

# Interactive Communication in Bilateral Trade

Jieming Mao ✉

Google Research, New York, NY, USA

Renato Paes Leme ✉

Google Research, New York, NY, USA

Kangning Wang ✉

Duke University, Durham, NC, USA

---

## Abstract

We define a model of interactive communication where two agents with private types can exchange information before a game is played. The model contains Bayesian persuasion as a special case of a one-round communication protocol. We define message complexity corresponding to the minimum number of interactive rounds necessary to achieve the best possible outcome. Our main result is that for bilateral trade, agents don't stop talking until they reach an efficient outcome: Either agents achieve an efficient allocation in finitely many rounds of communication; or the optimal communication protocol has infinite number of rounds. We show an important class of bilateral trade settings where efficient allocation is achievable with a small number of rounds of communication.

**2012 ACM Subject Classification** Theory of computation → Algorithmic game theory; Theory of computation → Solution concepts in game theory; Theory of computation → Exact and approximate computation of equilibria

**Keywords and phrases** Bayesian persuasion, bilateral trade, information design

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2022.105

**Related Version** *Full Version*: <https://arxiv.org/abs/2106.02150>

**Supplementary Material** *Software (Source Code)*: <https://gist.github.com/renatopl/72ee85d212d08ad9670977cc8ffa2afa>

*Software (Source Code)*: [https://colab.research.google.com/drive/1lQdRrZD7-aCuuYJM6bI1sf1fi\\_CSA4Fq?usp=sharing](https://colab.research.google.com/drive/1lQdRrZD7-aCuuYJM6bI1sf1fi_CSA4Fq?usp=sharing)

## 1 Introduction

We consider the situation where agents are allowed to have a conversation before playing a game. Unlike cheap talk [3], we assume agents are able to send verifiable signals. The ability to send verifiable signals is a new device introduced in the literature on Bayesian persuasion [29]. Without this power, various revelation-principle-type results show that anything that can be achieved through an interactive game/mechanism can also be achieved through direct revelation. In stark contrast to those, we will show that with verifiable signals, a (potentially very long) interactive protocol can lead to more efficient outcomes.

Consider the following scenario: Sally is a supplier of electronic components and Bob is a builder of computers. They have an ongoing business relationship and often negotiate the price of new custom parts. Since they come to the negotiation table over and over, they established a protocol which they follow during the negotiation. Through this protocol, Bob will convey information about his value for the new component and Sally will convey information about her cost for producing it. Each step of the communication is verifiable: Bob can for example show Sally a quote from a competing supplier, bounding his willingness to pay. Sally can tell Bob the cost of raw materials for building that component, which will refine Bob's information about Sally's production cost. Their goal is that if they talk long enough, they will be able to settle on a price in-between or figure out that no trade is possible (Sally's cost exceeds Bob's value).



© Jieming Mao, Renato Paes Leme, and Kangning Wang;  
licensed under Creative Commons License CC-BY 4.0

13th Innovations in Theoretical Computer Science Conference (ITCS 2022).

Editor: Mark Braverman; Article No. 105; pp. 105:1–105:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Informally, we will show that the optimal communication protocol will only end when agents reach an efficient outcome. To formalize this statement, we define a formal model of interactive communication, which contains the Bayesian persuasion model of Kamenica and Gentzkow [29] and Dughmi and Xu [23] as a special case of a one-round protocol. We define our protocol for a generic two-player game and study bilateral trade as a special case. Although the model can be generalized to multiple players, most applications we have in mind are about two agents talking to each other: buyer and seller in bilateral trade, sender and receiver in Bayesian persuasion or Alice and Bob in standard communication complexity.

## 1.1 Model of Interactive Communication

In this paper we study this question by first defining a communication protocol that takes place before Sally and Bob play a game where both players have private information and their payoffs depend on the combination of their private types. To make this concrete, we will use *bilateral trade* as the main running example. Sally (the seller) has a private cost  $\theta_S$  on the item being sold and Bob (the buyer) has a private value  $\theta_B$ . Sally's action in the game is to set a price  $p$ . If the price is below Bob's value the trade happens and Sally and Bob get utilities  $p - \theta_S$  and  $\theta_B - p$  respectively. Otherwise trade doesn't happen and both get zero utility.

Before Sally decides on the price, Bob may want to send a signal refining Sally's information in an attempt to persuade her to set a lower price. Since this is a game where both sides have private information, Bob's signal will depend on his information about Sally. Before Bob sends a signal, Sally may want to refine Bob's information about her in an attempt to persuade him to send a better signal. This can be taken one step further. Before Sally sends her signal, Bob can send a preliminary signal to influence how Sally will signal to influence Bob's subsequent signal.

The result can be interpreted as a communication protocol which consists of an alternating sequence of information refinements. In odd rounds, Bob sends a signal refining Sally's information and in even rounds Sally sends a signal refining Bob's information. We assume the same commitment structure as in Kamenica and Gentzkow [29]: agents truthfully communicate the realization of signals. This leads to a natural equilibrium notion which contains the persuasion scheme of [29] as a special case when the round complexity of the protocol is one.

To study how the round complexity of the protocol affects the efficiency of the game, we define the notion of *message complexity*. A game has message complexity  $t$  if players can't improve their payoffs after  $t$  rounds of communication. In other words, for any  $t' \geq t$  the payoffs the agents get with a  $t'$ -round protocol are the same as with a  $t$ -round protocol.

## 1.2 Efficient Communication in Bilateral Trade

Our main result is that for bilateral trade, if the message complexity is finite, then the allocation is efficient. This means that for any initial set of types and distributions in a bilateral trade game, one of the following must be true: either agents exchange information in a way that after finitely many rounds they can implement the efficient outcome (which is to trade whenever  $\theta_B > \theta_S$ ); or the efficiency of the allocation keeps improving the longer they talk.

For the case where the buyer has only two types in the support of his distribution, efficient allocation can always be achieved with two rounds of communication (regardless of how complicated the seller's type space is). If the buyer has three types in the support, we have

an example showing we need at least three rounds of communication. Finally, we conjecture that whenever type spaces are finite we can always achieve efficient allocation in finitely many rounds of communication.

### 1.3 A game with longer communication

We end the paper with an example of a game with longer communication. In the tradition of Cold War game theory we now have Sally (the Spymaster) and Bob (the Birdwatcher) be spies whose private types correspond to the countries they are actually serving. Sally and Bob don't know if they are friends or enemies. If they are friends (same type) their payoffs correspond to those of a cooperative game. If they are enemies (different types) their payoffs correspond to a zero-sum game. Bob and Sally can communicate before playing the game, but it is a delicate balance: both would like to know if their counterpart is their friend or enemy but both would also like to the other to believe they may be a friend when they are in fact enemies. We show that a long and gradual disclose of information can benefit both parties.

### 1.4 Techniques

From a technical standpoint, we build on top of two techniques introduced by Aumann and Maschler [4] that became standard in the toolbox of information design. The first is the observation that signals can be thought of as decomposition of a prior distribution into posteriors that are only required to average to the prior (“splitting lemma”). The second is the concavification of payoff functions, which is the observation that by signaling, the sender can replace their payoff function by its concave hull.

Our main technical innovation is to analyze the dynamic of *alternate concavification*, by which we mean the following: consider we start from payoff functions  $\pi_B(\mathcal{D}_S, \mathcal{D}_B)$  and  $\pi_S(\mathcal{D}_S, \mathcal{D}_B)$  that are defined as functions of Sally's information about Bob ( $\mathcal{D}_B$ ) and Bob's information about Sally ( $\mathcal{D}_S$ ). A message by Sally allows her to replace her payoff by its concave hull with respect to  $\mathcal{D}_S$ , smoothening out Bob's payoff correspondingly. Bob can similarly replace his payoff by its concave hull on  $\mathcal{D}_B$ . Note however, that whenever Sally concavifies her payoff it may cause Bob's payoff to be no longer concave and vice-versa. Our analysis will deal with understanding whether this procedure stabilizes after finitely many rounds (finite message complexity) or whether it goes on indefinitely. If it stabilizes, we are interested in understanding its properties.

In our analysis of bilateral trade, we will develop a higher-order version of the indifference argument of Bergemann, Brooks and Morris [10]. While [10] decomposes a prior distribution into posteriors for which the seller is indifferent about which price to set, we will decompose each prior into posteriors for which the buyer is indifferent between different ways to make the seller indifferent.

### 1.5 Implementation

In an online supplement<sup>1</sup> we provide an implementation of the alternating concavification procedure which takes a specification of a generic game with binary type spaces as an input and computes the payoffs after  $t$  rounds of communication together with the communication

---

<sup>1</sup> See code in Github (<https://gist.github.com/renatoppl/72ee85d212d08ad9670977cc8ffa2afa>) or Google Colab ([https://colab.research.google.com/drive/1lQdRrZD7-aCuuYJM6bI1sf1fi\\_CSA4Fq?usp=sharing](https://colab.research.google.com/drive/1lQdRrZD7-aCuuYJM6bI1sf1fi_CSA4Fq?usp=sharing)).

protocol. The computation is exact: it uses rational numbers (so there is no floating point precision issues) and computes a parametric concave hull so it doesn't need to rely on discretization.

## 1.6 Related Work

Our work is broadly situated in line of work in Economics studying how the information structure affects the outcome in auction and bargaining settings, which was initiated in Bergemann and Pesendorfer [13] and Esó and Szentes [27] and has been more recently explored in Bergemann et al [13, 11], Emek et al [26], Daskalakis et al [20] and Badanidiyuru et al [8]. In this line of work, the auction designer is more informed than the participants and must decide how much information to disclose as part of the design decision. We differ from this line of work in the sense that we assume that the information lies with the participants themselves and their decisions on how to disclose information affect their payoffs.

In that sense, our work is closer to the Bayesian persuasion model of Kamenica and Gentzkow [29, 28] and the algorithmic persuasion of Dughmi and Xu [24] and Dughmi et al [23]. In this setting there are two agents: *sender* and *receiver* where the sender is typically more informed than the receiver but the receiver is the one responsible for choosing an action in a base game. In contrast our agents both have uncertainty about each other. Kamenica and Gentzkow [29] briefly consider the setting where the receiver can also have private information, but only allow messages in one direction (from the sender to the receiver). Instead our paper considers both parties to have partial information and considers an interactive exchange of information. Doval and Ely [22] consider sequential disclosure of information that may depend on player behaviors, but it again differs from our setting in the sense that the information is initially held by the designer who then discloses it to agents. We refer to the excellent survey by Bergemann and Morris [12] for a unified treatment of those papers.

The notion of having a conversation before a game is played is the central talk in the *cheap talk* literature. In particular in “Long Cheap Talk”, Aumann and Hart [3] show how the set of equilibrium payoffs of a game can be expanded by an arbitrarily long conversation before the game is played. Their model assumes no commitment whatsoever (the players are free to send messages as they please) while in our model we have the same commitment structure as in the Bayesian persuasion literature, where the realization of signals are truthfully communicated.

The power of interactive communication has been extensively studied in communication complexity (Yao [33]). Nisan and Wigderson [32] show an exponential gap between  $k$  and  $(k - 1)$ -round communication complexity. Babai et al [5] show an exponential gap between simultaneous communication complexity and communication complexity in the multi-party number-on-forehead communication model. Similar demonstrations of the power of interactive communication have been shown when studying the communication complexity of mechanism design problems. Dobzinski, Nisan and Oren [21] study how the number of adaptive rounds affects welfare efficiency in communication protocols with polynomial communication complexity for combinatorial auctions. Subsequent work by Alon et al [1] and Assadi [2] provides tight bounds on the number of rounds necessary to obtain an efficient allocation. The motivation for interactive communication in those papers is a restriction in the number of bits used in each interaction. Hence the need for interactivity comes from algorithmic and not strategic considerations. In our paper, on the other hand, interactivity is driven by strategic considerations: agents will only reveal so much about their types until they can learn more about the other agent's type.

Our results also contribute to the line of work on bilateral trade started by Chatterjee and Samuelson [17] and Myerson and Satterthwaite [31]. Their message is in a sense the opposite of ours: interaction is not helpful (in the sense of the revelation principle) and efficient trade

can't be achieved by an incentive compatible mechanism. We show that communication is useful and leads to efficient allocation in many important cases. The main difference is that here we are giving the agents additional commitment power: they can credibly signal about their type, which is what drives the Bayesian persuasion literature. We refer the reader to Section 1C of Kamenica and Gentzkow [29] for an in-depth discussion of the source and motivation behind this additional commitment power.

There is an important line of work looking at the bilateral trade problem from the perspective of approximation algorithms. Blumrosen and Dobzinski [14] give the first approximation to efficiency in bilateral trade, later improved by Collini-Baldeschi et al [19] and Kang and Vondrak [30]. Recently, various new angles have been explored, such as multi-dimensional two-sided markets (Collini-Baldeschi et al [18], Cai et al [16]), gains from trade approximation (Brustle et al [15]), best of both worlds guarantees (Babaioff et al [6]), resource augmentation (Babaioff et al [7]), dynamic auctions (Balseiro et al [9]), and limited information (Dütting et al [25]).

## 2 Games with Interactive Communication

Our main objects of study are two-player games where both players (called Sally and Bob) have private types. Ex-ante (before types are revealed) the players can agree on a communication protocol to exchange information about their types. We will be interested in how the round complexity of the communication protocol can affect the outcome of the game. For simplicity we will restrict to games where only Sally has non-trivial actions.

### 2.1 Base game

We first define the *base game*, which is played after the communication protocol. In this game Bob has a type  $\theta_B \in \Theta_B$  and Sally has a type  $\theta_S \in \Theta_S$ . The types are drawn from independent known distributions  $\mathcal{D}_B$  and  $\mathcal{D}_S$  respectively. Only Sally has an action  $a_S \in A_S$  to choose. The utilities of both players are given by functions:

$$u_i : A_S \times \Theta_S \times \Theta_B \rightarrow \mathbb{R} \quad i \in \{S, B\}$$

As usual in games with private information, Sally knows her type but only knows the distribution over which Bob's type is drawn and vice versa. Her optimal strategy is rather simple:

$$a_S^*(\theta_S; \mathcal{D}_B) = \operatorname{argmax}_{a \in A_S} \mathbb{E}_{\theta_B \sim \mathcal{D}_B} [u_S(a, \theta_S, \theta_B)]$$

### 2.2 Example: Bilateral Trade

Bilateral trade will provide us the main running example. Sally the seller is trying to sell an item to Bob the buyer. The types of Sally and Bob are their values for the item. Sally's action is to choose a price to sell the good. The sets  $\Theta_S$ ,  $\Theta_B$  and  $A_S$  correspond to the non-negative real numbers  $\mathbb{R}_+$ . Upon setting a price  $a_S$ , Bob will buy whenever  $\theta_B \geq a_S$  leading to the following utilities:

$$u_B(a_S, \theta_S, \theta_B) = (\theta_B - a_S) \cdot \mathbf{1}\{\theta_B \geq a_S\} \quad u_S(a_S, \theta_S, \theta_B) = (a_S - \theta_S) \cdot \mathbf{1}\{\theta_B \geq a_S\}$$

Without any communication, the expected welfare is:

$$W(\mathcal{D}_S, \mathcal{D}_B) = \mathbb{E}_{\theta_S \sim \mathcal{D}_S} \mathbb{E}_{\theta_B \sim \mathcal{D}_B} [u_S(a_S^*(\theta_S; \mathcal{D}_B), \theta_S, \theta_B) + u_B(a_S^*(\theta_S; \mathcal{D}_B), \theta_S, \theta_B)]$$

## 105:6 Interactive Communication in Bilateral Trade

which is typically suboptimal when compared to the the welfare under efficient trade:

$$W^*(\mathcal{D}_S, \mathcal{D}_B) = \mathbb{E}_{\theta_S \sim \mathcal{D}_S} \mathbb{E}_{\theta_B \sim \mathcal{D}_B} [(\theta_B - \theta_S)^+] \quad (1)$$

If Bob doesn't buy, both players experience zero utility. So it is in the best interest of both Sally and Bob to make the trade happen whenever  $\theta_B > \theta_S$ . It is only natural for them to engage in *bargaining*, i.e., a conversation in which information is gradually revealed in a way that the price settles in at a point that hopefully leads to efficient trade. By gradual information revealing we mean gradually refining each other's priors to lead to an outcome that is improving for both parties.

### 2.3 Communication Protocol

A communication protocol will be an alternating sequence of information refinements: in odd rounds Bob will send a signal to refine Sally's information about his type and in even rounds Sally will refine Bob's information about her type.

#### 2.3.1 Notation

The protocol will be described as a sequence of messages by each player. We will adopt the unusual but convenient convention of using  $t$  to denote the  $t$ -to-last message. So  $m^1$  will denote the last message,  $m^2$  the penultimate message and so on. . . As we described before, odd messages will be sent by Bob and denoted  $m_B^t$ . Even messages will be sent by Sally and denoted  $m_S^t$ .

To describe the messages, we need an additional notation. Given any set  $\mathcal{X}$ , let  $\Delta(\mathcal{X})$  be the set of distributions over  $\mathcal{X}$ , and let  $\Delta(\Delta(\mathcal{X}))$  be the set of distributions over  $\Delta(\mathcal{X})$ . We will also define the following operator:

$$\mu : \Delta(\Delta(\mathcal{X})) \rightarrow \Delta(\mathcal{X})$$

that given a distribution over distributions returns a single distribution in the natural way: given  $M \in \Delta(\Delta(\mathcal{X}))$  build a distribution over  $\mathcal{X}$  by first sampling  $\mathcal{D} \sim M$  where  $\mathcal{D} \in \Delta(\mathcal{X})$  and then sampling an element of  $\mathcal{X}$  from  $\mathcal{D}$ . The resulting distribution over  $\mathcal{X}$  is  $\mu(M)$ .

#### 2.3.2 Information Refinement

We say that a distribution over distributions  $M \in \Delta(\Delta(\mathcal{X}))$  is an information refinement of a distribution  $\mathcal{D} \in \Delta(\mathcal{X})$  whenever  $\mu(M) = \mathcal{D}$ .

An information refinement is a convenient way to represent a signal. A signal about  $\mathcal{D}$  is a random variable  $Y$  that is correlated with a random variable  $X \sim \mathcal{D}$ . Upon observing  $Y$  one can perform a Bayesian update and obtain the distribution of  $X$  conditioned on  $Y$ . This induces a distribution over distributions  $M \in \Delta(\Delta(\mathcal{X}))$  s.t.  $\mu(M) = \mathcal{D}$ . Conversely given any  $M \in \Delta(\Delta(\mathcal{X}))$  such that  $\mu(M) = \mathcal{D}$  we can obtain a signal by letting  $Y$  represent the distribution sampled from  $M$  and  $X$  be an element in  $\mathcal{X}$  sampled from  $Y$  (here  $Y$  is both a random variable and a distribution in  $\Delta(\mathcal{X})$ ).

#### 2.3.3 Structure of the messages

Using the notation in the previous subsections we define messages in odd rounds as:

$$m_B^t : \Delta(\Theta_S) \times \Delta(\Theta_B) \rightarrow \Delta(\Delta(\Theta_B)) \quad \text{s.t.} \quad \mu(m_B^t(\mathcal{D}_S, \mathcal{D}_B)) = \mathcal{D}_B$$

A message will take as input the information Sally has about Bob and the information Bob has about Sally and output a refined version of Bob's information. Similarly in even rounds:

$$m_S^t : \Delta(\Theta_S) \times \Delta(\Theta_B) \rightarrow \Delta(\Delta(\Theta_S)) \quad \text{s.t.} \quad \mu(m_S^t(\mathcal{D}_S, \mathcal{D}_B)) = \mathcal{D}_S$$

A sequence of functions  $m_B^1, m_S^2, m_B^3, \dots, m_i^k$  (where  $i \in \{B, S\}$  depends on the parity of  $k$ ) describes an *alternating information refinement* protocol.

The protocol is executed then as follows: if the types are sampled from  $\mathcal{D}_B$  and  $\mathcal{D}_S$  respectively, we set  $\mathcal{D}_B^k = \mathcal{D}_B$  and  $\mathcal{D}_S^k = \mathcal{D}_S$ . Then for  $t = k, k-1, k-2, \dots, 1$ ,

$$\begin{aligned} \text{if } t \text{ is odd: } & \mathcal{D}_B^{t-1} \sim m_B^t(\mathcal{D}_S^t, \mathcal{D}_B^t) & \mathcal{D}_S^{t-1} &= \mathcal{D}_S^t \\ \text{if } t \text{ is even: } & \mathcal{D}_B^{t-1} = \mathcal{D}_B^t & \mathcal{D}_S^{t-1} &\sim m_S^t(\mathcal{D}_S^t, \mathcal{D}_B^t) \end{aligned}$$

Once the communication is over, Sally selects her optimal action  $a_S^*(\theta_S; \mathcal{D}_B^0)$  using the information available at that point.

### 2.3.4 Payoffs

We can now define recursively the payoffs of each agent. Let  $\pi_B^t(\mathcal{D}_S, \mathcal{D}_B)$  and  $\pi_S^t(\mathcal{D}_S, \mathcal{D}_B)$  be the expected payoffs obtained by Bob and Sally from a  $t$ -round protocol with initial information  $\mathcal{D}_S$  and  $\mathcal{D}_B$ . For  $t = 0$  we simply have the payoffs of the base game:

$$\pi_i^0(\mathcal{D}_S, \mathcal{D}_B) = \mathbb{E}_{\theta_S \sim \mathcal{D}_S, \theta_B \sim \mathcal{D}_B} [u_i(a_S^*(\theta_S; \mathcal{D}_B), \theta_S, \theta_B)] \quad i \in \{B, S\} \quad (2)$$

For  $t > 0$  and we have:

$$\begin{aligned} \pi_i^t(\mathcal{D}_S, \mathcal{D}_B) &= \mathbb{E}_{\mathcal{D}'_B \sim m_B^t(\mathcal{D}_S, \mathcal{D}_B)} [\pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}'_B)] & i \in \{B, S\} & \quad \text{for odd } t \\ \pi_i^t(\mathcal{D}_S, \mathcal{D}_B) &= \mathbb{E}_{\mathcal{D}'_S \sim m_S^t(\mathcal{D}_S, \mathcal{D}_B)} [\pi_i^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)] & i \in \{B, S\} & \quad \text{for even } t \end{aligned}$$

### 2.3.5 Solution Concept

We now define the notion of equilibrium of a communication protocol. In high level terms a protocol in equilibrium must satisfy two properties: (i) *voluntary communication* and (ii) *sub-protocol optimality*.

#### 2.3.5.1 Voluntary Communication

We say that a protocol satisfies *voluntary communication* if both agents weakly prefer communicating over skipping that round. Mathematically this can be stated as follows:

$$\pi_i^t(\mathcal{D}_S, \mathcal{D}_B) \geq \pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}_B), \forall i \in \{B, S\} \text{ and } \forall t > 0 \quad (3)$$

This means that the sender will never send a message decreasing their payoff. The receiver will refuse<sup>2</sup> to hear any message that decreases their payoff (by for example shutting their ears or deleting an email without reading).

---

<sup>2</sup> The important assumption here is that the receiver has a way to credibly "not listen". For example, whenever Bob tells Sally something, even if she decides not to use the information in subsequent rounds, Bob knows that she knows and that may prevent Bob from disclosing further information in the future. Sally must have a way to show Bob that she didn't update her prior based on that message.

We define the set of *voluntary refinements* as follows. For odd  $t$  define  $\mathbf{VR}_B^t(\mathcal{D}_S, \mathcal{D}_B)$  as the set of refinements  $M \in \Delta(\Delta(\Theta_B))$  with  $\mu(M) = \mathcal{D}_B$  such that

$$\mathbb{E}_{\mathcal{D}'_B \sim M}[\pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}'_B)] \geq \pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}_B), \quad \forall i \in \{B, S\}$$

Similarly in even rounds we define  $\mathbf{VR}_S^t(\mathcal{D}_S, \mathcal{D}_B)$  as the set of refinements  $M \in \Delta(\Delta(\Theta_S))$  with  $\mu(M) = \mathcal{D}_S$  such that

$$\mathbb{E}_{\mathcal{D}'_S \sim M}[\pi_i^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)] \geq \pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}_B), \quad \forall i \in \{B, S\}$$

### 2.3.5.2 Sub-protocol optimality

The second condition for equilibrium is that each message maximizes the payoff of the sender among all messages satisfying voluntary communication:

$$\begin{aligned} \pi_B^t(\mathcal{D}_S, \mathcal{D}_B) &= \max_{M \in \mathbf{VR}_B^t(\mathcal{D}_S, \mathcal{D}_B)} \mathbb{E}_{\mathcal{D}'_B \sim M}[\pi_B^{t-1}(\mathcal{D}_S, \mathcal{D}'_B)] && \text{for odd } t \\ \pi_S^t(\mathcal{D}_S, \mathcal{D}_B) &= \max_{M \in \mathbf{VR}_S^t(\mathcal{D}_S, \mathcal{D}_B)} \mathbb{E}_{\mathcal{D}'_S \sim M}[\pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)] && \text{for even } t \end{aligned} \quad (4)$$

### 2.3.5.3 Equilibrium selection

The equilibrium conditions (3) and (4) don't specify a unique protocol, since there may be multiple optimal information refinements for Sally leading to different utilities for Bob and vice-versa. For the remainder of the paper we will study the equilibrium in which each agent breaks ties in favor of the other one.<sup>3</sup>

With this tie-breaking in place, the values of  $\pi_i^t(\mathcal{D}_S, \mathcal{D}_B)$  in equilibrium are uniquely determined. From now on, whenever we refer to  $\pi_i^t(\mathcal{D}_S, \mathcal{D}_B)$  we will be referring to the values in equilibrium with that tie-breaking rule.

### 2.3.5.4 Note on Voluntary Communication

In settings such as binary type spaces ( $|\Theta_S| = |\Theta_B| = 2$ ) any refinement is voluntary. This is due to the fact that uncertainty can be represented by a single-parameter. We show this in the full version together with a discussion on Voluntary Communication.

## 2.3.6 Generic protocols

One could consider more general protocols, for example by allowing both agents to simultaneously send messages and by considering messages in a generic space instead of just information refinements. In the full version, we show any general protocol can be reduced to an alternating information refinement protocol. We give the intuition below.

Information refinement is enough since the only use of messages is for agents to perform a Bayesian update on their information about the other agent. Hence it is enough to reason about the information updates directly.

To see that alternating communication is without loss of generality, observe that only Sally has an action at  $t = 0$ . Hence there is no need for Sally to communicate at  $t = 1$  since

<sup>3</sup> Formally, this means that in odd rounds  $t$  Bob will choose a refinement  $M \in \mathbf{VR}_B^t(\mathcal{D}_S, \mathcal{D}_B)$  that lexicographically maximizes  $(\mathbb{E}_{\mathcal{D}'_B \sim M}[\pi_B^{t-1}(\mathcal{D}_S, \mathcal{D}'_B)], \mathbb{E}_{\mathcal{D}'_B \sim M}[\pi_S^{t-1}(\mathcal{D}_S, \mathcal{D}'_B)])$  and Sally will choose a refinement  $M \in \mathbf{VR}_S^t(\mathcal{D}_S, \mathcal{D}_B)$  that lexicographically maximizes  $(\mathbb{E}_{\mathcal{D}'_S \sim M}[\pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)], \mathbb{E}_{\mathcal{D}'_S \sim M}[\pi_B^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)])$ .



Bob has no subsequent action to take. Sally may very well stay silent at time  $t = 1$ . If that happens there is no need for Bob to say anything at  $t = 2$  since he can't influence any further message from Sally. Instead Bob can combine his communication at  $t = 2$  and  $t = 1$  and send it together at time  $t = 1$  staying silent at  $t = 2$ . The same argument can be applied recursively showing that we can obtain a protocol with the same effect by having Bob and Sally speaking in alternating rounds.

## 2.4 Message Complexity

By voluntary communication, the utilities of the agents are monotone along the protocol, i.e.,  $\pi_i^t(\mathcal{D}_S, \mathcal{D}_B) \geq \pi_i^{t-1}(\mathcal{D}_S, \mathcal{D}_B)$  for all  $i \in \{B, S\}$  and  $t > 0$ . This allows us to define *limit utilities* as follows:

$$\pi_i^\infty(\mathcal{D}_S, \mathcal{D}_B) = \lim_{t \rightarrow \infty} \pi_i^t(\mathcal{D}_S, \mathcal{D}_B).$$

With that definition, we can define the *message complexity* of a pair of distributions  $(\mathcal{D}_S, \mathcal{D}_B)$  as:

$$C(\mathcal{D}_S, \mathcal{D}_B) := \min \{t \in \mathbb{N} \cup \{\infty\} \mid \pi_S^t(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^t(\mathcal{D}_S, \mathcal{D}_B) = \pi_S^\infty(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^\infty(\mathcal{D}_S, \mathcal{D}_B)\}.$$

which corresponds to the minimum number of messages to achieve the best possible utilities.

Finally, if in equilibrium there is a message that moves from state  $(\mathcal{D}_S, \mathcal{D}_B)$  in period  $t$  to  $(\mathcal{D}'_S, \mathcal{D}'_B)$  in period  $t - 1$  we say that  $(\mathcal{D}'_S, \mathcal{D}'_B)$  is a *child* of  $(\mathcal{D}_S, \mathcal{D}_B)$  in period  $t$ .

## 3 Efficient Bilateral Trade

Bilateral trade will provide us with a concrete setting in which multiple rounds of interaction can lead to a better outcome – in particular the efficient allocation. We start with a numerical example.

### 3.1 Numerical example with one seller type

Recall that we are in the setting described in Section 2.2. Consider the setting where Sally has a single type  $\theta_s = 2$  and Bob has two possible types 3 and 6 with probability  $p = 1/3$  of having the high type. In the absence of any communication Sally will set a price<sup>4</sup> equal to 6 and sell with probability  $1/3$ , resulting in inefficient trade:

$$\pi_S^0 = \frac{4}{3} \quad \pi_B^0 = 0 \quad W^0 = \frac{4}{3} < W^* = 2$$

If we allow one round of communication, Bob can send a signal refining Sally's information about his type.<sup>5</sup> With probability  $8/9$  Bob will send a message that refines Sally's information to  $p = 1/4$ , where  $p$  is the probability of having the high type. With probability  $1/9$  Bob can refine Sally's information to  $p = 1$ . In the first case, Sally is indifferent between the two

<sup>4</sup> Posted-price mechanisms are optimal for Sally in general.

<sup>5</sup> As we discussed in Section 2.3.2 information refinements can always be interpreted as a Bayesian update of a signal correlated with the sender's type. In this particular case, Bob can accomplish this information refinement by sending one of two signals  $\{L, H\}$ . If Bob has the high type he can send the signal  $H$  with probability  $1/3$  and  $L$  with probability  $2/3$ , and whenever he has the low type he always sends  $L$ . One can readily verify that this leads to the refinement described.

## 105:10 Interactive Communication in Bilateral Trade

prices and may very well price at 3. In the second case, Sally will price at 6. In either case, she will sell with probability one, leading to efficient trade. The outcome is the following:

$$\pi_S^1 = \frac{4}{3} \quad \pi_B^1 = \frac{2}{3} \quad W^1 = 2 = W^* = 2$$

This is in fact a general phenomenon described in Bergemann, Brooks and Morris [10], who show that if there is no uncertainty about the seller's type, the buyer can always signal in order to extract the full surplus of the trade. The buyer does so by refining his distribution into revenue-equivalent distributions, i.e., distributions where the seller is indifferent between pricing at any point in the support. Below we state their result in our language, which will prove useful later.

► **Lemma 1** (Bergemann-Brooks-Morris [10]). *If  $\mathcal{D}_S$  has a single point in the support, then*

$$\pi_B^1(\mathcal{D}_S, \mathcal{D}_B) = W^*(\mathcal{D}_S, \mathcal{D}_B) - \pi_B^0(\mathcal{D}_S, \mathcal{D}_B) \quad \pi_S^1(\mathcal{D}_S, \mathcal{D}_B) = \pi_S^0(\mathcal{D}_S, \mathcal{D}_B).$$


### 3.2 Numerical example with two seller types

If Bob also has uncertainty about the Sally's type, then it is no longer possible to achieve efficiency by one round of signaling. Let's keep Bob's types as  $\Theta_B = \{3, 6\}$  with  $1/3$  probability on the high type. But now Sally's type is in  $\Theta_S = \{0, 2\}$  with  $1/2$  probability on the high type.

Without any communication, Sally sets price 6 whenever her type is 2 and 3 whenever her cost is 0, leading to the following outcome:

$$\pi_S^0 = \frac{13}{6} \quad \pi_B^0 = \frac{1}{2} \quad W^0 = \frac{8}{3} < W^* = 3$$

To describe the outcome with communication it is useful to use the diagram in Figure 1. Since we have two types for each agent, we can represent the distributions  $(\mathcal{D}_B, \mathcal{D}_S)$  by a pair  $(p, q) \in [0, 1]$  where  $p$  is the probability that Bob has the high type and  $q$  is the probability that Sally has the high type. Hence the initial distribution corresponds to the point  $(1/3, 1/2)$ .



■ **Figure 1** Each square corresponds to  $[0, 1]^2$ . A point  $(p, q)$  depicts a state where Bob has the high type with probability  $p$  and Sally has the high type with probability  $q$ . The initial state is  $(\frac{1}{3}, \frac{1}{2})$ . The arrows correspond to information refinements in the optimal 1-round (left) and 2-round (right) protocols. Red arrows correspond to Sally's refinement of Bob's information and blue arrows to Bob's refinement of Sally's information.

With one round of communication, the best thing that Bob can do is to refine<sup>6</sup> Sally's information  $p = 1/2$  with probability  $1/3$  and  $p = 1/4$  with probability  $2/3$  (see the diagram on the left in Figure 1). We will explain in the next subsection why this is the optimal choice for Bob. Now Sally's price depends both on her type and the information she has about Bob. Evaluating the four cases we get to the following outcome:

$$\pi_S^1 = \frac{13}{6} \quad \pi_B^1 = \frac{3}{4} \quad W^1 = \frac{35}{12} < W^* = 3$$

Bob's signal improves the efficiency of the allocation and extracts the additional efficiency as buyer surplus, but it is not quite enough to achieve full efficiency. With two rounds of communication, however, we obtain the efficient allocation. In the diagram on the right in Figure 1 we depict the optimal two-round protocol: first Sally sends a signal that refines Bob's signal to  $q = 0$  with probability  $1/4$  and  $q = 2/3$  with probability  $3/4$ . If Bob receives the  $q = 0$  signal, he stays silent since he knows Sally will already price at the low type. If Bob receives the  $q = 2/3$  signal, however, he refines Sally's information to  $p = 1/4$  with probability  $8/9$  and  $p = 1$  with probability  $1/9$ . This leads to the following outcome:

$$\pi_S^2 = \frac{9}{4} \quad \pi_B^2 = \frac{3}{4} \quad W^2 = 3 = W^* = 3$$

By revealing some information about her type, Sally incentivizes Bob to reveal more about his, leading to a distribution of information points in which efficient trade is possible. In the next section we show that this is a general phenomenon.

### 3.3 Nonstop Communication until Efficient Trade

Our main result is that agents don't stop talking until they reach an efficient outcome. Formally, we will show that if the agents have no further use for rounds of communication after a certain round (finite message complexity) then it must be because they have reached an efficient outcome.

► **Theorem 2 (Main Theorem).** *If the message complexity  $C(\mathcal{D}_S, \mathcal{D}_B) = t < \infty$  for bilateral trade, then the equilibrium allocation is efficient, i.e.,*

$$\pi_S^t(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^t(\mathcal{D}_S, \mathcal{D}_B) = W^*(\mathcal{D}_S, \mathcal{D}_B)$$

where  $W^*(\mathcal{D}_S, \mathcal{D}_B)$  is the welfare of the optimal allocation.

We will apply a recursive argument. First we show in Lemma 3 that if the allocation can't improve with any communication (zero message complexity) then the allocation must already be efficient. In the recursive step, we argue that if there are distributions  $(\mathcal{D}_S, \mathcal{D}_B)$  which both have finite message complexity and lead to an inefficient outcome, then it is possible to construct distributions  $(\mathcal{D}'_S, \mathcal{D}'_B)$  with strictly smaller message complexity that also lead to an inefficient outcome.

► **Lemma 3.** *If  $C(\mathcal{D}_S, \mathcal{D}_B) = 0$ , then the equilibrium allocation is efficient, i.e.,*

$$\pi_S^0(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^0(\mathcal{D}_S, \mathcal{D}_B) = W^*(\mathcal{D}_S, \mathcal{D}_B).$$

<sup>6</sup> The reader is invited to check that this is a valid refinement (as defined in Section 2.3.2) since  $\frac{1}{3} \cdot \frac{1}{2} + \frac{2}{3} \cdot \frac{1}{4} = \frac{1}{3}$ .

## 105:12 Interactive Communication in Bilateral Trade

**Proof.** Let  $F(\mathcal{D}_S)$  be the message where Sally reveals her true cost (full revelation). Then we know that:

$$\pi_S^1(\mathcal{D}_S, \mathcal{D}_B) = \pi_S^0(\mathcal{D}_S, \mathcal{D}_B) = \mathbb{E}_{\mathcal{D}'_S \sim F(\mathcal{D}_S)}[\pi_S^0(\mathcal{D}'_S, \mathcal{D}_B)] = \mathbb{E}_{\mathcal{D}'_S \sim F(\mathcal{D}_S)}[\pi_S^1(\mathcal{D}'_S, \mathcal{D}_B)]$$

where the first equality comes from the fact that  $\mathbf{C}(\mathcal{D}_S, \mathcal{D}_B) = 0$ , the second equality follows from the fact that  $\pi_S^0$  is linear in  $\mathcal{D}_S$  and the third follows from Lemma 1, since  $\mathcal{D}'_S$  has only a single point in its support. Hence Sally has a message in period  $t = 2$  that keeps her utility unchanged but improves Bob's utility if the allocation is inefficient. Therefore, the message complexity cannot be 0. ◀

Finally, we describe how the message complexity of a node in the protocol relates to the message complexity of its children. Recall the definition of a *child* from Section 2.4.

► **Lemma 4.** *If  $\mathbf{C}(\mathcal{D}_S, \mathcal{D}_B) = t$  with  $1 \leq t < \infty$ , then any child of  $(\mathcal{D}_S, \mathcal{D}_B)$  in period  $t$  must have message complexity at most  $t - 1$ .*

**Proof.** We prove by contradiction. Assume  $t$  is even (Sally's turn to communicate) and  $(\mathcal{D}'_S, \mathcal{D}_B)$  is a child of  $(\mathcal{D}_S, \mathcal{D}_B)$  with message complexity  $t' \geq t$ . Then there is some finite  $T$  with the same parity as  $t - 1$  (to make it the same player's turn) satisfying  $t \leq T \leq t' + 1$ ,

$$\pi_S^T(\mathcal{D}'_S, \mathcal{D}_B) + \pi_B^T(\mathcal{D}'_S, \mathcal{D}_B) > \pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B) + \pi_B^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)$$

by the definition of message complexity. Furthermore, by voluntary communication we have that:

$$\pi_S^T(\mathcal{D}'_S, \mathcal{D}_B) \geq \pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B), \quad \pi_B^T(\mathcal{D}'_S, \mathcal{D}_B) \geq \pi_B^{t-1}(\mathcal{D}'_S, \mathcal{D}_B).$$

Therefore, either  $\pi_S^T(\mathcal{D}'_S, \mathcal{D}_B) > \pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)$ , or  $\pi_S^T(\mathcal{D}'_S, \mathcal{D}_B) = \pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)$  and  $\pi_B^T(\mathcal{D}'_S, \mathcal{D}_B) > \pi_B^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)$ . In the first Sally could use her message of the  $t$ -th period in the  $(T + 1)$ -th period, obtaining:

$$\pi_S^{T+1}(\mathcal{D}_S, \mathcal{D}_B) \geq \mathbb{E}_{\mathcal{D}'_S \sim m_S^t(\mathcal{D}_S, \mathcal{D}_B)}[\pi_S^T(\mathcal{D}'_S, \mathcal{D}_B)] > \mathbb{E}_{\mathcal{D}'_S \sim m_S^t(\mathcal{D}_S, \mathcal{D}_B)}[\pi_S^{t-1}(\mathcal{D}'_S, \mathcal{D}_B)] = \pi_S^t(\mathcal{D}_S, \mathcal{D}_B).$$

Or in the second case, she can use the same construction to obtain an outcome that gives her the same utility while benefiting Bob, leading to  $\pi_S^{T+1}(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^{T+1}(\mathcal{D}_S, \mathcal{D}_B) > \pi_S^t(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^t(\mathcal{D}_S, \mathcal{D}_B)$ .<sup>7</sup> In either case we have a contradiction with the premise that  $\mathbf{C}(\mathcal{D}_S, \mathcal{D}_B) = t$ .

The same argument holds if  $t$  is odd and  $(\mathcal{D}_S, \mathcal{D}'_B)$  is a child with message complexity  $t' \geq t$  by replaying the argument reversing the roles of Sally and Bob. ◀

**Proof of Theorem 2.** Define the following sets of pairs of distributions:

$$\Gamma = \{(\mathcal{D}_S, \mathcal{D}_B) \mid \mathbf{C}(\mathcal{D}_S, \mathcal{D}_B) < \infty \text{ and } \pi_S^\infty(\mathcal{D}_S, \mathcal{D}_B) + \pi_B^\infty(\mathcal{D}_S, \mathcal{D}_B) < W^*(\mathcal{D}_S, \mathcal{D}_B)\}.$$

We want to show that  $\Gamma = \emptyset$ . If  $\Gamma$  is not empty, choose any pair  $(\mathcal{D}_S, \mathcal{D}_B)$  with smallest possible message complexity  $t = \mathbf{C}(\mathcal{D}_S, \mathcal{D}_B)$ . We know by Lemma 3 that  $t \geq 1$  hence state  $(\mathcal{D}_S, \mathcal{D}_B)$  must have an inefficient child  $(\mathcal{D}'_S, \mathcal{D}_B)$  or  $(\mathcal{D}_S, \mathcal{D}'_B)$  in period  $t - 1$ . By Lemma 4 this pair must have message complexity at most  $t - 1$  contradicting the minimality of  $(\mathcal{D}_S, \mathcal{D}_B)$ . ◀

<sup>7</sup> If Sally has a strategy to strictly improve her utility, we have  $\pi_S^{T+1}(\mathcal{D}_S, \mathcal{D}_B) > \pi_S^t(\mathcal{D}_S, \mathcal{D}_B)$ . Otherwise, Bob's utility must satisfy  $\pi_B^{T+1}(\mathcal{D}_S, \mathcal{D}_B) > \pi_B^t(\mathcal{D}_S, \mathcal{D}_B)$ . There is always  $\pi_S^{T+1}(\mathcal{D}_S, \mathcal{D}_B) \geq \pi_S^t(\mathcal{D}_S, \mathcal{D}_B)$  and  $\pi_B^{T+1}(\mathcal{D}_S, \mathcal{D}_B) \geq \pi_B^t(\mathcal{D}_S, \mathcal{D}_B)$  by voluntary communication.

### 3.4 Efficient Communication Conjecture

Theorem 2 leaves it possible that there are instances of bilateral trade where the allocation indefinitely improves with the increasing rounds of the communication protocol. This is an intriguing possibility, but based on evidence from computer simulations, we conjecture that an efficient allocation is always reached in finitely-many rounds of communication. We name it the “Efficient Communication Conjecture” (ECC) and formalize it below:

► **Efficient Communication Conjecture.** *In the bilateral trade setting, if  $\Theta_B$  and  $\Theta_S$  are finite types spaces, then for any distributions  $\mathcal{D}_B \in \Delta(\Theta_B)$  and  $\mathcal{D}_S \in \Delta(\Theta_S)$  the message complexity is finite:  $\mathbf{C}(\mathcal{D}_B, \mathcal{D}_S) < \infty$ .*

Resolving in either way would lead to an interesting result. Either we show that efficient allocation is always reached after finitely-many rounds or we exhibit the possibility where agents would prefer to have an infinitely-long conversation before trade.

In the Section 3.5 we make progress on this question by showing an important class of bilateral trade instances where an efficient allocation is achieved with two rounds of communication. Whenever Bob’s type space is binary ( $|\Theta_B| = 2$ ) and Sally’s type space is finite ( $|\Theta_S| < \infty$ ), we show that for any distributions  $\mathcal{D}_B$  and  $\mathcal{D}_S$  supported on  $\Theta_B$  and  $\Theta_S$  respectively, we have  $\mathbf{C}(\mathcal{D}_B, \mathcal{D}_S) \leq 2$ . In that case, we can completely describe the communication protocol.

Whenever Bob’s type space is larger ( $|\Theta_B| > 2$ ), two messages are no longer enough. Consider the following example where  $\Theta_B = \{3, 6, 12\}$  and  $\Theta_S = \{0, 2\}$ . Let  $\mathcal{D}_B$  be the uniform distribution over  $\Theta_B$  and  $\mathcal{D}_S$  be the distribution that puts  $1/5$  probability on the high type. The welfare of the efficient allocation is  $W^* = 33/5$ . In the table below we write the payoffs of Sally and Bob and the welfare for  $t = 0, 1, 2$ .

|   | $t = 0$ | $t = 1$ | $t = 2$ |
|---|---------|---------|---------|
| S | 58/15   | 58/15   | 62/15   |
| B | 8/5     | 12/5    | 12/5    |
| W | 82/15   | 94/15   | 98/15   |

Observe that the welfare is strictly smaller than  $W^*$ . By Theorem 2 we have that further rounds of communication must improve the welfare and therefore  $\mathbf{C}(\mathcal{D}_S, \mathcal{D}_B) \geq 3$ . In the full version, we present the details of this example.

### 3.5 Short protocol for binary buyer types

For an important class of bilateral trade instances we show that a short (two-round) communication protocol leads to efficiency. For this class, we will be able to fully characterize the communication protocol. The assumption will be that the buyer has a binary type space, i.e.  $|\Theta_B| = 2$ . In order to construct the communication protocol, we will develop a higher-order indifference argument. The original indifference argument in [10] decomposes a prior into posterior distributions where the seller is indifferent between different prices. Here: since the seller knows that the buyer will signal in order to make her indifferent, she will send a preliminary signal that makes the buyer indifferent among different ways to make the seller indifferent.

► **Theorem 5 (Efficient Communication for Buyer Binary Types).** *Let  $\Theta_B = \{v_1, v_0\}$  and  $\Theta_S = \{c_1, c_2, \dots, c_n\}$ . Then in two rounds of communication, Bob and Sally achieve an efficient allocation.*

## 105:14 Interactive Communication in Bilateral Trade

The rest of this section is dedicated to proving Theorem 5. We will use  $p$  to denote the probability that the buyer has the high type and let  $q$  be a vector in the  $n$ -simplex  $\Delta_n = \{q \in [0, 1]^n; \sum_i q_i = 1\}$  representing the probability that Sally has each type. This way, we can represent the pair  $(\mathcal{D}_B, \mathcal{D}_S)$  by  $(p, q)$ .

Without loss of generality we will also assume that:

$$c_1 < c_2 < \dots < c_n < v_0 < v_1$$

since any of Sally's type such that  $c_i > v_0$  is ignored by Bob since she will price at  $v_1$  or higher anyway if she has that type, regardless of what her information about Bob's type is.

To prove this theorem, we will first characterize  $\pi^0(p, q)$  and then obtain  $\pi^1(p, q)$  and  $\pi^2(p, q)$  using the equilibrium definition (equation (4)). We will view this as a process of alternating concavification: we alternatively replace the payoffs of Bob and Sally by their concave hulls.

### 3.5.1 Characterizing $\pi^0$

Without any communication, Sally's decision only depends on  $p$  and her type, and hence the payoffs will be linear in  $q$ , i.e.:

$$\pi_i^0(p, q) = \sum_{j=1}^n q_j \cdot \pi_i^0(p, e_j), \quad \forall (p, q) \in [0, 1] \times \Delta_n, i \in \{B, S\}$$

where  $e_j$  is the  $j$ -th unit vector. Therefore, we can focus on understanding the cases for  $q = e_j$ . If Sally has cost  $c_j$  and Bob has probability  $p$  of being the high type, Sally will price at  $v_1$  whenever:


$$p(v_1 - c_j) > v_0 - c_j$$

and at  $v_0$  otherwise. This gives the threshold  $p_j^*$  defined as follows:


$$p_j^* = \frac{v_0 - c_j}{v_1 - c_j}$$

It will be convenient in the following analysis to set  $p_0^* = 1$ .

We can plot the payoff curves for Bob and Sally at  $q = e_j$  in Figure 2. Before  $p_j^*$ , Sally prices at Bob's low type and get constant revenue equal to  $v_0 - c_j$ . After  $p_j^*$ , Sally prices at Bob's high type and his payoff is zero. For each  $q \in \Delta_n$ , Bob's utility is a combination of  $n$  curves like the ones in Figure 2 with breakpoints at  $p_i^*$ . See the left side of Figure 3.



■ **Figure 2** Payoffs of Bob (blue) and Sally (red) without communication for  $q = e_j$ .



■ **Figure 3** Bob’s best response in  $t = 1$  for a generic distribution  $q$ .

### 3.5.2 Characterizing $\pi^1$

From the previous subsection we understand the outcome if we reach the last round ( $t = 0$ ) for each pair  $(p, q)$ . From that we can infer what should be Bob’s response in the second-to-last round ( $t = 1$ ) for each pair  $(p, q)$ . By an information refinement, Bob can obtain a convex combination of payoffs  $\pi_B^0(p', q)$  as long as the distributions average to  $p$ . In other words, we can translate the equilibrium condition (4) as:

$$\pi_B^1(p, q) = \max_{\lambda} \sum_u \lambda_u \cdot \pi_B^0(p_u, q) \quad \text{s.t.} \quad \sum_u \lambda_u \cdot p_u = p, \quad \sum_u \lambda_u = 1, \quad \lambda_u \in [0, 1]$$

which means that  $\pi_B^1(p, q)$  is the concave hull of  $\pi_B^0(p, q)$  along the  $p$ -coordinate. The convex combination coefficients  $\lambda_u$  tell us how to obtain  $\pi_B^1(p, q)$  from  $\pi_B^0(p, q)$ .

For any fixed  $q$ , the function  $p \mapsto \pi_B^0(p, q)$  is piecewise linear with peaks at  $p_j^*$ . Bob’s signal at  $t = 1$  is obtained by taking the concave hull of his utility curve and signaling using the endpoints (right side of Figure 3). For the example in the figure, Bob will use the following strategy:

- Stay silent if  $p \leq p_4^*$ .
- If  $p \in [p_4^*, p_3^*]$ , refine  $p$  to the endpoints  $p_4^*$  and  $p_3^*$ .
- If  $p \in [p_3^*, p_1^*]$ , refine  $p$  to the endpoints  $p_3^*$  and  $p_1^*$ .
- If  $p \in [p_1^*, 1]$ , refine  $p$  to the endpoints  $p_1^*$  and 1.

### 3.5.3 Characterizing $\pi^2$

We will show that at  $t = 2$ , Sally has a signal that keeps Bob’s payoffs unchanged and captures the remaining surplus from the optimal allocation:

$$\pi_S^2(p, q) = W^*(p, q) - \pi_S^1(p, q) \quad \pi_B^2(p, q) = \pi_B^1(p, q) \tag{5}$$

To construct this signal, we will first define one *indifference distribution* for each possible support of Sally’s distribution. For each non-empty subset  $X = \{x(1), x(2), \dots, x(k)\} \subseteq [n]$  define the distribution  $q^X$  over  $[n]$  as follows (using  $p_0^* = 1$  and  $x(0) = 0$ ):

$$q_{x(i)}^X = \frac{(p_{x(i)}^*)^{-1} - (p_{x(i-1)}^*)^{-1}}{(p_{x(k)}^*)^{-1} - 1} \quad \text{for } i = 1, \dots, k \quad \text{and} \quad q_j^X = 0 \quad \text{for } j \notin X$$


The advantage of indifference distributions is that the peaks of  $\pi_B^0(p, q^X)$  are aligned (see Figure 4) and hence if Bob’s information about Sally is  $q^X$  at  $t = 1$ , then Bob’s optimal strategy is the following:

- If  $p \leq p_{x(k)}^*$ , Bob will stay silent.

## 105:16 Interactive Communication in Bilateral Trade

■ If  $p > p_{x(k)}^*$ , Bob will refine his signal to either 1 or  $p_{x(k)}^*$ . In either case, we are at an efficient allocation. If  $p \leq p_{x(k)}^*$ , then Sally prices at the lowest signal of Bob's support. If  $p = 1$  then Bob has the high type with probability 1 and Sally prices at that point. Hence:

$$\pi_S^1(p, q^X) + \pi_B^1(p, q^X) = W^*(p, q^X) \quad (6)$$



■ **Figure 4** Bob's best response in  $t = 1$  for an indifference distribution  $q$ .

With that we are ready to prove the statement of equation (5) which directly implies Theorem 5. We will decompose the proof in two lemmas:

► **Lemma 6.** *Let  $q$  be any distribution over  $[n]$ . Then there exists a decomposition of  $q$  into nested indifference distributions: there are subsets*

$$[n] \supseteq X_1 \supseteq X_2 \supseteq \dots \supseteq X_k \neq \emptyset$$

such that:

$$q = \sum_{i=1}^k \lambda_i q^{X_i} \quad \text{s.t.} \quad \sum_{i=1}^k \lambda_i = 1 \quad \text{and} \quad \lambda_i \geq 0$$

**Proof.** Let  $X$  be the support of  $q$  and define  $z = \min_{i \in X} q_i^X / q_i$ . Since both  $q$  and  $q^X$  have support  $X$ , the components  $q_i^X$  and  $q_i$  are strictly positive, hence  $z > 0$ . Also, observe that  $z = \min_i q_i^X / q_i \leq (\sum_i q_i^X) / (\sum_i q_i) = 1$  since both are distributions. If  $z = 1$  then  $q = q^X$  and we are done since  $q$  has a trivial decomposition into indifference distributions.

If  $z < 1$ , define  $q' \in \Delta_n$  as follows:

$$q' = \frac{q - zq^X}{1 - z}$$

To see that  $q'$  is a distribution, observe first that  $q_i - zq_i^X \geq 0$  by the definition of  $z$  and that:

$$\sum_i q'_i = \frac{1}{1 - z} \sum_i (q_i - zq_i^X) = \frac{1 - z}{1 - z} = 1$$

Finally, note that if  $i$  is the index such that  $z = q_i^X / q_i$  we have  $q'_i = 0$ . Hence we showed how to decompose  $q$  into a distribution of the form  $q^X$  and a distribution with support strictly contained in  $X$ . Applying the construction recursively for  $q'$  we obtain the decomposition in the statement. ◀

► **Lemma 7.** *Sally's signal that at  $t = 2$  refines her distribution into indifference distributions is her best response and satisfied equation (5).*



**Proof.** Let's assume without loss of generality that  $q$  has full support and that  $p > p_n^*$ . If  $p \leq p_n^*$ , then  $\pi_S^0(p, q) + \pi_B^0(p, q) = W^*(p, q)$  since Sally always prices at  $v_0$  so equation (5) holds trivially.

Consider the decomposition of Sally's original distribution  $q$  into nested indifference distributions in Lemma 6  $q = \sum_i \lambda_i q^{X_i}$ . Then:

$$\pi_S^2(p, q) \geq \sum_i \lambda_i \pi_S^1(p, q^{X_i}) = \sum_i \lambda_i [W^*(p, q^{X_i}) - \pi_B^1(p, q^{X_i})] = W^*(p, q) - \sum_i \lambda_i \pi_B^1(p, q^{X_i})$$

where the first equality comes from equation (6) and the second due to the fact that  $W^*(p, q)$  is linear in  $q$ .

To complete the proof, we need to show that  $\sum_i \lambda_i \pi_B^1(p, q^{X_i}) = \pi_B^1(p, q)$ . One of the directions is easy: if  $M$  is the optimal refinement for Bob in  $t = 1$  then:

$$\sum_i \lambda_i \pi_B^1(p, q^{X_i}) \geq \sum_i \lambda_i \mathbb{E}_{p' \sim M}[\pi_B^0(p, q^{X_i})] = \mathbb{E}_{p' \sim M}[\pi_B^0(p, q)] = \pi_B^1(p, q)$$

where the second inequality follows from the fact that  $\pi_B^0$  is linear in  $q$ . This inequality in particular, verifies that Sally's signal satisfies voluntary communication.

Showing the other direction is the delicate aspect of the proof and will use the fact that the sets  $X_i$  are nested. Start by choosing the smallest  $i$  such that

$$\min_{x \in X_i} p_x^* > p$$

If no such index exists, set  $i = k + 1$  and  $p_a^* = 1$ . Otherwise, we know by the assumption that  $p > p_n^*$  and hence  $i > 1$ . Let  $a = \max\{x; x \in X_i\}$  and  $b = \min\{x; x \in X_{i-1}\}$ . By the definition we have that  $p_b^* \leq p \leq p_a^*$ . The crucial observation is that the payoff functions  $p \mapsto \pi_B^1(p, q^{X_i})$  are linear for  $p \in [p_b^*, p_a^*]$ . The reason is that the  $\pi_B^1(p, q^{X_i})$  is continuous and piecewise linear with two pieces with a breakpoint at  $\min_{x \in X_i} p_x^*$ . Hence the interval  $[p_b^*, p_a^*]$  is always within the same linear piece. If  $p = z_a p_a^* + z_b p_b^*$  then:

$$\sum_i \lambda_i \pi_B^1(p, q^{X_i}) = \sum_i \lambda_i [z_a \pi_B^0(p_a^*, q^{X_i}) + z_b \pi_B^0(p_b^*, q^{X_i})] = z_a \pi_B^0(p_a^*, q) + z_b \pi_B^0(p_b^*, q) \leq \pi_B^1(p, q)$$

The first equality follows from the observation about the linearity of  $\pi_B^1(p, q^{X_i})$  in  $[p_b^*, p_a^*]$ , the second follows from the linearity of  $\pi_B^0(p, q)$  on  $q$  and the third is simply because the refining  $p$  to  $p_b^*$  and  $p_a^*$  is a valid strategy for Bob. ◀

#### 4 A Spy Game (with Longer Communication)

We conclude by showing an example of a game with a longer communication protocol. In the tradition of Cold War era game theory we will have Sally (the Spymaster) and Bob (the Birdwatcher) be spies. As with good spies, no one really knows which country they are actually serving. Their types  $\theta_S, \theta_B \in \{0, 1\}$  represent the country they are serving. If  $\theta_S = \theta_B$  then Sally and Bob are friends and should cooperate. If  $\theta_S \neq \theta_B$  they are enemies and they have opposing goals. In their meeting in West Berlin (Glienicke Bridge, more precisely), Sally must decide either to cooperate with Bob (C) or to expose him (E). The actions and types lead to the following utilities for Sally and Bob respectively in the base game:

|                          | C       | E        |
|--------------------------|---------|----------|
| $\theta_S = \theta_B$    | (1, 1)  | (-1, -1) |
| $\theta_S \neq \theta_B$ | (-2, 2) | (2, -2)  |

**105:18 Interactive Communication in Bilateral Trade**

If Sally and Bob are friends ( $\theta_S = \theta_B$ ) then they are playing a cooperative game: both get a payoff of 1 for cooperating and both get  $-1$  if Sally decides to expose Bob. If Sally and Bob are enemies ( $\theta_S \neq \theta_B$ ) they are playing a zero-sum game. If Sally cooperates with an enemy, she loses  $(-2)$  and Bob wins  $(+2)$ . If she exposes an enemy then Sally wins  $(+2)$  and Bob loses  $(-2)$ .


We analyze this game by following the same methodology as in Section 3.5 of taking alternate concave hulls in Bob's and Sally's dimensions (see implementation in footnote 1). Even though it is a game with binary types and only two actions, the communication protocol arising in equilibrium is quite complex. For example, the following table shows the payoffs of Sally and Bob from a protocol with  $t$  periods if they start from uniform distributions:

|   | $t = 0$ | $t = 1$ | $t = 2$ | $t = 3$ | $t = 4$ | $t = 5$ | $t = 6$ |
|---|---------|---------|---------|---------|---------|---------|---------|
| S | 0.5     | 0.5     | 0.722   | 0.722   | 0.738   | 0.769   | 0.801   |
| B | -1.5    | -0.166  | -0.166  | -0.107  | -0.107  | -0.077  | -0.075  |

In Figures 5 and 6 we depict the optimal protocol using the square diagrams similar to the ones introduced in Figure 1. The coordinates of each point  $(p, q) \in [0, 1]^2$  represent the probabilities that Bob ( $p$ ) and Sally ( $q$ ) have type 1.

We note that the optimal strategy can be represented simply by a partition of the square  $[0, 1]^2$  into sub-rectangles. Assume we are in state  $(p, q)$  in round  $t$  and its inside a sub-rectangle  $[p_0, p_1] \times [q_0, q_1]$ . If  $t$  is odd and Bob will refine  $(p, q)$  to  $(p_0, q)$  and  $(p_1, q)$ . If  $t$  is even, Sally will refine  $(p, q)$  to  $(p, q_0)$  and  $(p, q_1)$ .


At  $t = 0$  Sally's action varies whether her belief is in the intervals  $[0, 1/3]$ ,  $(1/3, 2/3)$  or  $[2/3, 1]$ . In the middle interval she always exposed Bob and in the first and third interval she will cooperate if Bob's likeliest type is the same as her type.



■ **Figure 5** Best response of Sally and Bob in the first few iterations.


What follows next is a delicate game where agents slowly reveal their types depending on how much information they have about other agents. For example, Bob's strategy in  $t = 1$  will be one of the following:

- If Bob has a more extreme information about Sally ( $q \leq 4/9$  or  $q \geq 5/9$ ) he will reveal more information about his type hoping to cooperate more frequently.
- If Bob is more uncertain about Sally's type ( $4/9 < q < 5/9$ ) Bob will reveal less information since he runs the risk of Sally having the opposite type and deciding to expose him when she would otherwise cooperate.



■ **Figure 6** Best response of Sally and Bob in iterations  $t = 3, 4, 5, 6$ .

In  $t = 2$  Sally is again in the same situation: the more certain she is about Bob’s type the more willing she is to tell her own type. Figure 6 shows the evolution of the game in each of the subsequent stages.



■ **Figure 7** Dynamics of the state  $(p, q)$  starting at  $(1/2, 1/2)$  for 5 periods.

Figure 7 shows how the state  $(p, q)$  of the game evolves starting from  $(1/2, 1/2)$  for 5 iterations. What we observe is that a long and gradual disclose of information benefits both parties.

**References**

- 1 Noga Alon, Noam Nisan, Ran Raz, and Omri Weinstein. Welfare maximization with limited interaction. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1499–1512, 2015. doi:10.1109/FOCS.2015.95.
- 2 Sepehr Assadi. Combinatorial auctions do need modest interaction. *ACM Trans. Econ. Comput.*, 8(1), March 2020. doi:10.1145/3381521.
- 3 Robert J Aumann and Sergiu Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660, 2003.
- 4 Robert J Aumann and Michael Maschler. *Repeated games with incomplete information*. MIT press, 1995.

- 5 László Babai, Anna Gál, Peter G. Kimmel, and Satyanarayana V. Lokam. Communication complexity of simultaneous messages. *SIAM J. Comput.*, 33(1):137–166, 2003. doi:10.1137/S0097539700375944.
- 6 Moshe Babaioff, Yang Cai, Yannai A Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 373–373, 2018.
- 7 Moshe Babaioff, Kira Goldner, and Yannai A Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2452–2471. SIAM, 2020.
- 8 Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2545–2563. SIAM, 2018.
- 9 Santiago R Balseiro, Vahab Mirrokni, Renato Paes Leme, and Song Zuo. Dynamic double auctions: Towards first best. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 157–172. SIAM, 2019.
- 10 Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.
- 11 Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017.
- 12 Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- 13 Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of economic theory*, 137(1):580–609, 2007.
- 14 Liad Blumrosen and Shahar Dobzinski. Reallocation mechanisms. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 617–617, 2014.
- 15 Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 589–590, 2017.
- 16 Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. On multi-dimensional gains from trade maximization. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1079–1098. SIAM, 2021. doi:10.1137/1.9781611976465.67.
- 17 Kalyan Chatterjee and William Samuelson. Bargaining under incomplete information. *Operations research*, 31(5):835–851, 1983.
- 18 Riccardo Colini-Baldeschi, Paul W Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. *ACM Transactions on Economics and Computation (TEAC)*, 8(1):1–29, 2020.
- 19 Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1424–1443. SIAM, 2016.
- 20 Constantinos Daskalakis, Christos Papadimitriou, and Christos Tzamos. Does information revelation improve revenue? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 233–250, 2016.
- 21 Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *the 46th annual ACM symposium on Theory of computing (STOC)*, 2014.
- 22 Laura Doval and Jeffrey C Ely. Sequential information design. *Econometrica*, 88(6):2575–2608, 2020.

- 23 Shaddin Dughmi, David Kempe, and Ruixin Qiang. Persuasion with limited communication. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 663–680, 2016.
- 24 Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. *SIAM Journal on Computing*, 50(3), 2021.
- 25 Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. *ACM Symposium on Theory of Computing, STOC'21*, 2020.
- 26 Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation (TEAC)*, 2(2):1–19, 2014.
- 27 Péter Eső and Balazs Szentes. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):705–731, 2007.
- 28 Matthew Gentzkow and Emir Kamenica. Costly persuasion. *American Economic Review*, 104(5):457–62, 2014.
- 29 Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- 30 Zi Yang Kang and Jan Vondrák. Fixed-price approximations to optimal efficiency in bilateral trade. *Available at SSRN 3460336*, 2019.
- 31 Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.
- 32 Noam Nisan and Avi Wigderson. Rounds in communication complexity revisited. *SIAM J. Comput.*, 22(1):211–219, February 1993. doi:10.1137/0222016.
- 33 Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing, STOC '79*, pages 209–213, New York, NY, USA, 1979. ACM. doi:10.1145/800135.804414.