Disclosing Preferences to Improve Recommendations

Amir Habibi (Humboldt University of Berlin)

August 2022

Introduction: Models of communication

'Economic models of communication have little to say about real conversations— dynamic exchanges in which people take turns.'
$\sim$ **Joel Sobel.**

Introduction: Models of communication

'Economic models of communication have little to say about real conversations— dynamic exchanges in which people take turns.'
$\sim$ **Joel Sobel.**

- A typical cheap talk game:

1. Message, $m \in \mathcal{M}$

Seller/Sender/Expert — Buyer/Receiver/DM

Private info: θ — 2. Action: $a \in A$

Payoff: $V(a, \theta)$ — Payoff: $U(a, \theta)$

Introduction: Back and forth cheap talk

- A modified game:

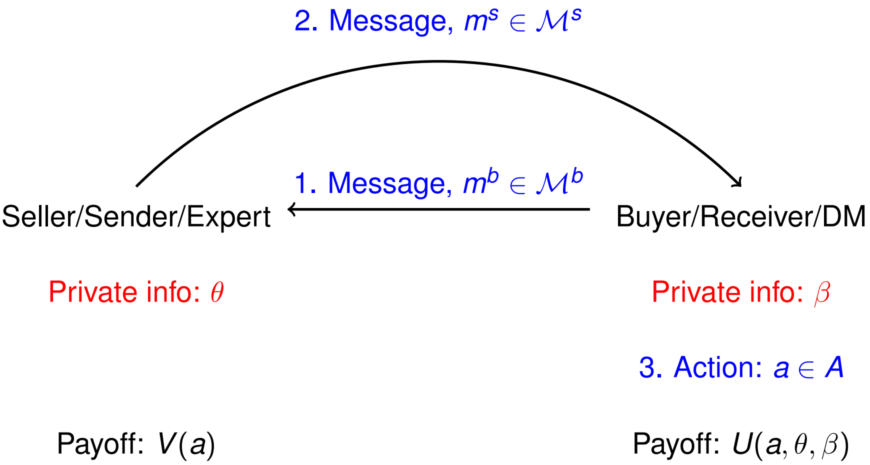

Model set up

Players

- A buyer/receiver/DM (she)
- A seller/sender/expert (he)

Model set up

Players

- A buyer/receiver/DM (she)
- A seller/sender/expert (he)

Information/states of the world

- There are two goods

Model set up

Players

- A buyer/receiver/DM (she)
- A seller/sender/expert (he)

Information/states of the world

- There are two goods
- The quality of goods is determined by a random variable $\theta \in \Theta$

Model set up

Players

- A buyer/receiver/DM (she)
- A seller/sender/expert (he)

Information/states of the world

- There are two goods
- The quality of goods is determined by a random variable $\theta \in \Theta$
- Buyer has a preference parameter given by $\beta \in B$

Model set up

Players

- A buyer/receiver/DM (she)
- A seller/sender/expert (he)

Information/states of the world

- There are two goods
- The quality of goods is determined by a random variable $\theta \in \Theta$
- Buyer has a preference parameter given by $\beta \in B$
- Players share a common prior, $\theta \sim G$ and $\beta \sim F$

Model set up

Model set up

Actions and timing.

1. The buyer privately learns the realisation of β, and the seller privately learns the realisation of θ

Model set up

Actions and timing.

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

Model set up

Actions and timing.

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

Model set up

Actions and timing.

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 $u_0 \sim U[0, 1]$

Model set up

Actions and timing.

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 $u_0 \sim U[0, 1]$
5. The buyer takes an action, $a \in \{a_0, a_1, a_2\}$: her outside option (a_0) or one of the two goods (a_1) and (a_2)

Model set up

Actions and timing.

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 $u_0 \sim U[0, 1]$
5. The buyer takes an action, $a \in \{a_0, a_1, a_2\}$: her outside option (a_0) or one of the two goods (a_1) and (a_2)
6. The players get their payoffs and the game ends

Model set up

Payoffs.

- The buyer's payoff:

$$U = \begin{cases} u_1(\theta, \beta) & \text{if } a = a_1 \\ u_2(\theta, \beta) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

Model set up

Payoffs.

- The buyer's payoff:

$$U = \begin{cases} u_1(\theta, \beta) & \text{if } a = a_1 \\ u_2(\theta, \beta) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

- The seller's payoff is **state-independent**:

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

Model set up

Equilibrium.

- **Solution concept**: perfect Bayesian equilibrium

Model set up

Equilibrium.

- **Solution concept**: perfect Bayesian equilibrium
- **Seller preferred equilibrium**: An equilibrium which maximises the seller's expected utility among the set of possible equilibrium payoffs

Model set up

Equilibrium.

- **Solution concept**: perfect Bayesian equilibrium
- **Seller preferred equilibrium**: An equilibrium which maximises the seller's expected utility among the set of possible equilibrium payoffs
- **Beneficial conversation equilibrium**: 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: $|\mathcal{M}^b| = 1$

Model set up

Equilibrium.

- **Solution concept**: perfect Bayesian equilibrium
- **Seller preferred equilibrium**: An equilibrium which maximises the seller's expected utility among the set of possible equilibrium payoffs
- **Beneficial conversation equilibrium**: 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: $|\mathcal{M}^b| = 1$

Question: When is there a beneficial conversation equilibrium?

Single attribute: Set-up

Single attribute: Set-up

Information and payoffs

- $\theta \in \{0, 1\}$, with $\Pr[\theta = 1] = 1/2$

Single attribute: Set-up

Information and payoffs

- $\theta \in \{0, 1\}$, with $\Pr[\theta = 1] = 1/2$
- $\beta = \beta_g \in [0, 1]$, with distribution F_g

Single attribute: Set-up

Information and payoffs

- $\theta \in \{0, 1\}$, with $\Pr[\theta = 1] = 1/2$
- $\beta = \beta_g \in [0, 1]$, with distribution F_g

$$U = \begin{cases} \beta_g \theta & \text{if } a = a_1 \\ (1 - \beta_g)(1 - \theta) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

Single attribute: Set-up

Information and payoffs

- $\theta \in \{0, 1\}$, with $\Pr[\theta = 1] = 1/2$
- $\beta = \beta_g \in [0, 1]$, with distribution F_g

$$U = \begin{cases} \beta_g \theta & \text{if } a = a_1 \\ (1 - \beta_g)(1 - \theta) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

- $\theta = 1$ means that good 1 has high quality and good 2 has low quality

Single attribute: Set-up

Information and payoffs

- $\theta \in \{0, 1\}$, with $\Pr[\theta = 1] = 1/2$
- $\beta = \beta_g \in [0, 1]$, with distribution F_g

$$U = \begin{cases} \beta_g \theta & \text{if } a = a_1 \\ (1 - \beta_g)(1 - \theta) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

- $\theta = 1$ means that good 1 has high quality and good 2 has low quality
- β_g represents the preference across goods

Single attribute: Simple example

Single attribute: Simple example

- F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$

Single attribute: Simple example

- F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$
- Suppose the seller used an information policy fully revealing θ

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

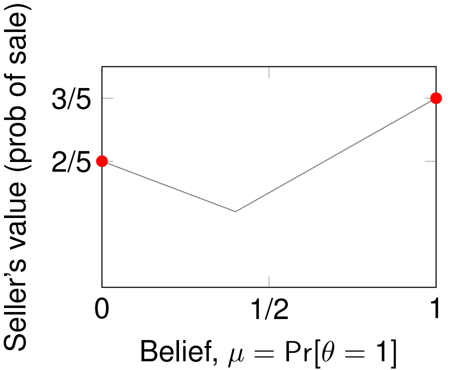

Single attribute: Simple example

- F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$

Single attribute: Simple example

- F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$
- Fully revealing θ is not an equilibrium

Single attribute: Simple example

- F_g satisfies the following: $\Pr[\beta_g = \frac{3}{5}] = 1$
- Fully revealing θ is not an equilibrium
- The following information policy is an equilibrium:

$$\Pr[\theta = 1 | m_1^s] = 2/3$$
$$\Pr[\theta = 1 | m_2^s] = 0$$

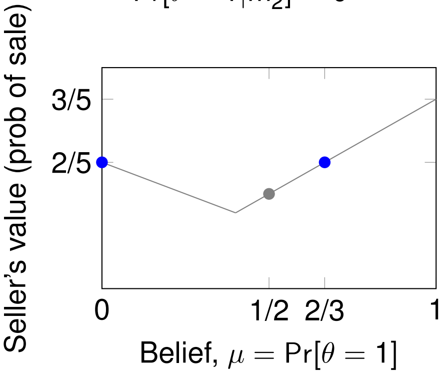

Single attribute: Securability

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)
- To achieve this in equilibrium, the seller can degrade the value from the posterior that achieves a higher payoff

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)
- To achieve this in equilibrium, the seller can degrade the value from the posterior that achieves a higher payoff
- Lipnowski and Ravid (2020) provide a general tool to find possible equilibrium payoffs in state-independent cheap talk games

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)
- To achieve this in equilibrium, the seller can degrade the value from the posterior that achieves a higher payoff
- Lipnowski and Ravid (2020) provide a general tool to find possible equilibrium payoffs in state-independent cheap talk games
  - make use of this in my model to solve the seller's problem

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)
- To achieve this in equilibrium, the seller can degrade the value from the posterior that achieves a higher payoff
- Lipnowski and Ravid (2020) provide a general tool to find possible equilibrium payoffs in state-independent cheap talk games
  - make use of this in my model to solve the seller's problem
  - can find (seller preferred) equilibrium payoff, and then find the seller's policy that achieves this

Single attribute: Securability

- In the example, the seller could secure a payoff of $2/5$ by fully revealing the state
  - a payoff is secured if it is the lowest payoff across messages sent (posterior beliefs induced)
- To achieve this in equilibrium, the seller can degrade the value from the posterior that achieves a higher payoff
- Lipnowski and Ravid (2020) provide a general tool to find possible equilibrium payoffs in state-independent cheap talk games
  - make use of this in my model to solve the seller's problem
  - can find (seller preferred) equilibrium payoff, and then find the seller's policy that achieves this
  - then can consider buyer incentives for communicating her preferences

Single attribute: Results

Proposition 1

With a single attribute, the (unique seller preferred) equilibrium is never a beneficial conversation equilibrium.

Single attribute: Results

Proposition 1

With a single attribute, the (unique seller preferred) equilibrium is never a beneficial conversation equilibrium.

- The buyer always wants to make the seller think that she has no preference for either good ($\beta_g = 1/2$)
  - doing so would mean the seller would fully reveal θ

Single attribute: Results

Proposition 1

With a single attribute, the (unique seller preferred) equilibrium is never a beneficial conversation equilibrium.

- The buyer always wants to make the seller think that she has no preference for either good ($\beta_g = 1/2$)
  - doing so would mean the seller would fully reveal θ
- This means the buyer cannot credibly disclose her preferences

Two attributes: Set-up

Information and payoffs

- $\theta = (\theta_1, \theta_2) \in \{0, 1\}^2$, with θ_1 and θ_2 drawn independently with $\Pr[\theta_i = 1] = 1/2$

Two attributes: Set-up

Information and payoffs

- $\theta = (\theta_1, \theta_2) \in \{0, 1\}^2$, with θ_1 and θ_2 drawn independently with $\Pr[\theta_i = 1] = 1/2$
- $\beta = \beta_a \in [0, 1]$, with distribution F_a

Information and payoffs

- $\theta = (\theta_1, \theta_2) \in \{0, 1\}^2$, with θ_1 and θ_2 drawn independently with $\Pr[\theta_i = 1] = 1/2$
- $\beta = \beta_a \in [0, 1]$, with distribution F_a

$$U = \begin{cases} \frac{1}{2}\big(\beta_a\theta_1 + (1 - \beta_a)\theta_2\big) & \text{if } a = a_1 \\ \frac{1}{2}\big(\beta_a(1 - \theta_1) + (1 - \beta_a)(1 - \theta_2)\big) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

Two attributes: Set-up

Information and payoffs

- $\theta = (\theta_1, \theta_2) \in \{0, 1\}^2$, with θ_1 and θ_2 drawn independently with $\Pr[\theta_i = 1] = 1/2$
- $\beta = \beta_a \in [0, 1]$, with distribution F_a

$$U = \begin{cases} \frac{1}{2}\big(\beta_a\theta_1 + (1-\beta_a)\theta_2\big) & \text{if } a = a_1 \\ \frac{1}{2}\big(\beta_a(1-\theta_1) + (1-\beta_a)(1-\theta_2)\big) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

- $\theta_j = 1$ means that for attribute j, good 1 has high quality and good 2 has low quality

Two attributes: Set-up

Information and payoffs

- $\theta = (\theta_1, \theta_2) \in \{0, 1\}^2$, with θ_1 and θ_2 drawn independently with $\Pr[\theta_i = 1] = 1/2$
- $\beta = \beta_a \in [0, 1]$, with distribution F_a

$$U = \begin{cases} \frac{1}{2}\big(\beta_a\theta_1 + (1-\beta_a)\theta_2\big) & \text{if } a = a_1 \\ \frac{1}{2}\big(\beta_a(1-\theta_1) + (1-\beta_a)(1-\theta_2)\big) & \text{if } a = a_2 \\ u_0 & \text{if } a = a_0 \end{cases}$$

- $\theta_j = 1$ means that for attribute j, good 1 has high quality and good 2 has low quality
- β_a represents the preference across attributes

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

- $\beta_a \in \{0, 1\}$ with $\Pr[\beta_a = 1] = p \in [0, 1]$

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

- $\beta_a \in \{0, 1\}$ with $\Pr[\beta_a = 1] = p \in [0, 1]$
- In equilibrium the seller fully reveals the state

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

- $\beta_a \in \{0, 1\}$ with $\Pr[\beta_a = 1] = p \in [0, 1]$
- In equilibrium the seller fully reveals the state
    - having revealed the quality of one attribute, the seller has no reason to not truthfully reveal the quality of the other attribute

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

- $\beta_a \in \{0, 1\}$ with $\Pr[\beta_a = 1] = p \in [0, 1]$
- In equilibrium the seller fully reveals the state
  - having revealed the quality of one attribute, the seller has no reason to not truthfully reveal the quality of the other attribute

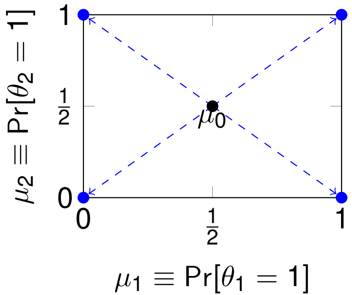

Two attributes: Simple example

Buyer only interested in one (unknown) attribute:

- $\beta_a \in \{0, 1\}$ with $\Pr[\beta_a = 1] = p \in [0, 1]$
- In equilibrium the seller fully reveals the state
  - having revealed the quality of one attribute, the seller has no reason to not truthfully reveal the quality of the other attribute

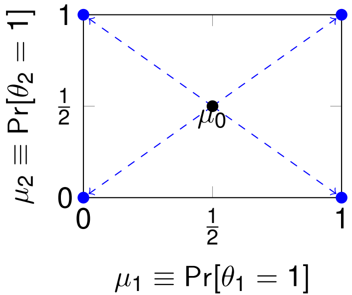

- There is **no benefit** from the buyer communicating her preferences (β_a)

Two attributes: Another example

Buyer potentially interested in both attributes:

- $\beta_a \in \left\{0, \frac{1}{2}, 1\right\}$ with $\Pr[\beta_a = 1] = \Pr[\beta_a = 0] = p \in (0, \frac{1}{2})$

Two attributes: Another example

Buyer potentially interested in both attributes:

- $\beta_a \in \left\{0, \frac{1}{2}, 1\right\}$ with $\Pr[\beta_a = 1] = \Pr[\beta_a = 0] = p \in (0, \frac{1}{2})$
- With no buyer communication, there is no longer an equilibrium in which the seller fully reveals the state

Two attributes: Another example

Buyer potentially interested in both attributes:

- $\beta_a \in \left\{0, \frac{1}{2}, 1\right\}$ with $\Pr[\beta_a = 1] = \Pr[\beta_a = 0] = p \in (0, \frac{1}{2})$
- With no buyer communication, there is no longer an equilibrium in which the seller fully reveals the state
  - suppose for attribute 1, the seller (truthfully) reveals that $\theta_1 = 1$
  - for attribute 2, the seller now has a strict preference for revealing that $\theta_2 = 1$

Two attributes: Another example

Buyer potentially interested in both attributes:

- $\beta_a \in \{0, \frac{1}{2}, 1\}$ with $\Pr[\beta_a = 1] = \Pr[\beta_a = 0] = p \in (0, \frac{1}{2})$
- With no buyer communication, there is no longer an equilibrium in which the seller fully reveals the state
  - suppose for attribute 1, the seller (truthfully) reveals that $\theta_1 = 1$
  - for attribute 2, the seller now has a strict preference for revealing that $\theta_2 = 1$
- With no buyer communication, the seller can only fully reveal the quality of one attribute and partially reveal for the other attribute

Two attributes: Another example

Buyer potentially interested in both attributes:

- $\beta_a \in \{0, \frac{1}{2}, 1\}$ with $\Pr[\beta_a = 1] = \Pr[\beta_a = 0] = p \in (0, \frac{1}{2})$
- With no buyer communication, there is no longer an equilibrium in which the seller fully reveals the state
  - suppose for attribute 1, the seller (truthfully) reveals that $\theta_1 = 1$
  - for attribute 2, the seller now has a strict preference for revealing that $\theta_2 = 1$
- With no buyer communication, the seller can only fully reveal the quality of one attribute and partially reveal for the other attribute
- There is a benefit from buyer communicating her preferences (β_a)
  - intuition: seller can provide more tailored recommendation for the buyer by providing information on buyer's preferred attribute

Two attributes: Results

Assumption 1

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

Proposition 2

With two attributes and no bias towards either good, there is a (seller preferred) 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 $\theta_j = 0$.*

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

Two attributes: Results

- In words:
  - buyer reveals which attribute she is most interested in

Two attributes: Results

- In words:
  - buyer reveals which attribute she is most interested in
  - seller fully reveals best good for that attribute and nothing about other attribute

Two attributes: Results

- In words:
  - buyer reveals which attribute she is most interested in
  - seller fully reveals best good for that attribute and nothing about other attribute
- Intuition:
  - this is an equilibrium: both buyer and seller follow equilibrium strategy
  - note given the information from the buyer, the seller cannot do better than to reveal information about the preferred attribute

Two attributes: Results

- In words:
  - buyer reveals which attribute she is most interested in
  - seller fully reveals best good for that attribute and nothing about other attribute
- Intuition:
  - this is an equilibrium: both buyer and seller follow equilibrium strategy
  - note given the information from the buyer, the seller cannot do better than to reveal information about the preferred attribute
  - an equilibrium in which the buyer requests (partial) information about both attributes is strictly worse for seller

Summary

- Study a back and forth cheap talk model with two-sided private information
  - very little research on this topic
- Application to buyer-seller both for online and offline interactions
  - relevant to debate on consumer privacy
- Key result: if an expert wants to convince a decision maker to take one of several non-default actions
  - single attribute: eliciting DM's preferences between options can only be harmful
  - multiple attributes: eliciting DM's preferences between different attributes is helpful for tailoring recommendations