2. Using a very similar argument to the proof of Lemma 2, there is no information policy that is strictly contained within the two regions \bar{M} .

Next, following a similar reasoning to the proof of Proposition 2, if the seller could choose an information structure over buyer's preferences, he would just learn whether β_a is above or below $\frac{1}{2}$. The logic is exactly in Proposition 2. I illustrate the equivalent figure below in Figure 9. Note that because of the restriction to values of β_a close to $\frac{1}{2}$ it means that either $v \left((1, \frac{1}{2} + \delta), \hat{F}_a \right)$ or $v \left((\frac{1}{2} + \delta', 1), \hat{F}_a \right)$




Figure 8: Set \bar{M} in (μ_1, μ_2) belief space. The blue dots represent the posteriors from the policy that secures $v\left(1, \frac{1}{2} + \delta\right), \hat{F}_a$.




Figure 9: Value functions: blue $v\left((1, \frac{1}{2} + \delta), \beta_a\right)$, yellow $v\left((\frac{1}{2} + \delta', 1), \beta_a\right)$, orange $v\left((1, 0), \beta_a\right)$, green $v\left((0, 1), \beta_a\right)$.

are optimal for the seller. Bringing these adaptations of Lemma 2 and Proposition 2 together gives the result in Proposition 3 below.

Proposition 3. *There exists an $\epsilon > 0$ such that if the support of F is such that $(\beta_a, \beta_g) \in [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]^2$, then the following result holds. If F_g is degenerate with a bias towards good 2 ($\beta_g < \frac{1}{2}$), and if Assumption 3 is satisfied then all (seller preferred) equilibria are beneficial buyer communication equilibria. 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_j^b , the seller sends one of two message m_1^s and m_2^s .
If $\theta_j = 1$ and $\theta_{-j} = 1$ the seller sends the message m_1^s ; if $\theta_j = 1$ and $\theta_{-j} = 0$

the seller mixes between the two messages; if $\theta_j = 0$ the seller sends m_2^s . The mixture probability is such that the posterior beliefs generated are as in Lemma 4 above.

Furthermore, the equilibrium above is unique if and only if $\Pr[\beta_a = \frac{1}{2}] = 0$.

If Assumption 3 is not satisfied then no equilibrium is a persuasive equilibrium with buyer communication.