The Pennsylvania State University

The Graduate School

Department of Economics

Essays in Game Theory and Political Economy

A Thesis in

Economics

by

R. Vijay Krishna

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

December 2004

The thesis of R. Vijay Krishna has been reviewed and approved* by the following:

Kalyan Chatterjee
Distinguished Professor of Economics
and Management Science
Thesis Co-advisor
Chairman of Committee

J. Tomas Sjöström
Professor of Economics
Thesis Co-advisor

James S. Jordan
Professor of Economics

Anthony M. Kwasnica
Assistant Professor of Business Economics

Robert C. Marshall
Professor of Economics
Head of the Department of Economics

*Signatures are on file in the Graduate School.

Abstract

Game Theory plays a fundamental role in the social sciences. In this dissertation, we present three essays—the first two dealing with the pure theory of games and the third with an application of game theory to political economy. In the first essay, we consider the effect of unbounded communication between agents. 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. Their research provokes (at least) three questions. The first is with regards to the structure of equilibrium payoffs when the player can communicate for large but finitely many periods. Is it true that the set of equilibrium payoffs stabilises (i.e. remains unchanged) if there are sufficiently many rounds of communication? The second is if the set of equilibria from communication which is unbounded but finite with probability one is the same as equilibria from communication which is just unbounded. The third question is whether any of these sets of equilibria are "simple" and if so, is there an algorithm to compute them. We show that in the context of finite Sender-Receiver games, the answer to all three 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.

In the second essay, we study the effect of communication in two-person games of incomplete information. We show that any rational mediated communication mechanism outcome satisfying a Nash domination condition can be implemented as the perfect Bayesian equilibrium of an extended communication game built from the original game and ends in finite time with probability 1.

In the third paper, we study a model of spatial voting where players have objectives different from the traditional Hotelling-Downs model. The players are rank maximisers and enter sequentially with a given order of entry and where k winners get a payoff v. If the players make their decisions and enter simultaneously, the resulting Nash equilibrium (with k winners) is the Greenberg-Shepsle k-equilibrium, which does not exist for all distributions of voters. We partition the set of distributions into 3 cells. The first of these cells consists of *regular* distributions, wherein the set of SPE is isomorphic to the set Greenberg-Shepsle equilibria in the sense that every SPE generates an outcome which is a k-equilibrium and every k-equilibrium is the outcome of an SPE. The second type of distributions are called *semi-clustered* distributions. These have Greenberg-Shepsle k-equilibria but also have SPE

that generate outcomes which are not k-equilibria. The remaining distributions are called *clustered* distributions. These distributions do not have k-equilibria. Nonetheless, they always have SPE with k entrants. We also give conditions to determine if a distribution has a k-equilibrium or not.

Contents

List of Figures vii

Acknowledgements viii

1 Overview 1

2 Extended Conversations in Sender-Receiver Games 2

2.1 Introduction 2
2.2 Related Literature 4
2.3 Model 6
2.4 The Mechanics of Cheap Talk 7
2.4.1 A Qualitative Discussion 7
2.4.2 The Mathematics of Cheap Talk 10
2.4.3 Game Theoretic Aspects of Cheap Talk 13
2.5 Long Conversations 15
2.6 (Trans)finite Conversations 22
2.7 Conclusion 25
2.8 Appendix 26
2.8.1 Bi-convexity and bi-martingales 27
2.8.2 A Lyrical Digression. 32
2.8.3 Proof of Lemma 2.5.14 32
2.8.4 Proof of Theorem 2.6.2 35
2.8.5 Graph of Payoffs in Sender-Receiver Games 36

3 Communication in Games of Incomplete Information: The Two Player Case 37

3.1 Introduction 37
3.2 Related Literature 40
3.3 The Model and Main Result 42
3.4 An Example 45
3.5 Proof of Main Theorem 48

3.6 Discussion 49

4 Voter Clustering and the Theory of Spatial Voting with Entry **52**

4.1 Introduction 52
4.2 The Model 56
4.2.1 The Environment 56
4.2.2 The Mechanism 57
4.3 Two Winners - A Complete Analysis 60
4.4 Examples 65
4.4.1 Regular Society 65
4.4.2 Semi-Clustered Society 66
4.4.3 Clustered Society 66
4.5 Concluding Remarks 67
4.5.1 On Polarised Societies 67
4.5.2 Multi-Dimensional Voting 68
4.5.3 Our Results 68
4.6 Appendix 69
4.6.1 Palfrey's (1984) dominant 2-candidate model. 69
4.6.2 Non-uniqueness of solutions to equations (4.1) and (4.2). 71
4.6.3 Analysis for arbitrary k. 71

Bibliography **75**

List of Figures

2.1 Illustrating signalling 8
2.2 Illustrating a Compromise 9
2.3 A bi-convex set that is not convex. 16
2.4 Example where $A = \text{bi-co}(A) \subsetneq A^*$ 17
2.5 Example 2.5.10 18
2.6 Receiver's payoffs given by upper envelope 18
2.7 Graph of $\text{proj}_{\mathbb{A}\times\mathbb{Q}}\, \mathscr{E}^+(q)$ — commonly known as the Sender's modified payoffs 19
2.8 $a_1 + a_2 = 4$ section of graph of Receiver's payoffs 22
2.9 $A = \{a_i\}_1^4 = \text{bi-co}(A) \subsetneq A^*$ 29

3.1 Example 46

Acknowledgements

I would like to thank my advisors Profs. Kalyan Chatterjee and Tomas Sjöström for their encouragement and advice. Prof. James S. Jordan provided immense support during the trying times that lead to the proof of Theorem A in chapter 2. I would also like to thank Françoise Forges and Roland Strausz for comments on earlier drafts. I would also like to thank my friends and colleagues, Siddhartha Bandhopadhyay, Stephen Binns and Abhiroop Mukhopahyay from whom I learned so much.

Chapter 1

Overview

Game theory is fundamental to analysis in modern social science. A game describes the strategic interaction between agents. Typically, this is a difficult problem when there are many agents. In chapters 2 and 3 of this thesis, we shall look at the simple case of interaction between two agents. Although this is conceptually very simple, the analysis of such situations is anything but.

In chapter 2, we shall be interested in the case where a Sender who has private information about the state of the world communicates with a Receiver who takes the only payoff relevant action. Such a situation is typical when an agent asks an expert for an opinion. For example, one may consider the situation where there is testimony before Congress. Another example is a restaurant recommendation by a tour guide to a tourist. Here, the interesting question is how much information can be transmitted when preferences are not aligned. If we only consider cheap talk, then the set of outcomes that can achieved is characterised by Aumann and Hart [3]. But their research provokes many questions. Namely, does the set of equilibria from finite conversations eventually stabilise? Are conversations that are finite a.s. the same as unbounded conversations? Are the sets of equilibrium outcomes of these different conversations simple? Is it the case that transfinite conversations do not expand the set of equilibrium outcomes? We show that all the questions are answered in the affirmative if a certain condition is satisfied by a game.

In chapter 3, we consider more complex unmediated mechanisms and show that the set of equilibrium outcomes that can be achieved approximates the set of mediated outcomes. Finally in chapter 4, we consider the political economy part of the thesis. Here, we consider the problem of sequential entry in a spatial voting model where politicians care about winning and entering the fray is costly. We characterise the outcomes possible in terms of the distribution voter preferences of the population.

Chapter 2

Extended Conversations in Sender-Receiver Games

2.1 Introduction

The problem of strategic interaction between two agents with private information has a central place in modern economic theory. The now classical method of solution is to use the Revelation Principle, which reduces the problem to a programming problem. Here, the players are assumed to be able to write an enforceable contract after which the players reveal their private information. If enforceable contracts cannot be written, the set of payoffs achievable can still be obtained provided there is a disinterested mediator who makes private recommendations based on the private messages that he receives. Unfortunately, in many applications, the players have neither the ability to commit to an enforceable contract nor access to a disinterested mediator. Nevertheless, it is known that if the agents are allowed to communicate, then there can be some useful transmission of information.

The usual assumption in the case of Sender-Receiver games – games where the Sender has some private information and the Receiver 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 only involves *only* one round of signalling by the Sender. In principle, there is no reason to assume *a priori* that there cannot be any future 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.

This mean that the natural objects to consider are conversations that do not have exogenous deadlines. The reader may feel that talk without a 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 [38] 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 [37], 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 §2.3.) 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 §2.2, we describe the relative position of our results with respect to the extant literature. In §2.3, we give a formal description of the model. In §2.4, we discuss the mechanics of information transmission in cheap talk games. In §2.4.1, we give a qualitative description of the workings of conversations and argue that one can restrict attention to the so-called canonical conversations. In §2.4.2, 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 §2.4.3, 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 §2.5. In §2.6, we consider both finite conversations and transfinite conversations and discuss the nature of our sufficient condition. §2.7 concludes and Appendix 2.8 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.

2.2 Related Literature

The theory of strategic information transmission when there are no signalling costs can be traced back to Crawford and Sobel [12], 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 [7] 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 [8] 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 [23] 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 [40]).

Extended conversations are described formally by Aumann and Hart in [3] who provide a complete characterisation of the set of equilibrium payoffs from unbounded coversations. An example that shows that extended talk can increase the set of equilibrium outcomes is due to Forges [16] 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 [18]. Another example due to Simon [39], 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 allow for Pareto improvements for all types of the Sender and the Receiver.

The standard 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 [31].

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 [6] 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 [43] 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 in Galois fields are extremely hard to compute.

2.3 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 $|T_S| = k$,[1] 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 $\Delta^{k-1} := \{p \in \mathbb{R}^k \ : \ \sum_1^k p_i = 1\}$ is the $(k-1)$-dimensional simplex) with $p \gg 0$ and $u_S : C_R \times T_S \to \mathbb{R}$ and $u_R : C_R \times T_S \to \mathbb{R}$ are utility functions[2] for the Sender and Receiver respectively. In other words, the only payoff 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 $|\mathsf{M}| > 1$). 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

[1] For any set A, we denote its cardinality by $|A|$.

[2] The 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.

message) while refraining from giving a formal definition. The length of the communication phase is denoted by $\mathscr{L}_{\mathrm{c}}$, a random variable that maps into the *class* of ordinals.[3] (Note that although $\mathscr{L}_{\mathrm{c}}$ can be *any* ordinal number (up to 2^{ω}), we will primarily be interested in the cases where $\mathscr{L}_{\mathrm{c}}$ 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 $\beta : C_R \to \mathbb{R}$ 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 $a : C_R \to \mathbb{R}^k$. We assume that there is perfect recall and that the model of communication is commonly known between the players.

2.4 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 §2.4.1 we shall informally describe the mechanics of cheap talk and in §2.4.2 we shall introduce the mathematics that describes cheap talk. In §2.4.3 we discuss the game theoretic aspects of long cheap talk.

2.4.1 A Qualitative Discussion

We shall assume here that each player has a finite message set, $\mathsf{M} := \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \ldots\}$. 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

[3]There are some complications that arise if we actually let $\mathscr{L}_{\mathrm{c}}$ take any value in the class of ordinals. Instead, we can assume that $\mathscr{L}_{\mathrm{c}}$ 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.

decision problem of the Receiver given his beliefs.

With only one round of communication, cheap talk is just signalling,[4] 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 [16] and [18].)

Example 2.4.1 *Signalling.*

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

Figure 2.1: Illustrating signalling

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 $\mathsf{M} := \{\mathsf{x}, \mathsf{y}\}$ 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 $q_{\mathsf{m}} :=$ prob(Sender is type $t_1|\mathsf{m}$) where $\mathsf{m} \in \mathsf{M}$. 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 is p_1 were positive, then $a^1_{\mathsf{x}} < a^1_{\mathsf{y}}$ (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. ✠

More generally, if each type of the Sender sends both messages with probability, it must be the case that for type j, $a^j_{\mathsf{x}} = a^j_{\mathsf{y}}$. In general, we have[5]

$$a^j_{\mathsf{x}} = a^j_{\mathsf{y}} \text{ implies } p_j \in [0,1]. \tag{2.1}$$

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

[5]Equations (2.1)–(2.3) represent the so-called *incentive compatibility* conditions in Forges [16].

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

$$a^j_{\mathsf{x}} \leqslant a^j_{\mathsf{y}} \text{ implies } p_j = 0 \text{ and} \tag{2.2}$$
$$a^j_{\mathsf{x}} \geqslant a^j_{\mathsf{y}} \text{ implies } p_j = 1. \tag{2.3}$$

In the case that $a^j_{\mathsf{x}} = a^j_{\mathsf{y}}$, 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, $a^j_{\mathsf{x}} \leqslant a^j_{\mathsf{y}}$. 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 §2.4.2 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 [8] illustrates this idea.

Example 2.4.2 *Compromising.*

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

Figure 2.2: Illustrating a Compromise

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,* $\{\mathsf{a},\mathsf{b}\}$. *Each player randomises uniformly over the two messages. If the messages sent are* $\{\mathsf{a},\mathsf{a}\}$ *or* $\{\mathsf{b},\mathsf{b}\}$, *then the Sender will reveal his type. If the outcome is* $\{\mathsf{a},\mathsf{b}\}$ *or* $\{\mathsf{b},\mathsf{a}\}$, *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.* ✠

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 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.

2.4.2 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 $\mathrm{co}(A)$, 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-dimensional[6] 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 $\mathbb{Z} := \mathbb{A} \times \mathbb{B} \times \mathbb{Q}$.[7] Recall that a *martingale* is a sequence of $\mathbb{Z}$-valued random variables $(\mathbf{z}_t)_0^\infty$ and Borel fields $(\mathscr{F}_t)_0^\infty$ such that for each $t = 0, 1, 2, \ldots$,

1. $\mathscr{F}_t \subset \mathscr{F}_{t+1}$ and $\mathbf{z}_t \in \mathscr{F}_t$;

2. $\mathsf{E}(|\mathbf{z}_t|) < \infty$;

3. $\mathbf{z}_t = \mathsf{E}(\mathbf{z}_{t+1}|\mathscr{F}_t)$, 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 $\mathbf{z}_t = (\mathbf{a}_t, \boldsymbol{\beta}_t, \mathbf{q}_t)$ such that $\mathbf{a}_{t+1} = \mathbf{a}_t$ when t is even, $\mathbf{q}_{t+1} = \mathbf{q}_t$ when t is odd and $(\mathbf{a}_0, \boldsymbol{\beta}_0, \mathbf{q}_0)$ is a constant[8] (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 §2.4.1 in the discussion of the incentive compatibility conditions (equations (2.1)–(2.3)), 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.

[6] k-dimensional because there are k types of the Sender.

[7] To 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 $[-2, 2]^k \times [-2, 2]$.

[8] This makes $\mathbf{z}_0$ the expectation of the di-martingale.

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 action[9] $y \in \Delta(C_R)$. A (mixed) strategy y is an equilibrium of the silent game $\Gamma(p)$ if

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

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

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, $\mathscr{E}(p)$ (cf. the discussion following equations (2.1)–(2.3)). These are precisely the virtual payoffs assigned to types which occur with probability 0. With the additional elements, we get $\mathscr{E}^+(p)$, the *modified equilibrium payoffs correspondence.* (The reasons for adding these additional points are discussed in greater detail in §2.4.3 below.) Once again, we will define $\mathscr{E}^+(p)$ formally. For each $p \in \Delta^{k-1}$, define $\mathscr{E}^+(p)$ as the set of all $(a, \beta) \in \mathbb{A} \times \mathbb{B}$ such that there exists $y \in \Delta(C_R)$ satisfying

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

The graph of the modified equilibrium correspondence is

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

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 $\operatorname{gr}\mathscr{E}^+$, 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 $(\Omega, \mathscr{F}, \mathsf{P})$ be a non-atomic probability space and let $\mathbb{N} := \{0, 1, 2, \ldots\}$ be the set of natural numbers. Also, let $(\mathbf{z}_n)_{n \in \mathbb{N}}$ be a sequence of random variables (taking values in a Euclidean space) and let $(\mathscr{F}_n)_{n \in \mathbb{N}}$ be a sequence of finite fields such that such that $(\mathbf{z}_n)$ is a martingale[10]with respect to $(\mathscr{F}_n)$.[11] Let $\mathbb{N}_\infty := \mathbb{N} \cup \{\infty\}$ and

[9]For 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.

[10]Standard 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.

[11]These ideas are described lucidly in §9.3 of [11], for example.

adjoin $\mathscr{F}_\infty := \bigvee_{n\in\mathbb{N}} \mathscr{F}_n$ (the minimal field containing $\mathscr{F}_n$ for each $n \in \mathbb{N}$) to $(\mathscr{F}_n)$. A random variable α is a *stopping time* if for every $n \in \mathbb{N}_\infty$, $\{\alpha = n\} \in \mathscr{F}_n$. A stopping time is *a.s. finite* if $\mathsf{P}(\alpha < \infty) = 1$. It is *a.s. bounded* if there exists $n_0 < \infty$ such that $\mathsf{P}(\alpha \leqslant n_0) = 1$.

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 lemma[12]) 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 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 $\operatorname{gr}\mathscr{E}^+$ as these are the only ones with the "correct" payoffs for the two players because $\operatorname{gr}\mathscr{E}^+$ represents the payoff the players can get in the action phase of the game. Now let $G := \operatorname{gr}\mathscr{E}^+$. 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*[13] in [2]. Nevertheless, they require proof and this is the substantial achievement recorded in [2] and [3].

2.4.3 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, $\mathscr{E}(p)$.[14]

[12]König's Lemma. *Every finitely generated tree with infinitely many points must contain at least one infinite branch.* (For a proof, see [42].)

[13]A *bi-martingale* is a di-martingale without the β-coordinate.

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

Consider the equilibrium payoff correspondence of the silent game, gr $\mathscr{E}$. 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, gr $\mathscr{E}^{+}$. In a signalling stage, the Receiver "promises" different payoffs to different types (cf. equations (2.1)–(2.3) 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 $\mathsf{M} = \{\mathsf{a}, \mathsf{b}\}$. 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 (2.1)). 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 that this is the situation considered in equations (2.2) and (2.3).) 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 $(\mathbf{a}, \boldsymbol{\beta}, \mathbf{p})$ 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 [23]. 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.

2.5 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 $(\Omega, \mathscr{F}, \mathsf{P})$ 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 $(\mathscr{F}_t)_0^\infty$ of subfields of $\mathscr{F}$, such that $(\mathbf{z}_t)$ is a martingale with respect to $(\mathscr{F}_t)$.
2. For each $t = 0, 1, \ldots$, 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 a $\mathbb{A}\times\mathbb{B}\times\mathbb{Q}$ 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.

Definition 2.5.1 *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 $A_a := \{q \in \mathbb{Q} \ : \ (a, q) \in \mathbb{A} \times \mathbb{Q}\}$ and $A_q := \{a \in \mathbb{A} \ : \ (a, q) \in \mathbb{A} \times \mathbb{Q}\}$ 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.

Example 2.5.2 *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 2.3 below, A is bi-convex but not convex.*

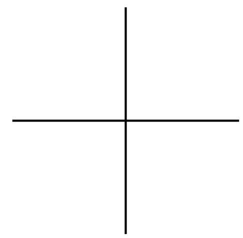

Figure 2.3: A bi-convex set that is not convex.

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

Definition 2.5.3 bi-co(A) *is the smallest bi-convex set containing* A.

Definition 2.5.4 $A^{\#} := \{z \in \mathbb{A} \times \mathbb{Q} : \exists$ *bi-martingale* $(\mathbf{z}_t)$ *with an a.s. finite stopping time* N *such that* $\mathbf{z}_N \in A$ $(a.s.)$ *and* $\mathbf{z}_0 = z$ $(a.s.)\}$.

Definition 2.5.5 $A^* := \{z \in \mathbb{A} \times \mathbb{Q} : \exists$ *bi-martingale* $(\mathbf{z}_t)$ *converging to* $\mathbf{z}_\infty$ *such that* $\mathbf{z}_\infty \in A$ $(a.s.)$ *and* $\mathbf{z}_0 = z$ $(a.s.)\}$.

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 $\text{proj}_{\mathbb{A}\times\mathbb{Q}}\, \mathscr{E}^+(q)$ (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, 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 [36]). The following example from [2] illustrates this.

Example 2.5.6 *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 2.4 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 2.8 below.*

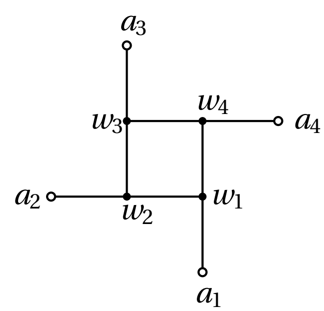

Figure 2.4: Example where $A = \text{bi-co}(A) \subsetneq A^*$

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.

Definition 2.5.7 *A* ***bi-convex combination*** *is a convex combination* $(a, q) = \sum_i \alpha_i (a_i, q_i)$ *(with* $\alpha_i \geqslant 0$ *and* $\sum_i \alpha_i = 1$*) where either* $a_1 = \cdots = a_n = a$ *or* $q_1 = \cdots = q_n = q$*.*

Definition 2.5.8 *A set* $A \subset \mathbb{A} \times \mathbb{Q}$ *is* ***generated*** *by* $A_0 \subset A$ *if for all* $a \in A$*, there exist* $\alpha_1, \ldots, \alpha_n$ *(with* $\alpha_i \geqslant 0$ *and* $\sum_i \alpha_i = 1$*) and* $a_1, \ldots, a_n \in A_0$ *such that* $a = \sum_1^n \alpha_i a_i$ *is a bi-convex combination.*

Definition 2.5.9 *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, $\text{proj}_{\mathbb{A}\times\mathbb{Q}}\, \mathscr{E}^+(q)$ is finitely generated. This is made precise in Lemma 2.5.11. Before proving the lemma, let us see an illustration of this for a game with two types.

Example 2.5.10 *Consider the following game (figure 2.5) due to Simon [39]. 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,* $\alpha, b, \ldots, g$*.*
The graph of the Receiver's payoffs are depicted in figure 2.6. 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.

		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 2.5: Example 2.5.10

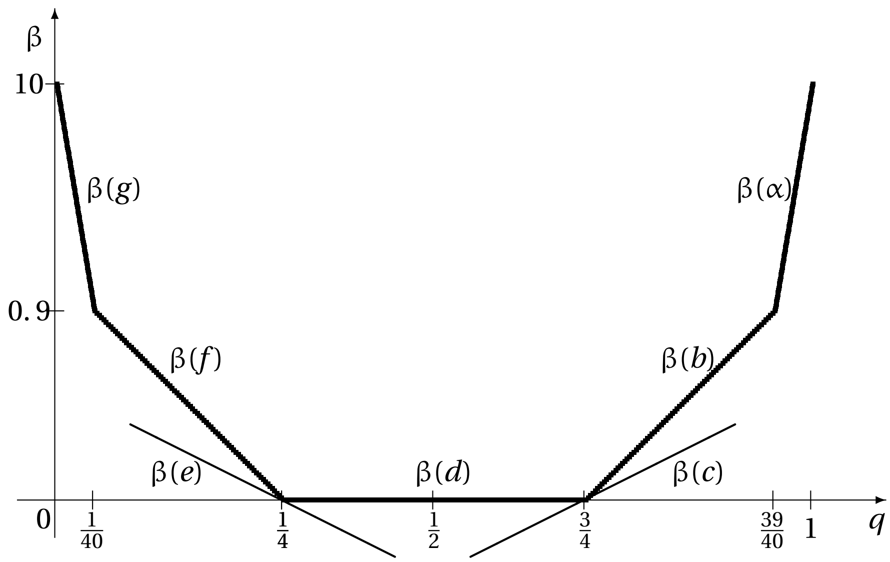

Figure 2.6: Receiver's payoffs given by upper envelope

Most interesting are the Sender's payoffs, for this set is finitely generated in the sense defined above. This is depicted in figure 2.7. 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 2.6). In figure 2.7, 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 Receier). 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 $q = \frac{1}{40}$.
✠

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

Lemma 2.5.11 *G is finitely generated.*

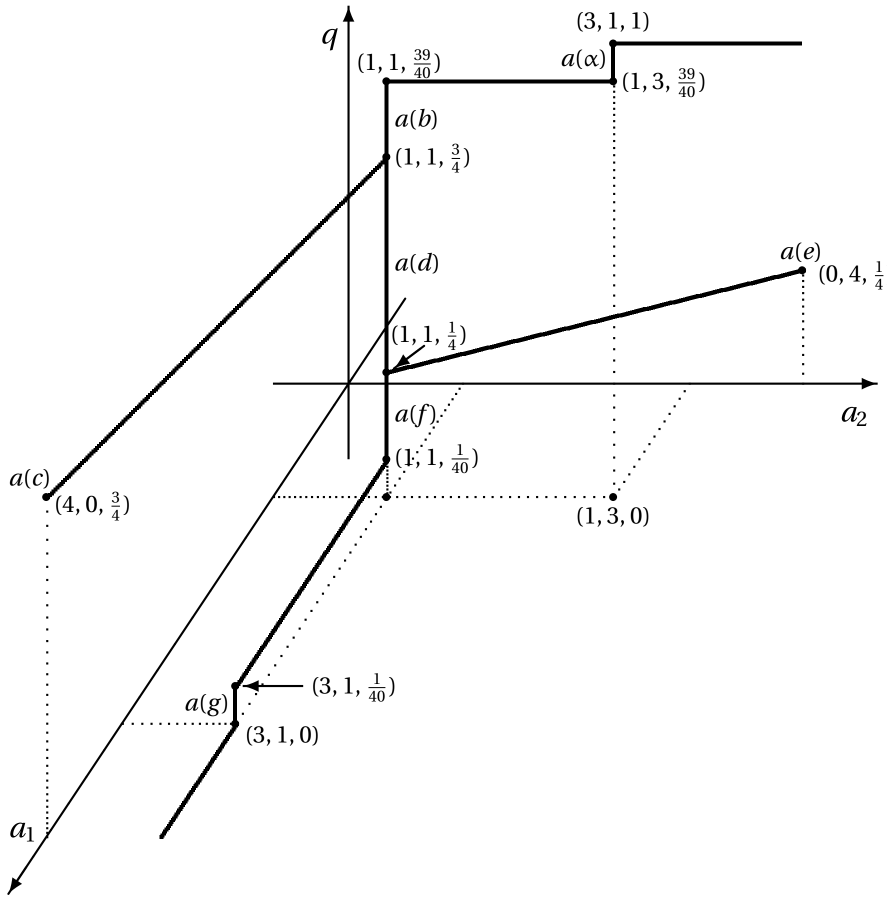

Figure 2.7: Graph of $\text{proj}_{\mathbb{A}\times\mathbb{Q}}\,\mathscr{E}^+(q)$ — commonly known as the Sender's modified payoffs

Proof. Let the Receiver's actions be $c_1, c_2, \ldots, c_n$ and β_i the expected payoff from action c_i. (Note that for each i, $\beta_i : \Delta^{k-1} \to \mathbb{B}$ 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 $F_i := \{q \in \Delta^{k-1} : \beta_i(q) > \beta_j(q) \quad \text{for } j \neq i\}$ which gives us the region of the probability space where the Receiver's best action is c_i. More generally, we can define $F_{i_1,\ldots,i_m} := \{q \in \Delta^{k-1} : \beta_{i_1}(q) = \cdots = \beta_{i_m}(q) > \beta_j(q) \quad \text{for } j \notin \{i_1, i_2, \ldots, i_m\}\}$ which gives us the region of the probability space where the Receiver's best actions are $c_{i_1}, c_{i_2}, \ldots, c_{i_m}$ and the Receiver is indifferent between these actions. Note that we do not require $F_{i_1,\ldots,i_m}$ 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, $F_{i_1,\ldots,i_m}$ 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 $F_{i_1,\ldots,i_m}$ can be written as a convex combination of these points). (This is a corollary of Theorem 19.1 in [36].)

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 $F_{i_1,\ldots,i_m}$. The Receiver is indifferent between actions $c_{i_1}, \ldots, c_{i_m}$ in this region. Thus, the (vectors of) Sender's payoffs are all convex combinations of $a(c_{i_1}), \ldots, a(c_{i_m})$ which is a finitely generated set. But this is 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 $\mathrm{proj}_{\mathbb{A}\times\mathbb{Q}}\,\mathrm{gr}\,\mathscr{E}$ 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 $\mathrm{proj}_{\mathbb{A}\times\mathbb{Q}}\,\mathrm{gr}\,\mathscr{E}^+$ is also finitely generated. □

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 $\{v_0, \ldots, v_q\}$ is an affine independent subset of some Euclidean space, then it spans the q-simplex $s := [v_0, \ldots, v_q] = \mathrm{co}(\{v_0, \ldots, v_q\})$. The **vertex set** of s is denoted by $\mathrm{Vert}(s) := \{v_0, \ldots, v_q\}$. 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 $\mathrm{Vert}(s \times t) := \mathrm{Vert}(s) \times \mathrm{Vert}(t)$. If σ is a bi-simplex, then a **face** of σ is a bi-simplex σ' with $\mathrm{Vert}(\sigma') \subset \mathrm{Vert}(\sigma)$.[15]

Definition 2.5.12 *A finite* ***bi-simplicial complex*** *K is a finite collection of bi-simplices in some product of Euclidean spaces such that:*
(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. Let G_0 be the finite set generating the modified graph of the payoffs, $\mathrm{gr}\,\mathscr{E}^+$ and let $W_0 := \mathrm{proj}_{\mathbb{A}}\, G_0 \times \mathrm{proj}_{\mathbb{Q}}\, G_0$.

> **Condition F [Finiteness Condition].** *For all $A \in 2^{W_0}$, $\text{bi-co}(A) = K$, where K is a bi-simplicial complex.*

Definition 2.5.13 *A Sender-Receiver game satisfies Condition F if the generating set of the graph of the Sender's payoffs (i.e. W_0 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.

Lemma 2.5.14 *Let $A \subset \mathbb{A}\times\mathbb{Q}$ be generated by a finite set A_0 and suppose W_0 satisfies Condition F. Then $A^* = A^\#$.*

[15]These definitions are the obvious analogues of standard notions in algebraic topology which can be found, for instance, in [10].

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 is 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 2.8 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 2.8).

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*[16] (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 $G = \mathrm{proj}_{\mathbb{A}\times\mathbb{Q}}\, \mathrm{gr}\, \mathscr{E}^+$ is the graph of the Sender's payoffs. In fact, we shall prove that $G^{\#} = G^*$. *Proof.* [Proof of Theorem A] From Lemma 2.5.11, we know that the modified graph of the Sender's payoffs is

[16]Given 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.

finitely generated. Therefore from Lemma 2.5.14, 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. □

Note that the proof of Lemma 2.5.14 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.

Example 2.5.15 *Consider once again the game introduced in Example 2.5.10. Below, in figure 2.8 is a picture of the section of the graph of the Receiver's payoffs with* $a_1 + a_2 = 4$. *(The coordinates in figure 2.8 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 2.7 that the graph of the Receiver's payoffs is biconvex, 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.

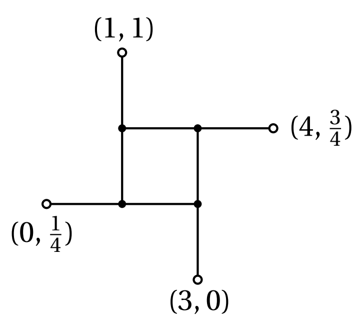

Figure 2.8: $a_1 + a_2 = 4$ section of graph of Receiver's payoffs

2.6 (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 §2.4.2, we let $G = \mathrm{proj}_{\mathbb{A}\times\mathbb{Q}}\, \mathrm{gr}\, \mathscr{E}^{+}$ be the graph of the Sender's (modified) payoffs and say that G is generated by G_0. We first consider finite canonical[17] 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 $\bigcup_{n\in\mathbb{N}} G_n =: \text{bi-co}(G_0) = \text{bi-co}(G)$.[18] We are now able to state our result on finite conversations.

Theorem 2.6.1 *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.* [Proof of Theorem 2.6.1] It will suffice to prove that there exists an N so that $G_N = \text{bi-co}(G_0)$. We know from Lemma 2.5.11 that G is finitely generated by G_0. By Condition F, it follows that 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 bi-co(G_0) is the smallest bi-convex set which contains G_0) which demonstrates our claim. $\square$

We now come to conversations that are longer that 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.

Theorem 2.6.2 *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 $G^{\#}$.*

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

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

Proof. See the appendix. □

Indeed, we can say even more.

Corollary 2.6.3 *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 2.6.2. □

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 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 2.6.2.

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 2.5.14 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 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.

2.7 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.

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*.[19] To understand the significance of this, consider first a set $A := \{x \in \mathbb{R}^n \ : \ f_1(x) \geqslant 0, \dots, f_\ell(x) \geqslant 0, g_1(x) = \cdots = g_m(x) = 0\}$ 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). The Aumann-Hart notions of 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 [9]) and stable equilibria which are set-valued as defined by Mertens in [29]. This is an attractive and desirable property as semi-algebraic sets are definable in a very precise sense (see [14]) 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

[19] A *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 [14].

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.

2.8 Appendix

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 2.5.14 and Theorem 2.6.2. We conclude the appendix with an axiomatisation of the graph of payoffs of the silent game ($\mathrm{gr}\,\mathscr{E}$) in Sender-Receiver games, thereby removing our problem from the corsettes of game theory and making it purely mathematical.

2.8.1 Bi-convexity and bi-martingales

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

The space $\mathscr{X} \times \mathscr{Y}$ inherits the product topology relative to the Euclidean spaces they live in. Let us denote this as $\mathscr{T}$. Let us denote the *relative* topology of a set $B \subset \mathscr{X} \times \mathscr{Y}$ by $\mathscr{T}_B$. Let us denote the *bi-relative* topology by $\mathscr{T}_{BR}$, where $\mathscr{T}_{BR}$ consists of all sets of the form $\{E \cap U : U \in \mathscr{T}\}$ and $E := \text{aff}(\text{proj}_{\mathscr{X}} B) \times \text{aff}(\text{proj}_{\mathscr{Y}} B)$. A point $z = (x, y) \in B$ is *bi-relatively interior* to B if z is interior in the topology $\mathscr{T}_{BR}$ (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 $\mathscr{X}$, a neighbourhood V of y, a collection of simplices $s_1, s_2, \ldots, s_m$ in $\mathscr{X}$ and a collection of simplices $t_1, t_2, \ldots, t_n$ in $\mathscr{Y}$ 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.

Proposition 2.8.1 (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.

Example 2.8.2 *Let $\mathscr{X} = \mathscr{Y} = [0, 1]$ 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 $\mathscr{X}$ and $\mathscr{Y}$ be compact, convex subsets of Euclidean spaces. Let $(\Omega, \mathscr{F}, \mathrm{P})$ be a non-atomic probability space. A sequence $(\mathbf{z}_t)_0^\infty := (\mathbf{x}_t, \mathbf{y}_t)_0^\infty$ of $\mathscr{X} \times \mathscr{Y}$-valued random variables is a *bi-martingale* if:

1. There exists a non-decreasing sequence $(\mathscr{F}_t)_0^\infty$ of subfields of $\mathscr{F}$, such that $(\mathbf{z}_t)$ is a martingale with respect to $(\mathscr{F}_t)$.

2. For each $t = 0, 1, \ldots$, 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 $\mathscr{X} \times \mathscr{Y}$. We will require the following definitions.

Definition 2.8.3 bi-co(A) *is the smallest bi-convex set containing* A.

Definition 2.8.4 $A^{\#} := \{z \in \mathscr{X} \times \mathscr{Y} : \exists$ *bi-martingale* $(\mathbf{z}_t)$ *with an a.s. finite stopping time* N *such that* $\mathbf{z}_N \in A$ $(a.s.)$ *and* $\mathbf{z}_0 = z$ $(a.s.)\}$.

Definition 2.8.5 $A^{*} := \{z \in \mathscr{X} \times \mathscr{Y} : \exists$ *bi-martingale* $(\mathbf{z}_t)$ *converging to* $\mathbf{z}_\infty$ *such that* $\mathbf{z}_\infty \in A$ *and* $\mathbf{z}_0 = z$ $(a.s.)\}$.

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

We will now define a process of separation. Let $B_0 := \mathscr{X} \times \mathscr{Y}$. Define inductively, $B_{\alpha+1} = \mathrm{ns_A}(B_\alpha)$ for every successor ordinal α and $B_\alpha = \bigcap_{\beta<\alpha} B_\beta$ for every limit ordinal α. This defines a non-increasing sequence[20] 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.

> **Theorem 2.8.6 (4.3 in [2])** *The limit set C satisfies $C = \mathrm{ns_A}(C)$ 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 $\mathrm{nsc}(B)$ $(\equiv \mathrm{nsc_A}(B))$ and a similar inductive process which will give us another limit set D. This gives us

> **Theorem 2.8.7 (4.7 in [2])** *The limit set D satisfies $D = \mathrm{nsc_A}(D)$ and is the largest such set. Also, $D = A^{*}$.*

Note that for each bi-convex set B that contains a set A, both $\mathrm{ns_A}(B)$ and $\mathrm{nsc_A}(B)$ are bi-convex sets. This is because if $z \in \mathrm{ns_A}(B)$ (say) then if must be the case that for all $f : B \to \mathbb{R}$ biconvex, $f(z) \leqslant \sup f(A)$. Now suppose $z, z' \in \mathrm{ns_A}(B)$, then for any $z'' := tz + (1-t)z'$ a bi-convex combination, we have $f(z'') \leqslant \max\{f(z), f(z')\} \leqslant \sup f(A)$ which implies that $z'' \in \mathrm{ns_A}(B)$, that is $\mathrm{ns_A}(B)$ is bi-convex. A similar argument shows that $\mathrm{nsc_A}(B)$ is also bi-convex. To demonstrate these ideas, let us reconsider example 2.5.6 above.

Example 2.8.8 *Let* $A := \{a_1 = (2/3, 0), a_2 = (0, 1/3), a_3 = (1/3, 1), a_4 = (1, 2/3)\}$ *and* $\mathscr{X} = \mathscr{Y} = [0, 1]$. *We claim that* $A^{\#} = A^{*} = \text{bi-co}(\{a_i, w_i\}_{i=1})$, *as illustrated in figure 2.9 below. By Lemma 2.8.11 below,* $\sup f(\{a_i, w_i\}_{i=1}^4) \geqslant \sup f(\text{bi-co}(\{a_i, w_i\}_{i=1}^4))$. *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*

[20]A sequence is a function whose domain is an ordinal (cf. §2.4 of [27]).

function f. Let us assume that w_1 is such that $f(w_1) \geqslant f(w_j)$ 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 an A. We can therefore restrict attention to separation by such functions. Thus, $f(w_1) \leqslant \frac{1}{2}[f(a_1) + f(w_2)]$ 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, 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 $A \subset \text{bi-co}(A) \subset A^\# \subset A^*$.

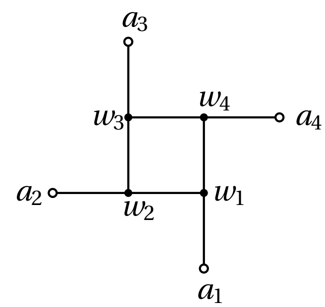

Figure 2.9: $A = \{a_i\}_1^4 = \text{bi-co}(A) \subsetneq A^*$

We will begin with some definitions of our own.

Definition 2.8.9 *A set $A \subset \mathscr{X} \times \mathscr{Y}$ is* ***generated*** *by $A_0 \subset A$ if for all $z \in A$, there exist $\alpha_1, \ldots, \alpha_n$ (with $\alpha_i \geqslant 0$ and $\sum_i \alpha_i = 1$) and $a_1, \ldots, a_n \in A_0$ such that $z = \sum_1^n \alpha_i a_i$ is a bi-convex combination.*[21]

Definition 2.8.10 *A set A is* ***finitely generated*** *if there exists a finite set A_0 which generates it.*

We shall denote the set of all bi-convex combinations of a set $A_0 \subset \mathscr{X} \times \mathscr{Y}$ 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.

Lemma 2.8.11 *Let $A \subset \mathscr{X} \times \mathscr{Y}$ and let $B = \text{bi-co}(A)$. Then for any bi-convex function $f : B \to \mathbb{R}$, $\sup f(B) = \sup f(A)$.*

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 $\sup f(A_0) \geqslant$

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

$\sum \alpha_i f(a_i) \geqslant f\left(\sum \alpha_i a_i\right) = f(a)$, where the second-to-last inequality is because f is bi-convex. Similarly, $\sup f(A_i) \geqslant \sup f(A_{i+1})$. Therefore, we get $\sup f(A_0) \geqslant \sup f(A_1) \geqslant \cdots \geqslant \sup f(A_i) \geqslant \cdots$. But (A_i) is an increasing sequence of sets with limit $B = \text{bi-co}(A_0)$, which implies that $\sup f(A_0) \geqslant \sup f(B)$. □

The Lemma above is equivalent to the observation that for any set A with $B = \text{bi-co}(A)$, $B = \text{ns}_{\text{A}}(B)$. 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.

Lemma 2.8.12 *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 $f : A^{\#} \to \mathbb{R}$ such that $f(z) > \sup f(A_0)$. (If such a point and corresponding function did not exist, then $A^{\#} = \text{ns}_{\text{A}_0}(A^{\#})$ which, by Theorem 2.8.6, 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 $f(z) > \sup f(A_0) \geqslant \sum \alpha_i f(a_i) \geqslant f\left(\sum \alpha_i a_i\right) = f(a)$, where the second-last inequality is because f is bi-convex. In other words, $f(z) > \sup f(A)$ which contradicts $A^{\#} = \text{ns}_{\text{A}}(A^{\#})$ which implies that such a z cannot exist. Thus, $A^{\#} \setminus A_0^{\#} = \varnothing$, i.e. $A^{\#} = A_0^{\#}$. □

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

Definition 2.8.13 *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 $y_1 = y_2$.*

Lemma 2.8.14 *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 2.8.14 and the Krein–Milman Theorem can be found in the Appendix to Chapter IV in [25].

> **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.

Definition 2.8.15 *Let $B \subset \mathscr{X} \times \mathscr{Y}$ 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.

Lemma 2.8.16 *Let $S \in \mathbb{R}^n$ be compact. Then*
(*i*) $\mathrm{co}(S)$ *has an extreme point; and*
(*ii*) *if $x \in \mathrm{co}(S)$ is an extreme point, then $x \in S$.*

Proof. (i) Since S is compact, it follows that $\mathrm{co}(S)$ is compact. By Lemma 2.8.14, $\mathrm{co}(S)$ has an extreme point.
(ii) Now $\mathrm{co}(S) = S \sqcup (\mathrm{co}(S) \setminus S)$, a disjoint union. Let $x \in \mathrm{co}(S)$ be an extreme point of $\mathrm{co}(S)$. If $x \in \mathrm{co}(S) \setminus S$, then x is a convex combination of points in S, contradicting its being an extreme point. Therefore, if x is an extreme point of $\mathrm{co}(S)$, $x \in S$. □

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

Lemma 2.8.17 *Let $A \subset \mathscr{X} \times \mathscr{Y}$ be compact. Then* $\text{bi-co}(A)$ *has a bi-extreme point in A.*

Proof. Recall that $\mathscr{X} \subset \mathbb{R}^m$ and $\mathscr{Y} \subset \mathbb{R}^n$. Let $\mathrm{co}(A)$ be the convex hull of A in $\mathbb{R}^{m+n}$. Then $\mathrm{co}(A)$ has an extreme point, by Lemma 2.8.14 above and if z is an extreme point of $\mathrm{co}(A)$, then $z \in A$. But this means that z cannot be written as a convex combination of elements in bi-co(A) (which lies in $\mathrm{co}(A)$) which implies that it cannot be written as a bi-convex combination of elements in bi-co(A) (as every bi-convex combination is also a convex combination). Thus, z is a bi-extreme point of bi-co(A). □

2.8.2 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.[22] 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 2.5.6.

Example 2.8.18 *Let* $A := \{a_1 = (2/3, 0), a_2 = (0, 1/3), a_3 = (1/3, 1), a_4 = (1, 2/3)\}$ *and* $\mathscr{X} = \mathscr{Y} = [0, 1]$. *Consider the bi-convex hull of* $\{a_i, w_i\}_{i=1}^4$*, as illustrated in figure 2.9. 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*[23] convex set is finitely generated, i.e. is the convex hull of finitely many points (cf. Theorem 19.1 in [36]). We may call a set *semi-bi-linear* if it is defined by finitely many bi-linear functions.[24] Such a bi-convex set need not be finitely generated as demonstrated in the example below.

Example 2.8.19 *Define* $A := \{(x, y) \in \mathbb{R}_+^2 : xy \geqslant 1, \quad (x-4)(y-4) \geqslant 1, \quad x \leqslant 4 \;\text{ and }\; y \leqslant 4\}$. *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.*

2.8.3 Proof of Lemma 2.5.14

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* $\text{bi-co}(E_0) = \langle E_0 \rangle_b^n$.

Lemma 2.8.20 *Let* $A_0 \subset \mathscr{X} \times \mathscr{Y}$ *be a finite set. Then* A_0 *satisfies Condition F iff* A_0 *satisfies Condition* H.

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

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

[24] A bi-linear function, also known as a *bi-affine* function is a function $f : B \to \mathbb{R}$ (where $B \subset \mathscr{X} \times \mathscr{Y}$ 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.

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 2.8.6 and 2.8.7 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.

Lemma 2.8.21 *Let $A_0 \subset \mathscr{X} \times \mathscr{Y}$ be finite, $A_0 \subset W_0$ also finite and $B = \text{bi-co}(W_0)$ and suppose $f : B \to \mathbb{R}$ 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, $\widehat{w} \in W_0$ such that $f(\widehat{w}) > \sup f(A_0)$.*

Proof. It follows from Lemma 2.8.17 that B has a bi-extreme point and that any bi-extreme point of B must lie in W_0. Note that if $f, g : B \to \mathbb{R}$ are bi-convex functions,[25] 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 2.8.11 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 $\widehat{w} \in W_0$ such that $f(\widehat{w}) = \max f(W_0) \geqslant f(z) > 0$. We shall now show that one can take such a $\widehat{w}$ 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 $\widehat{w}$ is not a bi-extreme point. Then we can write $\widehat{w}$ as a (non-trivial) bi-convex combination of points that are extreme in one of the (x- or y-) sections of $\widehat{w}$. Let us replace $\widehat{w}$ 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 process stops in finite time with some bi-extreme point taking the maximum value under f. □

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

Definition 2.8.22

1. *A set $D \subset \mathscr{X} \times \mathscr{Y} \times \mathbb{R}$ is **di-convex** if for each $x \in \mathscr{X}$ and each $y \in \mathscr{Y}$, the respective x- and y-sections $D_x := \{(y, \mu) \in \mathscr{Y} \times \mathbb{R} \ : \ (x, y, \mu) \in \mathscr{X} \times \mathscr{Y} \times \mathbb{R}\}$ and $D_y := \{(x, \mu) \in \mathscr{X} \times \mathbb{R} \ : \ (x, y, \mu) \in \mathscr{X} \times \mathscr{Y} \times \mathbb{R}\}$ are convex.*

[25] We shall assume that the domain of a bi-convex function is a bi-convex set.

2. *A convex combination* $(x, y, \mu) = \sum_i \alpha_i(x_i, y_i, \mu_i)$ *(with* $\alpha_i \geqslant 0$ *and* $\sum_i \alpha_i = 1$*) will be called* ***di-convex*** *if either* $x_1 = \cdots = x_n = x$ *or* $y_1 = \cdots = y_n = y$.

3. *The* ***di-convex hull*** *of a set A is the smallest di-convex set containing A.*

Lemma 2.8.23 *Let* $W_0 \subset \mathscr{X} \times \mathscr{Y}$ *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* $g : B \to \mathbb{R}$ *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 $D_0, \widetilde{D}_0 \subset \mathscr{X} \times \mathscr{Y} \times \mathbb{R}$ 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, $W_i := \langle W_0 \rangle_b^i = \text{proj}_{\mathscr{X} \times \mathscr{Y}} D_i$ and $\widetilde{W_i} := \langle \widetilde{W_0} \rangle_b^i = \text{proj}_{\mathscr{X} \times \mathscr{Y}} \widetilde{D}_i$. 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 $\left(z, \sum \alpha_j \min\{\mu : (z_j, \mu) \in \widetilde{D}_i\}\right) \in D_{i+1} \setminus \widetilde{D}_{i+1}$ 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

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

As each section of a di-convex set is convex, $g(z)$ is convex on each section (cf. Theorem 5.3 in [36]). Also, $g(z)$ is as required, by construction. □

We are now able to take our final steps.

Lemma 2.8.24 *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 $B \neq \mathrm{ns}_{\mathrm{A}_0}(B)$. (As mentioned above, we can take all functions to take the value 0 on A_0.) Then there exists $z \in B$ and $f : B \to \mathbb{R}$ bi-convex such that $f(z) > 0$. But by Lemma 2.8.21, there also exists a bi-extreme point of B, $\widehat{w} \in W_0$ such that $f(\widehat{w}) > 0$. Moreover, there exists a continuous function $g : B \to \mathbb{R}$ such that $g(z) > 0$ for all $z \in$ bi-co$(W_0) \setminus$ bi-co$(W_0 \setminus \{\widehat{w}\})$. □

Lemma 2.8.25 (Lemma 2.5.14 in the text) *Let $A \subset \mathscr{X} \times \mathscr{Y}$ 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 $B =$ bi-co(W_0) then $\mathrm{ns}_{\mathrm{A}_0}(B) \neq B$ implies that $\mathrm{nsc}_{\mathrm{A}_0}(B) \neq B$. 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 $(\mathrm{ns}_{\mathrm{A}_0}(B) \neq B) \to (\mathrm{nsc}_{\mathrm{A}}(B) \neq B)$. But this is not enough, as we have to show that the set A^* (and consequently $A^\#$) is obtained in this way. Let $W_0 : \{(x, y) \ : \ x \in \mathrm{proj}_{\mathscr{X}} A_0, y \in \mathrm{proj}_{\mathscr{Y}} A_0\}$ and let $W_{i+1} := \{w \in W_i \ : \ w$ is not a bi-extreme point of bi-co$(W_i)\}$. This process ends in finite time at W and $A^* = A^\# =$ bi-co(W). □

2.8.4 Proof of Theorem 2.6.2

Theorem 2.8.26 (Theorem 2.6.2 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 $A \subset \mathscr{X} \times \mathscr{Y}$ be a finite set and let $W := \mathrm{proj}_{\mathscr{X}} A \times \mathrm{proj}_{\mathscr{Y}} A$. Then, for all $C \subset W$ such that $C \supset A$, $\mathrm{proj}_{\mathscr{X}} C \times \mathrm{proj}_{\mathscr{Y}} C = W$. This is because for all $c = (c_x, c_y) \in C$, $c_x \in \mathrm{proj}_{\mathscr{X}} W$ and $c_y \in \mathrm{proj}_{\mathscr{Y}} W$.

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 $W_0 := \mathrm{proj}_{\mathbb{A}} G_0 \times \mathrm{proj}_{\mathbb{Q}} G_0$. Since Γ satisfies Condition F, it is the case that $G^\# = G^*$. Also, from the proof of lemma 2.5.14, we know that $G^* \subset$ bi-co(W_0) and that G^* is finitely generated by some $W' \subset W_0$ with $W' \supset G_0$. Therefore, $\mathrm{proj}_{\mathbb{A}} W' \times \mathrm{proj}_{\mathbb{Q}} W' = W_0$. But recall that W_0 satisfies Condition H, so that by the same argument as in the proof of lemma 2.5.14, 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 2.5.14 and the second equality is from page 178 of [2]. □

2.8.5 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 $a : K_n \to \mathbb{R}^n$ be a map associating an n-tuple of real numbers with each n-simplex and define

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

Lemma 2.8.27 *There exists a Sender-Receiver game Γ with graph of the payoff correspondence of the silent game* $\mathrm{gr}\,\mathscr{E}$ *such that* $\mathrm{proj}_{\mathbb{A}\times\mathbb{Q}}\,\mathscr{E}$ *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 $f : \mathbb{R}^n \to \mathbb{R}$ be given by $f(x) := \|x - \alpha\|^2$ where $\alpha \in \mathbb{R}^n$ is such that for all $v, w \in \mathrm{Vert}(K)$, $f(v) \neq f(w)$. 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 $F := \{(v, f(v)) : v \in \mathrm{Vert}(K)\}$ and let

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

It follows from Theorem 5.3 in [36] that g is a convex function on Δ^{n-1}. Moreover, it is polyhedral and for all $v \in \mathrm{Vert}(K)$, $f(v) = g(v)$. Also, every n-simplex in K is the projection of some lower face of $\mathrm{co}(F)$, which ensures that g induces the required simplicial decomposition. For each n-simplex σ, $g|_\sigma$ is a linear function which can be extend to the entire ambient space and denote the resulting linear function by β_{c_σ}. 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. $\square$

Chapter 3

Communication in Games of Incomplete Information: The Two Player Case

3.1 Introduction

Since the seminal paper of Crawford and Sobel [12], it has been well understood that communication in games of incomplete information can expand the set of equilibrium outcomes that can be achieved and provide the players with Pareto improving outcomes. In this paper, we examine the question of how much the set of equilibrium outcomes can be expanded via cheap communication procedures in two-person games of incomplete information.

The question is obviously moot if the players can write enforceable contracts wherein they report their private information to each other and subsequently take an action. If such incentive compatible contracts could be written, there would be nothing else left to do. We will refer to such situations as consisting of *perfect commitment.* The set of such contracts exhaust all the achievable outcomes. But in many economic situations of interest, this is not the case—players have unverifiable private information, there is no enforcement agency to which the players can appeal and they cannot commit to any course of action. Indeed, the Crawford-Sobel results are of such interest precisely because even though the players cannot commit to telling the truth, there is some sharing of information between a Sender (who has the private information) and a Receiver (who takes the only payoff relevant action).

In Bayesian games (i.e. games of incomplete information) there is another construct which is extremely useful. This is the notion of the mediated communication mechanism. The idea is that there is a disinterested party, known as the *mediator* to whom the players report their private information and the mediator then makes private, incentive compatible recommendations. It can be shown (see Myerson [31]) that even in instances where there is no commitment, the set of incentive compatible mediated solutions is exactly the same as the set of incentive compatible enforceable contracts in the case of perfect commitment. This is because although the players cannot commit, it is the disinterested mediator who makes the recommendations and his being disinterested guarantees that the players trust him. The players make honest reports to the mediator because they trust him to make recommendations according to the mediation plan. They then follow his recommendations because it is incentive compatible to do so. Following Myerson (pp. 261-262 in Chapter 6 of [31]), we say that "…incentive-compatible mediation plans [mediation communication mechanisms] are the appropriate generalisation of the concept of correlated equilibria in Bayesian games with communication."

Of course, this still leaves unresolved the case where there is no commitment and no disinterested mediator. In finite games of one-sided incomplete information, it is well known that extended cheap talk can increase the set of equilibrium outcomes. It is also known that the more general conversations that can be considered are conversations without an exogenous deadline. Indeed, there is an example (due to Simon [39]) of a game where finite conversations of any length neither transmit any information nor increase or improve on the babbling outcome. But unbounded conversations produces outcomes that are Pareto superior for both types of the player with the private information and the uninformed player.

It is known that in finite games of incomplete information with three or more players, there is a communication mechanism which implements any rational mediated outcome. In other words, the players can talk amongst themselves and achieve any outcome achievable in the presence of a mediator or if there was perfect commitment. A distinguishing feature of this mechanism is that it satisfies three important desiderata:

C1 The equilibria from the communication game induced by the protocol be perfect Bayesian;

C2 The communication be cheap, non-binding and non-verifiable and

C3 The mechanism be unmediated.

Before we discuss our desiderata, we shall briefly discuss the meaning of the terms used above. By cheap communication, we mean that communication is costless. By non-binding, we mean that by taking part in the mechanism, the players do not lose any strategic options that were available to them as part of the original Bayesian game. By non-verifiability, we mean that there can be no verification in the future of anything that was said in the past. In particular, there is no verifiability of past events. (Note that this strongly resembles our intuitive idea of what constitutes "talk.") Of course, the first requirement seems natural as we would not like our outcomes to depend on incredible threats. Hence, requirements C1 and C2 can be said to constitute a robust mechanism that represents our intuitive idea of cheap talk. But requirements C1 and C2 are not enough. These are satisfied by mediated communication mechanisms. Indeed, in the mechanisms, there is no cost to communicating with the mediator and the players still retain all of their strategic options when they are about to take an action. Thus, requirement C3 is crucial, that there be no other player.

Unfortunately, there is no protocol that satisfies the above three requirements in the two player case. This is made precise below. (The reader will find all the relevant definitions in §3.3.)

Theorem 0. There exists a Bayesian game with a mediated communication mechanism outcome which satisfies a Nash domination condition such that there is no communication mechanism satisfying requirements C1-C3 above which can approximate the mediated communication mechanism outcome.

Remark. It should be mentioned that in the sequel, we shall take all communication mechanisms to be of order type ω. For more on the so-called *transfinite conversations*, the reader is referred to Krishna [24].

Proof. We will demonstrate an example of a generic game with the property that no denumerable communication protocol will simulate the presence of a mediator. Indeed, in doing so, we will also not insist on requirement C1. It will suffice to provide a Sender-Receiver game as an example. (Recall that a Sender-Receiver game is a game where the Sender has some private information and the Receiver takes the only payoff relevant action.) By Theorem A of Aumann and Hart [3], any communication in a Sender-Receiver game can be thought of as a canonical conversation, where a canonical conversation is a conversation where in odd periods, the Sender sends a signal which could potentially give the Receiver some information about the Sender's type and in even periods, the players compromise about future courses of action via joint lotteries. But now consider the example of Forges [16] who gives an example of a Sender-Receiver game where the set of mediated communication mechanism outcomes is strictly larger than the set of outcomes from unbounded canonical conversations (i.e. conversations that have denumerable length). She also demonstrates the existence of a mediated communication mechanism outcome which cannot be approximated by unmediated cheap talk. □

As mentioned in the proof of Theorem 0, we do not get very far by relaxing requirement C1. Nevertheless, as we shall see below, we can do much better if we relax requirement C2. In our Main Theorem below, we show that any rational mediated communication mechanism outcome (see the definition in §3.3) that that satisfies a Nash domination condition can be implemented via cheap communication mechanism that satisfies requirements C1 and C3.

We are not stipulating that there somehow be any sort of commitment—all we are asking is that in the extensive form game which represents an instance of the mechanism, the players be allowed to verify their coordinates in the game-tree. (This will be made clear below, via an example in §3.4.)

From the point of view of applications, this raises some very interesting questions. It has long been thought that the lack of commitment in an economic environment poses some very real problems by reducing the set of achievable outcomes. But we have shown that this is just because all the possible communication mechanisms have not been explored. It should be emphasised that there is enforcement of any kind and the players' participation in the protocol (which induces the extended game) is completely voluntary. They do not have to take part in the protocol, but do so because it is in their best interests to do so. This demonstrates that the problem of lack of commitment can be greatly ameliorated with communication and more so with increasingly sophisticated communication mechanisms.

In fact, one could view our result as the culmination of what can be done without a mediator or equivalently, when there is no commitment. We show that given sufficient latitude in the choice of mechanism, there is no loss of efficiency at all. This implies that in any application where there is limited commitment, the choice of mechanism is not a benign choice. Indeed, we believe it is incumbent on the modeller to explain his choice of mechanism. For instance, consider Bester and Strausz [7] who study the case of limited commitment in principal agent problems. They characterise the incentive efficient equilibria with one round of signalling from the agent. But why should this be the appropriate mechanism to consider? Why not consider multiple rounds of cheap talk or infinitely many rounds of cheap talk, or a more sophisticated mechanism (such as the one we introduce in §§3.4 and 3.5)? We believe that these are important questions that the modeller should not shy away from.

The remainder of this paper is structured as follows. In §3.2 we describe the related literature and how it relates to the present paper. In §3.3 we introduce the model, all the relevant definitions and our Main Theorem. In §3.5 we provide a proof of the Main Theorem. We conclude with a brief discussion in §3.6.

3.2 Related Literature

The usefulness of cheap-talk in Bayesian games was first pointed out by Crawford and Sobel [12]. The idea of mediated solutions is discussed in great detail in Chapter 6 of Myerson [31]. A version of the mediated solution in extensive form games is described by Myerson in [30]. In Chapter 6 of [31], Myerson shows that noisy mechanisms can help agents achieve outcomes that are otherwise not attainable. He also alludes to an equivalence between noisy communication mechanisms in complete information games and noisy mechanisms in Sender-Receiver (and incomplete information) games. Although there are no theorems to this effect, the ideas presented therein have turned out to be immensely valuable. It is this equivalence which enables us to prove our Main Theorem in §3.5. We adapt a mechanism due to Ben-Porath [5] which uses urns and balls to implement correlated equilibria in a large class of bi-matrix games of perfect information to our environment. In the context of equivalence of noisy mechanisms in games of incomplete information and games of perfect information, we should mention a recent paper by Urbano and Vila [43] where it is shown that if the players have bounded computational capabilities, they can achieve any correlated equilibrium in bi-matrix games. Their result is based on the fact that certain algebraic operations (exponentiation and taking logarithms) in prime fields are extremely hard to compute (i.e. cannot be computed in polynomial time).

Barany [4] studies how a communication protocol can implement rational correlated equilibria in games of complete information. He shows that if there are at least four players, then any rational correlated equilibrium can be implemented via a communication protocol. Forges [17] extends this study to Bayesian games. She shows that when there are at least four players, every *ex-ante* mediated outcome can be implemented as a correlated equilibrium of a cheap talk extension of the original game. Now using the result of Barany, it is straightforward to show that every rational mediated outcome can be implemented as a Nash equilibrium. Here, ex-ante means that all the communication occurs before the players learn their types. Gerardi [19] extends this result to the standard interim case, where players communicate after learning their types. Unfortunately, none of these results use procedures that are sequentially rational.

The problem of sequential rationality is addressed by Ben-Porath [6] who also requires only three players. He shows that in any Bayesian game with at least three players, any rational mediated outcome satisfying certain individual rationality constraints can be implemented via a communication protocol which satisfies requirements (1) and (2) above and ends in finite time. (Requirements (1) and (2) were, in fact, adapted from [6].)

We should also mention a result due to Lehrer and Sorin [26]. They point out that in general, mediated solutions need not be stochastic. They introduce a deterministic mediated mechanism which can implement any rational mediated outcome (which, in general, is not deterministic). Their result applies to both complete information games (wherein they implement correlated equilibria and Bayesian games. It should be noted that the mechanism we use in our example and in the proof of our Main Theorem in §§3.4 and 3.5 respectively, seems as much a descendant of the Lehrer-Sorin mechanism as of the Ben-Porath mechanism.

There is also a literature which studies outcomes from pure cheap talk alone (i.e. communication which satisfies our requirement (2)). The main paper in this area is Aumann and Hart [3] who characterise the set of all Nash equilibrium outcomes from pure cheap talk in unbounded conversations Bayesian games with one-sided incomplete information. The example in Forges [16] is based on this theory. Krishna [24] studies a special class of such games, namely the Sender-Receiver games. He shows that if a Sender-Receiver game satisfies a certain condition, then, among other things, the set of equilibria from unbounded conversations is the same as the set of equilibria from conversations that are almost surely finite.

Pure cheap talk is used in a variety of applications and we shall only mention two instances. Prime among these is the Crawford-Sobel model. It and it's variants have been the workhorse in a large number of political economy models (see [22]). Another example of cheap talk is Matthews and Postlewaite [28], who show that with preplay communication in a standard k-double auction, the set of equilibrium outcomes is independent of the parameter k.

3.3 The Model and Main Result

A *game of incomplete information* (which we shall also refer to as a *Bayesian game*) is characterised by $\Gamma^b := (N, (C_i)_{i\in N}, (T_i)_{i\in N}, \mathsf{P}, (u_i)_{i\in N})$, where N is the set of players and for each player $i \in N$, the set of possible actions is C_i, the set of possible types is T_i, P is a probability measure on T and i's utility is given by u_i. If we let $C := \times_i C_i$ and $T := \times_i T_i$, then $u_i : C \times T \to \mathbb{R}$ is a von Neumann-Morgenstern utility function. We say that Γ^b is *finite* iff N, C_i and T_i are finite for each $i \in N$. (We will mainly be interested in the case where $N = 2$.)

Each player i first learns his own type and then forms beliefs over the types of the other players. His belief over the types of the other players is denoted by $\mathsf{p}_i : T_i \to \Delta(T_{-i})$ where $\Delta(T_i)$ is the set of probability distributions over T_{-i} and is calculated via Baye's rule so that

$$\mathsf{p}_i(t_{-i}|t_i) := \frac{\mathsf{P}(t)}{\sum\limits_{s_{-i}\in T_{-i}} \mathsf{P}(s_{-i}, t_i)}.$$

Thus, we can also say that beliefs are *consistent* since all the conditional beliefs p_i can be derived from a single probability measure P. Player i's strategy σ_i is a function $\sigma_i : T_i \to \Delta(C_i)$. A profile of strategies $\sigma := (\sigma_i)_{i\in N}$ is a *Bayesian Nash equilibrium* if for each player i, given his type $t_i \in T_i$, $\sigma_i(t_i)$ maximises his expected utility given his information and the other players' strategies. In other words,

$$\sum_{t_{-i}\in T_{-i}} \mathsf{p}(t_{-i}|t_i)\, u_i(\sigma_i(t_i), \sigma_{-i}(t_{-i}), t) \geqslant \sum_{t_{-i}\in T_{-i}} \mathsf{p}(t_{-i}|t_i)\, u_i(\widetilde{\sigma}_i(t_i), \sigma_{-i}(t_{-i}), t)$$

for all strategies $\widetilde{\sigma}_i \in \Delta(C_i)^{T_i}$.[1]

Now suppose there is a disinterested mediator[2] with whom the players can communicate. This gives us a Bayesian game with communication wherein the players communicate after they learn their types but before they choose an action. Let each player report his type confidentially to the mediator. The mediator, upon receiving messages from all the players, makes private probabilistic recommendations of actions to take to each player. Specifically, player i sends a message $m_i \in T_i$ to the mediator. Then the mediator makes a recommendation according to $\mu : T \to \Delta(C)$. Thus, a strategy for player i is a pair (m_i, δ_i) where m_i is the type report that player i sends to the mediator and $\delta_i : C_i \to \Delta(C_i)$ is such that player i, upon receiving the recommendation c_i, takes the action $\delta_i(c_i)$. We will require that $\sum_{\tilde{c} \in C} \mu(\tilde{c}|t) = 1$ and $\mu(c|t) \geqslant 0$ for all $c \in C$ and $t \in T$. Indeed, any such function $\mu : T \to \Delta(C)$ will be referred to as a *mediated communication mechanism.* Recall that such a mechanism still satisfies C1 and C2. We shall call a mediated communication mechanism *rational* if $\mu(c|t)$ is rational for each $c \in C$ and for every $t \in T$.

Let us consider the instance where each player reports his type honestly to the mediator and obeys the recommendation of the mediator. Then, player i's expected utility from the mediation mechanism μ is

$$U_i(\mu|t) := \sum_{t_{-i} \in T_{-i}} \sum_{c \in C} \mathsf{p}_i(t_{-i}|t_i)\mu(c|t)u_i(c, t).$$

It goes without saying that we are restricting attention to truth-telling strategies in the reporting stage precisely because of the Revelation Principle (see, for instance, [31]). The mediation plan μ therefore induces a communication game Γ^b_μ wherein each player first chooses a reporting strategy $m_i \neq t_i$ and a choice of action $\delta_i(c_i)$ upon receipt of the recommendation c_i. Thus a strategy is a pair (m_i, δ_i). The type sets are the same as in Γ^b and the utility functions in Γ^b_μ are derived in the obvious way. Now if player i were to use strategy (m_i, δ_i), his expected utility from the mechanism μ is given by

$$U_i(\mu, \delta_i, m_i|t_i) := \sum_{t_{-i} \in T_{-i}} \sum_{c \in C} \mathsf{p}_i(t_{-i}|t_i)\mu(c|t_{-i}, m_i)u_i(c_{-i}, \delta_i(c_i), t)$$

if all the other players remain honest and obedient. We will say that a mediated communication mechanism μ is *incentive compatible* iff being honest while reporting their type to the mediator and obedient while following the recommendation of the mediator is a Bayesian Nash equilibrium of the game with communication. Thus, μ is incentive compatible if, for all $i \in N$, $t_i \in T_i$, $m_i \in T_i$ (the report of player i) and $\delta_i \in \Delta(C_i)^{C_i}$ (the choice of action for

[1] If X and Y are sets, Y^X denotes the set of all mappings from X to Y.

[2] The discussion on mediated mechanisms is based on Chapter 6 of Myerson [31].

player i when the recommendation is c_i), it is the case that

$$U_i(\mu|t) \geqslant U_i(\mu, \delta_i, m_i|t_i).$$

The mediator represents, in an indirect way, commitment among the players. Although they cannot write binding contracts to reveal their information and act according to the resulting prescription (possibly because the information is not verifiable and there is no enforcement mechanism), they can nonetheless achieve the same outcomes through a mediator. Thus, the set of incentive compatible mechanisms represents all the outcomes that could potentially be realised as the equilibrium of any other mechanism or protocol.

Recall that an outcome function is a function $\psi : T \to \Delta(C)$. (Then $\psi(t)$ is a probability measure on C and $\psi(t)(c)$ is the probability of the outcome being c.) Consider now any other cheap communication extension of Γ^b represented by Γ^b_c. Here, Γ^b_c is an extensive form game where players communicate after learning their types. At some endogenously determined point in time, the players simultaneously choose their actions from the original game Γ^b. A strategy profile $\sigma = (\sigma_i)_{i\in N}$ in Γ^b_c induces an outcome function $\psi^\sigma : T \to \Delta(C)$. We will say that an outcome function ψ can be *implemented* via a communication mechanism if there exists a cheap communication extension Γ^b_c and a perfect Bayesian equilibrium σ which induces ψ. We will therefore call ψ (following Ben-Porath [6]) a *communication equilibrium outcome* if for all $i \in N$, $t_i, t'_i \in T_i$ and $\delta \in \Delta(C_i)^{C_i}$, it is the case that

$$\sum_{t_{-i}\in T_{-i}} \sum_{c\in C} \mathsf{p}(t_{-i}|t_i)\psi(t)(c)\,u_i(c,t) \geqslant \sum_{t_{-i}\in T_{-i}} \sum_{c\in C} \mathsf{p}(t_{-i}|t_i)\psi(t'_i, t_{-i})\,u_i(\delta(c_i), c_{-i}, (t_i, t_{-i})).$$

The right hand side of the equation above represent the two kinds of deviations that any player could have, namely by pretending to be another type t'_i or by choosing an action according to $\delta_i(c_i)$ instead of choosing c_i.

As mentioned above, the set of mediated communication mechanism outcomes encompasses everything that is achievable as the equilibrium outcome of some mechanism. We are interested in implementing these mediated outcomes via communication mechanisms that are unmediated. In other words, our goal is to find a mechanism so that for any mediated communication mechanism outcome, there is a perfect Bayesian equilibrium of the game induced by the mechanism that implements the mediated mechanism. We need one more definition before we continue.

Definition 3.3.1 *Let Γ^b be any Bayesian game. A mediated communication mechanism μ is* ***Nash dominating for player*** *m if there exists a Bayesian Nash equilibrium of Γ^b with payoffs α^j_m to type j of player m such that $U_m(\mu|t_j) > \alpha^j_m$ for all $t_j \in T_m$. We shall call such an equilibrium Nash dominated for player m.*

We are now ready to state our Main Theorem.

Main Theorem. Let Γ^b be any Bayesian game with two players and let ψ be an outcome function that is induced by some mediated communication mechanism μ. Suppose μ is Nash dominating for player m. Then, there exists a communication mechanism satisfying C1 and C3 that is cheap and non-binding that implements ψ with probability 1.

Remark. Note that the only criterion we have dropped is that the communication mechanism be non-verifiable.

We will need a few more ideas in order to move towards our end. We will use the idea of a *joint lottery*. A joint lottery simulates a public randomising device and works in a straightforward manner for lotteries of the form $(\lambda, 1-\lambda)$ where $\lambda \in \mathbb{Q} \cap [0,1]$. Consider the case where there are two players and suppose they wish to pursue a particular course of action called outcome 1 with (rational) probability λ and some other outcome, 2, with the complementary probability. To achieve this, we assume that $\lambda = \frac{p}{q}$ where $p, q \in \mathbb{Z}_+$ and are coprime. Now let each player have $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z} = \{0, \ldots, q-1\}$ as a message space. The players simultaneously send messages z_1 and z_2 respectively to each other from the message space with each message being drawn according to the uniform distribution. Then, we will stipulate that outcome 1 obtains if $z_1 + z_2 \pmod{q} \in \mathbb{Z}_p$. Note that $z_1 + z_2 \pmod{q}$ is uniformly distributed over $\mathbb{Z}_q$. It is now easy to see that outcome 1 obtains with probability $\lambda = \frac{p}{q}$ and it is also the case that neither player has an incentive to unilaterally choose the message according to some other distribution (which constitutes a change in strategy). We will also use the coding tools of envelopes with notes in them and larger envelopes which contain these envelopes. (These correspond respectively to the balls with messages in them and the urns containing the balls in Ben-Porath [5]. Our choice of this alternative terminology stems from æsthetic concerns as we will be using envelopes within envelopes (which we shall refer to as *nested envelopes*) and we find the image of urns within urns somewhat unappealing.)

3.4 An Example

We now consider an example that illustrates the workings of our mechanism. Consider the following game Sender-Receiver from §6.7 of Myerson [31]. 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 three possible actions, namely x, y and z. Consider the case where $q = \frac{1}{2}$.

We leave it to the reader to verify that with one round of cheap talk (as in Crawford-Sobel), there is no transmission of information and the only equilibrium is the so-called *babbling equilibrium* (also known as the equilibrium of the *silent game* in the terminology of Aumann and Hart [3]). In this equilibrium, the Sender's expected payoffs (for the respective types) is

		x	y	z
$[q]$	t_1	2,3	0,2	-1,0
$[1-q]$	t_2	1,0	2,2	0,3

Figure 3.1: Example

(0,2), while the Receiver gets an expected payoff of 2.

It turns out that any finite number of stages also do not increase the set of equilibrium outcomes. (For the initiated reader, this is because the graph of equilibrium payoffs of the silent game as a multi-function over the set of probability distributions over types, which in this case is $[0, 1]$, is *bi-convex*.) But there are mediated solutions that provide Pareto superior payoffs (in expectation) for both the Sender and the Receiver.

Mediated Communication Mechanism. Consider

$$\mu(x|t_1) = \frac{2}{3}, \quad \mu(y|t_1) = \frac{1}{3}, \quad \mu(z|t_1) = 0,$$

$$\mu(x|t_2) = 0, \quad \mu(y|t_2) = \frac{2}{3}, \quad \mu(z|t_2) = \frac{1}{3}.$$

which maximises the Receiver's payoffs. This gives expected payoffs

$$U_S(\mu|t_1) = 1\frac{1}{3}, \quad U_S(\mu|t_2) = 1\frac{1}{3}, \quad U_R(\mu) = 2\frac{1}{2}.$$

To see truthful reporting by the Sender, consider when t_1 lies. If he reports t_2 his payoff is

$$\frac{2}{3} \times 0 + \frac{1}{3} \times (-1) < 1\frac{1}{3}.$$

If the Sender is of type t_2, then by reporting t_1 he gets

$$\frac{2}{3} \times 1 + \frac{1}{3} \times 2 = 1\frac{1}{3} = U_s(\mu|t_2)$$

which is the same payoff as reporting truthfully. Thus, neither type has an incentive to report untruthfully. Now consider the Receiver's posterior probabilities upon receiving the various recommendations. If the Receiver gets the recommendation x, he knows that the true state of the world is t_1 as $\mu(x|t_2) = 0$. Conditional on his new posterior, his best action is the recommended action x. Similarly, for the case when he receives the recommendation z, his posterior probability assigns probability 1 to state t_2 which makes z the optimal action for him. If he receives y as a recommendation, he assigns probability $\frac{2}{3}$ to the state being t_2. But with these posteriors, $U_R(y|y) = 2 = U_R(z|y) > U_R(x|y)$, making the recommended y an optimal choice. We will now consider a cheap, unmediated communication mechanism which will implement the above mediated plan.

The Mechanism

Step 1 We shall assume that the players have access to sealable envelopes and pieces of paper on which they can write messages. Let the Receiver write messages and place each message in an envelope according to the mediated solution. This means that there are six envelopes, two of which contain the message x, three of which contain the message y and the last one contains the message z. The Receiver then places the envelopes in larger envelopes according to the type of the Sender. In particular, he takes two large envelopes and labels them t_1 and t_2. In the envelope labelled t_1, he places the two envelopes containing x and one containing y. In the envelope labelled t_2, he places all the other envelopes containing messages. Note that this mimics the distribution which is the mediated solution.

Step 2 We now come to the random monitoring stage. Our protocol requires the players to verify the contents of all the envelopes with some probability, ρ (which can be taken to be rational). This is accomplished by a joint lottery as described in §3.3. If the outcome of the lottery requires the verification of the contents of all the envelopes, then this is performed. If there is a deviation from the equilibrium by the Receiver, then the Sender decides to babble (which results in a payoff of 2 for the Receiver). If there is no deviation, then the Receiver repeats Step 1. If however, the outcome of the joint lottery was that there be no verification, the players move to Step 3.

Step 3 The Receiver then hands the two labelled envelopes to the Sender. The Sender privately takes a message out of the envelope corresponding to his type and hands it to the Receiver. The Receiver opens the envelope and takes the prescribed action. All the other envelopes remain unopened and are destroyed publicly.

We have not yet specified the random monitoring probability, ρ. We shall do so now. Let α_R denote the Receiver's payoffs from the babbling equilibrium and let W denote the maximal payoff available to him in the game. Also, let $U_R(\mu)$ be his expected payoff from following the mechanism. (Note that in this case, $\alpha_R = 2$, $W = 3$ and $U_R(\mu) = 2\frac{1}{2}$. Thus, $U_R(\mu) > \alpha_R$ which means that μ is Nash dominating for the Receiver.) Now if the Receiver deviates, his payoff is bounded above by

$$\rho\alpha_R + (1-\rho)W.$$

As mentioned above, his expected payoff from following the mechanism is $U_R(\mu)$. He will not deviate if

$$U_R(\mu) > \rho\alpha_R + (1-\rho)W. \tag{3.1}$$

But $U_R(\mu) > \alpha_R$, so there exists a $\overline{\rho}$ so that (3.1) is satisfied. We can then take any $\rho > \overline{\rho}$ such that $\rho \in \mathbb{Q} \cap [0, 1]$. It is easy to see that the mechanism ends in finite time with probability 1.[3] It is also easy to see that the equilibrium constructed is perfect Bayesian. In §3.5 below, we show that our protocol for the Sender-Receiver game is a special case of a general scheme.

3.5 Proof of Main Theorem

We shall now provide a generalisation of the protocol described in §3.4. Let us assume that the players wish to implement a mediated solution μ, as described in §3.3.

The Mechanism

Step 1 Suppose player 1 is of type $t_i \in T_1$ and player 2 is of type $\tau_j \in T_2$. Player 2 makes envelopes for this state ($t = (t_i, \tau_j)$) according to $\mu(t)$. In particular, suppose $\mu(c|t) = \frac{1}{2} = \mu(c'|t)$ (where $c = (c_1, c_2)$). Then player 2 makes two envelopes for state t. In the first envelope, say the envelope corresponding to the action c, there will be envelopes marked Player 1 and Player 2. The envelope Player 1 will have the action c_1 inside on a piece of paper and similarly for the other envelope. To recap, an envelope for state t has inside it more envelopes containing envelopes with messages. These are the *nested envelopes* referred to in §3.3.

Now player 2 makes envelopes for all states $t \in T$. But the labels on the envelope for a particular state are as follows: state $t = (t_i, \tau_j)$ will be marked t_i, m where $m \in |T_2|$ and there is a bijection $\xi : T_2 \to |T_2|$ such that $m = \xi(\tau_j)$. Note that the bijection $\xi \in \Xi$, the space of all bijections from T_2 to $|T_2|$, is chosen by player 2 at random according to the uniform distribution. Then the players proceed to Step 2.

Step 2 This is the random monitoring stage. The players conduct a joint lottery where with a probability ρ, the contents of all the envelopes are checked. If there is a deviation by player 2, i.e. if he has not placed the messages in the envelopes and the nested envelopes according to the mediated solution μ, then they play the Nash dominated equilibrium for player 2. If there is no deviation, they go back to Step 1. If the joint lottery does not require them to open all the envelopes, the players proceed to Step 3.

Step 3 Suppose that player 2's true type is τ_1. He now removes all the envelopes corresponding to states $t = (t_i, \tau_j)$ where $\tau_j \neq \tau_1$ and hands the remaining envelopes to player 1. In other words, he tells player 1 that he is of type $m = \xi(\tau_1)$, but since player 1 does

[3] The probability that it takes at least n iterations being ρ^n which becomes small very quickly as n becomes larger. Also, the protocol runs for an expected time of $\frac{\rho}{(1-\rho)^2}$.

not know ξ and all the ξ's are equally likely, his beliefs about player 2's type remain unchanged. Player 1 now picks the envelope corresponding to his type and picks one of the envelopes in it. This envelope contains the enveloped marked Player 1 and Player 2. He hands player 2 his envelope and they play the recommended action.

Remark. The reason for player 2 making envelopes corresponding to all states $t \in T$ is that if he just made it for all $t = (t_i, \tau_j)$ where τ_j is his true type and the players are required to inspect the contents of the envelopes in the monitoring stage, then player 1 can, in principle, infer the type of player 2.

All that remains to be done is to determine a monitoring probability, ρ. This will be done as in §3.4. Recall that α_2^j is the payoff to type τ_j of player 2. Now consider the case where player 2 of type τ_j deviates. Then, assuming that the maximum payoff in the game is W, his payoffs are bounded from above by

$$\rho\alpha_2^j + (1-\rho)W.$$

But his expected payoff from following the protocol is $U_2(\mu|\tau_j)$ and we require that

$$U_2(\mu|\tau_j) > \rho\alpha_2^j + (1-\rho)W. \tag{3.2}$$

But $U_2(\mu|\tau_j) > \alpha_2^j$ by assumption, therefore there exists a $\overline{\rho}$ such that (3.2) is satisfied and we can find a $\rho \geqslant \overline{\rho}$ such that $\rho \in \mathbb{Q} \cap [0,1]$. We now consider the issue of perfection. At each zero-probability node at which the players takes an action, their beliefs are exactly their priors. More precisely, the players' beliefs about each other's types are always updated using Baye's rule, and since deviations do not provide any information, the priors remain the same. Given an envelope with a recommendation, player 2 doesn't update his priors about the type of player 1 before he takes an action. Any updating to be performed is done by Baye's rule and is possible for any allocation of messages in envelopes in the various bigger envelopes by player 2. Thus, the equilibrium constructed is perfect Bayesian.

3.6 Discussion

In this paper, we have discussed the notion of a mediated solution and demonstrated that any mediated solution which satisfies a Nash domination condition can be implemented via a cheap communication mechanism. In this section, we will discuss some other ideas related to that of the mediator before we conclude.

In §§3.4 and 3.5, we adapted a mechanism due to Ben-Porath [5]. Our main technique is to use a *menu* of nested envelopes to simulate the role of the mediator. The key to a mediated mechanism is the observation that even upon receipt of the recommendation from the mediator, the players are unsure of the type of their opponent. This is so even after the players update their beliefs using Baye's rule. Thus, the essential role of the mediator is to move the players' posterior beliefs about each other's types around in an incentive compatible way. Thus, we can also view a mediator as representing a noisy communication medium. In other words, we can also think of the mediator as someone who introduces noise in the communication. Let us consider this more abstractly.

Let M_1 and M_2 represent the set of messages send to some machine (which could also be a disinterested third party, the mediator) and let X_1 and X_2 represent the set of messages that the players receive. Now define, $\mathsf{M} := \mathsf{M}_1 \times \mathsf{M}_2$ and $X := X_1 \times X_2$. Let $(X, \mathscr{X})$ be a measurable space and with a slight abuse of notation, let $\Delta(X)$ denote the space of probability measures on this measurable space. This is the abstract representation of any single stage communication mechanism where the players send messages and receive randomised messages after which they take an action. (It is straightforward to extend the definitions to multi-stage communication mechanisms, but we shall not do so here in the interests of notational simplicity.) This is made precise below.

Definition 3.6.1 *An* ***abstract communication mechanism*** *is a tuple* $(M, (X, \mathscr{X}), \nu)$ *where* $\nu : \mathsf{M} \to \Delta(X)$.

Thus, the signals that the players receive is a random variable drawn according to the probability measure $\nu(m)$ over $(X, \mathscr{X})$. If we assume that X is a topological space and $\mathscr{X}$ is the corresponding Borel σ-algebra, then the support of a measure P is a closed set $C =: \text{supp}(\mathsf{P})$ such that (i) $\mathsf{P}(C) = 1$ and (ii) $\mathsf{P}(C') < 1$ for all closed sets $C' \subset C$. The support exists if X is a sufficiently nice topological space, an assumption we shall make. This leads us to the following condition and definition.

Condition T. *For all* $m \in \mathsf{M}$, $\text{supp}(\nu(m))$ *is a singleton.*

Definition 3.6.2 *A communication mechanism is* ***noiseless***[4] *iff it satisfies Condition T. A communication mechanism that is not noiseless will be called* ***noisy***.

Remark. It is pertinent to point out Lehrer and Sorin's result at this stage. They show that by suitably expanding M_1 and M_2, any rational mediated solution can be achieved through a noiseless (deterministic) communication mechanism.

It is instructive to consider the role of the mediator here in the simple case of Sender-Receiver games, as in §3.4. Consider a mediated mechanism $\mu(\cdot|t)$. For each type t, the

[4]Lehrer and Sorin [26] call a noiseless mechanism *deterministic*. We use this alternative terminology to emphasise the distinction from noisy mechanisms.

Sender is promised that the outcome chosen will be drawn according to the probability distribution $\mu(\cdot|t)$. The Sender truthfully reveals his type t to the mediator if he believes that the outcome will, in fact, be generated by the said probability distribution. The mediator then make a recommendation to the Receiver by randomising over the elements in $\mathrm{supp}(\mu(\cdot|t))$ according to the probability distribution $\mu(\cdot|t)$ (with the recommendation being the outcome of the randomisation). Of course, the Sender has to truthfully reveal his type and the Receiver has to comply and take the recommended action, but this is what the incentive compatibility constraints guarantee. Thus, a mediated mechanism is an example of a noisy communication mechanism.

Another example of a noisy communication mechanism is given by the game in §3.4 above. Recall that there is no transmission of information even with unbounded cheap talk. However, Myerson shows that if the Sender has access to pigeons that get lost with some probability, then there can be some transmission of information. This is possible because when a pigeon doesn't arrive, the Receiver is not sure if it was because the Sender did not send the pigeon or because the pigeon got lost. Note that the message that the Receiver receives is the arrival of the pigeon and not the message the pigeon may bear. Myerson then shows that with more general noisy mechanisms (i.e. mediated mechanisms), even more equilibrium outcomes can be realised as communication equilibria. But a common feature of all these noisy mechanism is that they provide some degree of protection for the players. Specifically, in the game in §3.4, with just signalling, type t_1 of the Sender always wants to pretend to be type t_2. However, the introduction of a noisy communication channel provides him with more protection from exposure.

As mentioned before, our mechanism is a descendant of Ben-Porath's [5]. We modify his mechanism by allowing for a menu of nested envelopes from which one of the players chooses. As in [5], the envelopes with messages provide enough uncertainty to simulate the noise that a mediator brings. The main connection between our mechanism and the one used by Ben-Porath in [6] is the use of random monitoring. Otherwise, his mechanism ends in finite time and has no possibility of verification (as in requirement (2) above). We believe that with our Theorem 0 and Main Theorem and Ben-Porath's [6] result, we have a complete taxonomy of everything that can be achieved through unmediated talk. We also hope that our and the aforementioned results will persuade game theoretic modelers to pay more attention to the choice of mechanism in applications.

Chapter 4

Voter Clustering and the Theory of Spatial Voting with Entry

4.1 Introduction

The traditional Hotelling-Downs voting model has a continuum of voters on the unit interval, who vote for the candidate whose electoral position is closest to their own. With two candidates, the model predicts the candidates locating themselves at the (unique) median of the distribution. However, there are many criticisms of this electoral model, on both empirical and theoretical grounds. One of these criticisms is due to Smithies (1941), who points out that this prediction is extremely vulnerable to *successful* entry by a third candidate. As Palfrey (1984) says, "...the result is not robust to perturbations of the technical assumptions of the model." In particular, it is not robust to the possibility of entry.

To get around this, and shed some light on the stability of the two party system, Palfrey (1984) proposed the dominant 2-candidate equilibrium. This two period model has a symmetric, 'essentially' unimodal distribution of voters and 3 *vote maximising* candidates. The two *dominant* candidates (A and B) simultaneously choose positions in the first period, so that they may maximise their share of the votes with the (correct) premise that there will be entry in the next period. The third candidate (C) then chooses his location in keeping with his vote-maximising objectives. The resulting equilibrium is called the dominant 2-candidate equilibrium. In this equilibrium, the entrant will never beat the two established candidates who choose spatially differentiated positions. They play a Cournot-Nash equilibrium in the first period, but are Stackelberg leaders with respect to the entrant[1].

The model predicts that A and B will choose points $\hat{\theta}_A$ (given by the unique solution to $F(\hat{\theta}_A) = 1 - 2F(\frac{\hat{\theta}_A}{2} + \frac{1}{4})$) and $\hat{\theta}_B$ ($= 1 - \hat{\theta}_A$) with the entrant choosing $\hat{\theta}_C = 0.5$ (the me-

[1] For a brief, but more complete discussion of this model, see the Appendix, §A.

dian). By construction[2], the entrant will never get more votes than either of the two dominant candidates and consequently, never comes first, thereby establishing the dominance of the established candidates. Unfortunately, there is a problem with this formulation. It is that the entrant can get arbitrarily close to his best payoff (which he gets at the median), by entering arbitrarily close to, say, the left of candidate A. But by doing so, he comes outright second. So what does he do? Does he maximise votes, or does he also care about his electoral position (in which case, he sacrifices a few votes to come second)? The point is that one should either talk about vote maximising candidates or rank maximising candidates (or some explicitly specified combination thereof). Thus, the dominant 2-candidate model does not provide a convincing explanation for the stability of the two party system, which was the initial objective.

In their extremely influential paper Greenberg and Shepsle (1987) (henceforth referred to as GS) addressed the problems in Palfrey's formulation. They identified two broad classes of electoral systems (whose description we shall borrow in the sequel), the first of which is the *fixed-standard* method[3]. In this method, there is, as the name suggests, a fixed standard, and candidates who achieve it are elected, thus making it the appropriate system in which to study vote-maximising candidates. Under this system, with a given standard, we endogenously determine the number of winners (or more precisely, qualifiers). Thus, in the setting of a legislative body, we would endogenously determine its size. The second class of electoral systems mentioned by GS is the *fixed-number* system. Here, one fixes the number of winners (qualifiers), with relative performance being the criteria (thus making rank maximisation the appropriate objective). They then specifically considered the fixed-number electoral method and defined the notion of k-equilibrium. In a Greenberg-Shepsle k-equilibrium (henceforth referred to as the GS k-equilibrium), after the k candidates have chosen their positions, there is no incentive for further entry by any other candidate. In other words, no prospective entrant has any hope of beating any of the k incumbents (in terms of rank). They then demonstrated (under fairly standard assumptions) that for an arbitrary number of winners ($k \geqslant 2$) there exist societies (i.e. distributions of voters) for which there is no k-equilibrium. Thus, it would be appropriate to consider vote-maximising behaviour (as in Palfrey (1984)) only if we are studying fixed-standard electoral systems.

However, Greenberg and Shepsle left several questions unresolved. The first among these is, when does a GS k-equilibrium exist? Quite obviously, the next question is, what happens when there is no equilibrium? This is potentially a very interesting question. Unfortunately, the axiomatic framework of GS cannot provide us with an answer to the second question.

[2]See the Appendix, §A.

[3]An analysis of which is provided by Greenberg and Weber (1984). For a model with entry using the fixed-standard method see Weber (1998).

The negative results obtained by GS is symptomatic of the axiomatic method (as far as the theory of spatial voting is concerned), which pays inadequate attention to the roles of the specific institutions that may be involved in the phenomenon under study. To answer these questions, we need a more explicit mechanism. The non-cooperative approach provides us with precisely such a setting. In particular, it shows us how a particular set of institutions (rules) affect the outcome. (For an interesting discussion of why the study of institutions is important, see Riker (1980).) Consequently, we always have equilibria and are always able to make predictions. This is important because the theory of spatial voting with multiple candidates has a rich history of impossibility results. We demonstrate that this need not always be the case.

We therefore consider a non-cooperative model with sequential entry by the candidates. For generality, we consider a countably infinite set of potential candidates. There is a fixed cost c to entering the electoral fray. There is a cost of delay, d, a cost that has to be borne each period before entry, which reflects the fact that candidates who enter late have to make an additional effort to influence the same set of voters. Candidates are rank-maximisers and the payoff from being one of the first k candidates (in terms of vote share) is v, while all other candidates get a payoff of 0. This represents the fixed-number method.

With our non-cooperative method we are able to answer the following questions. The first of these is with regards to the equivalence of GS k-equilibrium and the SPE of our extensive form game. We show that if a GS k-equilibrium exists, then it is also the outcome of a SPE of the extensive form game. But this only shows that the set of GS k-equilibria is contained in the set of SPE[4]. We demonstrate that the equivalence holds only for a particular class of distributions[5], which we call *regular*. We then demonstrate the existence of distributions where the set of SPE outcomes is strictly larger than the set of GS k-equilibria. We call these distributions *semi-clustered*. This leaves the set of distributions for which there is no GS k-equilibrium. We find that the SPE for these distributions must necessarily entail strategic delay, i.e. there are some candidates who choose to not enter. These are referred to as *clustered distributions*. In these instances, entry occurs up to the point where the benefit of entry is just balanced by the cost of entry. There is no entry after this point (candidate). Contemporaneously, we provide necessary and sufficient conditions for the existence of a GS k-equilibrium. The conditions for a distribution to be semi-clustered are not easy to check. This is due to the sequential nature of the problem which forces one to look at many histories (and there are uncountably many of these). Nevertheless, we provide one set of sufficient conditions which indicates how to compute the entire set of necessary and sufficient

[4]We shall say that a GS k-equilibrium and an SPE are equivalent if both have the same outcome, in terms of candidate positioning.

[5]We shall refer, interchangeably, to both "societies" and "distributions of voters", with the understanding that both terms refer to the same object.

conditions.

The nomenclature used above refers to the fact that some distributions have sections of the population that are grouped together into "clusters". Each cluster has members whose ideal points are close to each other, i.e. there is some degree of homogeneity among members of a cluster. A clustered society has many such clusters, which are spatially separated. If these clusters are sufficiently far apart, then candidates cannot simultaneously service multiple clusters. This is the intuition behind the result obtained by Greenberg and Shepsle (1987). We provide a definition of clustering and show that a GS k-equilibrium exists if and only if society is not clustered. However, there may be small clusters that are sufficiently far apart so that there are Subgame Perfect outcomes which are not GS k-equilibria. These are the *semi-clustered* distributions. If clusters are sufficiently close together so as to not appear as inhomogeneous clumps, then society is regular and the set of Subgame Perfect outcomes is exactly the same as the set of GS k-equilibria (modulo permutations).

Regular distributions represent societies where there is a first mover advantage. Here, the first movers can block further entry. In contrast, if there is no GS k-equilibrium, the candidates may not want to enter too early. There is a tradeoff between entering early and ensuring that one cannot be defeated. In these clustered societies, the SPE have some waiting, with the last entrant being player N, after whom there is no entry because it is too costly. These distributions do not exactly have a last mover advantage, but there is some advantage to waiting. These equilibria have the same flavour as the (conjectured) equilibria in Martin Osborne's game of sequential entry,[6] but for entirely different reasons. Ours is not the first model to incorporate the non-cooperative approach. Nevertheless, it is substantively different from the models in the literature and with good reason. The main aim here is to describe the conditions under which we have stable outcomes (i.e. GS k-equilibria) and what happens if we don't . This is important because the GS model predicts and demonstrates the existence of instances where the model will not work. Our analysis suggests that if we incorporate the relevant institutions (i.e. rules of the game) in our analysis, then we can, in fact, make some reasonable predictions.

In §4.2 we outline the model underlying our analysis. In §4.3 we provide a partition of the set of all societies (into 3 cells) for the case of 2 winners. The restriction to the case of two winners is to keep the analysis simple. The non-cooperative approach, as developed below, will also shed some light on the set of possible outcomes when we consider voters distributed in multiple dimensions. There are impossibility results in multi-dimensional voting, (see, for example, Davis *et al.* (1972)) which are extensions from the theory of decision making by committees. It is easy to see that once we apply our sledgehammer (the non-cooperative approach), we get a host of SPE. Like Prescott and Visscher (1977), we also feel that the non-

[6]See Martin J. Osborne's website for details.

cooperative model of sequential entry is the appropriate model to look at since it is reasonable to assume that most candidates can expect someone to enter after they have entered. Another important feature of our model is that in every equilibrium outcome, the players are spatially separated. They space themselves so that no further entry is possible. This spatial differentiation is similar to Palfrey's equilibrium outcome. It should be noted that this is in contrast with other models of spatial voting, notably the classical Hotelling-Downs model and Osborne's model of sequential entry. It is pertinent to point out that any model which purports to explain the stability of multi-party systems *must* have spatially differentiated positions. We now proceed to the model.

4.2 The Model

In this section we formalise the model by making precise our assumptions.

4.2.1 The Environment

We first describe the environment in which our game is set.

Assumption 1 (Candidates) *The set of potential candidates is countably infinite.*

There is a continuum of voters with single-peaked, symmetric preferences. Their ideal points are distributed according to a cumulative distribution function $F(\cdot)$ (and associated density function $f(\cdot)$) defined on $[0, 1]$. $F(\cdot)$ has the following properties.

Assumption 2 (Society) $F(0) = 0$; $F(1) = 1$; $F(\cdot)$ *is strictly monotonic, continuous and differentiable on* $[0, 1]$.

Voters vote for the candidate closest to them. By assumption, there are no mass points. This means that the set of voters who are equidistant from any two candidates is of zero measure and thus inconsequential to the analysis.

Assumption 3 (Voter) $\mid x - x_i \mid < \mid x - x_j \mid \quad \forall j \neq i \rightarrow$ *voter with ideal position* x *votes for candidate* i.

Define $S_i(x_1, x_2, \ldots, x_m)$ by

$$S_i(x_1, x_2, \ldots, x_m) \quad := \quad \{x \in [0, 1] \ : \mid x - x_i \mid < \ \min\{\mid x - x_j \mid \ \forall j \neq i\}\}$$

$S_i(\cdot)$ denotes the set of voter ideal points that are closest to a_i. In the terminology of Greenberg-Shepsle, $S_i(\cdot)$ is i's *support.* Note that the set of voter ideal points that are indifferent between two or more candidates is of zero measure, because $F(\cdot)$ is assumed to be continuous. Now define the measure of the support, or the *base* for candidate i by

$$B_i(x_1, \ldots, x_m) := \int_{S_i(\cdot)} f(x)\, dx.$$

4.2.2 The Mechanism

We now describe the mechanism by which the game is played[7]. Candidates can choose to enter if they so desire. However, they cannot enter at a point previously chosen by another candidate.

Assumption 4 (Candidate's actions) *Candidates j chooses his action from his action set $A_j := \{Out\} \cup \{[0,1] \setminus \bigcup_{i=1}^{j-1} a_i\}$.*

Candidates choose their actions one at a time in a predetermined order. The predetermined order is meant to capture the existence of some already dominant candidates who may be more established and thus choose their positions before others.

Assumption 5 (Order of Entry) *There is a fixed order of entry, with one candidate having the option of entering in each period.*

Candidates get payoffs based on their rank.

Assumption 6 (Payoffs from Ranking) *The first k candidates get a payoff of $v < \infty$. All other candidates get a payoff of 0. If m $(< k)$ candidates are tied for the last position and there are n candidates who have more votes than them, then the first $k - n$ of these m candidates each gets v and the rest get 0.*

Candidates have to pay a fixed cost of entry c. This cost is represents the fact that there are costs to campaigning. There is also a per-period cost of delay, d, which attempts to capture the idea that candidates who enter late have to spend more effort to try and influence the same set of voters. The costs have to be paid *only if* the players choose to enter.

Assumption 7 (Costs – Fixed and Waiting) *There is a fixed cost of entry, $c \geqslant 0$. There is a per-period waiting cost, $d > 0$.*

To ensure that we do not have trivial counterexamples, we make the following two assumptions. The first guarantees us that at least $k > 1$ people can enter. The second tells us that the whole process of a countably infinite set of potential candidates making up their minds can be done in a finite amount of time and that we do, in fact, have an election.

[7] In the sequel, the word 'institution' is an allusion to the relevant mechanism.

Assumption 8 (Non-triviality I) *For fixed $k > 1$, there exists $N > k$ such that $c + Nd \leqslant v$ and $c + (N+1)d > v$.*

Assumption 9 (Non-triviality II) *All the candidates take their decisions in a finite amount of time after which the election is held.*

Assumption 1 is designed to help us do some comparative statics. Clearly, we don't need a countably infinite set, but we don't want an upper bound on the set of potential candidates either. Hence, Assumption 1. Assumptions 3 and 4 are standard assumptions regarding preferences and actions taken by voters. Assumption 5 is meant to reflect the fact that some candidates may be already established and get a priority in choosing their position. Weakening this assumption leads to some very interesting questions, but we shall choose to answer the simpler ones first. Assumption 6 reflects the fixed-number system. Assumption 7 is very important for the analysis of the problem. If there were no waiting cost d, then the game tree could have histories with no bound on the last entrant who enters in equilibrium. While this is an interesting technical detail, it does little more than obscure our main point. Also, our assumption reflects the thought that there is a point in time after which no further entry is possible and one does not expect entry to be possible infinitely close to this point. In light of our tie-breaking rule, Assumption 2 is not really necessary, but is nevertheless made because it simplifies our calculations. It also helps to reinforce the idea that to beat an incumbent, an entrant must get more votes than the incumbent.

The most important assumption here is Assumption 5 - the assumption of a fixed order of entry. This is crucial for the results. The assumption is supposed to reflect the idea that established candidates may be forced to move first or that (in an allusion to a more dynamic setting) established candidates may not be able to move too far away from their past positions. Nonetheless, our results will still hold if we randomize the order of entry. Clearly, endogenising the order of entry is much more satisfactory, but this will have to wait for another paper. The other question is whether our mechanism is, in fact, the best mechanism. It is important to note that our theory is a positive one and we would like to first have a consistent positive theory before we debate the normative aspects of the choice of mechanism.

Immediate consequences.
The non-cooperative approach has some rather straightforward and immediate consequences. Recall that Greenberg and Shepsle [20] showed that there exist societies with no k-equilibrium. In our model, this is equivalent to the existence of societies where there are gaps in entry. It is easy to see that we always have an SPE. Moreover, in every SPE there is entry with k candidates entering in equilibrium.

Theorem 4.2.1 *For every distribution of voters, the above game has a Subgame Perfect Equilibrium.*

Proof. The payoffs are fixed and is a maximum of $v < \infty$. Therefore, there exists a smallest natural number N (and candidate N) such that $c + (N + 1)d > v$. This means that in any continuation equilibrium, all candidates after candidate N will choose 'Out'. In other words, *every* Subgame Perfect Equilibrium (if one exists) has the *same* strategy for players $N + 1$ onwards. All that remains to be shown is that there does exist an SPE. Now, the last player who could enter in any equilibrium (candidate N) may be indifferent between some payoffs from some actions (indeed, since he gets either v or 0, this will probably be the case). Making a selection from candidate N's best response set and then proceeding inductively to the top of the tree will give us an SPE. Note that the best response set will always be non-empty because the set of payoffs (which for an entrant, minus the entry costs, is $\{0, v\}$) is finite. □

It is important to note that although the candidates' action sets are not compact, we still have an equilibrium. This is because the payoff space is finite and we can always map back into the strategy space in a way that enables the agent to make a best selection. More formally, for each player i and for every strategy profile of all the other players, there exists a partition of i's strategy set into 3 cells, which reflect his not entering, his entering and being among the top k candidates and his entering unsuccessfully. This ensures that he always has a best response.

Theorem 4.2.2 *In every SPE, there are k entrants.*

Proof. Clearly, there cannot be an SPE with more than k entrants (because someone would be getting a payoff of 0 while paying the entry costs – a negative net payoff). Suppose there is an SPE with only m $(0 \leqslant m < k)$ entrants. If player N does not enter in equilibrium, it is profitable for him to deviate by entering (because there is no continuation equilibrium in which there will be entry after him). Therefore, let us assume that player N does enter in equilibrium. Let Q $(Q \leqslant N)$ be the largest natural number such that player Q doesn't enter in equilibrium. Then, clearly, player Q can deviate and enter. But this is true for all $m < k$. But we have already established above that there cannot be SPE with more than k entrants. This means that every SPE must have exactly k entrants. □

We have thus demonstrated that the situation is not chaotic once we incorporate a more elaborate mechanism. What we haven't resolved is which set of distributions will have immediate entry and which set will entail some gaps in entry and their relation to GS k-equilibria. Thus what we are looking for is a taxonomy of the set of distribution in question on the basis of the existence of gaps in entry in equilibrium. In order to maintain notational simplicity, we shall provide such a taxonomy for the case of two winners, i.e. $k = 2$. The general case will be seen to follow immediately.

4.3 Two Winners - A Complete Analysis

In the present section we shall provide a complete analysis for the case of two winners, i.e. $k = 2$. Our immediate point of departure will be the (solution to the) question posed by Greenberg and Shepsle. For simplicity, we shall also assume that $c + (N+1)d > v$ for $N > 4$. In the notation above, this corresponds to the case when no player after player 4 finds it profitable to enter.

A k-equilibrium is $\widehat{x} = (\widehat{x}_1, \widehat{x}_2)$ such that these points are unbeatable. It must then be the case that no one can enter immediately to the left or right of these two points or somewhere in the interior of the interval $(\widehat{x}_1, \widehat{x}_2) \subset [0,1]$. We shall consider these two requirements separately.

We first require that nobody be able to enter arbitrarily close and beat the two entrants (who have chosen points x_1 and x_2). This means that for all $x \in [0, x_1)$,

$$F(x) \leq F\left(\frac{x_1 + x_2}{2}\right) - F(x_1)$$

But $F(\cdot)$ is continuous so we have

$$2F(x_1) \leq F\left(\frac{x_1 + x_2}{2}\right).$$

Similarly for x_1 (entry in $(x_2, x_2 + \varepsilon)$, small ε) we have

$$F\left(\frac{x + x_2}{2}\right) - F\left(\frac{x_1 + x}{2}\right) \leq F\left(\frac{x + x_2}{2}\right)$$

and as $x \searrow x_1$

$$F\left(\frac{x_1 + x_2}{2}\right) \leq 2F(x_1).$$

The above two inequalities give us our first necessary condition

$$2F(x_1) = F\left(\frac{x_1 + x_2}{2}\right) \tag{4.1}$$

Using the same ideas at the second point x_2 we get

$$F\left(\frac{x_1 + x_2}{2}\right) = 2F(x_2) - 1. \tag{4.2}$$

But how do we know that there exist solutions to the above two equations? The following Lemma establishes that fact.

Lemma 4.3.1 *There always exists a solution $\widehat{x} = (\widehat{x}_1, \widehat{x}_2)$ to equations (4.1) and (4.2).*

Proof. The Lemma is a trivial consequence of the fact that $F(\cdot)$ is continuous[8]. □

[8] See the appendix for the general argument.

The above Lemma states that we can always find points so that there is no entry in the immediate vicinity of these points[9]. However, this leaves open the possibility of entry in the interior of the interval $(\widehat{x}_1, \widehat{x}_2)$. In order to facilitate the analysis, we have the following definition.

Definition 4.3.2 (Clustered Distribution) *A distribution $F(\cdot)$ is said to be* **clustered** *iff for every ordered pair $(\widehat{x}_1, \widehat{x}_2)$ which is a solution to equations (4.1) and (4.2), there exists $x \in (\widehat{x}_1, \widehat{x}_2)$ such that*

$$F\left(\frac{x+\widehat{x}_2}{2}\right) - F\left(\frac{\widehat{x}_1+x}{2}\right) > \min\left\{F\left(\frac{\widehat{x}_1+x}{2}\right), 1 - F\left(\frac{x+\widehat{x}_2}{2}\right)\right\}.$$

This means that a sufficient condition for the SPE to have delay in entry is that the distribution be clustered. There are no points x_1 and x_2 which can successfully prevent further entry. In particular, for a distribution to be clustered means that for some point $x \in (\widehat{x}_1, \widehat{x}_2)$

$$\begin{aligned} & F\left(\frac{x+\widehat{x}_2}{2}\right) - F\left(\frac{\widehat{x}_1+x}{2}\right) &>& F\left(\frac{\widehat{x}_1+x}{2}\right) \\ \longleftrightarrow \quad & F\left(\frac{x+\widehat{x}_2}{2}\right) &>& 2F\left(\frac{\widehat{x}_1+x}{2}\right). \end{aligned} \tag{4.3}$$

Also,

$$\begin{aligned} & F\left(\frac{x+\widehat{x}_2}{2}\right) - F\left(\frac{\widehat{x}_1+x}{2}\right) &>& 1 - F\left(\frac{x+\widehat{x}_2}{2}\right) \\ \longleftrightarrow \quad & 2F\left(\frac{x+\widehat{x}_2}{2}\right) &>& 1 + F\left(\frac{\widehat{x}_1+x}{2}\right) \end{aligned} \tag{4.4}$$

Adding (4.3) and (4.4) we get

$$F\left(\frac{x+\widehat{x}_2}{2}\right) - F\left(\frac{\widehat{x}_1+x}{2}\right) > \frac{1}{3} \tag{4.5}$$

Note that (4.5) is always true for clustered distributions. Inequality (4.5) conveys the intuitive content of our definition of clustered distributions. It says that if there is a subinterval of $(\widehat{x}_1, \widehat{x}_2)$ of half the length but which has more than a third of the total votes, then the first two entrants cannot prevent further entry.

What happens if a distribution is clustered? Then it must be the case that there is no SPE in which the first two candidates can enter. That, in fact, was the point made by Greenberg and Shepsle. They showed that there always exist societies which are polarised, i.e. societies

[9] Note that the lemma does not say anything about the uniqueness of the solution. Once again, the reader is directed to the appendix.

for which there is no GS 2-equilibrium[10]. Thus, clustered societies are precisely the ones for which there is no GS 2-equilibrium. Before proceeding any further, let us first recall Greenberg and Shepsle's definition of k-equilibrium (using our terminology).

Definition 4.3.3 (Greenberg-Shepsle (1987), Definition 2) *A k-tuple $(x_1, \ldots, x_k)$ is a k-equilibrium iff for any other player choosing x_m, where $x_m \in [0,1] \setminus \bigcup_{i=1}^{k} x_i$, it is the case that*

$$B_m(x_1, \ldots, x_k, x_m) \leq \min_{1 \leqslant i \leqslant k} B_i(x_1, \ldots, x_k, x_m)$$

Theorem 4.3.4 *If a distribution is clustered, then there is no GS k-equilibrium.*

Proof. A straightforward consequence of the definitions. □

The above theorem actually says a lot about the existence of a GS 2-equilibrium. Since all the players move simultaneously in the GS world, either there is a 2-equilibrium or there isn't. It can be seen from the analysis above that if a society is *not* clustered, then there must exist a GS 2-equilibrium. Thus, the following, which is a strengthening of the statement above, is stated as a separate theorem to emphasise the fact that one can decide if a society has a 2-equilibrium or not.

Theorem 4.3.5 *If $F(\cdot)$ is* **not** *clustered, then there exists a GS k-equilibrium.*

As was mentioned in Theorems 4.2.1 and 4.2.2, in every SPE we must necessarily have k entrants. Moreover, in clustered distributions, it is clear that player N will enter. This gives us the following result.

Corollary 4.3.6 *If $F(\cdot)$ is clustered, then every SPE will have gaps in entry and in every SPE player N will enter.*

From the analysis above, it is clear that if society is *not* clustered, then there exist SPE which have immediate entry. But we cannot yet conclude that *every* SPE will have immediate entry, i.e. that society is, in fact, *regular* (a term we shall define below). This is because in every path of play, each players actions can be partitioned into two indifference classes, those in which he wins and those in which he loses. He is indifferent between members of each class, but his choice determines the actions of the players who move before him. Thus, his particular selection will determine the run of play and his partition is determined by the selections of players who move after him, i.e selections of paths further down in the game tree. It is now easy to see why clustered societies can have multiple subgame perfect equilibria.

[10]Recall that we refer to the Greenberg-Shepsle k-equilibrium as the GS k-equilibrium.

Let us assume that there exists an SPE where 1 is not supposed to enter. Let us further assume that 1 enters at $\widehat{x}_1$. In order for 2 and 3 (say) to be able to punish him, it must be the case that they get more votes than him and simultaneously prevent further entry. First, we need to determine the set of possible points which can support a punishment. Call this set $P(\widehat{x}_1) \subset [0,1]^2$. We can define this set as

$$\begin{aligned} P(\widehat{x}_1) \quad := \quad & \left\{ (x_2, x_3) \in [0,1]^2 \ : \ F\left(\frac{x_2 + x_3}{2}\right) - F\left(\frac{\widehat{x}_1 + x_2}{2}\right) > F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \right. \\ & \qquad \& \qquad\qquad \left. 1 - F\left(\frac{x_2 + x_3}{2}\right) > F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \right\}. \end{aligned}$$

Clearly, $P(\widehat{x}_1) \neq \varnothing$. Any punishment must lie in $P(\widehat{x}_1)$. Now we need to check if an optimal punishment actually exists. To prevent any further entry (say, by 4), we need to satisfy the following inequalities

$$F(x_2) - F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \leq \text{any 2 of} \left\{ F\left(\frac{\widehat{x}_1 + x_2}{2}\right), 1 - F\left(\frac{x_2 + x_3}{2}\right), F\left(\frac{x_2 + x_3}{2}\right) - F(x_2) \right\} \tag{4.6}$$

$$F\left(\frac{x_2 + x_3}{2}\right) - F(x_2) \leq \text{any 2 of} \left\{ F\left(\frac{\widehat{x}_1 + x_2}{2}\right), 1 - F\left(\frac{x_2 + x_3}{2}\right), F(x_2) - F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \right\} \tag{4.7}$$

$$F(x_3) - F\left(\frac{x_2 + x_3}{2}\right) \leq \text{any 2 of} \left\{ F\left(\frac{\widehat{x}_1 + x_2}{2}\right), 1 - F(x_3), F\left(\frac{x_2 + x_3}{2}\right) - F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \right\} \tag{4.8}$$

$$1 - F(x_3) \leq \text{any 2 of} \left\{ F\left(\frac{\widehat{x}_1 + x_2}{2}\right), F(x_3) - F\left(\frac{x_2 + x_3}{2}\right), F\left(\frac{x_2 + x_3}{2}\right) - F\left(\frac{\widehat{x}_1 + x_2}{2}\right) \right\}. \tag{4.9}$$

Note that right hand side of each inequality is a continuous function in x_2 and x_3. The solution of these inequalities guarantees us that players 2 and 3 cannot be punished by entry in their immediate vicinity. But this still leaves open the question of entry in the intervals $(\widehat{x}_1, x_2)$ and (x_2, x_3). That this can be prevented is represented by the inequalities below.
For all $x \in (\widehat{x}_1, x_2)$ it is the case that

$$F\left(\frac{x + x_2}{2}\right) - F\left(\frac{\widehat{x}_1 + x}{2}\right) \leq \text{any 2 of} \left\{ 1 - F\left(\frac{x_2 + x_3}{2}\right), F\left(\frac{x + x_2}{2}\right), F\left(\frac{x_2 + x_3}{2}\right) - F\left(\frac{x + x_2}{2}\right) \right\}. \tag{4.10}$$

Similarly, for all $x \in (x_2, x_3)$ we have

$$F\left(\frac{x+x_3}{2}\right) - F\left(\frac{x_2+x_3}{2}\right) \leq \text{any 2 of} \left\{1 - F\left(\frac{x+x_3}{2}\right), F\left(\frac{\widehat{x}_1+x_2}{2}\right), F\left(\frac{x_2+x}{2}\right) - F\left(\frac{\widehat{x}_1+x_2}{2}\right)\right\}. \tag{4.11}$$

We can write down similar inequalities for the case in which 1 enters at point $\widehat{x}_2$.

Definition 4.3.7 (Semi-clustered Distribution) *We say that $F(\cdot)$ is* **semi-clustered** *if $F(\cdot)$ is not clustered and if there exist x_2, x_3 such that $(x_2, x_3) \subset P(\widehat{x}_1)$ and x_2, x_3 satisfy inequalities (4.6)-(4.9) and inequalities (4.10) and (4.11).*

We can write down similar conditions for entry at $\widehat{x}_2$. This brings us to the following characterisation.

Lemma 4.3.8 *If $F(\cdot)$ is semi-clustered, then there exist GS k-equilibria and each GS k-equilibrium configuration can be supported as a Subgame Perfect Outcome with entry by the first k players*[11].

Proof. Since $F(\cdot)$ is not clustered, we have a GS k-equilibrium. Let us denote this equilibrium configuration by $(\widehat{x}_1, \widehat{x}_2)$. Let player 1 enter at $\widehat{x}_1$. Conditional on 1's entry at $\widehat{x}_1$ player 2 enters at $\widehat{x}_2$. By definition of the GS k-equilibrium, there can be no further entry. Thus, we have constructed a SPE whose outcome is the GS k-equilibrium. □

We can write similar conditions to check if it is possible for players 2 and 4 to prevent 1 from entering. We omit those inequalities as they are in the same spirit as the ones above. But it is important to notice that if neither 2 and 3 nor 2 and 4 can prevent 1 from entering, then it is the case that no larger set of players can prevent 1 from entering. Furthermore, it can be shown that there exist SPE where there are gaps in entry. A formal demonstration is avoided here (see the analysis above), but the reader is referred to the example in §4.4.2.

This brings us to the last set of distributions, namely the regular distributions. Our definition is tautological, but in light of the previous definitions, is the most appropriate.

Definition 4.3.9 (Regular Distribution) *A distribution $F(\cdot)$ is* **regular** *iff it is neither clustered nor semi-clustered.*

Regular distributions are very well behaved. These are the only distributions which have a one-to-one correspondence between the set of GS k-equilibria and the set of SPE.

[11] Here and below, take $k = 2$.

Lemma 4.3.10 *If $F(\cdot)$ is regular, then there exist GS k-equilibria and each GS k-equilibrium configuration can be supported as a Subgame Perfect Outcome with entry by the first k $(= k)$ players. Furthermore, every SPE entails entry by only the first k players with the corresponding Subgame Perfect Outcome being a GS k-equilibrium.*

Proof. That there exist GS k-equilibria that can be supported as SPE is obvious from the Lemma above. Furthermore, by virtue of $F(\cdot)$ being regular, there are no gaps in entry. This means that in any SPE only the first k players will enter (because no player can be optimally prevented from entering) and if this outcome path is subgame perfect, then there can be no further entry. But this means that this configuration is also a GS k-equilibrium. □

We have provided a complete taxonomy of the set of distributions by providing a partition consisting of 3 cells. In the next section, we provide some illustrative examples.

4.4 Examples

4.4.1 Regular Society

Example 4.4.1 (Regular Society) *Consider $k = 2$, uniform distribution of voters and $N \geqslant 4$. The unique subgame perfect outcome (modulo permutations) is player 1 choosing* $1/4$ *and player 2 choosing* $3/4$.

The uniform distribution is the prototypical example of the most well behaved distributions. As far as the question of timing of entry is concerned, everything is for the best in the world of regular societies; indeed, *tout est pour le mieux dans le meilleur des mondes possibles.* Voters are spread apart uniformly so that there is no role for further entry or for threats that might prevent candidates from entering. This is the unique solution to equations (4.1) and (4.2). It is easy to see that players 2 and 3 cannot prevent 1 from entering. Similarly, it can be checked that players 2 and 4 also cannot optimally prevent 1 from entering. This means that in any equilibrium, player 1 will always enter. Given this, it is immediate that player 2 will enter in such a manner as to prevent further entry. This is possible since the interval $(1/4, 3/4)$ is not clustered. Also, this is the unique GS k-equilibrium. Thus, the unique Subgame Perfect outcome is player 1 entering at $1/4$ and player 2 entering at $3/4$.

It is easy to see that the uniform distribution is not the sole example of a non-polarised society. Indeed, the structure of the requisite conditions is such that one can make minor perturbations to the distribution in the interior of the intervals $[0, 1/4)$, $(1/4, 3/4)$ and $(3/4, 1]$. These perturbations must be such that our society is not clustered or semi-clustered. This means that there are infinitely many distributions which are non-polarised and that there is a open set of regular distributions containing the uniform distribution.

4.4.2 Semi-Clustered Society

Example 4.4.2 (Semi-Clustered Society) *Consider $k = 2$, $N = 4$, but a symmetric distribution wherein $F(0) = 0, F(.095) = .17, F(.17) = .23, F(0.225) = .24, F(.25) = .25, F(.35) = .3, F(.375) = .4, F(.5) = .5$. On any interval, the density is constant.*

There are multiple subgame perfect outcomes here. The GS k-equilibrium configuration is $(0.25, 0.75)$. As was seen above, this can be supported as an SPE. However, there are SPE with gaps in entry.

1. *2 and 3 enter in equilibrium.* 2 and 3 can prevent 1 from entering at $.25$ by themselves taking the points .45 and .83. This prevents further entry and 1 comes third. Here, 2 and 3 get vote shares of 36% and 34% respectively, with 1 getting 30% of the votes. It is easily seen that player 4 cannot enter and beat any two of these three vote candidates. By symmetry, 1 can be prevented from entering at .75. Now that 1 is deterred from entering, 2 and 3 choose the points .25 and .75 which is an equilibrium.
2. *2 and 4 enter in equilibrium.* 1 can be deterred from entering. Let 2 enter at $x_2 = .375$. If 3 enters, it will be at the best response, which is to enter so that $1 - F(x_3) = F(x_3) - F((x_2 + x_3)/2)$. This is to prevent entry in the immediate vicinity of 3. But 4 can now enter at 0.65, thereby getting more votes than 3. Thus, 3 can be deterred from entering and only 2 and 4 enter in equilibrium. Note that this means there is a continuum of SPE and a corresponding continuum of subgame perfect outcomes.

4.4.3 Clustered Society

Example 4.4.3 (Greenberg-Shepsle (1987), Lemma 3) *For a distribution on $[0, 1]$, $F(\cdot)$ is such that $F(0.1) = F(0.4) = 0.3$ and $F(0.6) = F(0.9) = 0.7$. $F(\cdot)$ has constant density on the intervals where it has positive mass.*

It is easy to see that there are no points which are safe. Indeed, the symmetric solutions to equations (4.1) and (4.2) are $F^{-1}(1/4)$ and $F^{-1}(3/4)$. This interval is clustered. It should be mentioned that the distribution does not satisfy our assumptions (as it is not strictly monotonic). Nevertheless, we can take infinitesimal amounts of mass from the intervals with positive mass and add it uniformly to the intervals with zero mass. From the discussions above, it is clear that this can be done in a way that preserves the polar nature of the society.

Other equilibria. It is easy to see that in any equilibrium, player 4 (the last person) *must* enter (consider the contrary, the definition of non-polar societies, Lemmas 1 and 2, and player 4's possible actions). One would consequently expect a slew of equilibria. This indeed turns out to be the case. We list a couple of examples.

1. *3 and 4 enter in equilibrium.* Can we construct strategies to support this outcome? By the definition of polar societies, it is clear that if both 1 and 2 enter, then there is room for further entry and one of these two will get a negative payoff. Thus, one can conclude that there is no perfect outcome path where both these players enter. Now consider histories of length 2, where only one of them has entered. From the structure of the distribution it is easy to see that 3 and 4 can optimally punish this entrant.

2. *1 and 4 enter in equilibrium.* Let $x_1 = 0.5$. Consider $x_2 \in (0.4, 0.5)$. If $x_3 \in (0.5, 0.6)$ then $x_4 = 0.95$ ensures that this is suboptimal for 3. So 3 will not choose this. If $x_3 \in [0.9, 1]$ means $x_4 \in [0.4, x_2)$. This is suboptimal for 2. Similarly, $x_3 \in [0, 0.1]$ means $x_4 \in (0.5, 0.6]$. This is suboptimal for 2. So 2 will not choose $(0.4, 0.5)$ and by symmetry $(0.5, 0.6)$.
   Consider $x_2 = 0.4$. If $x_3 \in [0, 0.1]$ then $x_4 \in (0.5, 0.5]$; suboptimal for 2. Clearly, $x_3 \in (0.5, 0.6]$ is not optimal for one of 2 and 3, depending on who 4 chooses to punish. $x_3 \in [0.9, 1]$ is not good enough for 3. Similarly, for $(0.1, 0.4)$ and $(0.6, 0.9)$.
   1 and 4 can trust 2 to not enter in $[0, 0.1]$ because 3 can then enter at $x_3 = 0.6$ which is suboptimal for 2. The remaining intervals can be checked in a similar fashion. This means that there is no continuation equilibrium where both 2 and 3 enter. In particular, there is no point where 2 can enter. If 2 doesn't enter, neither will 3. Only 1 and 4 enter in equilibrium.

The above example provides us with the intuition for the cases in which there are gaps in entry, i.e. society is polarised. If the constituent groups are sufficiently far apart then the candidates who service one constituency cannot service another.

4.5 Concluding Remarks

4.5.1 On Polarised Societies

We have seen thus far that a complete analysis of the problem entails more than simply using the median-voter theorem. That the outcomes that societies are interested in could depend on more than the position of the median was pointed out by Esteban and Ray (1994). They rightly point out that not only are concentrations of the population important, but the spatial distance between these concentrations is of pivotal importance too. From the analysis above, it is clear that polarised societies are those in which there are concentrations of voters which potential candidates would like to service, but these concentrations are too far apart for a limited number of candidates to be able to service all of them while at the same time preventing further entry. This is precisely the Greenberg-Shepsle example. Semi-polarised societies do not have clusters of voters that have such a large spatial separation. Nevertheless, there are enough concentrations to have SPE with both delay and the lack thereof. Non-polarised societies do not have such problems. A leading example of this is the uniform distribution of voters.

However, our analysis points out that any study of the degree of polarisation should be crucially dependent on the set of possible outcomes. If, as is the case here, there are only a finite set of possible outcomes, then it is easy to see that any classification of distributions can have only finitely many types. Furthermore, it is not obvious that given a (minute) change in voter concentrations, polarisation either increases or decreases. The only reasonable way to order the set of distributions seems to be by saying that a distribution is more polarised than another if it is closer to being polar. One could measure this distance by using a metric on the space of distributions that we are looking at and thereby get a set of equivalence classes of distributions. Of course, this is being said with the benefit of hindsight.

4.5.2 Multi-Dimensional Voting

The theory of voting in multiple dimensions has thus far been axiomatic. Needless to say, there has been a preponderance of negative results. However, once one recognises the institutions in place and adopts a suitable non-cooperative approach, it is easy to see that there will always be equilibria. Whether these equilibria have gaps in entry or not is another question altogether which will be dealt with in a future paper. Once again, our sledgehammer provides us with answers.

4.5.3 Our Results

We partition the set of all distributions into 3 cells. The elements of the first of these cells are called regular. These distributions have GS k-equilibria and the set of GS k-equilibria and SPE are isomorphic with none of the SPE having any delay in entry. The second cell consists of semi-clustered distributions. These distributions also have GS k-equilibria. The set of GS k-equilibria is isomorphic (as above) to a strict subset of the set of SPE. The remaining SPE all have gaps in entry. The third cell consists of clustered distributions. These distributions do not have GS k-equilibria but nevertheless have SPE with entry (by k players). In every SPE here, there are gaps in entry and in every SPE player N (where N is the smallest natural number such that $c + (N + 1)d > v$) will always be the last player to enter. It is worthwhile to note that all our equilibrium outcomes are spatially separated. This is crucial in deterring further entry.

The model we describe has parameters that are not well defined in real political contests. Nevertheless, our attempt here is to try and explain the stability and persistence of multi-party (2 party) systems and to attempt to describe the possible outcomes when the aforementioned stability is absent. It is pertinent to ask if our model is plausible? In the United States, since 1900, there is only one instance of a presidential election in which a third party candidate defeated one of the two established parties. This was accomplished by Theodore Roosevelt who defeated William Taft. Roosevelt was, however, defeated convincingly by Woodrow Wilson, the other established party candidate. This was in the year 1912.

We have talked about established parties as parties who move first. But our model gives them the freedom to change their platform radically from one election to the next. Clearly, our model isn't dynamic in this respect, in that it doesn't account for "temporal inertia". Established candidates are like long run players. Their positions cannot be easily changed. If the concentration of the ideal points of the voters starts to change over time, it may well be the case that an established party becomes defunct. Examples of this sort of behaviour are not too hard to find, even in the history of the United States.

While we have provided answers to a few unanswered questions, there is a lot of work to be done. We have ignored the role of strategic voting (although this is probably insignificant in national elections) and disillusioned voters and how this may relate to why we look at the case of k winners (see, for example, Riker (1976)) for some k. Our analysis suggests that the way to proceed is by using non-cooperative models. More sophisticated models should synthesise the weakening of the assumptions of sincere voting, perfect information and the absence of network effects.

4.6 Appendix

4.6.1 Palfrey's (1984) dominant 2-candidate model.

The two dominant candidates are A and B. They choose (vote maximising) positions so that the entrant cannot get more votes than them. The distribution $F(\cdot)$ is symmetric, strictly increasing and twice differentiable on $[0, 1]$ and concave on $[0, 0.5]$. Let us look at a symmetric equilibrium where the dominant candidates choose their positions so that $\theta_B = 1 - \theta_A$. (Later on it will become obvious that this is the *unique* equilibrium). Clearly, a symmetric equilibrium (with a well defined best response) will have the entrant entering at the median. Suppose the entrant does enter at $\theta_C = 0.5$. Can this be supported as an equilibrium? We require that

$$\begin{aligned}\left[F\left(\frac{\theta_C+\theta_B}{2}\right)\right]-\left[F\left(\frac{\theta_A+\theta_C}{2}\right)\right] &= 2\left[F(\theta_C)-F\left(\frac{\theta_A+\theta_C}{2}\right)\right] \\ &\geq F(\theta_A) \qquad (4.12)\end{aligned}$$

which is the requirement that the entrant prefers the median to entering to the immediate left of candidate A. We also require that

$$2\left[F\left(\frac{1}{2}\right)-F\left(\frac{\theta_A+\theta_C}{2}\right)\right]\leq F\left(\frac{\theta_A+\theta_C}{2}\right) \tag{4.13}$$

which is the requirement that the entrant gets fewer votes than candidate A (and, by symmetry, candidate B). Satisfying these two inequalities guarantees us that the entrant can always get more votes by entering at $\theta_C=0.5$ than by entering to the left of A or to the right of B. Let us stipulate that θ_C is, in fact, 0.5. It is easy to see that (vote maximising) candidate A will choose θ_A so that (4.12) is binding and (4.13) is not (which would put A as close to the median as possible). This means that he chooses $\theta_A=\hat{\theta}_A$ such that

$$F(\hat{\theta}_A)=1-2\left[F\left(\frac{\hat{\theta}_A}{2}+\frac{1}{4}\right)\right].$$

Clearly, (4.13) is satisfied (and does not bind). We have established that the (vote maximising) entrant will enter somewhere in the interval $(\hat{\theta}_A,\hat{\theta}_B)$. Now for all $\theta_C\in(\hat{\theta}_A,\hat{\theta}_B)$, $\frac{\theta_C+\hat{\theta}_B}{2}-\frac{\hat{\theta}_A+\theta_C}{2}=\frac{\hat{\theta}_B-\hat{\theta}_A}{2}$. It is now easy to see that (by the unimodality of $F(\cdot)$) the entrant will choose the subinterval with the highest density, which corresponds to $\theta_C=0.5$. All that remains is to establish that there indeed exists a solution to the above (desired) equation and that it is unique.

Let $g(\theta_A):=2F^{-1}\left(\frac{1}{2}-\frac{F(\theta_A)}{2}\right)-\frac{1}{2}-\theta_A$. Then $\hat{\theta}_A$ solves the desired equation above *iff* $g(\hat{\theta}_A)=0$. Now $g(0)=\frac{1}{2}>0$ and $g(\frac{1}{2})=2F^{-1}(\frac{1}{4})-1<0$. This means that there exists $\hat{\theta}_A\in(0,0.5)$ such that $g(\hat{\theta}_A)=0$. Another important consequence of the unimodality is that the entrant will never enter to the immediate right of candidate A. In particular, this implies that

$$\begin{aligned}
2\left[F\left(\frac{1}{2}\right)-F\left(\frac{\theta_A+\frac{1}{2}}{2}\right)\right] &\geq F\left(\frac{1}{2}\right)-F(\theta_A)\\
\rightarrow \qquad F(\theta_A) &\geq \frac{1}{2}F\left(\frac{1}{2}\right)\\
&\geq \frac{1}{4}.
\end{aligned}$$

Now all that remains is to show that our solution is unique. We have

$$\begin{aligned}
g'(\theta_A) &= 2\frac{1}{f\left(\frac{1}{2}-\frac{F(\theta_A)}{2}\right)}\left(-\frac{1}{2}f(\theta_A)\right)-1\\
&= -\frac{f(\theta_A)}{f\left(\frac{1-F(\theta_A)}{2}\right)}-1\\
&< 0.
\end{aligned}$$

Clearly, $\hat{\theta}_A$ is the unique solution to $g(\theta_A) = 0$ which means that the posited equilibrium is indeed an equilibrium. It is easy to see from the symmetry of the problem that it is, in fact, the unique equilibrium.

4.6.2 Non-uniqueness of solutions to equations (4.1) and (4.2).

From equations (4.1) and(4.2) we can conclude that $F(x_2) - F(x_1) = 1/2$. We now write $x_2 = F^{-1}[F(x_1) + 1/2]$. Substituting in our equation (4.1) we get the map

$$g(x_1) := 2F(x_1) - F\left[\frac{x_1}{2} + \frac{1}{2}F^{-1}\left(F(x_1) + \frac{1}{2}\right)\right]$$

By construction, $x_1 \in [0, x_m]$ where $x_m = F^{-1}(1/2)$. Also, $g(0) < 0$ and $g(x_m) > 0$. By construction, it is seen that any zero of $g(\cdot)$ is a solution to our equations (with x_2 taken appropriately). Since $g(\cdot)$ is continuous, it has a zero in $(0, x_m)$. In fact, $g(\cdot)$ has *all* it's zeros in $(0, x_m)$. This means that a generic $g(\cdot)$ has an odd number of zeros. Is this odd number 1? Let us look at the slope of $g(\cdot)$ at its zero. We have

$$\begin{aligned} g'(x_1) &= 2f(x_1) - f(\cdot)\frac{1}{2}\left[1 + \frac{d}{dx}\left(F^{-1}\left(F(x_1) + \frac{1}{2}\right)\right)\right] \\ &= 2f(x_1) - \frac{1}{2}f\left(\frac{x_1 + x_2}{2}\right)\left(1 + \frac{f(x_1)}{f(x_2)}\right) \end{aligned}$$

where $x_2 = F^{-1}[F(x_1) + 1/2]$. It is easy to see that this slope need not always be positive. In particular, this means that the zero need not be unique. However, for generic distributions, $g(\cdot)$ has only finitely many zeros.

4.6.3 Analysis for arbitrary k.

We provide a sketch of the results that one might expect for the case of more than two winners. We begin by stating the conditions for our society to be polarised.

Let candidate i choose position θ_i. Without loss of generality, we may assume that $\theta_1 < \theta_2 < \cdots < \theta_k$. Then it must be the case that no other candidate can enter immediately to the left or right of any of the candidates. In particular, he must not be able to enter arbitrarily close to the first θ_1. For all $x \in [0, \theta_1)$ we have

$$F\left(\frac{x + \theta_1}{2}\right) \leq F\left(\frac{\theta_1 + \theta_2}{2}\right) - F\left(\frac{x + \theta_1}{2}\right)$$

Clearly, $B_{k+1}(\cdot)$ is increasing in x for $x \in [0, \theta_1)$. This means that the above conditions can be given by (in the limit as $x \nearrow \theta_1$)

$$\begin{aligned} F(\theta_1) &\leq F\left(\frac{\theta_1+\theta_2}{2}\right) - F(\theta_1) \\ \textit{i.e.} \qquad F(\theta_1) &\leq \frac{1}{2}F\left(\frac{\theta_1+\theta_2}{2}\right) \end{aligned} \tag{4.14}$$

Similarly, in the limit as $x \searrow \theta_1$, we have

$$\begin{aligned} F\left(\frac{x+\theta_2}{2}\right) - F\left(\frac{x+\theta_1}{2}\right) &\leq F\left(\frac{x+\theta_2}{2}\right) \\ \textit{i.e.} \qquad \frac{1}{2}F\left(\frac{\theta_1+\theta_2}{2}\right) &\leq F(\theta_1) \end{aligned} \tag{4.15}$$

From (4.14) and (4.15), we get the necessary condition

$$F(\theta_1) = \frac{1}{2}F\left(\frac{\theta_1+\theta_2}{2}\right) \tag{4.16}$$

Similarly, he must not be able to enter arbitrarily close to θ_2, either on the left or the right. As $x \nearrow \theta_2$ and $x \searrow \theta_2$, we have, respectively,

$$F(\theta_2) - F\left(\frac{\theta_1+\theta_2}{2}\right) \leq F\left(\frac{\theta_2+\theta_3}{2}\right) - F(\theta_2) \tag{4.17}$$

$$F\left(\frac{\theta_2+\theta_3}{2}\right) - F(\theta_2) \leq F(\theta_2) - F\left(\frac{\theta_1+\theta_2}{2}\right) \tag{4.18}$$

From (4.17) and (4.18) we get the necessary condition

$$2F(\theta_2) = F\left(\frac{\theta_1+\theta_2}{2}\right) + F\left(\frac{\theta_2+\theta_3}{2}\right) \tag{4.19}$$

It is easy to see that for the points $\theta_3, \ldots, \theta_{k-1}$, we have similar necessary conditions, i.e. for all i such that $2 \leqslant i < k$,

$$2F(\theta_i) = F\left(\frac{\theta_{i-1}+\theta_i}{2}\right) + F\left(\frac{\theta_i+\theta_{i+1}}{2}\right) \tag{4.20}$$

For the candidate on the extreme right, we have

$$F(\theta_k) - F\left(\frac{\theta_{k-1}+\theta_k}{2}\right) \leq 1 - F(\theta_k) \tag{4.21}$$

$$1 - F(\theta_k) \leq F(\theta_k) - F\left(\frac{\theta_{k-1}+\theta_k}{2}\right) \tag{4.22}$$

giving us the necessary condition

$$F\left(\frac{\theta_{k-1}+\theta_k}{2}\right) = 2F(\theta_k) - 1 \tag{4.23}$$

The question now is whether there exist points which satisfy the above k equations. The above necessary conditions represent a system of k equations in k variables. It turns out that the assumptions made on $F(\cdot)$ are sufficient to guarantee the existence of a solution. Of course, the above inequalities aren't the only ones that must be satisfied. We must also ensure, that picking a point arbitrarily close to one of the candidates doesn't get the challenger more votes than any of the other established candidates. These additional inequalities do not affect the following Lemma, but they tell us more about our solution[12].

Lemma 4.6.1 *For any distribution $F(\cdot)$ satisfying Assumption 1, there always exist points $(\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_k)$ satisfying equations (4.16), (4.19), (4.20) and (4.23).*

Proof. The proof is very intuitive, due to the especially simple structure of our equations.

From (4.16) we can solve for θ_2 in terms of θ_1. From (4.19) we can solve for θ_3 in terms of θ_1. Proceeding in this fashion using the equations represented by (4.20), we can write all the θ_i's in terms of θ_1, i.e. we can find points $\theta_1, \theta_2(\theta_1), \ldots, \theta_k(\theta_1)$ to satisfy equations (4.16), (4.19) and (4.20). This leaves us only equation (4.23) to satisfy.

Consider $\theta_1 = 0$. From above, this gives us $\theta_2 = \cdots = \theta_k = 0$. Clearly then,

$$F\left(\frac{\theta_{k-1} + \theta_k}{2}\right) \leq 2F(\theta_k) - 1.$$

Now, by construction, $\theta_1 < \theta_2 < \cdots < \theta_k$. This means that by increasing θ_1 from 0, we can eventually find a θ_1' such that $\theta_k(\theta_1') = 1$. (Note that for this value of θ_1, it is not the case that $\theta_i(\theta_1') = 1$ for $2 \leqslant i \leqslant k-1$.) It is now easy to see that

$$F\left(\frac{\theta_{k-1} + \theta_k}{2}\right) \geq 2F(\theta_k) - 1.$$

But $F(\cdot)$ is continuous by assumption. This means that there is a $\hat{\theta}_1$ (and consequently $\theta_i(\hat{\theta}_1)$ for $2 \leqslant i \leqslant k$) such that equations (4.16), (4.19), (4.20) and (4.23) are satisfied. Let us denote this solution by $(\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_k)$. $\square$

Definition 4.6.2 (Clustered Distribution) *The distribution $F(\cdot)$ is said to be* **clustered** *iff for every k-tuple $(\theta_1, \theta_2, \ldots, \theta_k)$ which solves equations (4.16)-(4.23), there exists $\theta \in (\theta_i, \theta_{i+1})$ for some $i \in \{1, 2, \ldots, k-1\}$ such that*

$$F\left(\frac{\theta + \theta_{i+1}}{2}\right) - F\left(\frac{\theta_i + \theta}{2}\right) \quad > \quad \min\left\{F\left(\frac{\theta_i + \theta}{2}\right), 1 - F\left(\frac{\theta + \theta_{i+1}}{2}\right), \min_{j \neq i, i+1} B_j(\theta_1, \ldots, \theta_k)\right\}.$$

[12] For $k = 2$, they tell us that $F(\theta_1) \leqslant \frac{1}{3}$ and $F(\theta_2) \geqslant \frac{2}{3}$. For $k = 3$, we get $F(\theta_1) \leqslant \frac{1}{4}$, $F(\theta_3) \geqslant \frac{3}{4}$ and $F(\theta_2) \leqslant 1$.

Theorem 4.6.3 *If $F(\cdot)$ is clustered, then there is no GS k-equilibrium.*

Proof. A trivial consequence of Lemma 4.6.1 and the definitions. $\square$

It is possible to provide a definition of semi-clustered societies, but we avoid needless notation in the hope that the spirit of our definitions have been conveyed above. We can similarly define regular distributions. From the arguments above, it is immediate that these distributions will have the same properties as the distributions in the case of 2 winners.

Bibliography

[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] Barany, Imré (1992). "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research, (17)*, 329–340.

[5] 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.

[6] Ben-Porath, Elchanan (2003). "Cheap Talk in Games with Incomplete Information," *Journal of Economic Theory, 108*, 45–71.

[7] Bester, Helmut and Roland Strausz (2001). "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica, 69(4)*, 1077–1098.

[8] Blume, Andreas and Joel Sobel (1995). "Communication-Proof Equilibria in Cheap-Talk Games," *Journal of Economic Theory, 65*, 359–382.

[9] Blume, Lawrence E. and William R. Zame (1994). "The Algebraic Geometry of Perfect and Sequential Equilibrium," *Econometrica, 62(4)*, 783–794.

[10] Bredon, Glen E. (1993). *Topology and Geometry, GTM 139*, Springer-Verlag, New York.

[11] Chung, Kai Lai (2001). *A Course in Probability Theory, Third Edition*, Academic Press, San Diego.

[12] Crawford, Vincent and Joel Sobel (1982). "Strategic Information Transmission," *Econometrica, 50,* 579–594.

[13] Davis, Otto, Morris DeGroot and Melvin Hinich (1972). "Social Preference Orderings and Majority Rule," *Econometrica, 40(2),* 147-157.

[14] Dries, Lou van den (1998). *Tame Topology and O-minimal Structures, London Mathematical Society Lecture Note Series, 248*, Cambridge University Press.

[15] Esteban, Joan-María and Debraj Ray (1994). "On the Measurement of Polarisation," *Econometrica, 62(4),* 819-851.

[16] Forges, Françoise (1990). "Equilibria with Communication in a Job Market Example," *Quarterly Journal of Economics, 105*, 375–398.

[17] Forges, Françoise (1990). "Universal Mechanisms," *Econometrica, (58),* 1341–1364.

[18] 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.

[19] Gerardi, Dino (2000). "Interim Pre-play Communication," mimeo, Northwestern University.

[20] Greenberg, Joseph and Kenneth Shepsle (1987). "The Effect of Electoral Rewards in Multi-Party Competition with Entry," *American Political Science Review, 81(2),* 525-537.

[21] Greenberg, Joseph and Shlomo Weber (1985). "Multi-Party Equilibria Under Proportional Representation," *American Political Science Review, 79*, 693-703.

[22] Grossman, Gene M. and Elhanan Helpman (2001). *Special Interest Politics,* MIT Press, Cambridge, Massachusetts.

[23] Hart, Sergiu (1985). "Nonzero-Sum Two-Person Repeated Games with Incomplete Information," *Mathematics of Operations Research, 10(1),* 117–153.

[24] Krishna, R. Vijay (2004). "Extended Conversations in Sender-Receiver Games," working paper, Pennsylvania State University.

[25] Lang, Serge (1993). *Real and Functional Analysis, Third Edition, GTM 142,* Springer-Verlag, New York.

[26] Lehrer, Ehud and Sylvain Sorin (1997). "One-Shot Public Mediated Talk," *Games and Economic Behavior, 20*, 131–149.

[27] LEVY, AZRIEL (1979). *Basic Set Theory*, Springer-Verlag, New York.

[28] MATTHEWS, STEVEN and ANDREW POSTLEWAITE (1989). “Preplay Communication in Two-Person Sealed-bid Double Auctions,” *Journal of Economic Theory, (48),* 239–263.

[29] MERTENS, JEAN-FRANÇOIS (1989). “Stable Equilibria–A Reformulation,” *Mathematics of Operations Research, 14(4),* 575–625.

[30] MYERSON, ROGER B. (1986). “Multistage Games with Communication,” *Econometrica, 54,* 323–358.

[31] MYERSON, ROGER B. (1991). *Game Theory: Analysis of Conflict,* Harvard University Press, Cambridge, Massachusetts and London, England.

[32] PALFREY, THOMAS R. (1984). “Spatial Equilibrium with Entry,” *Review of Economic Studies, 51*, 139-156.

[33] PRESCOTT, EDWARD and MICHAEL VISSCHER (1977). “Sequential Location among Firms with Foresight,” *Bell Journal of Economics, 8*, 378-393.

[34] RIKER, WILLIAM H. (1980). “Implications from the Disequilibrium of Majority Rule for the Study of Institutions,” *American Political Science Review, 74*, 432-446.

[35] RIKER, WILLIAM H. (1976). “The Number of Political Parties: A Reexamination of Duverger’s Law,” *Comparative Politics, 9*, 93-106.

[36] ROCKAFELLAR, R. TYRELL (1970). *Convex Analysis,* Princeton University Press.

[37] RUBINSTEIN, ARIEL (1982). “Perfect Equilibrium in a Bargaining Model,” *Econometrica, 50*, 97–110.

[38] RUBINSTEIN, ARIEL (1991). “Comments on the Interpretation of Game Theory,” *Econometrica, 59(4)*, 909–924.

[39] SIMON, ROBERT SAMUEL (2002). “Separation of Joint Plan Equilibrium Payoffs from the Min-max Functions,” *Games and Economic Behavior, 41*, 79–102.

[40] 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.

[41] SMITHIES, A. (1941). “Optimum Location in Spatial Competition,” *Journal of Political Economy*, 423-439.

[42] SMULLYAN, RAYMOND M. (1995). *First-Order Logic,* Dover Publications, New York.

[43] URBANO, AMPARO and JOSE E. VILA (2002). “Computational Complexity and Communication: Coordination in Two-Player Games,” *Econometrica, 70(5)*, 1893–1927.

[44] WEBER, SHLOMO (1998). “Entry Deterence in Eectoral Spatial Competition,” *Social Choice and Welfare, 15*, 31-56.

Vita

Date & Place of Birth 10th of August, 1976; Madras, India

Citizenship Indian

Education 2004, Ph.D. in Economics, Pennsylvania State University
1998, B.Tech. (Civil Engineering), Indian Institute of Technology, Madras, India

Fields Primary: Game Theory, Political Economy
Secondary: Monetary Theory, Mathematical Economics