



Edinburgh School of Economics
Discussion Paper Series
Number 126

Extended Conversations in Sender-Receiver Games

R Vijay Krishna (University of Edinburgh)

Date
November 2004

Published by

School of Economics
University of Edinburgh
30 -31 Buccleuch Place
Edinburgh EH8 9JT
+44 (0)131 650 8361

<http://www.ed.ac.uk/schools-departments/economics>



THE UNIVERSITY *of* EDINBURGH

Extended Conversations in Sender-Receiver Games

R. Vijay Krishna^{*†}

September, 2004

Abstract

Aumann and Hart (Econometrica, Nov. 2003) have shown that in games of one-sided incomplete information, the set of equilibrium outcomes achievable can be expanded considerably if the players are allowed to communicate without exogenous time limits and completely characterise the equilibria from such communication. Their research provokes (at least) four questions. (i) Is it true that the set of equilibrium payoffs stabilises (i.e. remains unchanged) if there are sufficiently many rounds of communication? (ii) Is the set of equilibria from communication which is unbounded but finite with probability one the same as equilibria from communication which is just unbounded? (iii) Are any of these sets of equilibria “simple” and if so, is there an algorithm to compute them? (iv) Does unbounded communication (of order type ω) exhaust all possibilities so that further communication is irrelevant? We show that in the context of finite Sender-Receiver games, the answer to all four is yes if the game satisfies a certain geometric condition. We then relate this condition to some geometric facts about the notion of bi-convexity and argue that if any of the questions has a negative answer then all three of the questions are likely to have a negative answer.

1 Introduction

Since the seminal paper of Crawford and Sobel [11], it has been well understood by game theorists and economists that cheap pre-play communication or cheap talk can expand the

^{*}This paper is based on Chapter 2 of my doctoral dissertation at The Pennsylvania State University. I would like to thank my advisors Profs. Kalyan Chatterjee and Tomas Sjöström for their support, Profs. Françoise Forges and Roland Strausz for very detailed comments on an earlier draft and Prof. Sergiu Hart for pointing out an error in an earlier version. I am also grateful to Prof. James Jordan for the numerous conversations which led to the proof of Theorem A. All errors that remain are my own.

[†]School of Economics, University of Edinburgh, William Robertson Building, 50, George Square, Edinburgh EH8 9JY, UK. *Email:* vijay.krishna@ed.ac.uk.

set of equilibrium outcomes that can be achieved by the players in a game. What is not well understood is the importance of the structure of the communication phase itself. In a wonderful new paper, Aumann and Hart [3] have shown that the Crawford-Sobel model of communication which consists of one round of signalling is significantly lacking in generality in that multiple rounds of communication can greatly increase the set of equilibrium outcomes. They then study unbounded communication in two-player games of one-sided incomplete information and completely characterise the set of equilibrium outcomes.

Their investigations provoke a number of interesting questions. First, Is it the case that the set of equilibria from finite conversations (we shall use conversation to indicate the kind of communication we have in mind) eventually stabilises if we consider sufficiently long conversations? Second, Is it the case that the equilibria from conversations that are finite almost surely are the same as the equilibria from unbounded conversations (which are considered by Aumann and Hart)? Third, Are these sets of equilibria “simple”?¹ And finally, Is the case that unbounded conversations (of order type ω) exhaust all possibilities so that longer questions are unnecessary? We find sufficient conditions for all of these questions to have an affirmative answer.

At this point it is pertinent to discuss the importance and relevance of the four questions posed above. These questions stem from the nature of the Aumann-Hart communication phase. Aumann and Hart consider communication of order type ω so that their games are of order type $\omega + 1$. But while their theory is conceptually simple, it is far from clear just what the practical implications of their theory are. In particular, one might think that just as in the case of most repeated games, it may be possible to approximate the equilibria from unbounded communication by considering games with communication that lasts sufficiently long. After all, finite games are things that are well understood. But as is well known, this is not the case (and an example of this is provided below).

Another possibility is our more mathematical conception of a game. If the players were to submit their strategies to a computing machine (a register machine, say), would the machine be able to tell us the outcome of the play? Unfortunately, we do not know the answer to this question. Of course, in this paper we do not answer this question but what we do achieve is provide sufficient conditions for a Sender-Receiver game so that the aforementioned machine is capable of providing us with the outcome.

To put matters in perspective, recall that the usual assumption in the case of Sender-Receiver games – games where the Sender has some private information and the Receiver

¹We shall call a set *simple* if the set can be built up from simple geometric structures of increasing dimension in finitely many steps. More technically, in the language of algebraic topology, a simple set has a simplicial decomposition.

takes the only payoff relevant action – is that there is one round of signalling after which the Receiver takes an action. However, a drawback of such a description of the state of affairs is precisely that it involves *only* one round of signalling by the Sender. In principle, there is no reason to assume *a priori* that there cannot be any further communication. Indeed, if opportunities for further communication did present themselves, the agents cannot commit to not using them. Assumptions to the contrary seem strange especially when we are assuming that the agents cannot make commitments.

We may therefore posit that the natural objects to consider are conversations that do not have exogenous deadlines. The reader may feel that talk without a finite deadline, which we shall refer to as *unbounded talk*, is unnatural. We would like to argue that it is talk with an exogenous deadline that is unnatural. We want to consider a scenario where the agents cannot commit to anything and are free to do as they please. There are no artificial restrictions on how long they can talk. The infinite horizon should be viewed only as the possibility of the players coming back and having another round of talk. It is but a thought process that gives both players real strategic alternatives which may not exist in the presence of deadlines. All that is required is that there be no exogenous deadline. A similar view has been taken by Rubinstein in [24] about the use of infinitely repeated games, where he says:

“By using infinite horizon games we do not assume that the real world is infinite. . . . An infinitely repeated game is meant to assist in analyzing specific situations where players examine a long-term situation without assigning a specific status to the end of the world.”

Another example where the absence of a deadline makes all the difference is the Rubinstein bargaining model [23], where the players could bargain forever but only play for one period in the unique equilibrium. Here too, it is the *possibility* of infinite interactions that leads (immediately) to the optimal outcome. It should be noted that Sender-Receiver games are special in that it is after the so-called *talking phase* that the Receiver takes an action. Thus, having unbounded talk could mean that the agents talk for infinitely many periods and the Receiver takes an action after that. Or it could mean that we only consider conversations that are unbounded but end in finite time with probability 1. Both possibilities will be examined in the sequel and we will demonstrate that in Sender-Receiver games such distinctions are irrelevant if the game satisfies a finiteness condition. This finiteness condition is related to geometric structure of the payoffs of the game and is not known to be violated by any game.

In this paper we restrict attention to communication in Sender-Receiver games where the Sender has finitely many types and the Receiver has finitely many actions. (Precise

definitions can be found in §3.) In our main result (Theorem A), we show that in Sender-Receiver games, the set of equilibrium outcomes that can be achieved with unbounded conversations is the same as the set of equilibrium outcomes that can be achieved with unbounded conversations that are finite with probability 1 when the game satisfies the aforementioned finiteness condition. This is a significant result for two reasons. The first reason is that the possibility of unbounded conversation usually provides a significant enlargement to the set of achievable outcomes. The second reason is that it tells us that our intuitive reasons for considering conversations with no deadline is without loss of generality as we don't have to consider conversations that literally last forever to exhaust all the possibilities for information transmission.

Following Aumann and Hart, there are some other questions that one may ask of cheap talk equilibria. One of these is whether the set of equilibria from finite conversations stabilise after some large finite number of stages. Another question is whether the sets of equilibria are simple. A third question is whether there can be any more information that can be transmitted after the players talk for infinitely many periods. It turns out that all of these questions are closely related and an affirmative answer to the first two questions and a negative answer to the third depends on whether some underlying sets are well behaved. We go on argue that all of these properties, in a sense, come together, hence representing a dichotomy of the good and the bad.

We achieve this by exploiting the special geometric structure of the graph of payoffs in Sender-Receiver games. Aumann and Hart characterise equilibria of games with extended talk (which they call *long cheap talk*) in terms of stochastic processes called *di-martingales*, that converge to limit distributions that lie (almost surely) in the graph of payoffs of the incomplete information game and show in [2] that the associated *bi-convex* set can be characterised in terms of separation by *bi-convex* functions. They also note that in general, the process of separation may be transfinite. In Sender-Receiver games, this graph essentially depends on only finitely many points – a fact which has significant consequences.

The remainder of the paper is structured as follows. In §2, we describe the relative position of our results with respect to the extant literature. In §3, we give a formal description of the model. In §4, we discuss the mechanics of information transmission in cheap talk games. In §4.1, we give a qualitative description of the workings of conversations and argue that one can restrict attention to the so-called canonical conversations. In §4.2, we describe the mathematical underpinnings of the Aumann-Hart theory which we shall use extensively in the sequel and also describe briefly the intuition behind the Aumann-Hart results. In §4.3, we discuss some of the more delicate game-theoretic aspects relating to the definitions of strategy and equilibrium concept used. We then prove our result on unbounded conversations in Sender-Receiver games (Theorem A) in §5. In §6, we consider both finite

conversations and transfinite conversations and discuss the nature of our sufficient condition. §7 concludes and Appendix A provides a complete and self-contained (and therefore, a necessarily concise) description of the mathematical concepts introduced in [2], missing proofs and an axiomatisation of the graph of the payoff correspondence of the silent game.

2 Related Literature

The theory of strategic information transmission when there are no signalling costs can be traced back to Crawford and Sobel [11], who consider a Sender-Receiver game with one round of communication where the preferences of the Sender and Receiver are not completely aligned. They show that as the preferences diverge, the amount of information transmitted decreases and for full revelation, their preferences must be perfectly aligned. There is also the literature on approaches to contracting with imperfect commitment (of which our games would be an extreme example). Here, Bester and Strausz [6] provide a characterisation of the set of incentive efficient equilibria when there is one round of communication. They convert the problem to a programming problem, albeit under the condition of *incentive efficiency*, which is stronger than the condition of incentive compatibility. They also deal with all intermediate cases of partial commitment and show that there is an equilibrium where all types signal truthfully with positive probability. The papers mentioned above represent one part of the literature on cheap talk, which is concerned with the expansion of the set of equilibria through talk. There is another part which is interested in refining the set of equilibria, for example Blume and Sobel [7] who consider the effects of introducing further communication possibilities on the set of equilibria and are interested in equilibria that are immune to such possibilities.

The study of extended conversations arose out of considerations stemming from the study of (undiscounted) repeated games of one-sided incomplete information. In that context, a complete characterisation is provided by Hart [15] using the concepts of bi-convexity and bi-martingales which are formally introduced by Aumann and Hart [2]. Cheap talk is conceptually simpler than repeated games because we do not have to keep track of infinite streams of payoffs, only of the probabilities. Another advantage in the case of cheap talk is that the issue of existence of Nash equilibria is settled trivially which is not the case in the case of repeated games (see [26]).

Extended conversations are described formally by Aumann and Hart in [3] who provide a complete characterisation of the set of equilibrium payoffs from unbounded conversations. An example that shows that extended talk can increase the set of equilibrium outcomes is due to Forges [13] who shows that unbounded conversations that are finite with probability 1 provide outcomes not achievable with finite conversations. Needless to say, her analysis

uses the concepts of bi-convexity and bi-martingales. She also provides a host of examples in [14]. Another example due to Simon [25], which we shall study in detail in the sequel, has the property that finite conversations of any length do not yield any equilibria other than the babbling equilibrium (which we shall call the equilibrium of the silent game), whereas infinite talk allows for Pareto improvements for all types of the Sender and the Receiver.

The benchmark for communication in games of incomplete information is the use of the disinterested mediator which entails the players sending their private information to the mediator and the mediator making incentive compatible probabilistic recommendations of actions to the players. This represents a *mediated solution* and any outcome that can be achieved by any communication mechanism can be realised as a mediated solution. For a thorough discussion of the so-called *Bayesian incentive compatible mechanisms*, the reader is referred to Chapter 6 of [21].

There is another literature which studies the expansion of the set of equilibria through unmediated communication mechanisms, but uses protocols in a manner akin to the computer science literature. Most notable among these is the paper Ben-Porath [5] who shows that in three person games of incomplete information, there exists a communication protocol which enables the players to implement any mediated solution. (By Theorems A and B of Aumann and Hart [3], there can be no such protocol in the case of two player games of incomplete information.) Another closely related paper is due to Urbano and Vila [28] who show that when players have bounded computational abilities, any correlated equilibrium in a two-person normal form game of complete information can be implemented via an appropriate protocol. They use a variant of the idea of public key encryption and rely on the fact that certain algebraic operations such as exponentiation and taking logarithms in prime fields are very complicated and very hard to compute.

3 Model

A Sender-Receiver game Γ can be characterised by $\Gamma := (T_S, C_R, p, u_S, u_R)$, where T_S is a finite set of types of the Sender with $|T_S| = k$,² C_R is the set of actions of the Receiver (which is assumed to be finite), $p \in \Delta^{k-1}$ is a probability vector representing the Receiver's prior beliefs about the Sender's type (where $\Delta^{k-1} := \{p \in \mathbb{R}^k : \sum_{i=1}^k p_i = 1\}$ is the $(k-1)$ -dimensional simplex) with $p \gg 0$ and $u_S : C_R \times T_S \rightarrow \mathbb{R}$ and $u_R : C_R \times T_S \rightarrow \mathbb{R}$ are utility functions³ for the Sender and Receiver respectively. In other words, the only payoff

²For any set A , we denote its cardinality by $|A|$.

³The utility functions u_S and u_R are extended to mixed strategies by linearity. For notational simplicity, we shall define utility functions only over pure strategies in the sequel, with the understanding that they can be extended to mixed strategies in the obvious way.

relevant action is taken by the Receiver.

We will now consider the extended Sender-Receiver game—the game with communication. (Since the point of a Sender-Receiver game is to model communication, we shall refer to a Sender-Receiver game with communication as a Sender-Receiver game in the sequel.) The sequence of events is as follows. First, Nature picks the Sender's type at random. As in [3], we shall call this the *information phase*. This is followed by a *communication phase*. The communication phase consists of players sending *verbal* messages to each other from a finite message set M (where $|M| > 1$). Here, we shall take verbal to mean that there is no notion of verifiability (i.e. it is not possible to determine which type of a player sent a particular message) while refraining from giving a formal definition. The length of the communication phase is denoted by \mathcal{L}_c , a random variable that maps into the *class* of ordinals.⁴ (Note that although \mathcal{L}_c can be *any* ordinal number (up to 2^ω), we will primarily be interested in the cases where \mathcal{L}_c is either less than ω (i.e. finite) a.s. or is a constant taking the value ω , the latter case being the one considered in Aumann and Hart [3].) Finally, there is an *action phase*, where the Receiver takes an action. At the end of of the communication phase, the Receiver has some posterior beliefs (possibly different from the priors) about the Sender's type. Given these beliefs, he takes an action. Let $\beta : C_R \rightarrow \mathbb{R}$ be his expected payoffs at the end of the communication phase from each action taken. In what follows, we shall also denote the vector of payoffs that the Sender gets by $a : C_R \rightarrow \mathbb{R}^k$. We assume that there is perfect recall and that the model of communication is commonly known between the players.

4 The Mechanics of Cheap Talk

In this paper, we consider plain talk, also called *cheap talk* because it is costless, unmediated, non-verifiable and non-binding. In this section we shall describe the so-called *canonical conversations* and the underlying mathematical notions following Aumann and Hart in [3]. In §4.1 we shall informally describe the mechanics of cheap talk and in §4.2 we shall introduce the mathematics that describes cheap talk. In §4.3 we discuss the game theoretic aspects of long cheap talk.

⁴There are some complications that arise if we actually let \mathcal{L}_c take any value in the class of ordinals. Instead, we can assume that \mathcal{L}_c takes values in the *set* of ordinals less than, say, 2^ω , where ω is the first infinite ordinal. Of course, we can take 2^ω to be as large as we wish.

4.1 A Qualitative Discussion

We shall assume here that each player has a finite message set, $M := \{a, b, c, \dots\}$. Cheap talk consists of two components. The first involves signalling by the Sender, where the only message sent is by the Sender and the second, called compromising, involves simultaneous messages sent by both players. As there is no substantive information to be conveyed by the Receiver, we can ignore stages where messages are sent by the Receiver alone. In communication games there always exists a *babbling equilibrium*, an equilibrium where there is no substantive communication. More precisely, the Sender's message does not depend on his type and the Receiver's action depends only on his priors. Following Aumann and Hart [3], we too shall refer to these as the equilibria of the *silent game*. (Thus, an equilibrium always exists, regardless of the length of the game.) Note that the silent game in [3] is actually a game, as both players (potentially) have an action to take, whereas our silent game is just the decision problem of the Receiver given his beliefs.

With only one round of communication, there are two possibilities for cheap talk⁵, one where both the players speak and another where only the Sender speaks. The latter is just signalling, with the only message being sent by the Sender. As mentioned above, there is always the babbling message where the priors on the Sender's type remains unchanged after the message. Now suppose that there exist messages that convey some information, i.e. the posteriors based on the message are different from the priors. In this case, any probability that is a convex combination of the original prior and the posteriors mentioned above can be achieved by appropriately adjusting the probabilities with which the different types send messages. Thus, signalling serves the purpose of convexifying across probabilities. Let us consider an example to make this clear. (The arguments below are adapted from [13] and [14].)

4.1 Example. Signalling.

		c_1	c_2
$[\frac{1}{2}]$	t_1	0,2	1,1
$[\frac{1}{2}]$	t_2	0,0	2,2

Figure 1: Illustrating signalling

⁵Indeed, in what follows, all communication will be assumed to be costless. We shall therefore refer to the signalling and compromising stages of the communication process to differentiate between the two possible forms of communication – both of which are cheap.

Consider the game where the Sender has two types, t_1 and t_2 , each occurring with probability $\frac{1}{2}$. The Receiver has two actions c_1 and c_2 . If there is only one round of signalling, it is easy to see that the Sender will always pretend to be the type t_2 . Let us describe the strategies more precisely. Suppose the possible messages are $M := \{x, y\}$ with generic message being denoted by m . A strategy for the Sender can be described as follows: if type t_i , send x with probability p_i . Then, contingent on the message, the posterior probabilities that the Receiver will have are $q_m := \text{prob}(\text{Sender is type } t_1 | m)$ where $m \in M$. The payoffs the (two types of the) Sender can expect from the message m is a_m . In our game, let us suppose that $p_2 = 0$. Then, it follows that $p_1 = 0$. The reason is that if p_1 were positive, then $a_x^1 < a_y^1$ (because the Receiver now plays c_1 with positive probability which is not good for t_1), thus it is not optimal for type 1 to send x with positive probability. \blackbox

More generally, if each type of the Sender sends both messages with probability, it must be the case that for type j , $a_x^j = a_y^j$. In general, we have⁶

$$a_x^j = a_y^j \text{ implies } p_j \in [0, 1]. \quad (1)$$

If the payoffs to the different messages are not the same, then we get

$$a_x^j \leq a_y^j \text{ implies } p_j = 0 \text{ and} \quad (2)$$

$$a_x^j \geq a_y^j \text{ implies } p_j = 1. \quad (3)$$

In the case that $a_x^j = a_y^j$, we see the first instance of the so-called *martingale* property, i.e. regardless of the message sent by type j , the expected payoff remains the same. The martingale property is essential to the analysis and in longer conversations, there will be instances where, say, $a_x^j \leq a_y^j$. Here, we shall assign the type whose posterior payoff is 0, a *virtual* payoff which is high enough so that the payoff from each signal is the same whereby the martingale property is restored. (See §4.2 for a more geometric description and [3] for a thorough discussion of the issues involved.) Note also that the martingale property can be viewed as convexification in probabilities because if two posteriors can be reached from the two messages for a given prior, then any convex combination of those posteriors can also be taken as the prior probability. In other words, using these messages, we can find signalling probabilities for the two types if the Receiver has different priors as long as the priors are convex combinations of the posterior probabilities.

Aumann and Hart show that the use of simultaneous messages can also be extremely useful. These do not convey information, but are seen as a compromise between the Sender and the Receiver about future courses of action. A simple example of a game from Blume and Sobel [7] illustrates this idea.

		c_1	c_2
$[\frac{1}{2}]$	t_1	1,2	0,0
$[\frac{1}{2}]$	t_2	0,1	1,2

Figure 2: Illustrating a Compromise

4.2 Example. Compromising.

Consider the game where the Sender has two types, t_1 and t_2 , each occurring with probability $\frac{1}{2}$. The Receiver has two actions c_1 and c_2 .

There are two signalling equilibria in this game. There is the babbling equilibrium where the Receiver takes the action c_1 with his expected payoff given by $\beta(c_1) = 1\frac{1}{2}$ and the Sender's payoffs are $a(c_1) = (1, 0)$ (representing the payoffs of each type). There is also a fully revealing (separating) equilibrium where each type reveals the truth. Here, strategies are such that the Sender of type t_1 sends message x and the Sender of type t_2 send message y . Conditional on x , the Receiver plays c_1 and upon receipt of y , the Receiver plays c_2 . This equilibrium gives payoffs $(1, 1)$ to the two types of the Sender and 2 to the Receiver. Note that type t_1 of the Sender is indifferent between revealing the truth and just babbling, but the Receiver is not. Now consider players sending each other messages simultaneously before the signalling period. Let us consider just two messages, $\{a, b\}$. Each player randomises uniformly over the two messages. If the messages sent are $\{a, a\}$ or $\{b, b\}$, then the Sender will reveal his type. If the outcome is $\{a, b\}$ or $\{b, a\}$, the Sender will babble. Thus, we have an equilibrium with expected payoffs $(1, \frac{1}{2})$ to the two types of the Sender and $1\frac{3}{4}$ to the Receiver.

Such a construction is called a *joint lottery* and is used to model the outcome of a compromise, which is supposed to be random. The important feature here is that no player can unilaterally alter the probability distribution over the set of outcomes, hence the term joint lottery. The joint lottery helps us convexify between payoffs, *without* affecting the Receiver's beliefs about the Sender. The convexification helps the players achieve a payoff that was otherwise not achievable. \boxtimes

Thus, simultaneous messages or joint lotteries serve the purpose of convexifying across payoffs. Now consider a conversation of some, possibly unbounded length. It is clear that the conversation can only involve either the Sender sending a message (signalling) or the two players sending simultaneous messages (joint lotteries). We can convert this into a conversation where, say, the odd periods involve signalling and the even periods involve

⁶Equations (1)–(3) represent the so-called *incentive compatibility* conditions in Forges [13].

joint lotteries. This is called a *canonical conversation* in [3], where it is shown that a payoff from a conversation is an equilibrium if and only if it is the payoff to a canonical conversation. Of course, this is not trivial and demonstrating it takes some work. For a fuller discussion with all the relevant details, the reader is referred to [3]. Note also that if the original conversation is finite, then the associated canonical conversation is also finite and has at most twice the length. Clearly, if the original conversation has infinite length, then the associated canonical conversation is also infinite but has the same cardinality, ω .

4.2 The Mathematics of Cheap Talk

To get a better understanding of the mathematics underlying the theory of equilibria of cheap talk games, we first consider martingales that have limits in a particular set. For concreteness, let $A \subset \mathbb{R}^n$ for some n . Consider the set of all expectations of bounded martingales whose limits are almost surely in A . This is nothing but $\text{co}(A)$, the convex hull of A . This is because we can think of a martingale as the splitting of a particle where the center of mass is fixed and therefore the limit distribution is nothing but the limit cloud that forms. At each point in time, the starting point lies in the convex hull of the limit cloud. Thus, the starting points of the splitting process is the convex hull of A . We shall now describe how this idea relates to cheap talk equilibria.

Due to the Aumann-Hart theorem (Theorem A in [3]) that any cheap talk equilibrium can be viewed as a canonical equilibrium, it suffices to restrict attention to these equilibria. Let us now consider stochastic processes in $\mathbb{A} \times \mathbb{B} \times \mathbb{Q}$, where \mathbb{A} is a k -dimensional⁷ compact, convex subset of Euclidean space representing the payoffs that the Sender can have, \mathbb{B} is a non-degenerate interval in which the Receiver gets payoffs and $\mathbb{Q} = \Delta^{k-1}$ represents the set of all probabilities over the Sender's type and let $\mathbb{Z} := \mathbb{A} \times \mathbb{B} \times \mathbb{Q}$.⁸ Recall that a *martingale* is a sequence of \mathbb{Z} -valued random variables $(\mathbf{z}_t)_0^\infty$ and Borel fields $(\mathcal{F}_t)_0^\infty$ such that for each $t = 0, 1, 2, \dots$,

1. $\mathcal{F}_t \subset \mathcal{F}_{t+1}$ and $\mathbf{z}_t \in \mathcal{F}_t$;
2. $E(|\mathbf{z}_t|) < \infty$;
3. $\mathbf{z}_t = E(\mathbf{z}_{t+1} | \mathcal{F}_t)$, a.s.

(Here, we follow the notation in [3] and represent random variables in bold type.) A *di-martingale* is a \mathbb{Z} -valued martingale $\mathbf{z}_t = (\mathbf{a}_t, \beta_t, \mathbf{q}_t)$ such that $\mathbf{a}_{t+1} = \mathbf{a}_t$ when t is even,

⁷ k -dimensional because there are k types of the Sender.

⁸To avoid complications in the sequel, we will actually take \mathbb{A} and \mathbb{B} to be a little bigger than the payoffs that the players get. For example, if all the payoffs to the Sender and Receiver in the game are in $[-1, 1]^k$ and $[-1, 1]$ respectively, we may take $\mathbb{A} \times \mathbb{B}$ to be $[-2, 2]^k \times [-2, 2]$.

$\mathbf{q}_{t+1} = \mathbf{q}_t$ when t is odd and $(\mathbf{a}_0, \beta_0, \mathbf{q}_0)$ is a constant⁹ (i.e. is deterministic) with all the equalities holding almost surely.

From the definition above, it is seen that the q and a coordinates split alternately. This corresponds to signalling and compromising via joint lotteries. As mentioned above, since signalling and joint lotteries perform the mathematical task of convexifying across different sections and convexifying can be thought of as a splitting process, it makes sense to think of canonical equilibria in terms of di-martingales. (That the martingale property holds in the compromise stage of the communication is clear. As was argued in §4.1 in the discussion of the incentive compatibility conditions (equations (1)–(3)), the martingale property also holds in the signalling stages.) Now consider the Receiver's action at the end of the communication stage, regardless of its length. At the end of all the communication he might have new beliefs about the Sender's type and takes a corresponding optimal action. This determines his expected payoff and the Receiver's payoff. If we think of all possible posterior beliefs, this is the graph of the silent game (i.e. the game without any communication, where the Receiver takes an action based solely on his priors). Let us make this precise.

For each $p \in \Delta^{k-1}$, consider the Sender-Receiver game Γ with priors p , represented by $\Gamma(p)$. A mixed strategy of the Receiver is a mixed action¹⁰ $y \in \Delta(C_R)$. A (mixed) strategy y is an equilibrium of the silent game $\Gamma(p)$ if

$$\beta(y) := \sum_{t \in T_S} p_t u_R(y, t) = \max_{\tilde{y} \in \Delta(C_R)} \sum_{t \in T_S} p_t u_R(\tilde{y}, t).$$

We now let $\mathcal{E}(p)$ be the set of equilibrium payoffs in $\Gamma(p)$ which is easily seen to be non-empty for each p . Thus, for each p , $\mathcal{E}(p)$ consists of pairs (a, β) where for each p , there is only one β but multiple values of a . Denote by $\text{gr } \mathcal{E}$, the graph of the *payoff correspondence* $\mathcal{E}(p)$.

In order to ensure that the canonical conversations can be represented by di-martingales, we may have to add a few more elements to the equilibrium payoff correspondence of the silent game, $\mathcal{E}(p)$ (cf. the discussion following equations (1)–(3)). These are precisely the virtual payoffs assigned to types which occur with probability 0. With the additional elements, we get $\mathcal{E}^+(p)$, the *modified equilibrium payoffs correspondence*. (The reasons for adding these additional points are discussed in greater detail in §4.3 below.) Once again, we will define $\mathcal{E}^+(p)$ formally. For each $p \in \Delta^{k-1}$, define $\mathcal{E}^+(p)$ as the set of all $(a, \beta) \in \mathbb{A} \times \mathbb{B}$

⁹This makes \mathbf{z}_0 the expectation of the di-martingale.

¹⁰For any finite set F , the set of probability distributions on F is $\Delta(F) := \{x \in \mathbb{R}_+^{|F|} : \sum_{f \in F} x_f = 1\}$, the standard $(|F| - 1)$ -dimensional simplex.

such that there exists $y \in \Delta(C_R)$ satisfying

$$\begin{aligned}\beta &= \sum_{t \in T_S} p_t u_R(y, t) = \max_{\tilde{y} \in \Delta(C_R)} \sum_{t \in T_S} p_t u_R(\tilde{y}, t), \\ a_t &\geq u_S(y, t), \quad \forall t \in T_S; \text{ and} \\ a_t &= u_S(y, t) \text{ if } p_t > 0, \text{ for all } t \in T_S.\end{aligned}$$

The graph of the modified equilibrium correspondence is

$$\text{gr } \mathcal{E}^+ := \{(a, \beta, p) \in \mathbb{A} \times \mathbb{B} \times \Delta^{k-1} : (a, \beta) \in \mathcal{E}^+(p)\}.$$

This is the last element that we need to complete the picture. We now consider the expectations of all di-martingales whose limits are almost surely in $\text{gr } \mathcal{E}^+$, a section (corresponding to the original priors) of which gives us the cheap talk equilibria.

We need a few more definitions to exploit the above ideas fully. We first recall the notion of *stopping time* of a martingale. Let (Ω, \mathcal{F}, P) be a non-atomic probability space and let $\mathbb{N} := \{0, 1, 2, \dots\}$ be the set of natural numbers. Also, let $(z_n)_{n \in \mathbb{N}}$ be a sequence of random variables (taking values in a Euclidean space) and let $(\mathcal{F}_n)_{n \in \mathbb{N}}$ be a sequence of finite fields such that (z_n) is a martingale¹¹ with respect to (\mathcal{F}_n) .¹² Let $\mathbb{N}_\infty := \mathbb{N} \cup \{\infty\}$ and adjoin $\mathcal{F}_\infty := \bigvee_{n \in \mathbb{N}} \mathcal{F}_n$ (the minimal field containing \mathcal{F}_n for each $n \in \mathbb{N}$) to (\mathcal{F}_n) . A random variable α is a *stopping time* if for every $n \in \mathbb{N}_\infty$, $\{\alpha = n\} \in \mathcal{F}_n$. A stopping time is *a.s. finite* if $P(\alpha < \infty) = 1$. It is *a.s. bounded* if there exists $n_0 < \infty$ such that $P(\alpha \leq n_0) = 1$.

We will be interested in martingales that have almost surely finite stopping times. It should be noted that while a stopping time that is a.s. bounded is a.s. finite, the converse is not necessarily true. Our assumption that the fields are finite also ensure that a stopping time that is a.s. bounded is also everywhere bounded which implies (by König's lemma¹³) that it is everywhere finite (see §4 of [2]). Another point to note is that the assumption of a non-atomic probability space is made to allow the representation of a tree for each martingale in question.

Let $G \subset \mathbb{Z}$, where \mathbb{Z} is defined as above. Then the *di-convex hull* of G is the set of expectations of all finite di-martingales whose limits are almost surely in G . $G^\#$ is the set of expectations of all di-martingales that stop in finite time almost surely and whose limits are almost surely in G . Finally, G^* is the set of all di-martingales whose limits are almost surely

¹¹Standard definitions of martingale require neither a non-atomic probability space, nor finite fields and the definitions of martingale and stopping times are valid without them.

¹²These ideas are described lucidly in §9.3 of [10], for example.

¹³König's Lemma. *Every finitely generated tree with infinitely many points must contain at least one infinite branch.* (For a proof, see [27].)

in G . The set G^* is also called the *di-span* of G . We are interested in di-martingales whose limits are almost surely in $\text{gr } \mathcal{E}^+$ as these are the only ones with the “correct” payoffs for the two players because $\text{gr } \mathcal{E}^+$ represents the payoff the players can get in the action phase of the game. Now let $G := \text{gr } \mathcal{E}^+$. This leads to the following results.

1. The \mathbf{q}_0 -section of the di-convex hull of G represents the set of all payoffs from finite (i.e. almost surely bounded) conversations.
2. The \mathbf{q}_0 -section of $G^\#$ represents the payoffs attainable with conversations that are almost surely finite.
3. The \mathbf{q}_0 -section of G^* represents the payoffs attainable with unbounded conversations.

These results make intuitive sense and follow immediately from Theorem B in [3] and the characterisation of *bi-martingales*¹⁴ in [2]. Nevertheless, they require proof and this is the substantial achievement recorded in [2] and [3].

4.3 Game Theoretic Aspects of Cheap Talk

We shall now discuss the issue of strategy and solution concept. But before that, let us first revisit our reasons for modifying the payoff correspondence, $\mathcal{E}(p)$.¹⁵

Consider the equilibrium payoff correspondence of the silent game, $\text{gr } \mathcal{E}$. By allowing the types of the Sender with zero probability (i.e. on the boundary of Δ^{k-1}) to get *more*, we get the modified equilibrium payoff correspondence, $\text{gr } \mathcal{E}^+$. In a signalling stage, the Receiver “promises” different payoffs to different types (cf. equations (1)–(3) above). It follows from Theorem B in [3], that we can restrict attention to the case where all communication is done with only two messages, so that $M = \{a, b\}$. Here, two cases arise. Consider first the case where type j of the Sender is at a signalling node v where the Sender of type j sends each message with positive probability. Let w and w' be the nodes that are the descendants of v (i.e. have v as their common predecessor). Then it must be that the payoff to type j of the Sender is such that $a_j|_w = a_j|_{w'}$; for if $a_j|_w > a_j|_{w'}$ (say), then the Sender of type j would pick the node w with probability 1. This means that $a_j|_w = a_j|_{w'} = a_j|_v$ because v is the average of the payoffs at w and w' (as in equation (1)). Thus, if each type j of the Sender sends both messages a and b with positive probability, we get a martingale.

However, if a particular type is to be assigned probability zero after the signalling (i.e. picks one of the messages for sure), how much should the Receiver promise him? (Recall

¹⁴A *bi-martingale* is a di-martingale without the β -coordinate.

¹⁵Our discussion here is perfunctory at best. For a fuller treatment of the issues involved, the reader is referred to [3].

that this is the situation considered in equations (2) and (3).) In other words, what should the continuation payoff be at a node that occurs with probability 0? Indeed, it is in this case that the martingale property may not hold. In other words, it may be the case that $a_j|_v = a_j|_w > a_j|_{w'}$ where under the strategy in question, node w' occurs with probability 0. At the node w' the Receiver cannot detect a deviation, so $(\mathbf{a}, \beta, \mathbf{p})$ is not a di-martingale. To resolve this issue, Aumann and Hart propose raising the payoff at w' so that the expectation of the process a_j is still a martingale. It goes without saying that these modifications have to be made on the entire martingale. For further conceptual and technical details, the reader is once again referred to [3]. We should mention this method was first used by Hart in [15]. Of course, the zero probability situations may arise in other ways too. For this, the reader is referred, once again, to [3].

The other issue that we have ignored is a description of strategies and the solution concept used. In our informal description, we have described behavioural strategies for the players. But a strategy is supposed to tell us what each player will do at each node and all we permit at the start of play is a mixture over these strategies. But these are not easily defined in infinite games, as we cannot have independent randomisations at all the nodes where we have described behavioural strategies (as there are uncountably many of these). Nevertheless, Aumann [1] has shown that we can still define strategies consistently and allow for mixtures in a manner analogous to the finite case. Such a procedure has been used in [3] and we shall not repeat it here.

The solution concept that we use is *Bayesian Nash Equilibrium*. Nevertheless, following Aumann and Hart [3] we put some restrictions on the off the equilibrium behaviour. The main issue here is that the Receiver cannot tell when the Sender deviates. If the Sender does deviate, how do we define payoffs at nodes in the game tree that were not supposed to be reached? But recall that at each node, the goal is to define for each player, an expected payoff from going through that node. The expedient method is therefore to define the payoff of type j of the Sender at node v to be the expected payoff at v if from then on, the Sender plays a best response to the Receiver's strategy.

5 Long Conversations

In this section we shall prove the first of our main results, Theorem A, where we show that the set of equilibria in cheap talk games with unbounded conversations is the same as the equilibria where the conversations are unbounded but end in finite time with probability 1. For this we shall need a little geometric machinery. As before, this machinery comes from [2].

Let \mathbb{A} and \mathbb{Q} be compact, convex subsets of Euclidean spaces (of possibly different dimensions). Let (Ω, \mathcal{F}, P) be a non-atomic probability space. A sequence $(\mathbf{z}_t)_0^\infty := (\mathbf{a}_t, \mathbf{q}_t)_0^\infty$ of $\mathbb{A} \times \mathbb{Q}$ -valued random variables is a *bi-martingale* if:

1. There exists a non-decreasing sequence $(\mathcal{F}_t)_0^\infty$ of subfields of \mathcal{F} , such that (\mathbf{z}_t) is a martingale with respect to (\mathcal{F}_t) .
2. For each $t = 0, 1, \dots$, either $\mathbf{a}_t = \mathbf{a}_{t+1}$ or $\mathbf{q}_t = \mathbf{q}_{t+1}$ (a.s.).
3. \mathbf{z}_0 is constant (a.s.).

Thus, a bi-martingale is nothing but a di-martingale with the β coordinate missing and we shall restrict our attention to the set of expectations of bi-martingales with limits almost surely in the graph of the Sender's payoffs in the silent game. We can do this for the following reason: Consider an $\mathbb{A} \times \mathbb{B} \times \mathbb{Q}$ valued di-martingale (\mathbf{z}_t) with $\mathbf{z}_\infty \in A \times B \times Q$ a.s., where $A \times B \times Q$ is a measurable subset of $\mathbb{A} \times \mathbb{B} \times \mathbb{Q}$. This now defines an $\mathbb{A} \times \mathbb{Q}$ valued bi-martingale, (\mathbf{y}_t) with $\mathbf{y}_t \in A \times Q$ a.s. This bi-martingale stops if and only if the associated di-martingale stops. The if part is clear. To see the only if part, note that if the bi-martingale stops, then the β coordinate cannot move on its own in either the odd or even periods. Thus, the di-martingale must stop too. We shall be interested in the sets of expectations of bi-martingales that converge to certain sets and shall make extensive use of the following notions.

5.1 Definition. A set $A \subset \mathbb{A} \times \mathbb{Q}$ is **bi-convex** if for each $a \in \mathbb{A}$ and each $q \in \mathbb{Q}$, the respective a - and q -sections $A_a := \{q \in \mathbb{Q} : (a, q) \in A\}$ and $A_q := \{a \in \mathbb{A} : (a, q) \in A\}$ are convex.

It is easy to see that every convex set is bi-convex, but a bi-convex need not be convex. The following example illustrates this idea.

5.2 Example. Let $\mathbb{A} = \mathbb{Q} = [0, 1]$ and $A \subset \mathbb{A} \times \mathbb{Q}$ where $A := \{(a, \frac{1}{2}) : a \in [0, 1]\} \cup \{(\frac{1}{2}, q) : q \in [0, 1]\}$. As is easily seen from figure 3 below, A is bi-convex but not convex.

As with the case of di-martingales, the following definitions are immediate. Let A be a measurable subset of $\mathbb{A} \times \mathbb{Q}$.

5.3 Definition. $bi-co(A)$ is the smallest bi-convex set containing A .

5.4 Definition. $A^\# := \{z \in \mathbb{A} \times \mathbb{Q} : \exists \text{ bi-martingale } (\mathbf{z}_t) \text{ with an a.s. finite stopping time } N \text{ such that } \mathbf{z}_N \in A \text{ (a.s.) and } \mathbf{z}_0 = z \text{ (a.s.)}\}$.

5.5 Definition. $A^* := \{z \in \mathbb{A} \times \mathbb{Q} : \exists \text{ bi-martingale } (\mathbf{z}_t) \text{ converging to } \mathbf{z}_\infty \text{ such that } \mathbf{z}_\infty \in A \text{ (a.s.) and } \mathbf{z}_0 = z \text{ (a.s.)}\}$.




Figure 3: A bi-convex set that is not convex.

The splitting process analogy is equally useful in the bi-convex case. This shows that $A^\#$ and A^* are bi-convex sets. The reason for looking at bi-martingales rather than the original di-martingales is that these have a slightly simpler geometric structure and there is no loss of generality in doing so. From the game theoretic point of view, once the bi-martingale hits $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \mathcal{E}^+(q)$ (which is what A is supposed to represent), there is no more information to be transmitted. Once we know the posterior beliefs of the Receiver, we can compute his optimal action and his unique expected payoff resulting from an optimal action. This enables us to trace back and talk about the original di-martingale. The major difference between convexity and bi-convexity is that in the bi-convex case, $\text{bi-co}(A) \subset A^\# \subset A^*$ with the inclusions being strict in general. In the convex case, the three are the same, namely the convex hull of A , a result which follows immediately from Carathéodory's Theorem (cf. §17 in [22]). The following example from [2] illustrates this.

5.6 Example. Let $A := \{a_1 = (2/3, 0), a_2 = (0, 1/3), a_3 = (1/3, 1), a_4 = (1, 2/3)\}$ and $\mathbb{A} = \mathbb{Q} = [0, 1]$. Then, $\text{bi-co}(A) = A$, but $A^\# = A^* = \text{bi-co}(\bigcup_{i=1}^4 \{a_i, w_i\})$ is much bigger, as illustrated in figure 4 below. For a demonstration of the fact that $\text{bi-co}(A) \subsetneq A^\# = A^*$, the reader is referred to Example 2.5 in §2 of [2] or Appendix A below.




Figure 4: Example where $A = \text{bi-co}(A) \subsetneq A^*$

We shall look at a special class of sets and the bi-convex sets that they generate. These are what we call *finitely generated* sets and are defined below.

5.7 Definition. A **bi-convex combination** is a convex combination $(a, q) = \sum_i \alpha_i (a_i, q_i)$ (with $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$) where either $a_1 = \dots = a_n = a$ or $q_1 = \dots = q_n = q$.

5.8 Definition. A set $A \subset \mathbb{A} \times \mathbb{Q}$ is **generated** by $A_0 \subset A$ if for all $a \in A$, there exist $\alpha_1, \dots, \alpha_n$ (with $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$) and $a_1, \dots, a_n \in A_0$ such that $a = \sum_1^n \alpha_i a_i$ is a bi-convex combination.

5.9 Definition. A set A is **finitely generated** if there exists a finite set A_0 which generates it.

The reason for looking at finitely generated bi-convex sets is that in Sender Receiver games, the (modified) graph of Sender's payoffs, $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \mathcal{E}^+(q)$ is finitely generated. This is made precise in Lemma 5.11. Before proving the lemma, let us see an illustration of this for a game with two types.

5.10 Example. Consider the following game (figure 5) due to Simon [25]. The Sender has two types, t_1 and t_2 with the prior probability of t_1 being $q \in [0, 1]$. The Receiver has seven possible actions, α, b, \dots, g .

		α	b	c	d	e	f	g
$[q]$	t_1	1, 10	1, 1	$4, \frac{1}{2}$	1, 0	$0, -\frac{3}{2}$	1, -3	3, -354
$[1 - q]$	t_2	3, -354	1, -3	$0, -\frac{3}{2}$	1, 0	$4, \frac{1}{2}$	1, 1	1, 10

Figure 5: Example 5.10

The graph of the Receiver's payoffs are depicted in figure 6. For illustrative clarity, the Receiver's payoffs from any action are not drawn for all possible values of the Sender's type. Nevertheless, the Receiver's payoffs from an action can be extended to the entire probability space as they are linear in the probabilities of the types of the Sender.

Of greater interest, however, is the graph of the Sender's payoffs, for this set is finitely generated in the sense defined above. This is depicted in figure 7. We shall demonstrate below that this is not a coincidence, but is a general property of Sender-Receiver games.

Consider action f . It is optimal when for the Receiver when he believes the Sender to be of type 1 with probability $q \in [\frac{1}{40}, \frac{1}{4}]$, (as is seen in figure 6). In figure 7, this means that the two types of the Sender get the vector of payoffs given by $a(f)$ whenever the Sender is of type 1 with probability $q \in [\frac{1}{40}, \frac{1}{4}]$ (which means that f is the optimal action for the Receiver). But when $q = \frac{1}{40}$, for example, action g is also optimal for the Receiver. Hence the payoff correspondence for the Sender consists of all convex combinations of $a(f)$ and $a(g)$ at $q = \frac{1}{40}$. \blackboxtimes

Let $G := \text{proj}_{\mathbb{A} \times \mathbb{Q}} \text{gr } \mathcal{E}^+$, that is G is the graph of the Sender's modified payoffs from a Sender-Receiver game. Then,




Figure 6: Receiver's payoffs given by upper envelope

5.11 Lemma. G is finitely generated.

Proof. Let the Receiver's actions be c_1, c_2, \dots, c_n and β_i the expected payoff from action c_i . (Note that for each i , $\beta_i : \Delta^{k-1} \rightarrow \mathbb{B}$ is a linear function and that the Receiver's payoffs are given by the upper envelope of these n linear functions, i.e. the graph of $\beta(q) := \max_i \{\beta_i(q)\}$.) Let $F_i := \{q \in \Delta^{k-1} : \beta_i(q) > \beta_j(q) \text{ for } j \neq i\}$ which gives us the region of the probability space where the Receiver's best action is c_i . More generally, we can define $F_{i_1, \dots, i_m} := \{q \in \Delta^{k-1} : \beta_{i_1}(q) = \dots = \beta_{i_m}(q) > \beta_j(q) \text{ for } j \notin \{i_1, i_2, \dots, i_m\}\}$ which gives us the region of the probability space where the Receiver's best actions are $c_{i_1}, c_{i_2}, \dots, c_{i_m}$ and the Receiver is indifferent between these actions. Note that we do not require F_{i_1, \dots, i_m} to be non-empty.

Finite Partition of Probability Space. The F 's defined above partition the probability space Δ^{k-1} into finitely many regions where the Receiver has a set of best actions. Now, F_{i_1, \dots, i_m} is defined as the intersection of finitely many (bounded) affine sets and is therefore finitely generated (i.e. there exists a finite set so that each point in F_{i_1, \dots, i_m} can be written as a convex combination of these points). (This is a corollary of Theorem 19.1 in [22].)

Lifting to Sender's Payoffs. We shall now show that the Sender's payoffs are finitely generated. Consider the region of the probability space given by the set F_{i_1, \dots, i_m} . The Receiver is indifferent between actions c_{i_1}, \dots, c_{i_m} in this region. Thus, the (vectors of) Sender's payoffs are all convex combinations of $a(c_{i_1}), \dots, a(c_{i_m})$ which is a finitely generated set. But this is




Figure 7: Graph of $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \mathcal{E}^+(q)$ — commonly known as the Sender's modified payoffs

true in any cell in our partition, and there are only finitely many regions according to our partitions all of which are finitely generated, so that $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \text{gr } \mathcal{E}$ is finitely generated. It is now a simple matter to add the additional payoffs at the boundaries of the simplex Δ^{k-1} to show that $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \text{gr } \mathcal{E}^+$ is also finitely generated. \square

We shall now come to a statement of the finiteness condition which is essential to our making progress. Before we state the condition, we shall make some definitions. Recall that if $\{v_0, \dots, v_q\}$ is an affine independent subset of some Euclidean space, then it spans the q -simplex $s := [v_0, \dots, v_q] = \text{co}(\{v_0, \dots, v_q\})$. The *vertex set* of s is denoted by $\text{Vert}(s) := \{v_0, \dots, v_q\}$. If s is a q -simplex and t is an r -simplex, then we shall call $s \times t$ a $q \times r$ -bi-simplex and we shall denote its vertex set by $\text{Vert}(s \times t) := \text{Vert}(s) \times \text{Vert}(t)$. If σ is a bi-simplex, then a *face* of σ is a bi-simplex σ' with $\text{Vert}(\sigma') \subset \text{Vert}(\sigma)$.¹⁶

5.12 Definition. A *finite bi-simplicial complex* K is a finite collection of bi-simplices in some product of Euclidean spaces such that:

¹⁶These definitions are the obvious analogues of standard notions in algebraic topology which can be found, for instance, in [9].

- (i) if $\sigma \in K$, then every face of σ also belongs to K ;
(ii) if $\sigma, \tau \in K$, then $\sigma \cap \tau$ is either empty or a common face of σ and of τ .

We now introduce a finiteness condition.

Condition F [Finiteness Condition]. Given a set W_0 , then for all $A \in 2^{W_0}$, $bi-co(A) = K$, where K is a bi-simplicial complex.

If we let G_0 be the finite set generating the modified graph of the payoffs, $gr \mathcal{E}^+$ then $W_0 := \text{proj}_{\mathbb{A}} G_0 \times \text{proj}_{\mathbb{Q}} G_0$.

5.13 Definition. A Sender-Receiver game satisfies Condition F if the generating set of the graph of the Sender's payoffs, G_0 , is such that $W_0 = \text{proj}_{\mathbb{A}} G_0 \times \text{proj}_{\mathbb{Q}} G_0$ (as defined above) satisfies Condition F.

It should be mentioned that there is no example of a Sender-Receiver game that does not satisfy Condition F. In the appendix, we shall relate Condition F to a purely geometric condition on the bi-convex hulls of finite sets.

Theorem A below relies on the following lemma.

5.14 Lemma. Let $A \subset \mathbb{A} \times \mathbb{Q}$ be generated by a finite set A_0 and suppose W_0 satisfies Condition F. Then $A^* = A^\#$.

Proof. This is a key lemma and we shall informally describe why one might expect it to be true. For a formal proof, the reader can be referred to the Appendix. Analogous to the convex case, in the bi-convex case we can define *bi-convex* functions. These are functions defined on a bi-convex domain that are convex on every a - and q -section of the domain. Just as in the convex case, we can talk about the set of points that can be separated from a given set via bi-convex functions. Indeed, this can be defined as an inductive process starting from a large enough set (so that things remain interesting). By Theorem 4.3 in [2] (see Appendix A below), for a set $A \subset \mathbb{A} \times \mathbb{Q}$, $A^\#$ is the limit of an inductive process of separation by bi-convex functions. Now, if we restrict ourselves to separation by bi-convex functions that are *continuous* on A , then the resultant limit set defined is A^* (Theorem 4.7 in [2], also described in Appendix A).

At this point, it may be pertinent to point out the difficulties that arise. In the convex case, results on separation stem from the separating hyperplane theorem which essentially says that when considering separation properties, we may restrict attention to separation by linear functionals. This is a global result, as linear functionals defined on a convex set can be extended to the entire space. Unfortunately, there is no such global analogue in the bi-convex case. Instead, we obtain a global result by considering separation at a local level (by piecewise bi-affine functions) and performing this separation across the entire space.

More specifically, it turns out that for finitely generated sets, one can restrict attention to separation from the finite generating set. We then show that by taking a large enough finite set and considering its bi-convex hull, the process of separation (by both bi-convex functions and bi-convex functions continuous on A) is equivalent to removing *bi-extreme points*¹⁷ (points which are not non-trivial bi-convex combinations of other points) of the bi-convex hull of the larger finite set. The process ends in finite time as there are only finitely many points to consider and our result follows immediately. \square

We are now ready to state our main result.

Theorem A. *Let Γ be a finite Sender-Receiver game that satisfies Condition F. Then the set of equilibria with unbounded conversations is the same as the set of equilibria with unbounded conversations that end in finite time with probability one.*

In other words, the \mathbf{q}_0 -section of $G^\#$ is the same as the \mathbf{q}_0 -section of G^* , where $G = \text{proj}_{\mathbb{A} \times \mathbb{Q}} \text{gr } \mathcal{E}^+$ is the graph of the Sender's payoffs. In fact, we shall prove that $G^\# = G^*$.

Proof of Theorem A. From Lemma 5.11, we know that the modified graph of the Sender's payoffs is finitely generated. Therefore from Lemma 5.14, the set of equilibrium payoffs from unbounded conversations is the same as the set of equilibrium payoffs from unbounded conversations that end in finite time with probability one. \square

Note that the proof of Lemma 5.14 is actually constructive, in that we provide an algorithm to compute the set of equilibria. But this is merely a consequence of Condition F. In other words, the set of equilibria can be decomposed into a bi-simplicial complex (as it is the section of a bi-simplicial complex). We shall now look at an example which illustrates what this set of equilibria might look like.

5.15 Example. Consider once again the game introduced in Example 5.10. Below, in figure 8 is a picture of the section of the graph of the Receiver's payoffs with $a_1 + a_2 = 4$. (The coordinates in figure 8 represents the pair (a_1, q) , which automatically gives the value of a_2 at that point to be $4 - a_1$.) It is easily seen from figure 7 that the graph of the Receiver's payoffs is bi-convex, which means that there can be no information transmitted with finite conversations of any length. But the moment we consider infinite conversations, we get Pareto superior payoffs for both types of the Sender and the Receiver. For instance, there is an equilibrium where the two types of the Sender get expected payoffs $(2, 2)$ and the Receiver gets expected payoff 5.

¹⁷Given a convex set C , $x \in C$ is an *extreme point* of C if $x = \lambda y_1 + (1 - \lambda)y_2$ with $\lambda \in (0, 1)$ implies that $y_1 = y_2 = x$. The definition above is the appropriate analogue for the bi-convex case.




Figure 8: $a_1 + a_2 = 4$ section of graph of Receiver's payoffs

6 (Trans)finite Conversations

In this section we consider some important questions that are related to the length of potential conversations. In particular, we shall consider finite and transfinite conversations. We shall first provide definitions and results and then conclude the section with a discussion of the results.

As in §4.2, we let $G = \text{proj}_{\mathbb{A} \times \mathbb{Q}} \text{gr } \mathcal{E}^+$ be the graph of the Sender's (modified) payoffs and say that G is generated by G_0 . We first consider finite canonical¹⁸ conversations. Let $\langle G \rangle_b$ be the set of all *bi-convex combinations* of G and define $G_{n+1} := \langle G_n \rangle_b$. It follows from (a modification of) Theorem B in [3] that the set of equilibrium payoffs (for the Sender) from finite conversations that last n -periods is the \mathbf{q}_0 -section of G_n . Then, the set of equilibrium payoffs from all finite conversations is the \mathbf{q}_0 -section of $\bigcup_{n \in \mathbb{N}} G_n =: \text{bi-co}(G_0) = \text{bi-co}(G)$.¹⁹ We are now able to state our result on finite conversations.

6.1 Theorem. *Let Γ be a Sender-Receiver game that satisfies Condition F. Then there exists an N such that for all $n > N$, the \mathbf{q}_0 -section of G_n is identical to the \mathbf{q}_0 -section of G_N .*

In other words, the set of equilibria from finite conversations eventually stabilises if we allow the players to talk for large but finite periods. This is a direct consequence of Condition F as we demonstrate below.

Proof of Theorem 6.1. It will suffice to prove that there exists an N so that $G_N = \text{bi-co}(G_0)$. We know from Lemma 5.11 that G is finitely generated by G_0 . By Condition F, it follows that $\text{bi-co}(G_0)$ is a bi-simplicial complex. It is now immediate that there exists an N such that $\text{bi-co}(G_0) = G_N$. Also, $\text{bi-co}(\text{bi-co}(G_0)) = \text{bi-co}(G_0)$ (because $\text{bi-co}(G_0)$ is the smallest bi-convex set which contains G_0) which demonstrates our claim. \square

¹⁸In fact, we shall only consider canonical conversations. We shall therefore drop the qualifier "canonical" in what follows.

¹⁹See the Appendix for a proof that for any set A_0 , $\text{bi-co}(A_0) := \bigcup_{n \in \mathbb{N}} \langle A_n \rangle_b$.

We now come to conversations that are longer than ω . For concreteness, we consider only conversations of length $\omega + \omega$. Let the set of equilibrium payoffs for the Sender from conversations that are of length ω be denoted by the \mathbf{q}_0 -section of G^* and the set of equilibrium payoffs from conversations of length $\omega + \omega$ by the \mathbf{q}_0 -section of $(G^*)^*$. At this point, it may be useful to clarify the meaning of conversations of length $\omega + \omega$. Recall that conversations of length ω mean that the players communicate till infinity and the Receiver then takes an action. This makes the game tree of order type $\omega + 1$. But after a conversation of length ω , the players may choose to have another round of conversation of length ω (or less). This seems natural especially since we are assuming that the players cannot commit to not talking any more and are free to do as they please. Thus, if a conversation of length ω makes sense, then so does a conversation of length $\omega + \omega$. The question is, does this further expand the set of equilibria? Well, not if the game satisfies Condition F .

6.2 Theorem. *Let Γ be a Sender-Receiver game that satisfies Condition F . Then the \mathbf{q}_0 -section of G^* is identical to the \mathbf{q}_0 -section of $(G^*)^*$ which is the same as the \mathbf{q}_0 -section of $G^\#$.*

Proof. See the appendix. □

Indeed, we can say even more.

6.3 Corollary. *Let Γ be a Sender-Receiver game that satisfies Condition F . Then the set of equilibria from conversations that are finite with probability 1 (the \mathbf{q}_0 -section of $G^\#$) is the same as the set of equilibria from conversations that are of order type $n\omega$ for all $n \in \omega$.*

Proof. Same as the proof of theorem 6.2. □

Once again, we see that Condition F is the central ingredient to our result. The reason Condition F plays such a pivotal role has to do with the inhomogeneous nature of bi-convexity. As is demonstrated in the Appendix, bi-convexity is very different from convexity in that a number of theorems from the convex case do not carry over to the bi-convex case. A good number of these theorems have to do with the structure of convex sets which implies that for a given set A , the topological properties of $\text{bi-co}(A)$, $A^\#$ and A^* can be very strange. Recall that for a set A_0 , we denoted the set of bi-convex combinations of A by $\langle A_0 \rangle_b$ and defined $A_{n+1} := \langle A_n \rangle_b$. Now, there exist sets such that for all $n \in \mathbb{N}$, $\langle A_n \rangle_b \subsetneq \langle A_{n+1} \rangle_b$. In other words, the bi-convex hull of such a set can only be obtained after an infinitary process. Thus, the notion of bi-convexity is not finitistic (while convexity is). Therefore, by imposing Condition F , we automatically exit the realm of the infinite where things can be very strange and enter a realm where things are well behaved.

But that is not all. Aumann and Hart (§5 of [2]) give an example of a set A such that $A^\# \subsetneq A^*$. The set A in their example is piecewise algebraic, but it does not come from a

Sender-Receiver game (as it is not finitely generated). Also, while A is piecewise algebraic, neither $A^\#$ nor A^* are. Moreover, they also give an example of a set A where $A^* \subsetneq (A^*)^*$ which shows the relevance of Theorem 6.2.

But what about the converse? Just how strong is our Condition F ? The problem we face is that we do not have an example of a game that does not satisfy Condition F . Essentially, what we are doing in the proof of Lemma 5.14 is that we take the finitely many points and remove the ones that are bi-extreme with respect to some bi-convex set. As everything satisfies Condition F , we are fine. But if Condition F fails, then our proof does not go through. But what is worse, it is then also possible that $\text{bi-co}(G)$ does not stabilise after some large but finite number of rounds. Also, we would lose the niceness of the set of equilibrium payoffs.

7 Conclusion

In this paper, we consider some questions regarding the set of equilibria in Sender-Receiver games with extended conversations. The question of whether unbounded conversations are the same as conversations that are finite with probability 1 is important from the point of view of applications. After all, we are trying to understand the advantages of not placing deadlines on conversations when there may mutually beneficial exchanges (figuratively speaking) that are possible. We find a sufficient condition for the above question to have an affirmative answer. Surprisingly enough, the sufficient condition is closely related to the questions of whether the equilibria from finite talk eventually stabilise and whether transfinite talk may be useful and also ensures that the set of equilibria can be built up from geometric objects of increasing dimension in finitely many steps.

The importance of the set of equilibria being simple (in particular, a bi-simplicial complex) cannot be overstated. When a game satisfies Condition F , we have shown that the set of equilibria in unbounded talk is actually *semi-algebraic*.²⁰ To understand the significance of this, consider first a set $A := \{x \in \mathbb{R}^n : f_1(x) \geq 0, \dots, f_\ell(x) \geq 0, g_1(x) = \dots = g_m(x) = 0\}$ where the f_i 's and g_i 's are “nice” functions, say polynomials. Upon performing elementary logical, algebraic and topological operations like taking unions, convexifying, projecting etc., we get new sets. If we keep performing these operations, two scenarios emerge. In the first, the process stabilises after some number of steps and in the second, ever more complicated sets arise (e.g. the Cantor sets and Borel sets of arbitrarily high complexity). Semi-algebraic sets are of the first kind. In particular, every semi-algebraic set

²⁰A *semi-algebraic* set is defined by finitely many polynomial inequalities (with the polynomials being over any field). When the field in question is real, closed, these sets are some of the simplest sets that one can define and permit quantifier elimination. For a further description of these ideas, the reader is referred to [12].

has a simplicial decomposition. The Aumann-Hart notions of $\text{bi-co}(A)$, $A^\#$ and A^* are of the second kind. They are extremely complicated objects in general and even taking sections of these objects can only be done in principle. In other words, there may not be a decision procedure to determine if a point lies in the set.

Nevertheless, equilibria in finite games are semi-algebraic, i.e. belong to the first scenario described above. Moreover, all the popular refinements in use also give rise to semi-algebraic sets. This includes all the refinements that are defined pointwise (e.g. sequential equilibrium and perfect equilibrium - see Blume and Zame [8]) and stable equilibria which are set-valued as defined by Mertens in [19]. This is an attractive and desirable property as semi-algebraic sets are definable in a very precise sense (see [12]) and this means that not only can the game theorist compute the set of equilibria, but the players can too. We would like equilibria of unbounded conversations to have the same property, especially since we are restricting attention to finite games.

Our inability to say more about the when Condition F holds can be traced back to the peculiar geometry of bi-convex sets and the fact that the bi-convex hull of a set has an internal representation which is necessarily the limit of a countable process. Nevertheless, we have made some progress with regards to the effective length of conversations needed to achieve certain equilibrium outcomes. It has been an open question, at least since the publication of [2], whether there exists a game where communication which is unbounded but finite with probability one and communication which is unbounded has different sets of equilibria. We take a small step towards answering that question by identifying exactly what must happen in order that the two sets be different. Moreover, we present the dichotomous nature of sets of equilibria from finite and transfinite talk. We can either expect all of them to behave nicely or it is probably the case that none of them will.

From a purely mathematical point of view, another important question arises (which is closely related to Condition F). Namely, are their algebraic conditions on a finite (or semi-algebraic) set which will ensure that the bi-convex hull of the set will be achieved in finitely many iterations? This characterisation may not be possible in general, but we hope that the extra structure that the graph of Sender-Receiver games possess will be of some use.

The present paper also represents a departure from previous work in that it points out the algebraic difficulties that must be overcome to say more about the problem at hand. It should be noted that this is not a topological problem as bi-convexity is not preserved under continuous transformations (which means that we cannot bring to bear powerful methods from algebraic topology). (It may be shown that the any group of transformations which preserves bi-convexity and the origin is a subgroup of $GL(m, \mathbb{R}) \times GL(n, \mathbb{R})$.) But we are hopeful that a clarification of the issues at hand will make further progress possible.

A Miscellany

In this appendix, we tie up a number of loose ends in the paper. We first provide an alternate characterisation of bi-convex sets which entails providing a self-contained description of the concepts of bi-convexity and bi-martingales. It is therefore written in greater generality than the text above. It should be noted that many of the propositions below will seem obviously true. While this may be so, we still provide proofs and give some examples of statements that seem “obviously true” (because of the analogue in the convex case) but are not in the bi-convex case. We then provide proofs of Lemma 5.14 and Theorem 6.2. We conclude the appendix with an axiomatisation of the graph of payoffs of the silent game ($\text{gr } \mathcal{E}$) in Sender-Receiver games, thereby removing our problem from the corsettes of game theory and making it purely mathematical.

A.1 Bi-convexity and bi-martingales

Let us recall the definitions in Aumann and Hart [2]. Let \mathcal{X}, \mathcal{Y} be compact convex subsets of finite dimensional Euclidean spaces. Let $B \subset \mathcal{X} \times \mathcal{Y}$. Let $B_x := \{y \in \mathcal{Y} : (x, y) \in B\}$ and $B_y := \{x \in \mathcal{X} : (x, y) \in B\}$. B is *bi-convex* if for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, B_x and B_y are convex sets. Let $f : B \rightarrow \mathbb{R}$. The function f is *bi-convex* if for all $x \in \mathcal{X}$, $f(x, \cdot)$ is a convex function on B_x and if for all $y \in \mathcal{Y}$, $f(\cdot, y)$ is a convex function on B_y .

The space $\mathcal{X} \times \mathcal{Y}$ inherits the product topology relative to the Euclidean spaces they live in. Let us denote this as \mathcal{T} . Let us denote the *relative* topology of a set $B \subset \mathcal{X} \times \mathcal{Y}$ by \mathcal{T}_B . Let us denote the *bi-relative* topology by \mathcal{T}_{BR} , where \mathcal{T}_{BR} consists of all sets of the form $\{E \cap U : U \in \mathcal{T}\}$ and $E := \text{aff}(\text{proj}_{\mathcal{X}} B) \times \text{aff}(\text{proj}_{\mathcal{Y}} B)$. A point $z = (x, y) \in B$ is *bi-relatively interior* to B if z is interior in the topology \mathcal{T}_{BR} (not the topology relative to B). A point $z = (x, y) \in B$ is *locally bi-simplicial at z* if there exists a neighbourhood U of x in \mathcal{X} , a neighbourhood V of y , a collection of simplices s_1, s_2, \dots, s_m in \mathcal{X} and a collection of simplices t_1, t_2, \dots, t_n in \mathcal{Y} such that (putting $s = \bigcup_{i=1}^m s_i$ and $t = \bigcup_{i=1}^n t_i$), $s \times t \subset B$ and $(U \times V) \cap B = (U \times V) \cap (s \times t)$. We then have the following proposition.

A.1 Proposition (Propositions 3.6 and 3.7 in [2]). *Let f be a bi-convex function on a bi-convex set B , and let $z \in B$.*

- (i) *If z is a bi-relatively interior point of B , then f is lower-semi-continuous at z .*
- (ii) *If B is locally bi-simplicial at z , then f is upper-semi-continuous at z .*

The following example also from [2] illustrates these ideas beautifully.

A.2 Example. Let $\mathcal{X} = \mathcal{Y} = [0, 1]$ and let $B := \{(t, t) : 0 < t < 1\}$. Then every point of B is relatively interior but none is bi-relatively interior. Similarly, every point of B is locally simplicial but none is locally bi-simplicial. Also, *any* function on B is bi-convex.

Let \mathcal{X} and \mathcal{Y} be compact, convex subsets of Euclidean spaces. Let (Ω, \mathcal{F}, P) be a non-atomic probability space. A sequence $(\mathbf{z}_t)_0^\infty := (\mathbf{x}_t, \mathbf{y}_t)_0^\infty$ of $\mathcal{X} \times \mathcal{Y}$ -valued random variables is a *bi-martingale* if:

1. There exists a non-decreasing sequence $(\mathcal{F}_t)_0^\infty$ of subfields of \mathcal{F} , such that (\mathbf{z}_t) is a martingale with respect to (\mathcal{F}_t) .
2. For each $t = 0, 1, \dots$, either $\mathbf{x}_t = \mathbf{x}_{t+1}$ or $\mathbf{y}_t = \mathbf{y}_{t+1}$ (a.s.).
3. \mathbf{z}_0 is constant (a.s.).

Let A be a measurable subset of $\mathcal{X} \times \mathcal{Y}$. We will require the following definitions.

A.3 Definition. *bi-co(A) is the smallest bi-convex set containing A.*

A.4 Definition. $A^\# := \{z \in \mathcal{X} \times \mathcal{Y} : \exists \text{ bi-martingale } (\mathbf{z}_t) \text{ with an a.s. finite stopping time } N \text{ such that } \mathbf{z}_N \in A \text{ (a.s.) and } \mathbf{z}_0 = z \text{ (a.s.)}\}$.

A.5 Definition. $A^* := \{z \in \mathcal{X} \times \mathcal{Y} : \exists \text{ bi-martingale } (\mathbf{z}_t) \text{ converging to } \mathbf{z}_\infty \text{ such that } \mathbf{z}_\infty \in A \text{ and } \mathbf{z}_0 = z \text{ (a.s.)}\}$.

For a given bi-convex set B that contains A , the set $ns_A(B)$ consists of all points of B that cannot be separated from A by any bi-convex function, i.e. for all $z \in ns_A(B)$ and any bounded, bi-convex function $f : B \rightarrow \mathbb{R}$, $f(z) \leq \sup f(A)$.

We will now define a process of separation. Let $B_0 := \mathcal{X} \times \mathcal{Y}$. Define inductively, $B_{\alpha+1} = ns_A(B_\alpha)$ for every successor ordinal α and $B_\alpha = \bigcap_{\beta < \alpha} B_\beta$ for every limit ordinal α . This defines a non-increasing sequence²¹ of sets (B_α) with limit $C := B_\gamma$ for some ordinal γ . By Zorn's Lemma, it follows that C is well defined. We then have the following proposition.

A.6 Theorem (4.3 in [2]). *The limit set C satisfies $C = ns_A(C)$ and is the largest such set. Also, $C = A^\#$.*

If we take the set A to be closed (as is the case in all of our applications), we can define a similar notion of separation in terms of bi-convex functions which are continuous at every point of the set A , which we shall denote by $nsc(B)$ ($\equiv nsc_A(B)$) and a similar inductive process which will give us another limit set D . This gives us

²¹A sequence is a function whose domain is an ordinal (cf. §2.4 of [18]).

A.7 Theorem (4.7 in [2]). *The limit set D satisfies $D = \text{nsc}_A(D)$ and is the largest such set. Also, $D = A^*$.*

Note that for each bi-convex set B that contains a set A , both $\text{ns}_A(B)$ and $\text{nsc}_A(B)$ are bi-convex sets. This is because if $z \in \text{ns}_A(B)$ (say) then it must be the case that for all $f : B \rightarrow \mathbb{R}$ bi-convex, $f(z) \leq \sup f(A)$. Now suppose $z, z' \in \text{ns}_A(B)$, then for any $z'' := tz + (1-t)z'$ a bi-convex combination, we have $f(z'') \leq \max\{f(z), f(z')\} \leq \sup f(A)$ which implies that $z'' \in \text{ns}_A(B)$, that is $\text{ns}_A(B)$ is bi-convex. A similar argument shows that $\text{nsc}_A(B)$ is also bi-convex. To demonstrate these ideas, let us reconsider example 5.6 above.

A.8 Example. Let $A := \{a_1 = (2/3, 0), a_2 = (0, 1/3), a_3 = (1/3, 1), a_4 = (1, 2/3)\}$ and $\mathcal{X} = \mathcal{Y} = [0, 1]$. We claim that $A^\# = A^* = \text{bi-co}(\{a_i, w_i\}_{i=1}^4)$, as illustrated in figure 9 below. By Lemma A.11 below, $\sup f(\{a_i, w_i\}_{i=1}^4) \geq \sup f(\text{bi-co}(\{a_i, w_i\}_{i=1}^4))$. We can therefore restrict attention to the w_i 's. Now suppose we could separate one of the points w_i with some bi-convex function f . Let us assume that w_1 is such that $f(w_1) \geq f(w_j)$ for $j = 2, 3, 4$. Now note that if g is a bi-convex function then $f := \max\{g - \sup g(A), 0\}$ is also bi-convex and takes value 0 on A . We can therefore restrict attention to separation by such functions. Thus, $f(w_1) \leq \frac{1}{2} [f(a_1) + f(w_2)]$ which is impossible as $f(a_1) = 0$. Thus, we cannot separate any of the points, $(w_i)_1^4$. That this indeed is A^* is now easy to see. (Consider separation by bi-convex functions of the form $h(x, y) := \max\{x - x_0, 0\} \max\{y - y_0, 0\}$.)

Now, in game-theoretic terms, the set A is the graph of the modified payoffs of the silent game, $\text{bi-co}(A)$ is the set of payoffs that one can achieve with finite conversations, $A^\#$ is the set of payoffs achievable with unbounded conversations that are finite with probability 1 and A^* is the set of payoffs achievable with unbounded conversations. It follows from the definitions that $A \subset \text{bi-co}(A) \subset A^\# \subset A^*$.

We will begin with some definitions of our own.

A.9 Definition. A set $A \subset \mathcal{X} \times \mathcal{Y}$ is **generated** by $A_0 \subset A$ if for all $z \in A$, there exist $\alpha_1, \dots, \alpha_n$ (with $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$) and $a_1, \dots, a_n \in A_0$ such that $z = \sum_1^n \alpha_i a_i$ is a bi-convex combination.²²

A.10 Definition. A set A is **finitely generated** if there exists a finite set A_0 which generates it.

²²A convex combination $(x, y) = \sum_i \alpha_i (x_i, y_i)$ (with $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$) will be called *bi-convex* if either $x_1 = \dots = x_n = x$ or $y_1 = \dots = y_n = y$.




Figure 9: $A = \{a_i\}_1^4 = \text{bi-co}(A) \subsetneq A^*$

We shall denote the set of all bi-convex combinations of a set $A_0 \subset \mathcal{X} \times \mathcal{Y}$ by $\langle A_0 \rangle_b$. Thus, if we let $A_1 := \langle A_0 \rangle_b$ and $A_{i+1} := \langle A_i \rangle_b = \langle A_0 \rangle_b^{i+1}$, the bi-convex hull of A_0 is given by $\text{bi-co}(A_0) = \bigcup_i A_i$. To see that this characterisation is correct, let $B = \bigcup_i A_i$ and consider two points $z_1, z_2 \in B$. Let $z_1 \in A_m$ and $z_2 \in A_n$ and suppose $m < n$. This implies that $z_1 \in A_n$ and every bi-convex combination of z_1 and z_2 is in A_{n+1} , i.e. in B . Thus, B is a bi-convex set and the result is immediate. We now note the following remark.

A.11 Lemma. *Let $A \subset \mathcal{X} \times \mathcal{Y}$ and let $B = \text{bi-co}(A)$. Then for any bi-convex function $f : B \rightarrow \mathbb{R}$, $\sup f(B) = \sup f(A)$.*

Proof. Let us take $A_0 := A$ and as above, let $A_1 := \langle A_0 \rangle_b$ and $A_{i+1} := \langle A_i \rangle_b$. For any $a \in A_1$, we can write $a = \sum \alpha_i a_i$, a bi-convex combination where $a_i \in A_0$. This implies $\sup f(A_0) \geq \sum \alpha_i f(a_i) \geq f(\sum \alpha_i a_i) = f(a)$, where the second-to-last inequality is because f is bi-convex. Similarly, $\sup f(A_i) \geq \sup f(A_{i+1})$. Therefore, we get $\sup f(A_0) \geq \sup f(A_1) \geq \dots \geq \sup f(A_i) \geq \dots$. But (A_i) is an increasing sequence of sets with limit $B = \text{bi-co}(A_0)$, which implies that $\sup f(A_0) \geq \sup f(B)$. \square

The Lemma above is equivalent to the observation that for any set A with $B = \text{bi-co}(A)$, $B = \text{ns}_A(B)$. In other words, $B \subset A^\#$, a fact which follows immediately from the definitions of B and $A^\#$. As mentioned above, we shall be concentrating on the generators of sets. The lemma below says that this is without loss of generality.

A.12 Lemma. *Let A be generated by A_0 . Then*

- (i) $\text{bi-co}(A_0) = \text{bi-co}(A)$ and
- (ii) $A_0^\# = A^\#$.

Proof. (i) By definition, $\text{bi-co}(A_0) \subset \text{bi-co}(A)$. If the inclusion is strict, then there exists a bi-convex set B' such that $A_0 \subset B'$ and $A \not\subseteq B'$. But this is impossible as every point in A is a bi-convex combination of points in A_0 .

(ii) Suppose $A_0^\# \subsetneq A^\#$, then there exists $z \in A^\#$ and a bi-convex function $f : A^\# \rightarrow \mathbb{R}$ such that $f(z) > \sup f(A_0)$. (If such a point and corresponding function did not exist, then $A^\# = \text{ns}_{A_0}(A^\#)$ which, by Theorem A.6, contradicts $A_0^\#$'s maximality.)

Now, for any $a \in A$, we know that $a = \sum \alpha_i a_i$, a bi-convex combination with $a_i \in A_0$. But $f(z) > \sup f(A_0) \geq \sum \alpha_i f(a_i) \geq f(\sum \alpha_i a_i) = f(a)$, where the second-last inequality is because f is bi-convex. In other words, $f(z) > \sup f(A)$ which contradicts $A^\# = \text{ns}_A(A^\#)$ which implies that such a z cannot exist. Thus, $A^\# \setminus A_0^\# = \emptyset$, i.e. $A^\# = A_0^\#$. \square

Before introducing an important concept regarding the geometry of bi-convex sets, we recall some analogous ideas in the convex case.

A.13 Definition. An *extreme point* of a convex set S in a vector space E is a point $x \in S$ such that for all $y_1, y_2 \in S$ such that $x = ty_1 + (1 - t)y_2$ with $0 < t < 1$ implies $y_1 = y_2$.

A.14 Lemma. Let S be a non-empty, compact, convex subset of E , a vector space (over the reals). Then there exists an extreme point of S .

Of course, we are assuming that E^* is a vector space of linear maps of E into \mathbb{R} (not necessarily the dual of E) and that E^* separates E , that is, if $x \in E$, then there exists $\lambda \in E^*$ such that $\lambda(x) \neq 0$. Also, E^* is endowed with the weak topology, i.e. the coarsest topology which makes all the $\lambda \in E^*$ continuous. The other important property of convex sets is that they can be expressed as the convex hulls of their extreme points. This is the celebrated Krein–Milman theorem. The proofs of both Lemma A.14 and the Krein–Milman Theorem can be found in the Appendix to Chapter IV in [16].

Krein–Milman Theorem. Let K be a convex, compact set in a vector space E . Let S be the set of extreme points of K . Then K is the smallest closed convex set containing all the elements of S .

Bi-convex sets also have points analogous to extreme points in the convex case. These are what we shall call *bi-extreme points*, defined below.

A.15 Definition. Let $B \subset \mathcal{X} \times \mathcal{Y}$ be a bi-convex set. A point $z \in B$ is a *bi-extreme point* of B if for all z_1, z_2 such that $z = tz_1 + (1 - t)z_2$ (with $0 < t < 1$) is a bi-convex combination implies $z_1 = z_2$.

We shall prove below that if B is the bi-convex hull of a compact set then it has a bi-extreme point. We shall first prove a useful lemma.

A.16 Lemma. Let $S \in \mathbb{R}^n$ be compact. Then

- (i) $\text{co}(S)$ has an extreme point; and
- (ii) if $x \in \text{co}(S)$ is an extreme point, then $x \in S$.

Proof. (i) Since S is compact, it follows that $\text{co}(S)$ is compact. By Lemma A.14, $\text{co}(S)$ has an extreme point.

(ii) Now $\text{co}(S) = S \sqcup (\text{co}(S) \setminus S)$, a disjoint union. Let $x \in \text{co}(S)$ be an extreme point of $\text{co}(S)$. If $x \in \text{co}(S) \setminus S$, then x is a convex combination of points in S , contradicting its being an extreme point. Therefore, if x is an extreme point of $\text{co}(S)$, $x \in S$. \square

We now show that the bi-convex hull of compact set has a bi-extreme point.

A.17 Lemma. Let $A \subset \mathcal{X} \times \mathcal{Y}$ be compact. Then $\text{bi-co}(A)$ has a bi-extreme point in A .

Proof. Recall that $\mathcal{X} \subset \mathbb{R}^m$ and $\mathcal{Y} \subset \mathbb{R}^n$. Let $\text{co}(A)$ be the convex hull of A in \mathbb{R}^{m+n} . Then $\text{co}(A)$ has an extreme point, by Lemma A.14 above and if z is an extreme point of $\text{co}(A)$, then $z \in A$. But this means that z cannot be written as a convex combination of elements in $\text{bi-co}(A)$ (which lies in $\text{co}(A)$) which implies that it cannot be written as a bi-convex combination of elements in $\text{bi-co}(A)$ (as every bi-convex combination is also a convex combination). Thus, z is a bi-extreme point of $\text{bi-co}(A)$. \square

A.2 A Lyrical Digression

It may be pertinent to point out that the bi-convex case is sufficiently inhomogeneous, in the sense that some fundamental theorems from convexity theory do not carry over to the bi-convex case.²³ We first note that there is no analogue to the Krein–Milman Theorem in the bi-convex case. An example will illustrate this fact. Let us reconsider the set first encountered in Example 5.6.

A.18 Example. Let $A := \{a_1 = (2/3, 0), a_2 = (0, 1/3), a_3 = (1/3, 1), a_4 = (1, 2/3)\}$ and $\mathcal{X} = \mathcal{Y} = [0, 1]$. Consider the bi-convex hull of $\{a_i, w_i\}_{i=1}^4$, as illustrated in figure 9. Clearly, the (only) bi-extreme points of this bi-convex set are $A = \{a_i\}_1^4$. But $\text{bi-co}(A) = A$ and not the larger bi-convex set.

As a second example, consider the fact that a *polyhedral*²⁴ convex set is finitely generated, i.e. is the convex hull of finitely many points (cf. Theorem 19.1 in [22]). We may call a set *semi-bi-linear* if it is defined by finitely many bi-linear functions.²⁵ Such a bi-convex set need not be finitely generated as demonstrated in the example below.

A.19 Example. Define $A := \{(x, y) \in \mathbb{R}_+^2 : xy \geq 1, (x-4)(y-4) \geq 1, x \leq 4 \text{ and } y \leq 4\}$. It is easy to see that A is bi-convex and both the functions $f(x, y) := xy$ and $g(x, y) := (x-4)(y-4)$ are bi-affine. But A is *not* finitely generated. Indeed, its boundary has curvature, which is not the case for a polyhedral set.

²³Carathéodory's Theorem also does not hold, as is demonstrated in [2].

²⁴A polyhedral convex set (or equivalently, a *semi-linear* convex set), is defined by finitely many linear (weak) inequalities.

²⁵A bi-linear function, also known as a *bi-affine* function is a function $f : B \rightarrow \mathbb{R}$ (where $B \subset \mathcal{X} \times \mathcal{Y}$ is bi-convex) such that $f(x, \cdot)$ is an affine function on B_x and $f(\cdot, y)$ is an affine function on B_y , where B_x and B_y are x - and y -sections of B .

A.3 Proof of Lemma 5.14

We now relate the discussion above to the finitely generated sets that are our primary concern. We shall be interested in the following property of finite sets in product spaces.

Condition H. For any set $E \subset A_0$, there exists an n such that $\text{bi-co}(E_0) = \langle E_0 \rangle_b^n$.

A.20 Lemma. Let $A_0 \subset \mathcal{X} \times \mathcal{Y}$ be a finite set. Then A_0 satisfies Condition F iff A_0 satisfies Condition H.

Proof. Straightforward. □

We are interested in the expectations of bi-martingales with limits in a finite set A_0 . But these sets of expectations can also be represented as the limits of a separation process (see Theorems A.6 and A.7 above). The following lemma shows that if any separation takes place (from a finite set) in the bi-convex hull of a finite set, then there is also separation of a bi-extreme point. This is made precise below.

A.21 Lemma. Let $A_0 \subset \mathcal{X} \times \mathcal{Y}$ be finite, $A_0 \subset W_0$ also finite and $B = \text{bi-co}(W_0)$ and suppose $f : B \rightarrow \mathbb{R}$ is a bi-convex function. Also suppose that W_0 satisfies Condition H. Then there exists $z \in B$ with $f(z) > \sup f(A_0)$ iff there exists a bi-extreme point of B , $\hat{w} \in W_0$ such that $f(\hat{w}) > \sup f(A_0)$.

Proof. It follows from Lemma A.17 that B has a bi-extreme point and that any bi-extreme point of B must lie in W_0 . Note that if $f, g : B \rightarrow \mathbb{R}$ are bi-convex functions,²⁶ then $\max\{f, g\}$ is also a bi-convex function. This is because the maximum of two convex functions is convex and f and g are convex on every section. Now suppose g is a bi-convex function. Then $f := \max\{g - \max g(A_0), 0\}$ is a bi-convex function that separates everything that g does (i.e. $f(z) > 0$ iff $g(z) > 0$) and takes the value 0 on all of A_0 . We can thus restrict attention to f such that $f(A_0) = \{0\}$.

As $B = \text{bi-co}(W_0)$, by Lemma A.11 it is the case that $\max f(W_0) = \sup f(B)$. This implies that if there is $z \in B$ with $f(z) > 0$, then there exists $\hat{w} \in W_0$ such that $f(\hat{w}) = \max f(W_0) \geq f(z) > 0$. We shall now show that one can take such a \hat{w} to be a bi-extreme point of B .

Let n be such that $W_n = \langle W_0 \rangle_b^n = B = \text{bi-co}(W_0)$ and suppose \hat{w} is not a bi-extreme point. Then we can write \hat{w} as a (non-trivial) bi-convex combination of points that are extreme in one of the (x - or y -) sections of \hat{w} . Let us replace \hat{w} in W_{n-1} with these points. Note that all these points must also take the maximum value under f . If none of these points are bi-extreme in B , then proceed inductively and replace points in W_{n-2} and so on. This

²⁶We shall assume that the domain of a bi-convex function is a bi-convex set.

process stops in finite time with some bi-extreme point taking the maximum value under f . \square

We shall recall a few more definitions (from [3]).

A.22 Definition.

1. A set $D \subset \mathcal{X} \times \mathcal{Y} \times \mathbb{R}$ is **di-convex** if for each $x \in \mathcal{X}$ and each $y \in \mathcal{Y}$, the respective x - and y -sections $D_x := \{(y, \mu) \in \mathcal{Y} \times \mathbb{R} : (x, y, \mu) \in D\}$ and $D_y := \{(x, \mu) \in \mathcal{X} \times \mathbb{R} : (x, y, \mu) \in D\}$ are convex.
2. A convex combination $(x, y, \mu) = \sum_i \alpha_i (x_i, y_i, \mu_i)$ (with $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$) will be called **di-convex** if either $x_1 = \dots = x_n = x$ or $y_1 = \dots = y_n = y$.
3. The **di-convex hull** of a set A is the smallest di-convex set containing A .

A.23 Lemma. Let $W_0 \subset \mathcal{X} \times \mathcal{Y}$ be finite, $B = \text{bi-co}(W_0)$, $\widehat{w} \in W_0$ be a bi-extreme point of B and suppose W_0 satisfies Condition H. Then there exists a continuous, bi-convex function $g : B \rightarrow \mathbb{R}$ such that (i) $g(\widehat{w}) = 1$, (ii) $g(z) > 0$ for all $z \in \text{bi-co}(W_0) \setminus \text{bi-co}(W_0 \setminus \{\widehat{w}\})$ and (iii) $g(z) = 0$ for all $z \in \text{bi-co}(W_0 \setminus \{\widehat{w}\})$.

Proof. Let us define

$$g(w) = \begin{cases} 1 & \text{if } w = \widehat{w}; \\ 0 & \text{otherwise.} \end{cases}$$

Let us also define $D_0, \widetilde{D}_0 \subset \mathcal{X} \times \mathcal{Y} \times \mathbb{R}$ by $D_0 = \{(w, g(w)) : w \in W_0\}$ and $\widetilde{D}_0 = \{(w, g(w)) : w \in W_0 \setminus \{\widehat{w}\}\}$. We shall denote the set of di-convex combinations of a set D by $\langle D \rangle_d$. This permits us to define inductively, $D_{i+1} = \langle D_i \rangle_d$ and $\widetilde{D}_{i+1} = \langle \widetilde{D}_i \rangle_d$. Note that D_0 and \widetilde{D}_0 are compact (closed and bounded), which implies that for each i , D_i and \widetilde{D}_i are also compact. Note also that for each i , $W_i := \langle W_0 \rangle_b^i = \text{proj}_{\mathcal{X} \times \mathcal{Y}} D_i$ and $\widetilde{W}_i := \langle \widetilde{W}_0 \rangle_b^i = \text{proj}_{\mathcal{X} \times \mathcal{Y}} \widetilde{D}_i$. We shall demonstrate our claim by induction.

Induction hypothesis. For all $z \in \widetilde{W}_i$, $\min\{\mu : (z, \mu) \in \widetilde{D}_i\} = 0$, for all $z \in W_i \setminus \widetilde{W}_i$, $\min\{\mu : (z, \mu) \in \widetilde{D}_i\} > 0$ and $\widehat{w} \in W_i \setminus \widetilde{W}_i$ such that $\min\{\mu : (\widehat{w}, \mu) \in \widetilde{D}_i\} = 1$.

Now suppose $z \in \widetilde{W}_{i+1}$. Then $z = \sum \alpha_j z_j$ a bi-convex combination where $z_j \in \widetilde{W}_i$ for each j . Then, $\min\{\mu : (z, \mu) \in \widetilde{D}_{i+1}\} = 0$. Also, if $z \in W_{i+1} \setminus \widetilde{W}_{i+1}$, then $z = \sum \alpha_j z_j$, a bi-convex combination with some $z_k \in W_i \setminus \widetilde{W}_i$ and the corresponding $\alpha_k > 0$. Therefore $(z, \sum \alpha_j \min\{\mu : (z_j, \mu) \in \widetilde{D}_i\}) \in D_{i+1} \setminus \widetilde{D}_{i+1}$ and $\min\{\mu : (z, \mu) \in D_{i+1} \setminus \widetilde{D}_{i+1}\} > 0$. As \widehat{w} is a bi-extreme point, it cannot be written as a bi-convex combination of other point so that $\min\{\mu : (\widehat{w}, \mu) \in \widetilde{D}_{i+1}\} = 1$.

By Condition H , we know that there is an n such that $D = \langle D_0 \rangle_d^n$ is the di-convex hull of D_0 . Let us now define

$$g(z) = \min\{\mu \in \mathcal{X} \times \mathcal{Y} \times \mathbb{R} : (z, \mu) \in D\}.$$

As each section of a di-convex set is convex, $g(z)$ is convex on each section (cf. Theorem 5.3 in [22]). Also, $g(z)$ is as required, by construction. \square

We are now able to take our final steps.

A.24 Lemma. *Let A_0 be finite, $A_0 \subset W_0$ finite and $B = \text{bi-co}(W_0)$ and suppose W_0 satisfies Condition H . Then $B \neq \text{ns}_{A_0}(B)$ implies $B \neq \text{nsc}_{A_0}(B)$.*

Proof. Suppose $B \neq \text{ns}_{A_0}(B)$. (As mentioned above, we can take all functions to take the value 0 on A_0 .) Then there exists $z \in B$ and $f : B \rightarrow \mathbb{R}$ bi-convex such that $f(z) > 0$. But by Lemma A.21, there also exists a bi-extreme point of B , $\hat{w} \in W_0$ such that $f(\hat{w}) > 0$. Moreover, there exists a continuous function $g : B \rightarrow \mathbb{R}$ such that $g(z) > 0$ for all $z \in \text{bi-co}(W_0) \setminus \{\hat{w}\}$. \square

A.25 Lemma (Lemma 5.14 in the text). *Let $A \subset \mathcal{X} \times \mathcal{Y}$ be finitely generated and suppose Condition H holds. Then $A^* = A^\#$.*

Proof. Let A_0 generate A . Consider W_0 finite such that $A_0 \subset W_0$ and note that $A_0^\# = A^\#$. If $B = \text{bi-co}(W_0)$ then $\text{ns}_{A_0}(B) \neq B$ implies that $\text{nsc}_{A_0}(B) \neq B$. But in the proof of the lemma above, we proved this by constructing a function that was continuous everywhere on the domain. Thus, we actually proved $(\text{ns}_{A_0}(B) \neq B) \rightarrow (\text{nsc}_A(B) \neq B)$. But this is not enough, as we have to show that the set A^* (and consequently $A^\#$) is obtained in this way. Let $W_0 : \{(x, y) : x \in \text{proj}_{\mathcal{X}} A_0, y \in \text{proj}_{\mathcal{Y}} A_0\}$ and let $W_{i+1} := \{w \in W_i : w \text{ is not a bi-extreme point of } \text{bi-co}(W_i)\}$. This process ends in finite time at W and $A^* = A^\# = \text{bi-co}(W)$. \square

A.4 Proof of Theorem 6.2

A.26 Theorem (Theorem 6.2 in the text). *Let Γ be a Sender-Receiver game that satisfies Condition F . Then the \mathbf{q}_0 -section of G^* is identical to the \mathbf{q}_0 -section of $(G^*)^*$ which is the same as the \mathbf{q}_0 -section of $G^\#$.*

Proof. First, a simple observation. Let $A \subset \mathcal{X} \times \mathcal{Y}$ be a finite set and let $W := \text{proj}_{\mathcal{X}} A \times \text{proj}_{\mathcal{Y}} A$. Then, for all $C \subset W$ such that $C \supset A$, $\text{proj}_{\mathcal{X}} C \times \text{proj}_{\mathcal{Y}} C = W$. To see this, first note that for all $c = (c_x, c_y) \in C$, $c_x \in \text{proj}_{\mathcal{X}} W$ and $c_y \in \text{proj}_{\mathcal{Y}} W$. Thus, $C \subset W$. Furthermore, $W = \text{proj}_{\mathcal{X}} A \times \text{proj}_{\mathcal{Y}} A \subset \text{proj}_{\mathcal{X}} C \times \text{proj}_{\mathcal{Y}} C$. Therefore, $\text{proj}_{\mathcal{X}} C \times \text{proj}_{\mathcal{Y}} C = W$.

Now, let G be the graph of the modified payoffs correspondence of the Sender-Receiver game, Γ and let G_0 be the finite set that generates G . As above, let $W_0 := \text{proj}_{\mathbb{A}} G_0 \times \text{proj}_{\mathbb{Q}} G_0$. Since Γ satisfies Condition F , it is the case that $G^\# = G^*$. Also, from the proof of lemma 5.14, we know that $G^* \subset \text{bi-co}(W_0)$ and that G^* is finitely generated by some $W' \subset W_0$ with $W' \supset G_0$. Therefore, $\text{proj}_{\mathbb{A}} W' \times \text{proj}_{\mathbb{Q}} W' = W_0$. But recall that W_0 satisfies Condition H , so that by the same argument as in the proof of lemma 5.14, it follows that $(G^*)^\# = (G^*)^*$. From [2] (page 178), we know that $(G^\#)^\# = G^\#$. Therefore, we have $G^* = G^\# = (G^\#)^\# = (G^*)^\# = (G^*)^*$, where the first and third equalities are lemma 5.14 and the second equality is from page 178 of [2]. \square

A.5 Graph of Payoffs in Sender-Receiver Games

We now present an axiomatisation of the fundamental unit of our analysis—the graph of payoffs of the silent game. For simplicity, we shall only consider generic Sender-Receiver games and we shall also avoid the additional notational burden of modifying the payoff correspondence. (Taking care of both of these situations is straightforward.)

Let K be a simplicial complex so that $|K| = \Delta^{n-1} \subset \mathbb{R}^n$ (where n is the number of types of the sender and denote by K_n , the n -simplices in K). Let $\sigma \in K_n$ be an n -simplex and $a : K_n \rightarrow \mathbb{R}^n$ be a map associating an n -tuple of real numbers with each n -simplex and define

$$G_0 := \bigcup_{\sigma \in K_n} \left(\bigcup_{q \in \text{Vert}(\sigma)} (q, a(\sigma)) \right).$$

A.27 Lemma. *There exists a Sender-Receiver game Γ with graph of the payoff correspondence of the silent game \mathcal{E} such that $\text{proj}_{\mathbb{A} \times \mathbb{Q}} \mathcal{E}$ is generated by G_0 .*

The idea is that each action of the Receiver is optimal for some beliefs over types. Since the expected payoff from an action is linear in probabilities, this induces a simplicial decomposition of Δ^{n-1} . But over each simplex, the Sender's payoff is constant.

Proof. For each $\sigma \in K_n$, assign the Receiver an action c_σ with payoffs to (each type of) the Sender $a(c_\sigma)$. Now to define the Receiver's payoffs so that it induces the required simplicial decomposition. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be given by $f(x) := \|x - \alpha\|^2$ where $\alpha \in \mathbb{R}^n$ is such that for all $v, w \in \text{Vert}(K)$, $f(v) \neq f(w)$. Such an α can always be chosen and is, in fact, generic (by the Baire Category Theorem). Note that f is convex which means that f restricted to Δ^{n-1} is also convex. Now let $F := \{(v, f(v)) : v \in \text{Vert}(K)\}$ and let

$$g := \min\{\mu : (x, \mu) \in \text{co}(F)\}.$$

It follows from Theorem 5.3 in [22] that g is a convex function on Δ^{n-1} . Moreover, it is polyhedral and for all $v \in \text{Vert}(K)$, $f(v) = g(v)$. Also, every n -simplex in K is the projection of some lower face of $\text{co}(F)$, which ensures that g induces the required simplicial decomposition. For each n -simplex σ , $g|_{\sigma}$ is a linear function which can be extended to the entire ambient space and denote the resulting linear function by $\beta_{c_{\sigma}}$. Let e_i be the vertex in Δ^{n-1} which assigns probability 1 to type i of the Sender.

We now assign the Receiver the payoff $\beta_{c_{\sigma}}(e_i)$ when the Receiver takes the action c_{σ} and the type of the Sender is i . This gives us all the components of the game Γ . \square

References

- [1] Aumann, Robert J. (1964). “Mixed and Behaviour Strategies in Infinite Extensive Games,” M. Dresher *et. al.* (eds.), *Advances in Game Theory (Annals of Mathematics Study, 52)*, Princeton University Press, 627–650.
- [2] Aumann, Robert J. and Sergiu Hart (1986). “Bi-convexity and Bi-martingales,” *Israel Journal of Mathematics, 54(2)*, 159–180.
- [3] Aumann, Robert J. and Sergiu Hart (2003). “Long Cheap Talk,” *Econometrica, 71(6)*, 1619–1660.
- [4] Ben-Porath, Elchanan (1998). “Communication without Mediation: Expanding the Set of Equilibrium Outcomes by “Cheap” Pre-play Procedures,” *Journal of Economic Theory, 80*, 108–122.
- [5] Ben-Porath, Elchanan (2003). “Cheap Talk in Games with Incomplete Information,” *Journal of Economic Theory, 108*, 45–71.
- [6] Bester, Helmut and Roland Strausz (2001). “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica, 69(4)*, 1077–1098.
- [7] Blume, Andreas and Joel Sobel (1995). “Communication-Proof Equilibria in Cheap-Talk Games,” *Journal of Economic Theory, 65*, 359–382.
- [8] Blume, Lawrence E. and William R. Zame (1994). “The Algebraic Geometry of Perfect and Sequential Equilibrium,” *Econometrica, 62(4)*, 783–794.
- [9] Bredon, Glen E. (1993). *Topology and Geometry, GTM 139*, Springer-Verlag, New York.
- [10] Chung, Kai Lai (2001). *A Course in Probability Theory, Third Edition*, Academic Press, San Diego.
- [11] Crawford, Vincent and Joel Sobel (1982). “Strategic Information Transmission,” *Econometrica, 50*, 579–594.
- [12] Dries, Lou van den (1998). *Tame Topology and O-minimal Structures, London Mathematical Society Lecture Note Series, 248*, Cambridge University Press.
- [13] Forges, Françoise (1990). “Equilibria with Communication in a Job Market Example,” *Quarterly Journal of Economics, 105*, 375–398.

-
- [14] Forges, Françoise (1994). “Non-Zero Sum Repeated Games and Information Transmission,” in *Essays in Game Theory in Honor of Michael Maschler*, 65–95, ed. Nimrod Megiddo, Springer-Verlag, New York.
- [15] Hart, Sergiu (1985). “Nonzero-Sum Two-Person Repeated Games with Incomplete Information,” *Mathematics of Operations Research*, 10(1), 117–153.
- [16] Lang, Serge (1993). *Real and Functional Analysis, Third Edition, GTM 142*, Springer-Verlag, New York.
- [17] Lehrer, Ehud and Sylvain Sorin (1997). “One-Shot Public Mediated Talk,” *Games and Economic Behavior*, 20, 131–149.
- [18] Levy, Azriel (1979). *Basic Set Theory*, Springer-Verlag, New York.
- [19] Mertens, Jean-François (1989). “Stable Equilibria—A Reformulation,” *Mathematics of Operations Research*, 14(4), 575–625.
- [20] Myerson, Roger B. (1986). “Multistage Games with Communication,” *Econometrica*, 54, 323–358.
- [21] Myerson, Roger B. (1991). *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts and London, England.
- [22] Rockafellar, R. Tyrell (1970). *Convex Analysis*, Princeton University Press.
- [23] Rubinstein, Ariel (1982). “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50, 97–110.
- [24] Rubinstein, Ariel (1991). “Comments on the Interpretation of Game Theory,” *Econometrica*, 59(4), 909–924.
- [25] Simon, Robert Samuel (2002). “Separation of Joint Plan Equilibrium Payoffs from the Min-max Functions,” *Games and Economic Behavior*, 41, 79–102.
- [26] Simon, R., S. Spiez, and H. Torunczyk (1995). “The Existence of Equilibria in Certain Games, Separation for Families of Convex Functions and a Theorem of Borsuk-Ulam Type,” *Israel Journal of Mathematics*, 92, 1–21.
- [27] Smullyan, Raymond M. (1995). *First-Order Logic*, Dover Publications, New York.
- [28] Urbano, Amparo and Jose E. Vila (2002). “Computational Complexity and Communication: Coordination in Two-Player Games,” *Econometrica*, 70(5), 1893–1927.