ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ

Αλγοριθμικές πτυχές διαμόρφωσης απόψεων σε κοινωνικά δίκτυα

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

της

ΧΡΙΣΤΙΑΝΑΣ ΛΥΜΟΥΡΗ

Επιβλέπων: Δημήτριος Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΜΩΝ
Αθήνα, Ιούλιος 2015

Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών
Εργαστήριο Λογικής και Επιστήμης Υπολογισμών

Αλγοριθμικές πτυχές διαμόρφωσης απόψεων σε κοινωνικά δίκτυα

Διπλωματική Εργασία

της

ΧΡΙΣΤΙΑΝΑΣ ΛΥΜΟΥΡΗ

Επιβλέπων: Δημήτριος Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 21η Ιουλίου 2015.

(Υπογραφή)	*(Υπογραφή)*	*(Υπογραφή)*
....	....	....
Δημήτριος Φωτάκης	Νικόλαος Παπασπύρου	Αριστείδης Παγουρτζής
Επίκουρος Καθηγητής	Αναπληρωτής Καθηγητής	Αναπληρωτής Καθηγητής
Ε.Μ.Π.	Ε.Μ.Π.	Ε.Μ.Π.

Αθήνα, Ιούλιος 2015

(Υπογραφή)

....

Χριστιανα Λυμουρη
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Ευχαριστίες

Ολοκληρώνοντας αυτή την προσπάθεια, θα ήθελα να ευχαριστήσω τους κ. Ζάχο, ομότιμο καθηγητή Ε.Μ.Π., κ. Παγουρτζή, αναπληρωτή καθηγητή Ε.Μ.Π. και κ. Φωτάκη, επίκουρο καθηγητή Ε.Μ.Π., για τη στήριξή τους τόσο σε ακαδημαϊκό, όσο και σε προσωπικό επίπεδο.

Ιδιαίτερα θα ήθελα να ευχαριστήσω τον επιβλέποντα της διπλωματικής μου, κ. Φωτάκη, για τη συνεχή καθοδήγηση και ενθάρρυνση που μου παρείχε κατά την προσπάθεια αυτή. Ο χρόνος που διέθεσε, καθώς και η εμπιστοσύνη που μου έδειξε για την εκπόνηση της διπλωματικής μου εργασίας, ήταν καθοριστικής σημασίας για την πορεία μου.

Τέλος, θα ήθελα να ευχαριστήσω την οικογένειά μου και τους φίλους μου για την διαρκή υποστήριξη τους καθ΄ όλη τη διάρκεια της προπτυχιακής μου ζωής.

Περίληψη

Σκοπός της παρούσας διπλωματικής εργασίας είναι η μελέτη της διαμόρφωσης απόψεων σε κοινωνικά δίκτυα. Οι δυναμικές γνωμών είναι μια πολύπλοκη διαδικασία που ασχολείται με το πώς ένα άτομο ενσωματώνει την επιρροή των απόψεων των φίλων του, για το σχηματισμό μιας αναθεωρημένης άποψης. Εξαιτίας της τεράστιας επιρροής των κοινωνικών δικτύων σε πολλές πτυχές της καθημερινής μας ζωής, όλο και περισσότερες μελέτες προσπαθούν να μοντελοποιήσουν με υπολογιστικές μεθόδους τις δυναμικές διαμόρφωσης γνωμών σε κοινωνικά δίκτυα. Στην παρούσα εργασία, μελετάμε κυρίως non-Baysian μοντέλα στα οποία τα άτομα ανανεώνουν τις γνώμες τους, βασιζόμενοι στην αρχική τους γνώμη και λαμβάνοντας υπόψιν τους τις γνώμες αυτών με τους οποίους αλληλεπιδρούν. Ειδικότερα, μελετάμε στο πλαίσιο της θεωρίας παιγνίων, το μοντέλο του Kleinberg και επεκτείνουμε το μοντέλο αυτό σε περισσότερες από μια διαστάσεις. Μελετάμε εάν είναι δυνατή μια κατάσταση ισορροπίας για το παίγνιο, καθώς και πως αυτή εξαρτάται από τη δομή του δικτύου και από τις αρχικές γνώμες τον ατόμων που το αποτελούν. Επιπλέον, μελετάμε την ταχύτητα σύγκλισης σε μια κατάσταση ισορροπίας και υπολογίζουμε την απόκλιση της κατάστασης ισορροπίας από τη βέλτιστη λύση, η οποία προκαλείται από την ορθολογική συμπεριφορά των ιδιοτελών ατόμων που απαρτίζουν το δίκτυο.

Λέξεις Κλειδιά

Αλγοριθμική Θεωρία Παιγνίων, Κοινωνικά Δίκτυα, Δυναμική Διαμόρφωση 'Αποψης

Abstract

In this thesis we deal with opinion dynamics in social networks. Opinion dynamics is a complex procedure that involves a cognitive process of how a person integrates influential opinions to form a revised opinion. Nowadays, because of the amazing power of social networks, more and more studies tend to model opinion dynamics by developing computational methods. In this thesis, we are particularly interested in non-Baysian models in which individuals iteratively update their opinions, based on their initial opinions and observing the opinions of their neighbors. Precisely, we examine in the context of strategic game theory, the Kleinberg's model and propose an extended version of it in many dimensions. We study whether an equilibrium can emerge as a result of individuals' interactions and how such equilibrium possibly depends on the network structure and the individuals' initial opinions. We also study the convergence speed to such equilibrium and measure the inefficiency in equilibrium situations that is provoked from rational behavior of self-interested individuals.

Keywords

Algorithmic game theory, Social Networks, Opinion dynamics

Contents

List of Figures

Chapter 1

Introduction

> "*To speak of social life is to speak of the association between people - their associating in work and in play, in love and in war, to trade or to worship, to help or to hinder. It is in the social relations men establish that their interests find expression and their desires become realized.*"
>
> —Peter M. Blau, *Exchange and Power in Social Life*, 1964.

After all, everyone has an opinion. Whether it is politics, art or religion, we all feel compelled to have a say and defend our positions, either through arguments or by pure dismissal of alternative points of view. We develop our initial opinions by perceiving pieces of reality and interpreting the facts through our ideas and emotions. However, others develop different opinions, even for the same facts. As a result, after discussions, opinions tend to influence each other and evolve over time.

A definition arising from the above is the following. Opinions are subjective perspectives that evolve over time. And they evolve as the result of the continuous exchange of information among individuals [54]. The impact of the individuals' opinions on each other forms a network, and as the time progresses, their opinions change by taking into account the attitudes and opinions of others with whom they interact [36]. This is evident in many fields, such as advertising with social cues [7, 57] and diffusion of political views [63]. Therefore, a central question to be answered is how an individual's opinion is influenced by others' opinions and whether this interaction process leads to a socially beneficial stable situation of the network. In this direction, models of social networks and opinion dynamics are structures made up of individuals who interact. Such models are studied in order to explain the confidence or influence flow in populations without relying on detailed social psychological findings.

Since there is a nearly limitless set of situations in which people are influenced by others and new ideas and new behaviors spread through a population, opinion dynamics is of high interest in many areas including politics, as in voting prediction [13]; physics, as in spinning particles [14]; sociology, as in the diffusion of innovation [92], the electronic exchange of personal information [66], language change [24, 89]; and finally economics, [88]. Opinion Dynamics aims to understand how social opinions evolve and converge by defining different interaction mechanisms from individual levels. Next, we discuss briefly about the study of opinion dynamics.

1.1 Opinion dynamics in a nutshell

Over the past few decades prevalent computational models have been focusing on how opinion evolves in a social environment. To begin with, efforts to describe social structures with mathematical models followed suit in the mid 1900s, as the field of sociometry started to develop with the work of the social scientist J. Moreno [68] and the anthropologist J. Barnes [12]. The study of opinion dynamics and social networks goes back to the early work by French [35], who explored the patterns of interpersonal relations and agreements that can explain the influence process in groups of individuals. Subsequently, Harary provided a necessary and sufficient condition to reach a consensus in French's model of power networks [47].

From the aspect of how we define opinion variables mathematically, there are continuous opinion models [27, 97, 61, 49, 34] and discrete opinion models [87, 86, 40, 38, 26, 60, 65, 37]. In continuous opinion models, opinion variables are usually real-valued scalars within a given interval, or sometimes real-valued vectors [34]. These models could help to understand the evolution of opinion in social scenarios such as the decision-making of train fare in public hearing. In discrete opinion models, opinions could be some given values. The most studied discrete opinion model is binary opinions or binary choices, since it could be considered simply to give an yes/no, or agree/disagree to any proposal or suggestion. Either continuous or discrete, the general framework of opinion models is as follows: A distribution of opinions exists in a given society and at each time step, some groups of individuals interact according to some dynamical rule, and they generally tend to find an agreement or compromise among the group.

A popular opinion update rule is the non-Bayes method of averaging neighbors' opinions [2]. This method provides a good approximation to the behavior of a large population without relying on detailed social psychological findings. The work provided by French [35] became the pioneer of such models. In his work, individuals form their new opinions by averaging the opinions of other people with whom they have communicated directly. Following French's model, several non-Bayesian approaches were developed around the idea that "the opinions of each agent evolve dynamically over time as a function of their neighbors' opinions". Much works in this area is based on classical models of interacting particle systems motivated by statistical mechanics [59]. For instance, we have the first uses of the Ising model to describe the behavior of laborers in a strike [39], the voter model, where only one neighbor affects the agent at a time [50, 21], and continuous opinion models like the famous DeGroot model [28].

In social interaction networks, a common way of defining neighboring relation is based on bounded confidence, a label coined by Krause [56]. In such models an individual interacts only with those whose opinions are close enough to his/her own. However, this is not the case for the majority of social interactions. Therefore, many models such as [16], begin with an arbitrary network $G(V, E)$, in which nodes represent the individuals and edges represent their relationships and interactions.

To recapitulate, although a wide range of opinion dynamic models has been studied, every model has three key components.

- Priors. Any model of opinion formation has to start with some type of prior opinions for an individual or group of individuals.

- Sources of information. An individual updates his/her opinion based on new information that he/she receives. This might come from his/her own experiences, from observing others, or from communicating directly with others.
- Method of information processing. This part pertains the process of how the individual combines his/her priors and the information he/she receives.

1.2 Games of opinion formation

Despite these significant findings, to fully understand opinion dynamics through social processes is never a trivial job. Game theory on the other hand, analyzes models of fully rational agents acting in their own best interests, models of collusion and cooperation between agents, and even behavioral models based on experiments with actual human subjects. In other words, game theory is concerned with situations in which decision-makers interact with one another, and in which the happiness of each participant with the outcome depends not just on his/her own decisions but on the decisions made by everyone else who is involved in the game.

Game theory, may be said to begin with the work of von Neumann [71] and von Neumann and Morgenstern [93]. The next major development was John Nash's modication [70] of the von Neumann and Morgenstern approach. Nash formally defined an equilibrium of a noncooperative game to be a profile of strategies, one for each player in the game, such that each player's strategy maximizes his expected utility payoff against the given strategies of the other players. Among the major contributors to game theory following the work of Nash were Reinhard Selten and John Harsanyi. Selten [82] showed that for many games, Nash equilibrium can sometimes generate too many equilibria, whereas Harsanyi [48] introduced the formal modeling of uncertainty into game theory by introducing Bayesian game models and opening the door for models of incomplete information.

Nowadays, game theory is not limited by the mathematical perception. Rather, it is studied in a wider spectrum; economical, physical, anthropological, sociological, philosophical, polythological, legislative and biological. This helps systematize the perception of strategic situations and analyze the dependence of strategies and the assumptions connected with them.

But how can we model individuals' interactions and opinion formation as a game? The game theory framework [23, 44, 74] is an important tool in the study of opinion formation and has been used to model the interactions between individuals [64]. In fact, the literature on social network games, studies settings in which players are located on a network. Each player is associated with a node in the network, with the edges representing the relations between them. Players play a fixed game with their neighbors in the network.

More precisely, game theory assumes a rational behavior, which improves the player's situation. Furthermore, a game by definition, should guarantee a result (cost). Such a cost is strictly connected with the level of player's satisfaction from participating in the game and is represented by a numerical value. Thus, we represent opinions with real values and we assume that each player updates his/her opinion in order to minimize a personal cost that reflects the cost of disagreement with his/her intrinsic opinion, as well the cost of disagreement with his/her neighbors' opinions.

What is more, a highly important component of a game is strategy. A strategy is a set of rules, that is the point of reference for every possible situation during a game. When we study opinion formation in social networks from a game theoretic perspective, we define a player's expressed opinion as his/her strategy at each time step. Thus, in the game theoretic approach, individuals are called the players of the game and their opinion formation is expressed through cost functions. Rather than simply expessing an intrinsic opinion on a subject, individuals stratigically express opinions that may differ from what they really believe.

In this thesis, we merge the computational approaches to learning in games and learning in social networks into one framework. The models that we examine, allow us to study and make conjectures about social phenomena in which individuals update their opinion by observing their neighbors' opinions. In other words, we examine how opinions form by interpreting repeated averaging as a decentralized dynamics for selfish players. In fact, we observe miopically each individual's rational and selfish behavior, and compare the final situation with the optimal one, in which we confront players as a group and the sum of the players' costs is minimized. We therefore adopt a social network structure, in which individuals only communicate directly with their neighbours, but still influence indirectly the whole network. In such games, consensus is not reached in general, as players do not compromise further when this diminishes their utility. Rather, these dynamics will converge to an equilibrium in which players disagree; hence we study the final opinions of the game, answer the question of how quickly equilibria are reached by decentralized dynamics, and bound the cost of disagreement.

1.3 Outline of the thesis

With such background in mind, the goal of this thesis is to model the evolution of opinions in society. In fact, we answer fundamental questions about opinion formation within the framework of game theory. For this reason, firstly we illustrate how game theory can be used to model individuals' behaviors in social networks and analyze the corresponding equilibria and efficiency of such games and secondly we present the theory and the main contributions to the study of social networks. Combining the above theories we study a major recent development in the modeling of opinion formation, came about with the work of Kleinberg and al. [16]. Finally, we propose a model that captures the individuals' consistency in their opinions, by extending the Kleinberg's model in multiple dimensions.

More precisely, the organization of the thesis can be split into three main parts. The first part consists of the chapters 2 to 4 and provides some background on social networks as well as on the main ideas and results of algorithmic game theory, that help us develop a deep perspective on the properties of opinion dynamics. Specifically, chapter 2 serves as a preliminary chapter and it introduces basic game-theoretic definitions that are used in the following chapters. Furthermore, in chapter 3 we discuss in depth about smoothness theory, a novel method that is used to measure the inefficiency of a game. Finally, chapter 4 presents some of the models that have been used to study social networks. This chapter examines how network models are used to predict the choice of behaviors and opinions by people that are connected by a network.

In the second part of the thesis, which consists of the chapter 5, we study in depth a specific social network model, introduced by Kleinberg and al. [16]. More specifically, in a social network with natural payoffs, by interpreting the repeated averaging of opinions as best-response dynamics, we are able to study the opinion formation of the network. Thus, using the theory and the tools that we describe in the first part, we examine the existence of a stable situation for this game and we characterize this situation as well as the convergence time. Finally, using the smoothness theory, described in chapter 3, we are able to study the cost of disagreement in this model, relative to the social optimum.

In chapter 6, which is the last part of the thesis, we extend Kleinberg's model in a multi-dimensional setting. Precisely, a large body of literature in public opinion suggests that people tend to be consistent in their opinions on a wide variety of issues. Thus, the model we propose in chapter 6, encompasses the natural correlation of an individual's opinions on different issues. To illustrate what we mean by this consistency, consider the following two-dimensional structure of policy preferences. One dimension is defined by issues like abortion and women's rights, while the second is defined by religious issues. It can be clearly seen that an individual's opinions on the above issues are not independent. The above observations are captured in the model we propose and thus we study its properties in depth; First, we examine the existence of a stable situation for this game and subsequently we characterize this situation as well as the convergence time. Finally, using the smoothness theory, described in chapter 3, we study the cost of disagreement in this model relative to the social optimum.

Chapter 2

Nash Equilibrium

2.1 Introduction to Strategic Games

In the introduction, we discussed about Algorithmic Game Theory, a field which provides a natural framework in which to talk about situations when a collection of individuals interact strategically–in other words, with each trying to optimize an individual objective function. Along the way of this chapter, we will introduce some notions that occupy central positions in the general area of game theory and are useful tools for our further analysis on social networks in the next chapters.

First, let's describe in more detail how a game is defined; a game is a model designed to explain some regularity observed in auctions, Internet, social networks and more. That is, every individual -which we call a player of the game- chooses his strategy at each time step based on the knowledge he has on others' individuals' decisions. Below we give a more formal definition of a game.

Basic Ingredients of a Game A game is any situation with the following three aspects:

- There is a set $\langle N \rangle$ of individuals, whom we call 'the players'.
- Each player has a set of options for how to behave; we will refer to these as the player's possible strategies and symbolize it as $\langle X_i \rangle$.
- For each choice of strategies, each player i receives a payoff (or cost) c_i that can depend on the strategies selected by everyone. The payoffs will generally be numbers, with each player preferring smaller payoffs to larger payoffs in cost minimization games and the reverse in cost maximization games.

Next, to fix ideas, we describe what is perhaps the most well-known and well-studied game.

Example 2.1.1 (Prisoners' Dilemma)**.** *Two prisoners are on trial for a crime and each one faces a choice of confessing to the crime or remaining silent. If they both remain silent, the authorities will not be able to prove charges against them and they will both serve a short prison term of 2 years for minor offenses. If only one of them confesses, his term will be reduced to 1 year and he will be used as a witness against the other, who in turn will get a sentence of 5 years. Finally if they both confess, they both will get a small break for cooperating with the authorities and will have to serve prison sentences of 4* *years each (rather than 5). Clearly, there are four total outcomes depending on the choices made by each of the two prisoners. We can succinctly summarize the costs incurred in these four outcomes via the following two-by-two matrix.*

		Prisoner 1	
		Confess	*Silent*
Prisoner 2	*Confess*	*(4,4)*	*(1,5)*
	Silent	*(5,1)*	*(2,2)*

It should be clear that in correspondence to the above description, the players on our example are the two prisoners. Each of them can decide either to confess or to remain silent; these are the two possible strategies for each player. Concerning each players' payoff, it depends from the strategy of both prisoners, being 4 if they both confess, 2 if they both remain silent and 1 and 5 if they confess and remain silent respectively.

Now, its time to make some underlying assumptions. First, we assume everything that a player cares about is summarized in the player's payoffs. However, it is not required that players care only about personal rewards. For example, a player may care about both his own benefits, and the other player's benefit. If so, then the payoffs should reflect this; once the payoffs have been defined, they should constitute a complete description of each player's happiness with each of the possible outcomes of the game.

Second, we assume of rational decision-makers, that each player chooses the best available action, which in general depends on the other's strategies. Therefore, each player actually succeeds in selecting the optimal strategy and have some knowledge about the others' strategies. This knowledge is usually based on previous behaviors. We will refer to this as *best response*: the best choice of one player, given a belief about what the other player will do.

Having these two assumptions in mind, can we predict players' behaviors? Let's go back to the Prisoners' dilemma example described above and try to predict first prisoner's strategy if he knew what the second prisoner decided to do. If the second prisoner decided to:

- <u>confess</u>, then the first prisoner would have a payoff of 5 if he remained silent and a payoff of 4 if he would also confess. Thus, he obviously would prefer to confess in this case.

- <u>remain silent</u>, then the first prisoner would have a payoff of 2 if he remained silent and a payoff of 1 if he would decide to confess. Thus, he prefers to confess in this case too.

Therefore, the first prisoner would decide to confess no matter what the second prisoner decides to do. We expect exactly the same from the second prisoner, since the game is symmetric. This is refereed as *strictly dominant strategy*, since each player has a strictly better choice to make without carrying about the others' players' decisions. A weaker definition is that of *dominant strategies*, where a strategy is at least as good as any other possible strategy of a player. We say then that a dominant strategy is a best response to others' players' decisions, whereas a strictly dominant strategy is a strict best response.

It is useful now, before introducing the setting of the Nash Equilibrium, to discuss the above results. To begin with, it can be clearly seen that the situation where both prisoners confess is stable, as no one has an intensive to deviate from his choice. It is also worth noting that despite this, it would be better for both of them if none of them confess. However, we know that this cannot be achieved by rational play.

Finally, clearly these results depend on our assumptions that each player cares about minimizing his personal prison term. If the payoffs were aligned in a different way, stable stations could be different. For example, suppose that they care equally for their personal and the other's prison term. In this case, the situation where both players remain silent, is also stable. Therefore, we can conclude that multiple stable situations can exist in games. However, is it also possible for none to exist?

The next sectors of this chapter shed light to the above issues. More precisely, in section 1.2, we introduce the pure Nash equilibrium concept. In the section 1.3 we describe a class of strategy games where pure Nash equilibrium is always guaranteed, the so-called potential games. However, the lack of existence of pure Nash equilibrium on other games leads us to other equilibria concepts; this is the main setting of section 1.4. Finally, in section 1.5 we present two widely used methods of how to evaluate a game from the perspective of an outside observer who cares for all players equally.

2.2 Pure Nash Equilibrium

Let's start with the analysis of another well-known game called *the battle of the sexes.*

Example 2.2.1. *Consider that two players, a boy and a girl, are deciding on how to spend their evening. They both consider two possibilities: going to a football game or going to the theater. The boy prefers football and the girl prefers the theater, but they both would like to spend the evening together rather than separately. Here we express the players' preferences via payoffs as follows.*

		Boy	
		football	*theater*
Girl	*football*	(3,5)	(1,1)
	theater	(2,2)	(5,4)

Clearly, the two solutions where the two players choose different games are not stable - in each case, either of the two players can improve their payoff by switching their action. On the other hand, the two remaining options, both attending the same game, whether it is football or theater, are both stable solutions; the boy prefers the first and the girl prefers the second.

Thus, this is another example where more than one stable solution exist. Furthermore, games rarely possess dominant strategy solutions and the more complicated a game is, the more difficult is to predict the final solution. However, even if no dominant strategies exist, we expect that each player's strategy decision is a best response to the other players' choice. This idea is captured by the notion of pure Nash equilibria, introduced by J. Nash [70]. The best-response requirement seems reasonable and very intuitive, since if one player would not choose his/her best response answer, he/she would like to deviate from this strategy in order to maximize his/her payoff. Therefore, the Nash equilibrium captures the notion of a stable solution, from which no single player can individually improve his/her welfare by deviating.

Mathematically, x is a Nash equilibrium if for every player $i \in N$,

$$c_i(x_i, x_{-i}) \geq c_i(y_i, x_{-i}), \forall x_i \in S_i$$

where x_i denotes the strategy of player i, x_{-i} represents the strategies of all players other than player i and c_i is the payoff function of player i. Finally S_i is the set of all available strategies of i.

However, not all games have a pure Nash equilibrium. To clear this consider the following example of the famous "rock-paper-scissors" game.

Example 2.2.2 (rock-paper-scissors game). *It is a two-player game with the payoffs given in the following table.*

Player 1

Player 2

	Rock	*Paper*	*Scissors*
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

One can easily confirm that the above game has no stable situation. In each situation, at least one of the two players has an intensive to deviate from his strategy in order to maximize his payoff.

The next section is concerned with a special tool for detecting games with pure Nash equilibria: the so-called potential functions.

2.3 Potential Games

Representing multi-player games with large player populations in the normal form is undesirable for both practical and conceptual reasons. On the practical side, the number of parameters that must be specified grows exponentially with the size of the population. On the conceptual side, the normal form may fail to capture structure that is present in the strategic interaction, and which can aid understanding of the game and computation of its equilibria.

On the other hand, consider for example a graph representing a social network, where each individual is represented by a node and his/her neighbors are the people he/she interacts with. Individual decisions often depend on the relative proportions of neighbors taking actions, as in deciding on whether to buy a product, change technologies, learn a language and so forth. This can result in multiple equilibrium points. For instance, some people may be willing to adopt new technology only if others do, and so it would be possible for nobody to adopt it, or for some fraction to adopt it.

This setting belongs to a class called graphical games, a representation of multi-player games meant to capture and exploit locality of direct influences. It should be clear that examining the existence of pure Nash equilibria by examining one by one the strategies of the players is impossible in such games. Rather, any quantity that strictly decreases with a strict decrease on any player's payoff, and which can only decrease a finite number of times, would be exactly what we need. With this in mind, we try to formulate a measure that has this property. The measure will not necessarily have a strong intuitive meaning, but this is fine as long as we need it as a tool. This measure is provided by potential functions, an elegant an economical way to summarize the information concerning pure Nash equilibria into a single function.

More precisely, Shapley and al. [67] introduced several classes of potential games. A common feature of these classes is the existence of a real-valued function on the strategy space that incorporates information about the strategic possibilities of all players simultaneously. In this section we review the results concerning exact potential games. Exact potential games are characterized from a real valued function on the strategy space which exactly measures the difference in the payoff that accrues to a player if he/she unilaterally deviates. Below we give a formal definition of an exact potential game.

Theorem 2.3.1. *A strategic game G is an exact potential game if there exists a function $P : X \to \mathbb{R}$ such that for all players $i \in N$, for all strategies $x_{-i} \in X_{-i}$ and all $x_i, y_i \in X_i$:*

$$c_i(x_i, x_{-i}) - c_i(y_i, x_{-i}) = P(x_i, x_{-i}) - P(y_i, x_{-i})$$

The function P is called an (exact) potential function for G.

Let's return to the Prisoners' dilemma example described in the beginning of the chapter, to fix the above ideas. The possible outcomes were the following: if they both remain silent, they will both serve a short prison term of 2 years. If only one of them confesses, his term will be reduced to 1 year and he will be used as a witness against the other, who in turn will get a sentence of 5 years. Finally if they both confess, they both will have to serve prison sentences of 4 years each. Thus, the Potential function of the game is given by the formula below:

$$P(p1, p2) = \begin{cases} 4 & \text{if both confess} \\ 6 & \text{if both remain silent} \\ 5 & \text{if exactly one of them confesses} \end{cases}$$

This implies the following result.

Corollary 2.3.1. *If a strategic game G has an exact potential function P, then the Nash equilibria of G and the game obtained by replacing each payoff function by the potential P, coincide.*

It can be clearly seen that if a game is potential, there are many potential functions that could be applied. As we already discussed, a potential function captures the changes on each player's payoff. Therefore, if P_1 is a potential function for a strategic game G, the function $P_2 = P_1 + c$, where c is an arbitrary constant, is also a potential function of the game. For instance, in the previous example of Prisoners' dilemma another potential function could be the following:

$$P(p1, p2) = \begin{cases} 8 & \text{if both confess} \\ 10 & \text{if both remain silent} \\ 9 & \text{if exactly one of them confesses} \end{cases}$$

In fact, the relationship between potential functions is more strict, as they must differ only by a constant factor c. This leads us to the following theorem.

Theorem 2.3.2. *Let G be a game with exact potential functions P and Q. Then, $P-Q$ is a constant function.*

Proof. Let $y, z \in X$ be two different sets of strategies. We know from the definition that if P and G are potential functions of G, then the following equation must hold.

$$P(z) - P(y) = Q(z) - Q(y)$$

Thus $P - Q = c$, where c is an arbitrary constant. □

Consider now a game with a finite set of strategies X for the players, called finite game. If this game is also exact potential game, then it has at least one pure Nash equilibrium. To see why this exists, consider the following. Since X is finite, $argmax_{x\in X}P(x)$ is a non-empty set (if we refer to cost minimization game, the corresponding $argmin_{x\in X}P(x)$ is non-empty). However, regarding the corollary 2.3.1, all members of this set are pure Nash equilibria.

Corollary 2.3.2. *Let G be a finite exact potential game. Then G has at least one pure Nash equilibrium.*

Now let $G = \langle N, (X_i)_{i\in N}, (c_i)_{i\in N}\rangle$ be a strategic game as defined before. A path in the strategy space X is a sequence $\gamma = (x_1, x_2, ...)$ of elements $x_k \in X$, such that for all $k = 1, 2, ..$ the strategy combinations x_k and x_{k+1} differ in exactly one, say the $i - th$ coordinate. A finite path $\gamma = (x_1, x_2, ..., x_k)$ is called *close* if $x_1 = x_k$. It is a *simple closed path* if it is closed and apart from the initial and final point of the path, all strategy combinations are different. The number of distinct strategy combinations in a simple closed path is called the *length* of the path.

Definition 1. *Let $c = (c_i)_{i\in N}$ be the vector of payoff functions and $\gamma = (x_1, x_2, ..., x_k)$ be a finite path. Define $I(\gamma, c) = \sum_{m=1}^{k-1}[c_{i(m)}(x_{m+1}) - c_{i(m)}(x_m)]$, where $i(m)$ is the unique deviating player at step m.*

Shapley and al. [67] proved that if G is an exact potential game, $I(\gamma, c) = 0$, for all closed paths γ. Next, we illustrate this via an example and afterwards the theorem and the corresponding proof are formally presented.

Example 2.3.1. *Consider a two-player strategic game with the payoffs given in the table below.*

	s_3	s_4
s_1	(0,2)	(2,3)
s_2	(2,5)	(4,6)

A potential function for the game is given by $P(s_1, s_3) = 0, P(s_1, s_4) = 1, P(s_2, s_3) = 2, P(s_2, s_4) = 3$. $\gamma = ((s_1, s_3), (s_1, s_4), (s_2, s_4), (s_2, s_3), (s_1, s_3))$ is a closed and simple path with length 4. Notice that $I(\gamma, c) = (3 - 2) + (4 - 2) + (5 - 6) + (0 - 2) = 0$

Theorem 2.3.3. *Let $G = \langle N, (X_i)_{i\in N}, (c_i)_{i\in N}\rangle$ be a strategic game. The following claims are equivalent:*

1. *G is an exact potential game.*

2. *$I(\gamma, c) = 0$ for all closed paths γ.*

3. *$I(\gamma, c) = 0$ for all simple closed paths γ.*

4. *$I(\gamma, c) = 0$ for all closed paths γ of length 4.*

Proof.

- 1. $\rightarrow$ 2.First, we prove the one direction from 1. to 2. It is obvious that the same proof exists for the 1. $\rightarrow$ 3. and 1. $\rightarrow$ 4., as they are special cases of 2. Let P be an exact potential of the game and $\gamma = (x_1, ...x_k)$ a closed path. Then $I(\gamma, c) = I(\gamma, P) = P(x_k) - P(x_1) = 0$.

- For the opposite direction, 2. $\rightarrow$ 1., fix a stragy vector $x \in X$ and take $P(x) = 0$. Let $y \in X, y \neq x$ and $\gamma = (x_1, ..., x_k)$ be a path from x to y (meaning $x_1 = x, x_k = y$). Also, define $P(y) = I(\gamma, c)$. Next, we show that P is a well defined, exact potential function of the game.

  For this, let $\gamma' = (y_1, ..., y_m)$ be a different path from x to y. To proof our claim, $P(y)$ should be equal to $I(\gamma', c)$. However, this follows from the path $\gamma'' = (x_1, ...x_k, y_{m-1}, ...y_1)$ which is a closed path and $I(\gamma'', c) = I(\gamma, c) - I(\gamma', c) = 0$

  Of course, we have not done yet, since we need to show that:

  $$P(y_i, x_{-i}) - P(z_i, x_{-i}) = c_i(y_i, x_{-i}) - c_i(z_i, x_{-i}), i \in N, y_i, z_i \in X_i \text{ and } x_{-i} \in X_{-i}.$$

  Let $\gamma = (y_1, ..., y_k)$ be a path from x to $y = (y_i, x_{-i})$ and $\gamma' = (z_1, ..., z_m)$ a path from x to $z = (z_i, x_{-i})$. Consider the closed path $\gamma'' = (y_1, ..., y_{k-1}, y_k, z_m, z_{m-1}, ...z_1)$. By definition,

  $$I(\gamma'', c) = I(\gamma, c) + c_i(z) - c_i(y) - I(\gamma', c) = P(y) + c_i(z) - c_i(y) - P(z).$$

- 4. $\rightarrow$ 2. Assume $I(\gamma, c) = 0$ for all simple closed paths of length 4. Suppose there is a closed path $\gamma = (x_1, ...x_k)$ such that $I(\gamma, c) \neq 0$ and let γ have a minimal length (≥ 5) among closed paths with this property. Since $i(1)$ deviates at the first step does not make two consecutive deviations: $m \in \{3, ..., k-1\}$. Define the closed path $m' = (x_1, ..., x_{m-1}, y_m, x_{m+1}, ...x_k)$ in such way that deviations of players i(m-1) and i(m)=i(1) are reversed, i.e.,

  $$y_{i_m} = \begin{cases} x_{i_{m+1}} & \text{if i=i(m)=i(1)} \\ x_{i_{m-1}} & \text{otherwise} \end{cases}$$

  The simple closed path $\nu = (x_{m-1}, x_m, x_{m+1}, y_m, x_{m-1})$ of length 4 satisfies $I(\nu, c) = 0$, so $I(\gamma, c) = I(m', c)$ but in the closed path m' player $i(1)$ deviates one step earlier than in γ. Continuing in this way one finds a closed path τ of the same length as γ with $I(\gamma, c) = I(\tau, c) \neq 0$ in which i(1) deviates in two consecutive steps, contradicting the minimality on γ.

□

By know, we have analyzed several properties of games with pure Nash equilibrium and especially the so called potential games. However, we implied that there are games where pure Nash equilibrium does not exist. This motivates the introduction of equilibrium concepts that are much more permissive than pure Nash equilibria. The next section is concerned about this issue and presents other types of equilibria.

2.4 The need for other equilibria concepts

As we showed before, there are games where no pure Nash equilibrium exists. To illustrate this claim, let's start this section with the following example.

Example 2.4.1 (Matching Pennies). *Two payers, each having a penny, are asked to choose from among two strategies - heads (H) and tails (T). The row player wins if the two pennies match, while the column player wins if they do not match, as shown by the following payoff matrix, where 1 indicates win and -1 indicated loss.*

1 \ 2	*Head*	*Tail*
Head	(1,-1)	(-1,1)
Tail	(-1,1)	(1,-1)

Just as in "rock-paper-scissors" game, it is easy to see that this game has no pure Nash equilibria. If both show heads or both show tails, the penny of player 2 goes to player 1, otherwise the penny of player 1 goes to player 2. No matter what strategy pair is chosen, there will always be a player with an incentive to deviate. Instead, it seems best for the players to randomize in order to thwart the strategy of the other player.

More precisely, the absence of pure Nash equilibria, as described above, is usually solved by introducing a larger strategy space: instead of just choosing one of the (pure) strategies, a player may choose each of his/her pure strategies with a certain probability, a so-called mixed strategy.

Except the reasonable explanation of mixed strategies, there is another reason they are worth studying; Nash [70] proved that under this extension, every game with a finite number of players, each having a finite set of strategies, has a Nash equilibrium.

Theorem 2.4.1. *Every game with a finite set of players and finite set of strategies has a Nash equilibrium of mixed strategies.*

Now, we need to understand how they evaluate the random outcome. Would they prefer a choice that leads to a small positive payoff with high probability, or is it better to have a small loss with high probability, and a large gain with small probability? For the notion of mixed Nash equilibrium, we assume that players are risk-neutral; that is, they act to maximize the expected payoff.

Let's return to the matching pennies example and try to figure the probability distribution each player selects. First, it can be clearly seen that due to the symmetric structure of the game, the distributions will be the same for the two players. Second, if each player picks each of his/her two strategies with probability 1/2, then we obtain a stable solution in a sense.

It is important to note that a game may have both pure-strategy and mixed- strategy equilibria. As a result, one should first check all pure outcomes (given by pairs of pure strategies) to see which, if any, form equilibria. Then, to check whether there are any mixed-strategy equilibria, we need to see whether there are mixing probabilities that are best responses to each other.

For instance, in two players games, if there is a mixed-strategy equilibrium, then we can determine player 2's strategy from the requirement that player 1 randomizes. Player

1 will only randomize if his/her pure strategies have equal expected payoff. This equality of expected payoffs for player 1 gives us one equation which we can solve to determine his/her strategy. The same process gives an equation to solve for determining player 2's strategy.

However, we introduced a restriction which we did not discussed until know. More specifically, we required that the games are finite with a finite number of strategies. Therefore, games with an infinite number of players, or games with a finite number of players who have access to an infinite strategy set may not have Nash equilibria. This leads to a further relaxation of the Nash equilibrium notion, introduced by Aumann [5], called correlated equilibrium. The main idea of correlated equilibrium is the following; In contrast to the Nash equilibrium where each player chooses his/her stragy independently, in a correlated equilibrium a coordinator can choose strategies for the players. However, the chosen strategies have to be stable, so that each player find it in his/her interest to follow the recommended strategy.

Formally, a correlated equilibrium is a probability distribution over strategy vectors. Let $p(s)$ denote the probability of strategy vector s, where we will also use the notation $p(s) = p(s_i, s_{-i})$ when talking about a player i. The distribution is a correlated equilibrium if for all players i and all strategies $s_i, s'_i \in S_i$, we have the inequality

$$\sum_{s_{-i}} p(s_i, s_{-i})c(s_i, s_{-i}) \geq \sum_{s_{-i}} p(s_i, s_{-i})c(s'_i, s_{-i})$$

In other words, if player i receives a suggested strategy s_i, the expected profit of the player cannot be increased by switching to a different strategy $s_i \in S_i$.
Next, we illustrate the above with an example.

Example 2.4.2 (Traffic lights)**.** *Consider that there are two players drive up to the same intersection at the same time. If both attempt to cross, the result is a fatal traffic accident. The game can be modeled by a payoff matrix given below.*

		2: *Cross*	2: *Stop*
1	*Cross*	*(-100,-100)*	*(1,0)*
1	*Stop*	*(0,1)*	*(-1,-1)*

In this game there are two pure Nash equilibria, where one of the two cars crosses and of course there is also a mixed strategy giving small possibility to each car to cross. But lets think about what would happen in the real world; a natural solution would be to add a traffic light, a fair randomized device that tells one of the players to go and the other one to wait.

The benefits of the traffic light introduction, make this seem as a natural solution. More precisely, the negative outcomes are completely avoided and the fairness is achieved.

We could use the same idea in less evident examples, such as the battle of the sexes discussed before, where we could have a situation were the boy and the girl flip a coin and depending on how the coin comes up they go either to the theater or to the football game. It can be clearly seen that in this case also, none of the girl and the boy would want to deviate from this correlated equilibrium, as if the other one would follow it, he or she would end up with a lower payoff.

Before closing this paragraph, it is useful to discuss the above results. To begin with, Nash equilibria are special cases of correlated equilibria, where the distribution over S is the product of independent distributions for each player. In other words, the correlated equilibrium is a generalization of the Nash equilibrium, allowing the probabilities on the space of strategy profiles to be correlated arbitrarily.

Furthermore, another reason of why correlated equilibria are so important is that the sum of social welfare in a correlated equilibrium solution, can exceed the corresponding sum of a Nash equilibrium solution. But how do we measure social welfare? This is the main setting of the last section of the chapter.

2.5 Pareto Optimality and Social Optimality

So far we examined the properties of strategic games and presented the evolution of such games from the players' perspectives. In this section, we examine these games from the perspective of an outside observer who watches the game and judges the result. More precisely, as we show, players act selfishly in order to maximize their payoff function. However, this does not mean that their actions are "good for society". But how could someone evaluate an outcome in this way? There are two candidates that we present below.

2.5.1 Pareto-Optimality

To give some intuition, lets start with the following concept; consider a game where there is an outcome o, which is at least as good as any other outcome o' for all players and there is also at least one player who strictly prefers o to o'. In this case, it seems reasonable to say that o is better than o'. Technically, we say that o *Pareto-dominates* o'. Next, we give a formal definition of Pareto optimality.

Definition 2 (Pareto-Optimum). *A set of strategies is Pareto-optimal if there is no other set of strategies in which all players receive payoffs at least as high, and at least one player receives a strictly higher payoff.*

Let's return to the battle of the sexes example, to illustrate the above definition. Remember that both the boy and the girl prefer to spend the evening together and the boy receives a payoff of 5 and the girl a payoff of 3 if they go to watch football, whereas the boy receives a payoff of 4 and the girl a payoff of 5 if they go to the theater. It can be clearly seen that both the situations where they go together to the theater or to watch football are Pareto-optimal. Therefore, one can assume that more than one Pareto-optimal outcomes are possible.

Now, lets think about the prisoners' Dilemma example discussed above. In this game there are four possible outcomes; If both the prisoners remain silent, they will both serve a short prison term of two years for minor offenses. If only one of them confesses, his term will be reduced to 1 year and he will be used as a witness against the other, who in turn will get a sentence of 5 years. Finally if they both confess, they both will get a small break for cooperating with the authorities and will have to serve prison sentences of 4 years each (rather than 5). There is something striking about the outcomes of this game; the only outcome that is not Pareto-optimal is the one corresponding to the unique Nash equilibrium.

2.5.2 Social Optimality

A stronger condition that is even more intuitive than Pareto-optimality is the so-called *social optimality.*

Definition 3 (Social Optimum)**.** *In a cost-maximization game, a set of strategies is a social welfare maximizer (or socially optimal), if it maximizes the sum of the players' payoffs.*

Clearly, in cost-minimization games, a set of strategies that minimizes the sum the players' payoffs is socially optimal. Of course, adding up the payoffs is not always the appropriate way to evaluate a game. To illustrate the above, return again to the prisoners' dilemma example. It can be clearly seen that the social optimum situation is the one where both players remain silent with a total cost of 2+2=4 years (instead of 8 and 6 in the other cases).

Now, before closing this section, it is useful to make some remarks. More precisely, one can notice that a game can have more than one social optimal outcomes and that at least one social optimal outcome exists in every game. Furthermore, outcomes that are socially optimal must also be Pareto-optimal: if such an outcome was not Pareto-optimal, there would be a different outcome in which all payoffs were at least as large, and one was larger - and this would be an outcome with a larger sum of payoffs. On the other hand, a Pareto-optimal outcome need not be socially optimal.

Finally, it is not the case that Nash equilibria are at odds with the goal of social optimality in every game. This triggered much research on how to compare a final situation of a game (if one exists) with a desirable social optimal solution and this is the main concept analyzed in the next chapter.

Chapter 3

Smooth Games and the Price of Anarchy

3.1 The Inefficiency of equilibria

In the previous chapter of this thesis, we analyzed strategic games from the players' perspective. More precisely, players act selfishly in order to minimize their personal cost. However, it became crystal clear that the outcome provoked from rational behavior of self-interested players may be in contrast with a centrally desirable outcome. In this chapter we examine how much an equilibria situation deviates from the optimal one.

To begin with, in order to measure the inefficiency of a game, we need to specify several factors of the game. First, we need to determine which solutions are the optimal ones. A possible candidate is a Pareto optimal solution, which discussed in the previous chapter. More precisely, a Pareto optimal solution, is a solution that is at least as good as any other solution for all players and strictly better for at least one player of the game. For instance, recall the prisoners' dilemma example. In this game, the only pure Nash equilibria is the situation where both prisoners confess and suffer a cost of 4. One can observe that this situation is Pareto inefficient, since there is another outcome in which all the players receive a lower cost. This qualitative perspective is appropriate in games where the "cost" or the "payoff" of a player is an abstract quantity that only expresses player's preferences among possible outcomes. However, costs and payoffs have concrete interpretations in many games, such as time delay and money.

Therefore, we can use the social optimum solution, which provides a more quantitative framework. More precisely, for the rest of our analysis we use the sum of the players' costs as an *objective function* that expresses numerically the social welfare of the game. Notice that the pure Nash Equilibrium on prisoners' dilemma example does not minimize either this function.

Introducing this objective function helps us evaluate an outcome of a game and approximate optimality. More precisely, in a cost-minimization game, an optimal solution is an outcome that minimizes the objective function of the game, whereas an optimal outcome should maximize the objective function in a cost-maximization game. One can consider games where equilibrium solution and optimal solution coincide. However, this is not the case for the majority of strategic games because of rational behavior of self-interested players. In the rest of this chapter we discuss methods that measure in what senses the equilibria outcomes approximate the social optimum outcome. All of these methods quantify the extent to which a given equilibrium outcome approximates an optimal one according to the ratio between the objective function values of the two outcomes.

Before we continue, it is useful to present some assumptions. First, we consider only non-negative objective functions, so this ratio is always non-negative. Second, we interpret the ratio $c/0$ as 1 if $c = 0$ and as $+\infty$ if $c > 0$. Third, this ratio is at least 1 for cost-minimization games and at most 1 for cost-maximization games. In either case, a value close to 1 indicates that the given outcome is approximately optimal.

To fix ideas, let's return again to the prisoners dilemma example. The sum of the players' costs in the Nash equilibrium is 8, whereas the minimum sum of costs is 4. Therefore, the corresponding ratio for the equilibrium outcome is 2. This notion of approximation, is discussed later in this chapter.

However, as we discussed in the previous chapter, there are games with more than one pure Nash equilibrium. In such games, it is not clear which equilibrium should be compared to an optimal outcome. In sections 3.2 and 3.3, we discuss the two most popular approaches. In section 3.4, we describe a novel method that let us identify sufficient upper bounds for the inefficiency of pure Nash equilibria and thus, approximate optimality.

3.2 The Price of Stability

To begin with, we are interested in identifying games where the equilibria solution is a good approximation of the optimal one. By this, we mean that the ratio between the objective function values of the equilibrium and optimal outcomes is close to 1. However, in a game where multiple equilibria exist, which one should be compared to the optimal outcome? A first thought could be to choose randomly one of the pure Nash equilibrium outcomes and compare it to the optimal one. But this would cause a very large ratio and thus selfish behavior could not be characterized as benign in such games.

The *price of stability* is a measure of inefficiency designed to differentiate between games in which all equilibria are inefficient and those in which some equilibria are inefficient. Formally, the price of stability of a game is the ratio between the best objective function value of one of its equilibria and that of an optimal outcome.

Of course, a bound on the price of stability, which ensures only that one equilibrium is approximately optimal, provides a considerable weak guarantee that the game is near optimal. Nevertheless, there are two reasons to study the price of stability. First, in some applications, a non-trivial bound is possible only for the price of stability. Second, the price of stability has a natural interpretation in many network games, where the best equilibrium is an obvious solution to propose.

On the other hand, it can be clearly seen that price of stability is not suitable in games where no pure Nash equilibrium exists. For instance, since a mixed-strategy Nash equilibrium might randomize only over outcomes that are not (pure-strategy) Nash equilibria, it is not clear how to interpret it as a single proposed outcome for future use by selfish players.

Now, lets examine an example to fix ideas. Recall the battle of the sexes example, discussed in the previous chapter. In this example, a boy and a girl, are deciding on how to spend their evening. They both consider two possibilities: going to a football game or going to the theater. The boy prefers football and the girl prefers the theater, but they both would like to spend the evening together rather than separately. The payoffs for both players are summarized in the next array.

		Boy	
		football	theater
Girl	football	(3,5)	(1,1)
	theater	(2,2)	(5,4)

Clearly, there are two equilibrias in this game; the situations where both doing the same activity, whether it is football or theater. In the case where both attend football, the value of the social welfare function is 8, whereas the corresponding value when both attend theater is 9. Furthermore, the social optimal solution, which is the one that maximizes the objective function is 9. Therefore, the price of stability of the game is 1

$$PoS = \frac{\text{value of the best pure Nash equilibrium}}{\text{value of the optimal solution}} = 1$$

Therefore, in the battle of the sexes game, despite the fact that the players act selfishly, a stable situation can coincide with the optimal solution.

By now, we only quantified the inefficiency of the best equilibrium of a game. Other interesting approaches are the analysis of a "typical" equilibrium and of the worst equilibrium of a game. Concerning the "typical" equilibrium, it is difficult to be defined in a meaningful and analytically tractable way and therefore, has not yet been used successfully to study the inefficiency on a game. On the other hand, the worst equilibrium of the game, is the most popular measure of the inefficiency of equilibria and it is discussed briefly in the next section.

3.3 The Price of Anarchy

The *price of anarchy* [55], the most popular measure of the inefficiency of equilibria, resolves the issue of multiple equilibria by adopting a worst-case approach. Precisely, the price of anarchy of a game is defined as the ratio between the worst objective function value of an equilibrium of the game and that of an optimal outcome. Note that both the price of anarchy and the price of stability of a game are defined with respect to a choice of objective function and a choice of equilibrium concept. In other words, they still exist even if we choose another objective function, such as the maximum payoff among the players' payoffs.

To get an idea of the price of anarchy, let us consider the battle of the sexes example discussed above. Clearly, the worst pure Nash equilibrium of the game is when both the boy and the girl decide to attend football. Then, while the price of stability is 1, the price of anarchy is

$$PoS = \frac{\text{value of the worst pure Nash equilibrium}}{\text{value of the optimal solution}} = \frac{8}{9}$$

It can be clearly seen that for a game with multiple equilibria, a price of stability is at least as close to 1 as its price of anarchy and it can be much closer. On the other hand, a bound on the price of anarchy provides a stronger guarantee than a bound on the price of stability about the efficiency of the game. Finally, in games where only one pure Nash equilibrium exist, the price of stability and the price of anarchy coincide.

Now, as we discussed before, the motivation behind studying the price of anarchy is to quantify the increase of the social cost due to selfish behavior. Furthermore, when we examined price of stability in section 3.2, we considered that it is more appropriate for pure equilibria concepts. However, in the price of anarchy case, we are not able to make such a consideration.

If one wants to study a distributed system in which agents repeatedly perform improvement steps until they reach a Nash equilibrium, then pure equilibria are the right solution concept. However, upper bounds about the price of anarchy for mixed equilibria are more robust than upper bounds for pure Nash equilibria, as mixed equilibria are more general than pure ones. In this chapter, we consider both of these equilibrium concepts; in section 3.4, we analyze smoothness arguments for pure Nash equilibria concept, whereas in section 3.5 we extend this theory to more general equilibrium concepts.

3.4 Smoothness Theory

A rigorous way to argue that a system with self-interested participants has good performance is to prove that every possible outcome of the system has an objective function value close to that of an optimal outcome. Thus, in this section, we identify a sufficient condition for an upper bound on the price of anarchy of pure Nash equilibria of a game for the welfare objective function. This condition encodes a canonical proof template for deriving such bounds. We call such proofs "smoothness arguments", as introduced by T. Roughgarden [78]. Next we give the formal definition of smooth games.

Definition 4 (Smooth games). *A cost-minimization game is (λ, μ)-smooth if for every two outcomes $\vec{s}$ and $\vec{s}^*$,*

$$\sum_{i=1}^{k} C_i(s_i^*, \vec{s}_{-i}) \leq \lambda C(\vec{s}^*) + \mu C(\vec{s}) \tag{3.1}$$

There is an analogous definition on cost-maximization games, where the first part of the above inequality, is greater or equal to the second. It can be clearly seen from the above definition that smoothness controls the cost of a set of "one-dimensional perturbations" of an outcome, as a function of both the initial outcome $\vec{s}$ and the perturbations $\vec{s}^*$. Precisely, if we start with a strategy profile $\vec{s}$, and consider for each player i the cost of i if he/she unilaterally deviates to his/her strategy in profile $\vec{s}^*$, then the sum of all such costs is upper-bounded jointly using costs $\vec{s}$ and $\vec{s}^*$.

To make sense out of this condition, we next show how a smoothness argument implies price of anarchy bounds, in a canonical way. More specifically, we claim that if a game is (λ, μ)-smooth, with $\lambda \geq 0$ and $\mu \leq 1$, then each of its pure Nash equilibria $\vec{s}$ has cost at most $\frac{\lambda}{1-\mu}$ times that of an optimal solution $\vec{s}^*$. The formal proof of the claim follows.

Proof.

$$\begin{aligned} C(s) &= \sum_{i=1}^{N} C_i(\vec{s}) \\ &\leq \sum_{i=1}^{N} C_i(s_i^*, \vec{s}_{-i}) \\ &\leq \lambda C(\vec{s}^*) + \mu C(\vec{s}) \end{aligned}$$

Clearly, the first inequality arises from the equation by applying the Nash equilibrium condition, once to each player i with the hypothetical deviation s_i^*. The second inequality follows from the definition of smooth games in 3.1. Now, by rearranging it follows that $C(\vec{s}) \leq \frac{\lambda}{1-\mu} C(\vec{s}^*)$. □

It is useful now to make some remarks before we continue. First, the above proof uses the Nash equilibrium hypothesis, but only to justify why each player i selects his/her strategy s_i rather than his/her strategy s_i* in the optimal outcome. Second, definition 3.1 is sufficient for the last line of this proof, but it insists on more than what is needed; it demands that the inequality holds for every outcome $\vec{s}$ and not only for Nash equilibria. This is the basic reason why smoothness arguments imply worst-case bounds beyond the set of pure Nash equilibria.

We say that a price of anarchy bound is *robust*, if it is the best upper bound provable via a smoothness argument.

Definition 5 (Robust price of anarchy). *The robust price of anarchy of a cost-minimization game is*

$$inf\{\frac{\lambda}{1-\mu} : (\lambda, \mu) \ such\ that\ the\ game\ is\ (\lambda, \mu) - smooth\}.$$

Smoothness arguments are a class of upper bound proofs for the POA of pure Nash equilibria that are confined to use the equilibrium hypothesis in a minimal way. Thus, if we care only about the price of anarchy of pure Nash equilibria, then we are free to establish an upper bound using any argument that we please.

On the other hand, we can make some relaxations on the definition 3.1 that will allow us compute more robust bounds for some games. Precisely, we define the framework of *local smoothness*, which provides a sufficient condition for a game to have a bounded price of anarchy. This framework refines the smoothness paradigm for games with convex strategy sets, by requiring the inequality to hold only for nearby pairs of outcomes, rather than for all pairs of outcomes. Next, we define the local smooth games as follows.

Definition 6 (Local Smooth Games). *A cost-minimization game is locally (λ, μ)-smooth with respect to the outcome $\vec{s}^*$ if for every outcome $\vec{s}$,*

$$\sum_{i=1}^{N} [c_i(\vec{s}) + \triangledown_i c_i(\vec{s})^T (s_i^* - s_i)] \leq \lambda C(\vec{s}^*) + \mu C(\vec{s}) \tag{3.2}$$

where $\triangledown_i c_i(\vec{s})$ denotes the gradient of c_i with respect to s_i.

Comparing the equations 3.1 and 3.2, it becomes crystal clear that the rough intuition behind local smoothness is to require the constraint 3.1 only for outcomes $\vec{s}^*$ that are arbitrarily close to $\vec{s}$. Since dropping constraints increases the set of feasible values for λ and μ, this idea has the potential to yield improved upper bounds on price of anarchy.

Since the smoothness theory is very general, there is a wide range of applicability. Next, in order to illustrate how smoothness arguments work, we examine in depth an example on congestion games with affine cost functions. The price of anarchy in such games was first studied by Awerbuch and al. [6] and Christodoulou and Koutsoupias [20].

Example 3.4.1 (Congestion Games with Affine Cost Functions). *Congestion games are a class of strategic games first proposed by Rosenthal [77]. In such games we define a set of N players and a set of resources E. Players choose strategies with their strategies sets* *being* $S_1, S_2, ..., S_N \subseteq 2^E$*. For each element* $e \in E$*, there is a cost function* $c_e : Z \to \mathbb{R}$*. We assume that every cost function is affine, meaning that* $c_e(x) = a_e x + b_e,\ a_e, b_e \geq 0\ \forall e \in E$*.*

To illustrate congestion games better, consider a traffic network, where E *is the edge set of the network and the strategies of a player correspond to paths between its source and sink vertices. Given a strategy profile* $\vec{s} = (s_1, s_2, ...s_N)$*, with* $s_i \in S_i$ *for each* i*, we say that* $x_e = |\{i : e \in s_i\}|$ *is the load induced on* e *by* $\vec{s}$*. In other words,* x_e *is the number of players that use* e *in* $\vec{s}$*. The cost to player* i *is defined as* $c_i = \sum_{e \in s_i} c_e(x_e)$*. Thus, the objective function of the game is defined as follows.*

$$C(\vec{s}) = \sum_{i=1}^{N} c_i(\vec{s}) = \sum_{e \in E} c_e(x_e) x_e$$

We claim that every congestion game with affine costs is $(\frac{5}{3}, \frac{1}{3})$*-smooth and hence has a robust price of anarchy of at most 5/2.*

Proof. We begin our proof with an inequality noticed by Christodoulou and Koutsoupias [20]:

$$y(z+1) \leq \frac{5}{3}y^2 + \frac{1}{3}z^2$$

Thus, for all $a, b \geq 0$ and non-negative integers y, z, the following inequalities hold.

$$ay(z+1) \leq a\frac{5}{3} + a\frac{1}{3}z^2 \tag{3.3}$$

$$by \leq b\frac{5}{3} + b\frac{1}{3}z \tag{3.4}$$

and thus

$$ay(z+1) + by \leq \frac{5}{3}(ay+b)y + \frac{1}{3}(az+b)z \tag{3.5}$$

Let $\vec{s}$ be a strategy profile on Nash equilibrium and $\vec{s^*}$ on the optimal outcome. From definition for each player i,

$$c_i(\vec{s}) \leq c_i(s_i^*, s_{-i}) \leq \sum_{e \in S_i^*} (c_e(x_e + 1)) = \sum_{e \in S_i^*} (a_e(x_e + 1) + b_e)$$

since the number of players using e in the outcome (s_i^*, s_{-i}) is at most one more than that in $\vec{s}$. Therefore, summing up all for all i,

$$\begin{aligned}
\sum_i c_i(\vec{s}) &\leq \sum_i c_i(s_i^*, s_{-i}) && (3.6)\\
&\leq \sum_i \sum_{e \in S_i^*} (c_e(x_e + 1)) && (3.7)\\
&= \sum_{e \in E} (c_e(x_e + 1) x_e^*) && (3.8)\\
&= \sum_{e \in E} (a_e(x_e + 1) + b_e) x_e^* && (3.9)\\
&\leq \sum_{e \in E} \frac{5}{3}(a_e x_e^* + b_e) x_e * + \sum_{e \in E} \frac{1}{3}(a_e x_e + b_e) x_e && (3.10)\\
&= \frac{5}{3} C(\vec{s^*}) + \frac{1}{3} C(\vec{s}) && (3.11)
\end{aligned}$$

where 3.11 follows from 3.5, with x_e^* and x_e playing the roles of y and z respectively. This proof implies that every congestion game with affine cost functions has an upper bound of $\frac{5}{2}$ on the price of anarchy of pure Nash equilibria. This fact was first proved independently in [20] and [6], along with matching worst-case lower bounds. □

Despite the above analysis, we have discussed so far a number of reasons why pure Nash equilibria might not occur. Precisely, the players might be playing a game in which computing a pure Nash equilibrium is a computationally intractable problem [32], or even worse, a game in which a pure Nash equilibrium does not even exist. These restrictions motivate worst-case performance bounds that apply to as wide a range of outcomes as possible, and under minimal assumptions about how players play and coordinate in a game. This is the main topic of the next section of the chapter where we present a theorem that extends every smoothness argument to more general equilibrium concepts.

3.5 Extensions to other equilibrium concepts

As we discussed in chapter 2, there are several reasons to study more permissive equilibrium concepts than pure Nash equilibria. First, while there are games where no pure Nash equilibria exist, every finite game has at least one mixed Nash equilibrium [70]. Second, while computing Nash equilibrium is a computationally intractable problem in general [25, 31], this is not the case in the correlated equilibrium computation [43]. These lead us to try to extend the smoothness arguments to more general equilibrium concepts.

In this section, we prove an extension theorem that claims the following; every price of anarchy bound for pure Nash equilibria that follows from a smoothness argument, extends automatically to the more general equilibrium concepts discussed in chapter 2, and to the coarse-correlated equilibria - an even more general concept. The drawback of the local smoothness arguments is that the corresponding price of anarchy extends automatically to the correlated equilibria of a game, but not necessarily to the coarse-correlated equilibria [79].

We begin our analysis with a brief presentation of these equilibrium concepts, in order to recall their main properties.

mixed-Nash equilibria

A set $(\sigma_1, \sigma_2, ..\sigma_N)$ of independent probability distributions over strategy sets- one per player of a cost-minimization game- is a *mixed Nash equilibrium* of the game if no player can decrease its expected cost under the product distribution $\sigma = \sigma_1 \times \sigma_2 \times ... \times \sigma_N$ via a unilateral deviation. In other words,

$$E_{\vec{s}\sim\sigma}[c_i(\vec{s})] \leq E_{\vec{s}_{-i}\sim\sigma_{-i}}[c_i(s_i', \vec{s}_{-i})] \tag{3.12}$$

Clearly, every pure Nash equilibrium is a mixed Nash equilibrium and not conversely, since many finite games do not have pure Nash equilibrium, where a mixed-Nash equilibrium always exist.

correlated equilibria

A *correlated equilibrium* [5] of a cost-minimization game G is a joint probability distribution σ over the outcomes of G with the property that

$$E_{\vec{s}\sim\sigma}[c_i(\vec{s})|s_i] \leq E_{\vec{s}\sim\sigma}[c_i(s_i', \vec{s}_{-i})|s_i] \tag{3.13}$$

for every i and $s_i, s_i' \in S_i$. In other words, if player i receives a suggested strategy s_i, the expected profit of the player cannot be increased by switching to a different strategy $s_i \in S_i$.

coarse-correlated equilibria

Finally, a *coarse-correlated equilibrium* [69] of a cost-minimization game is a probability distribution σ over outcomes that satisfies

$$E_{\vec{s}\sim\sigma}[c_i(\vec{s})] \leq E_{\vec{s}\sim\sigma}[c_i(s_i', \vec{s}_{-i})] \tag{3.14}$$

for every i and $s_i' \in S_i$. While a correlated equilibrium protects against deviations by a player aware of its recommended strategy, a coarse-correlated equilibrium is only constrained by player deviations that are independent of the sampled outcome.

We now give our extension theorem for smoothness arguments described before.

Theorem 3.5.1 (Extension Theorem)**.** *For every cost-minimization game G that is (λ, μ)-smooth with respect to every outcome s^* of G, for every coarse correlated equilibrium σ of G the following holds:*

$$E_{\vec{s}\sim\sigma}[C(\vec{s})] \leq \frac{\lambda}{1-\mu} C(s^*) \tag{3.15}$$

Proof. Let G be an (λ, μ)-smooth cost-minimization game, σ a coarse-correlated equilibrium and s^* an outcome of G. We write:

$$\begin{aligned}
E_{\vec{s}\sim\sigma}[C(\vec{s})] &= E_{\vec{s}\sim\sigma}[\sum_{i=1}^{N} c_i(\vec{s})] && (3.16)\\
&= \sum_{i=1}^{N} E_{\vec{s}\sim\sigma}[c_i(\vec{s})] && (3.17)\\
&\leq \sum_{i=1}^{N} E_{\vec{s}\sim\sigma}[c_i(s_i^*, \vec{s}_{-i})] && (3.18)\\
&= E_{\vec{s}\sim\sigma} \sum_{i=1}^{N} [c_i(s_i^*, \vec{s}_{-i})] && (3.19)\\
&\leq E_{\vec{s}\sim\sigma}[\lambda C(\vec{s^*}) + \mu C(\vec{s})] && (3.20)\\
&\leq \lambda C(\vec{s^*}) + \mu E_{\vec{s}\sim\sigma}[C(\vec{s})] && (3.21)
\end{aligned}$$

where the equality 3.16 follows from the definition of the objective function, equalities 3.17, 3.19 and 3.21 follow from linearity of expectation, inequality 3.18 follows from the definition of a coarse-correlated equilibrium 3.14 and inequality 3.20 follows from the definition of smooth games. □

As a last part of this chapter, we next prove that if a game is locally (λ, μ)-smooth with respect to an optimal outcome with $\mu \leq 1$, then the expected cost of every correlated equilibrium (and hence every pure and mixed Nash-equilibrium) is at most $\frac{\lambda}{1-\mu}$ times that of an optimal outcome.

Theorem 3.5.2 (Extension Theorem on Local Smoothness)**.** *Let σ be a correlated equilibrium of a cost-minimization game. If the game is locally (λ, μ)-smooth with respect to the outcome $\vec{s}^*$, then*

$$E_{\vec{s}\sim\sigma}[C(\vec{s})] \leq \frac{\lambda}{1-\mu} C(s^*) \tag{3.22}$$

Proof. In the first part of the proof, lets assume that

$$E_{\vec{s}\sim\sigma}[\bigtriangledown_i c_i(s)^T (s_i^* - s_i)] \geq 0, \forall i. \tag{3.23}$$

Bearing this in mind, we can complete the proof as follows.

$$\begin{aligned} & E_{\vec{s}\sim\sigma}[C(\vec{s})] \\ & \leq \sum_{i=1}^{N} E_{\vec{s}\sim\sigma}[c_i(\vec{s}) + \bigtriangledown_i c_i(s)^T (s_i^* - s_i)] \\ & \leq E_{\vec{s}\sim\sigma}[\lambda C(\vec{s}^*) + \mu C(\vec{s})] \end{aligned} \tag{3.24}$$

Now, we only need to prove 3.23. For this, suppose for contradiction that

$$E_{\vec{s}\sim\sigma}[\bigtriangledown_i c_i(s)^T (s_i^* - s_i)] < 0$$

We can define $s_\epsilon := ((1-\epsilon)s_i + \epsilon s_i^*, s_{-i})$. Since stragy sets are convex, s_ϵ is a well-defined stragy for every ϵ between 0 and 1.
In the limit as ϵ goes to zero, $E_{\vec{s}\sim\sigma}[\frac{1}{\epsilon}(c_i(s_\epsilon) - c_i(s))]$ tends to $E_{\vec{s}\sim\sigma}[\bigtriangledown_i c_i(s)^T (s_i^* - s_i)]$
This can be explained by the convergence theorem. Precisely, since cost functions are continuously differentiable with bounded derivatives, there is some $M < \infty$ so that for all $\vec{s} \in S$ we have $|\frac{1}{\epsilon}(c_i(s_\epsilon) - c_i(s))| < M$. Hence, $lim_{\epsilon\to 0} \int \frac{1}{\epsilon}(c_i(s_\epsilon) - c_i(s)) d\sigma(\vec{s}) = \int \bigtriangledown_i c_i(s)^T (s_i^* - s_i) d\sigma(\vec{s})$.
However , $E_{\vec{s}\sim\sigma}[\bigtriangledown_i c_i(s)^T (s_i^* - s_i)]$ is strictly negative by assumption and thus there is a sufficiently small $\epsilon > 0$ such that $E_{\vec{s}\sim\sigma}[c_i(\vec{s}_\epsilon)] < E_{\vec{s}\sim\sigma}[c_i(\vec{s})]$ which contradicts the assumption that σ is a correlated equilibrium.

□

To recapitulate, local-smoothness arguments and smoothness arguments are applied in a wide range of games. In the chapters 5 and 6 we use the local smoothness theory in order to give robust price of anarchy bounds for a class of opinion formation games introduced by Kleinberg and al. [16].

Chapter 4

Introduction to Social Networks

4.1 Modeling Social Networks

Social networks play a central role in the spread of information and the formation of opinions. In fact, there is a nearly limitless set of situations in which people are influenced by others such as, in the opinions they hold, the technologies they use, their political positions, and many others. The countless ways in which people are affected by others, make it critical to understand the process by which new ideas and new behaviors spread through a population.

Precisely, there has been a systematic study of such processes since the second half of the 20th Century. The initial research on this topic was empirical [80, 53, 22, 76], however since the 1970s, economists and mathematical sociologists such as Schelling(1978) and Granovetter(1978), began formulating basic mathematical models for the mechanisms by which ideas and behaviors diffuse through a population. However, while the literature on networks has been thriving in sociology for many decades, it has emerged in economics primarily over the past 10 to 15 years. Its explosion in computer science has been rapid and has taken place mostly during the last decade.

As there is a variety of networks of relationships, there is no single way of modeling social networks that encompasses all applications. Therefore, in this chapter, we present some of the game-theoretic and algorithmic grounding of the area and discuss the basic models. More specifically, in the section 4.2 of the chapter, we present some of the fundamentals of how networks are represented. This will provide us the basic background to discuss different types of models in section 4.3 of the chapter.

4.2 Representation of Social Networks

In this section, we present some of the fundamentals of how networks are represented and the basics for the language of networks. In other words, we focus on some representations that serve us a useful basis for capturing the models we discuss next. This involves basic the representation of the individuals that consist the network as well as the representation of their personal relationships.

Players

When we examine a social network, we study the interactions between selfish individuals that can be either people, countries, firms or organizations. The individuals that are involved in a social network, are represented by a set of nodes $N = \{1, 2, ..., n\}$ and we refer to them as "individuals","players","vertices", "agents" or simply as "nodes", depending on the setting.

Representation of the Social Network

The canonical form of a network is an undirected graph, in which two nodes are either connected or not connected. For such graphs it cannot be that one node is related to another without the second being related to the first. This is a natural representation for many social and economic relationships, such as friendships, alliances and acquaintances. However, there are situations that are better modeled as directed networks, in witch the relationship between two nodes is not necessarily two-way. For instance, a network that keeps track of which web page have links to which others would naturally take the form of a directed graph. This distinction results in some basic differences to the modeling and analysis of the network and therefore different conclusions arise. From now on, the default is that the network is undirected, unless it is mentioned explicitly as 'directed'.

It is also possible for the entries of the graph to take on more than two values, in order to capture the intensity level of relationships. We refer to such graphs as *weighted graphs.* Otherwise, it is standard to use the values of either 0 and 1, and in this case the graphs are called *unweighted.* Next, we give a formal definition of graphs that represent networks.

Definition 7 (Graphs)**.** *A graph (N, g) consists of a set of nodes $N = \{1, ..., n\}$ and a real valued matrix $g_{n \times n}$, where g_{ij} represents the relation between i and j. This matrix is often referred as an adjacency matrix, since it indicates which nodes are adjacent to one another. In most of the cases, N is fixed or given. Therefore, we often refer to g, as being a network or a graph.*

It is useful to mention that there are more general graph structures that can represent the possibility of multiple relationships between different nodes. These graphs are sometimes referred to as a *multiplex networks.* To fix ideas, a social network is represented in the simple graph that follows, with the individuals being the nodes of the graph and their personal relationship being reflected by the edges of the graph.

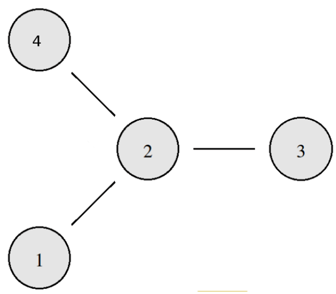

Figure 4.1: A simple network

Next, we define the neighborhood of an individual i and the degree of a node, which are fundamental for our further analysis.

Definition 8 (Neighborhood). *The neighborhood of a node i, is the set of nodes that i is linked to.*

$$N_i = \{j : g_{ij} = 1\}$$

For instance, in the network 4.2, the neighborhood of individual 2 is the set $N_2 = \{1, 3, 4\}$.

Definition 9 (Degree). *The degree of a node i is the number of links that involve this node, which is the cardinality of the node's neighborhood.*

$$d_i = \#\{j : g_{ij} = 1\} = \#N_i$$

It is useful to notice that in the case of a directed graph the above calculation is the node's in-degree.

Now, we are able to discuss about very basic models of decision-making by individuals.

4.3 Models of Network Formation

We will consider specifically on how new behaviors, practices, opinions, conventions, and technologies spread from person to person through a social network, as people influence their friends to adopt new ideas. Our understanding of how this process works is built on a long history of empirical work in sociology known as the *diffusion of innovations* [22, 76, 83].

Precisely, given the role of social networks in the formation of opinions and beliefs, and the subsequent shaping of behaviors, it is critical that we have a thorough understanding of how the structure of such networks affects learning and the diffusion of information. So far, various models have been proposed and in this chapter we explore two different types of models. The first is a *Bayesian learning model*, in which individuals observe actions and results experienced by their neighbors and process the information in a sophisticated manner. Such models provide information about whether individuals in a society come to hold a common belief or remain divided in opinions. The second model is based on a form of updating, in which individuals exchange information with their neighbors over time and then update their opinions by taking a weighted average of what they hear. This type of models is quite tractable and allows us to incorporate rich network structures and develop a deep perspective on how quickly individuals learn and whether initially diverse information scattered through the society can reach a consensus. There is an active line of recent work that considers processes by which a group of people in a social network can arrive at a shared opinion through a form of repeated averaging [29, 3, 51, 46].

For further background on social networks, a great many surveys are available. These sources include the general-interest books by Watts [96], Barabasi [11], and Gladwell [45]; the more technical books of Wasserman and Faust [94], Scott [81], Buchanan [17], Dorogovtsev and Mendes [95] and Watts [96]; and survey articles by Newman [72, 73], Albert and Barabasi [4], Strogatz [84] and Dorogovtsev and Mendes [30].

Bayesian Learning

The main idea behind Bayesian learning models is that an individual who is doing significantly worse than a neighbor must realize this over time and change his strategy in order to do as well as his neighbor. Thus, all connected neighbors must end up with the same limiting payoffs. It can be clearly seen that there are situations where a cascade is developed. This happens when people abandon their own information in favor of inferences based on earlier people's actions. This terminology comes from the work of Banerjee [10] as well as from work by Bikhchandani, Hirshleifer, and Welch [15, 98].

Before we continue, it is useful to remark some interesting observations. First, although individuals may end up with the same limiting payoffs, this does not imply that they all learn to take the best possible action. Furthermore, individuals in a cascade are imitating the behavior of others, but it is not mindless imitation. Rather, it is the result of drawing rational inferences from limited information.

The following learning setting is based on that studied by Bala and Goyal [8]. There are n individuals connected in an undirected social network. In each period $t \in \{1, 2, ...\}$, the individuals simultaneously choose among a finite set of actions. The payoffs to the actions are random and the individuals are all faced with the same set of possible actions. They have the same preferences and face the same form of uncertainty about the actions. At each time step, each individual observes not only his/her outcome, but also his/her neighbors' choices and outcomes.

To illustrate this better, consider a two choices setting, let's say A and B and suppose that A results in a payoff of 1 per period, whereas B pays 2 with probability p and 0 with probability $1 - p$. An individual is willing to maximize the expected sum of discount payoffs,

$$E\left[\sum_t \delta^t \pi_{it}\right]$$

where $\delta \in (0, 1)$ is a discount parameter and π_{it} is the payoff that i receives at time t. Clearly, if $p > \frac{1}{2}$ then every individual would prefer to choose B, while if $p < \frac{1}{2}$ every individual would prefer to choose A. However, p is unknown to the individuals and can take on a finite set of values $p \in \{p_1, ..., p_k\}$. Let individual i begin with a prior μ_i over this set, such that $\mu_i(p_k)$ is the probability that i initially assigns to p_k being the probability that action B pays 2.

The learning in such an environment can be quite complicated. For instance, seeing that a neighbor chooses an action B might indicate that the individual's neighbors have had good outcomes from B in the past. Thus, beyond simply seeing actions and outcomes, an individual can make inferences about outcomes of indirect neighbors by observing the action choices of neighbors. Such full Bayesian learning is explored in the context of three link networks [19, 41], but quickly becomes intractable in larger networks. Instead, Bala and Goyal [9] examined a limited form of Bayesian updating, where individuals only process the information from actions and outcomes and ignore any indirect information that might be gleaned from the action sequences of neighbors.

In this version of the model, observing a 0 or 2 must obviously indicate that an individual took action B and a payoff of 1 indicates that the individual took action A. However, Bala and Goyal [9] proved that with probability 1, there is a certain time after which all individuals settle down to play the same action. Note that this does not imply that the individuals will eventually choose the action with the higher payoff - this depends on the prior and chance.

Before closing our discussion about Bayesian learning it is interesting to present the following variation. Acemoglu and al. [1] extended the framework to consider a social network environment in which individuals have stochastic neighborhoods. Their main result is that asymptotic learning occurs even with bounded beliefs for stochastic topologies such that there is an infinitely growing subset of agents who are probabilistically "well informed" with respect to whom the rest of the agents have expanding observations.

The DeGroot Model

In a network context, DeGroot [28] provides the most influential non-Bayesian framework; individuals observe signals just once and communicate with each other and update their beliefs via a weighted and possibly directed trust matrix.

In the DeGroot model each individual starts with $p_i(0)$ lying in the interval $[0,1]$. Precisely, we write an initial n-dimensional vector of probabilities $p(0) = (p_1(0), ..., p_n(0))$, where each p_i can be thought of as the probability that the individual might engage in a given activity, vote for a specific party and many others. We let T be a matrix indicating the interaction patterns, where every individual updates his/her belief as a weighted sum of the beliefs. Clearly, T is a (row) stochastic matrix, where the entry T_{ij} represents the weight or trust that an individual i places on individual's j belief. Therefore, believes' updates correspond to the following formula.

$$p(t) = Tp(t-1) = T^t p(0)$$

It can be clearly seen that we have a Markov chain here defined by T. Thus, the following theorem holds as a result of Markov chain theory.

Theorem 4.3.1. *T converges for every initial conditions if and only if every set of nodes that is strongly connected and closed is aperiodic.*

However, while convergence is a useful property, it does not imply that a consensus is reached. We say that a group of individuals N reaches a consensus under T for initial belief vector $p(0)$ if $\lim_{t\to\infty} p_i(t) = \lim_{t\to\infty} p_j(t)$ for all $i, j \in N$. The following result indicates that convergence is to a consensus.

Theorem 4.3.2. *Under T, any strongly connected and closed group of agents reaches a consensus for every initial vector of beliefs if and only if it is aperiodic.*

Next, we present a simple example to fix ideas.

Example 4.3.1 (Updates in the DeGroot Model)**.** *Suppose that there are three individuals with their initial vector of beliefs being $p(0) = (1,0,0)$. Let the updating matrix be the following.*

$$T = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix}$$

The corresponding social graph looks as follows.

Clearly, individual 1 puts equal weight (1/3 each) on individuals 2 and 3 when forming his/her belief. Individual 2 weights his/her own belief slightly more, but completely discounts individual 3. Individual 3 meanwhile, puts the most weight/trust in his/her own belief (3/4) and then puts a weight of 1/4 on individual 2. Thus,

$$p(1) = Tp(0) = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix}$$

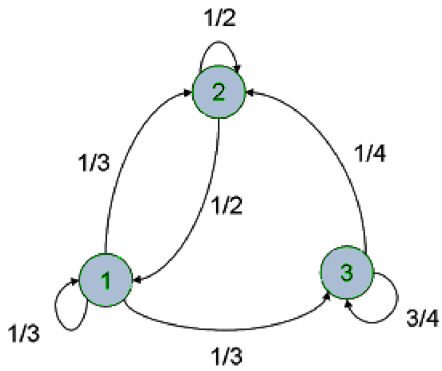

Figure 4.2: Degroot

$$p(2) = Tp(1) = \begin{pmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & 1/2 & 0 \\ 0 & 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1/3 \\ 1/2 \\ 0 \end{pmatrix} = \begin{pmatrix} 5/8 \\ 5/12 \\ 1/8 \end{pmatrix}$$

And iteratively,

$$p(t) = Tp(t-1) = T^t p(0) \rightarrow \begin{pmatrix} 3/11 \\ 3/11 \\ 3/11 \end{pmatrix}$$

Based on the DeGroot model, many variations have been studied in the literature of social networks. Next, we finish our discussion about the DeGroot model by presenting some variations of it. Precisely, DeMarzo, Vayanos and Zwiebel [29] examine a variation of the Degroot model where an individual might place more (or less) weight on his/her own belief over time. This is related to a variation of the DeGroot's model of Chatterjee and Seneta [18]. Furthermore, there have been examined several models in which updating always mixes in some weight on an individual's initial beliefs. These have been studied from Friedkin and Johnsen [36] and Kleinberg and al. [16]. It is useful to notice that in such models consensus may never be reached. Last but not least, Krause [56] and Hegselmann and Krause [49] introduced variations of the DeGroot model in which an individual pays attention only to other individuals whose beliefs do not differ much from his own.

So far, we discussed about social networks and several concepts on algorithmic game theory. In the next chapter we analyze in depth several properties of the variation of the DeGroot model that introduced by Kleinberg and al. [16]. Finally, in the last chapter of this thesis, we propose and analyze a generalization of this model that extends the setting in many dimensions.

Chapter 5

Kleinberg-Bindel Model

5.1 Introduction

When people are connected to a network, it becomes possible for them to influence each others' behavior and opinion. For example, a discussion with friends can influence someone's opinion on which movies to see, on a possible holiday destination or even on the political situation of the country.

Since the early 40's, understanding such interactions and predicting how specific opinions spread throughout social networks has triggered vast research by economists, sociologist, psychologies and physicists [80, 53, 22, 76]. However, algorithmic game theory provides a natural framework for modeling such interactions between individuals and the networks they generate.

In this direction, in order to study how consensus is formed when individual opinions are updated, using the average of the neighborhood, DeGroot [28] introduced the following model. A set of individuals are connected by a social network. Each individual has a prior belief on a subject that he/she updates over time by averaging his/her opinion and his/her neighbors' opinions. DeGroot's model triggered much research in the opinion formation games, although it is very naive in the representation of individuals' interactions. More precisely, despite the influention of others, people tend to be more persistent to their initial opinion. Furthermore, they tend to evaluate others' opinions, being more influenced by people they appreciate or people they are closer with. These two remarks have been incorporated in the model introduced by Kleinberg and al. [16]. This model is a variation of DeGroot's model and can be described as follows.

In Kleinberg's model, individuals (represented by nodes) are connected in a weighted undirected graph $G(V, E)$, as it is assumed that the relationships between two individuals are symmetric. The neighbors of an individual are the persons he/she interacts with, while the extent of a neighbor's influence is reflected by the edge weight $w_{i,j} = w_{j,i}$. In contrast to the DeGroot model, each individual i, holds a persistent internal opinion equal to a real number s_i, which might for example represent a position on a political spectrum, or a probability that i assigns to a certain belief. In addition to his/her internal opinion, each individual holds an external expressed opinion z_i, that he/she updates at each time step, influenced by his/her neighbor's external opinions and his/her initial internal opinion.

This behavior can be expressed by a cost function that each individual holds and tries to minimize. In other words, an individual is continually prepared to improve his/her personal cost in response to changes made by his/her neighbors. Quantifying the above, an individual updates his/her external opinion z_i at each time step, in order to minimize the following quadratic function:

$$c_i = (z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j}\,(z_i - z_j)^2$$

As a result, each individual updates each of his/her opinions to the average:

$$z_i = \frac{s_i}{1 + \sum_{j \in N(i)} w_{i,j}} + \frac{\sum_{j \in N(i)} w_{ij} z_j}{1 + \sum_{j \in N(i)} w_{i,j}}$$

We refer to this process as *best response dynamics*, as each individual's updates are based on his/her best response to the current situation.

In the next paragraphs of this chapter, we analyze in depth the properties of this model; In section 5.2 we examine whether an equilibrium exists for the multi-dimensional opinion formation game. The positive results produce some further questions about the characterization, the convergence time and the optimality of the equilibrium which are discussed in the sections 5.3, 5.4 and 5.5 respectively.

5.2 Existence of a Nash Equilibrium

We are particularly interested in stable solutions, where the best response of each agent is to stay put. Thus, a very natural question rises from the above: Does this process of repeating updates ever stops? In chapter 2, we defined the Nash Equilibrium as a situation where no individual has an incentive to deviate. However, as described above, individuals act selfishly in order to minimize their personal cost, so there is no evidence that these updates end up to an equilibrium.

Moreover, as the repeating updates of an individual's opinion affect society, it is useful now to remind the definition of *social optimum* as introduced in 2.5.2. We say then, that a solution is social optimum if it minimizes the total cost to all individuals. One could understand this solution as if there was a supervisor who views all agents as equals and thus evaluates the quality of the solution by summing up their personal costs.

At this point, one could easily understand that there is no proof that an equilibrium on our game even exists. Hence, in the rest of this section we focus on a proof of an equilibrium existence.

As the optimal solution for our game is the one that minimizes the total cost to all individuals, it would be ideal if we could prove that the social optimum is a stable state for our game. Unfortunately, as the following example illustrates, social optimum could be an unstable solution for the opinion formation game.

Example 5.2.1. *Consider the simple instance of a two node path, as seen in the figure below, in which the individuals have internal opinions 0 and 1 respectively and the weight of their connecting edge is $\frac{1}{2}$. One can easily prove that social optimum is $SC(\vec{y}) = \frac{3}{9}$, which is minimized by the following vector of opinions: $\vec{y} = (\frac{1}{3}, \frac{2}{3})$. However, both individuals would deviate from these opinions, as they do not minimize individuals' personal costs.*

(a) *internal opinions* (b) *optimal solution*

Figure 5.1: A two-node example

The previous example illustrates that the Nash Equilibrium, if exists, could be worse than the social optimum solution. Thus, we must not confuse the social optimum with the global minimum of the game, as social optimum does not necessarily decrease when an individual updates his/her opinion. The next step then, is to try to formulate a function that strictly decreases when an individual reduces his/her personal cost.

In chapter 2, we discussed about potential games, where the incentive of all players to change their strategy can be expressed using a single global function called the potential function. In potential games, the existence of a Pure Nash Equilibrium is reduced to the existence of a local minimum of the potential function. According to the above analysis, our intuition suggests that the potential function should have the following property:

- If an individual i reduces his cost by $c_i - c_i'$, the same reduction should be made in the value of the potential function.

Following this property, we calculate formally the potential function of our game.

Theorem 5.2.1. *The opinion formation game described above, is potential with the following potential function:*

$$P = \sum_i (z_i - s_i)^2 + \sum_i \sum_{j \in N(i)} \frac{1}{2} w_{i,j} (z_i - z_j)^2$$

Proof. If P is the potential function, since the incentives of all players are mapped into one function the following must be applied:

$$\frac{\partial P(z_i, z_{-i})}{\partial z_i} = \frac{\partial c(z_i)}{\partial z_i} = 2z_i - 2s_i + \sum_{j \in N(i)} 2w_{i,j}(z_i - z_j), \forall i \in V$$

For the rest of the proof, we consider the complete graph, where $w_{ij} = 0$ if nodes i and j are not linked in the initial graph. Then, for $i = 1$,

$$\begin{aligned} \frac{\partial P(z_1, z_{-1})}{\partial z_1} &= 2z_1 - 2s_1 + \sum_{j \neq 1} 2w_{1,j}(z_1 - z_j) \implies \\ P(z_1, z_{-1}) &= {z_1}^2 - 2z_1 s_1 + \sum_{j \neq 1} 2w_{1,j}(z_1 - z_j) + C(z_{-1}) \end{aligned} \tag{5.1}$$

Furthermore:

$$\frac{\partial P(z_1, z_{-1})}{\partial z_2} = 2z_2 - 2s_2 + \sum_{j \neq 2} 2w_{2,j}(z_2 - z_j) \tag{5.2}$$

By (5.1),(5.2), it is implied that

$$C(z_{-1}) = {z_2}^2 - 2z_2 s_2 + \sum_{j \neq 1,2} w_{2,j}(z_2^2 - 2z_2 z_j) + C(z_{-\{1,2\}})$$

And by iterating this process $\forall i \in V$, we have

$$\begin{aligned} P(z_i, z_{-i}) &= \sum_i ({z_i}^2 - 2z_i s_i) + \sum_i \sum_{j \neq i} \frac{1}{2} w_{i,j} (z_i - z_j)^2 + C \\ &= \sum_i \left((z_i - s_i)^2 - {s_i}^2 + \sum_{j \in N(i)} (z_i - z_j)^2 \right) + C \end{aligned}$$

For $C = \sum_i {s_i}^2$,

$$P(z_i, z_{-i}) = \sum_i (z_i - s_i)^2 + \sum_i \sum_{j \in N(i)} \frac{1}{2} w_{i,j} (z_i - z_j)^2$$

Which is a potential function, since $\forall i \in V$,

$$\begin{aligned} P(z_i, z_{-i}) - P(z_i', z_{-i}) &= (z_i - s_i)^2 + \sum_{j \in N(i)} (z_i - z_j)^2 \\ &- (z_i' - s_i)^2 - \sum_{j \in N(i)} (z_i' - z_j)^2 \\ &= c_i(z_i, z_{-i}) - c_i(z_i', z_{-i}) \end{aligned}$$

□

At this point, we have proven that our game is potential and as we already pointed out, the set of minimizers of the potential function refines the set of Nash equilibria. But can this set be empty for the opinion formation game? The potential function we computed before, is a continuous and bounded function. Thus, a global minimum exists and our game has at least one Nash Equilibrium. Furthermore, one can easily prove that it is also a convex function and thus there are no local minima for the opinion formation game.

Corollary 5.2.1. *The opinion formation game has a unique Pure Nash Equilibrium.*

5.3 Characterization of the Equilibrium

As proven in the previous section, the opinion formation game has a unique pure Nash Equilibrium to which it converges. However, beyond knowing whether individuals' opinions converge to a stable solution, we are interested in characterizing the stable solution. In fact, we want to ascertain how each individual in the social network influences the equilibrium solution, as well as how his/her opinion in the stable state is affected by previous external opinions of his/her neighbors.

This context has been examined in depth by Ghaderi and al. [42]. Regarding the simple version of the Degroot model [28], where the individuals do not persist to their initial opinions, dynamics converge to consensus and the convergence issues are already well understood in the context of consensus and distributed averaging [75, 52, 33, 62, 85, 58, 90, 91].

Before examining the equilibrium state formally, lets figure out the properties of such a stable solution. First, our intuition suggests that the equilibrium state is influenced by the weights that individuals put one another. To illustrate this, consider the trivial example where each individual puts zero weights on his/her neighbors' opinions. In this case, the game is always in a stable state, as individuals have no incentive to deviate from their initial internal opinions. At the other side of the spectrum, as we already discussed, if individuals do not put any weight on their initial opinions, they will finally reach a consensus from which no one will want to deviate. By this, we understand that the more weight the individuals put each other and to their initial opinions, the more non-trivial is the convergence to a stable solution. Next we prove this intuition formally.

At this point we remind that an individual i updates his/her opinion to the average:

$$z_i = \frac{s_i}{1 + \sum_{j \in N(i)} w_{i,j}} + \frac{\sum_{j \in N(i)} w_{ij} z_j}{1 + \sum_{j \in N(i)} w_{i,j}}$$

For the rest of the proof, in order to have a global aspect of the network, we will write the opinions in a matrix form. For this, we define a matrix $A_{n \times n}$, such that:

$$A_{ij} = \begin{cases} \frac{w_{i,j}}{1+\sum_{j \in N(i)} w_{i,j}} & \forall (i,j) \in E \\ 0 & \text{otherwise} \end{cases}$$

One can easily verify that matrix $A_{n \times n}$, is an irreducible sub-stochastic matrix with the row-sum of the elements of each row being less than one.

Furthermore, we define a matrix $B_{n \times n}$ as follows:

$$B_{ij} = \begin{cases} \frac{1}{1+\sum_{j \in N(i)} w_{i,j}} & \text{if i=j} \\ 0 & \text{if } i \neq j \end{cases}$$

Thus, we can write the individual opinions in a matrix form:

$$\vec{z}(t+1) = A\vec{z}(t) + B\vec{s} \tag{5.3}$$

And by iterating the above equation, arises the following formula:

$$\vec{z}(t+1) = A^t \vec{s} + \sum_{k=0}^{t-1} A^k B \vec{s}$$

One can keep track on the equilibrium in which the game converge, by increasing the above formula to its limit as $t \to \infty$. The results are described in the following theorem.

Theorem 5.3.1. *The pure Nash Equilibrium of the opinion formation game is:*

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} A^k B \vec{s} = (I - A)^{-1} B \vec{s}$$

Proof. Let $\rho_1(A) = max_i|\lambda_i(A)|$ denote the spectral radius of A. By the Peron-Ferobenius theorem, the largest eigenvalue of A should be real $1 > \lambda_1 > 0$ and $\rho_1(A) = \lambda_1$. Hence, $\lim_{t \to \infty} A^t = 0$ and

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} A^k B \vec{s} = (I + A^1 + ...) B \vec{s} = (I - A)^{-1} B \vec{s}$$

□

However, the above form does not give any insight on how the equilibrium depends on the graph structure. Therefore, next we describe the equilibrium in terms of explicit quantities that depend on the graph structure. For this, we define the graph $\hat{G}(\hat{V}, \hat{E})$ as follows: on the initial graph $G(V, E)$, we connect a new vertex v_i to each $i \in V$ with a weight of 1, corresponding to the internal opinion of i.

As a result, $\hat{G}(\hat{V}, \hat{E})$ is an undirected weighted graph with weights

$$w_{i,j},\ \forall (i,j) \in E \text{ and } w_{i,u_i} = 1,\ \forall (i, u_i) \in \hat{E} \setminus E$$

Next, we define $w_i = \sum_{(i,j)\in\hat{E}} w_{i,j}$ as the weighted degree of node i. By this, we have

$$w_i = \begin{cases} \sum_{(i,j)\in E} w_{i,j} + 1 & \text{if } i \in V \\ 1 & \text{if } i \in \hat{V} \setminus V \end{cases}$$

Lets consider now the random walk $Y(t)$ over $\hat{G}$ where the probability of transition from vertex i to vertex j is $P_{ij} = \frac{w_{ij}}{w_i}$. Assuming that the walk starts from some initial vertex $Y(0) = i \in V$, we define τ_j as the first hitting time to vertex j, for any $j \in \hat{V}$. Next follows a theorem that presents a different aspect from the one discussed above for the equilibrium state that the opinion formation game converges.

Theorem 5.3.2. *Based on the random walk over the graph $\hat{G}$, the equilibrium state where the opinion formation game ends is:*

$$z_i(\infty) = \sum_{j\in V} P_i(\tau = \tau_{u_j}) \times \vec{s}$$

for all $i \in V$, where $P_i(\tau = \tau_k), k \in \hat{V} \setminus V$, is the probability that the random walk hits vertex k first among vertices in $\hat{V} \setminus V$, given that the random walk starts from i.

This theorem offers a very simple perspective of the network, as it treats an individual's internal opinion like an additional neighbor. Thus, if an individual i puts more weight to his/her internal opinion, the probability to hit first the corresponding u_i node is very high. Furthermore, the more weight an individual i puts on a neighbor j, the higher the probability hitting first the corresponding u_j node. Next, we construct a proof for the above.

Proof. The transition probability matrix of the random walk defined above is given by

$$P = \begin{pmatrix} A_{n\times n} & B_{n\times n} \\ I_{n\times n} & 0 \end{pmatrix}$$

Apparently, when the walk reaches u_i, it returns to its corresponding individual i with probability 1. Nonzero elements of A correspond to transitions between vertices of V. Elements of B correspond to transitions from an individual to his/her corresponding u_i. For each vertex $i \in V$, let $F_{ij} = P_i(\tau = \tau_{u_j})$ be the probability that the random walk hits u_j first, among vertices, given that the random walk starts from i. Then, we can write the following recursive formula for F_{ij} probabilities:

$$F_{ij} = B_{ij} + \sum_{k\in V} A_{ik} F_{kj}$$

Hence, writing the above in a matrix form,

$$\vec{F} = B + A\vec{F} \implies \vec{F} = (I - A)^{-1} B$$

and the equilibrium at each node $i \in V$ is a convex combination of initial opinions of individuals, where $z_i(\infty) = \sum_{j\in V} F_{ij} s_j$. $\square$

5.4 Convergence time

Until now, we have proven that the opinion formation game has a unique Nash Equilibrium given by the following formula:

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} A^k B\vec{s} = (I - A)^{-1} B\vec{s}$$

However, while eventual convergence to the equilibrium is important, it is also important to know how quickly individuals reach this stable solution. In other words, we need to answer the following question: How many rounds of opinion updates are needed in order to reach the equilibrium?

At this point, there is no evidence that the opinion formation game converges in a reasonable time. In fact, we know that the potential function decreases with a decrease in an individual's personal cost, but when individuals update their opinions simultaneously the potential function does not necessarily decrease. In the rest of this section, we will examine how quickly do the individuals reach the equilibrium, by using elements of Linear Algebra.

First, in order to develop the basic intuitions and begin quantifying the speed of convergence, it is useful to start with the case of two individuals, as that is particularly transparent and actually provides the basis for a general analysis. Our intuition suggests that the speed of the convergence depends on the weight they put each other. An abstract proof for this observation is that, the more heavy is their link, the more they try to fit their opinions closer at each time step, so the convergence will be fast. Conversely, if each individual pays little attention to the other, the other's updates will occur slight changes to his/her opinion and the convergence will be relatively slower.

To develop these ideas more formally, we need to know how the vector of opinions $\vec{z}(t)$, differ from the opinions in stable time $\vec{z}(\infty)$, at each time step t. For this let $\vec{e}(t) = \vec{z}(t) - \vec{z}(\infty)$ be the error vector. Confirming our previous observations, the error vector depends only from array A as $e(t) = Ae(t-1)$.

Proof. From the definition of $e(t)$, we have:

$$\begin{aligned} e(t) &= A^t\vec{s} + \sum_{k=0}^{t-1} A^k B\vec{s} - \sum_{k=0}^{\infty} A^k B\vec{s} \\ &= A^k\vec{s} - \sum_{k=t}^{\infty} A^k B\vec{s} \\ &= A^t(\vec{s} - \sum_{k=0}^{\infty} A^k B\vec{s}) \\ &= Ae(t-1) \end{aligned}$$

□

Now, in the case of two connected individuals the array A has the following form:

$$A = \begin{pmatrix} 0 & \frac{w_{ij}}{1+w_{ij}} \\ \frac{w_{ij}}{1+w_{ij}} & 0 \end{pmatrix}$$

In order to see how A^t evolves over time, it is useful to rewrite A using what is known as its diagonal decomposition. Since A is symmetric, its eigenvalues are real, it is diagonalizable, and the sets of right and left eigenvectors are the same. Let u be the matrix of left-hand eigenvectors of A and Λ be the diagonal matrix with the eigenvalues of A on its diagonal. Thus, we know that

$$uA = \Lambda u$$

Furthermore, as we know the matrix u^{-1} is the right-hand eigenvector of A. Thus, we also have the following:

$$Au^{-1} = u^{-1}\Lambda$$

and finally we can write A on its diagonal decomposition

$$A = u^{-1}\Lambda u$$

It follows that the powers of A can be written as:

$$A^2 = u^{-1}\Lambda u u^{-1}\Lambda u = u^{-1}\Lambda^2 u$$

$$A^3 = u^{-1}\Lambda u u^{-1}\Lambda^2 u = u^{-1}\Lambda^3 u$$

and generally

$$A^t = u^{-1}\Lambda^t u.$$

So the convergence of A and thus the convergence of the two-individuals formation game, is directly related to the convergence of Λ.

The above analysis helps us develop an intuition for the converge on more complex networks where multiple individuals exist and prove the main theorem of this section. Let $\tilde{\pi} = [\frac{w_i}{Z} : i \in V]^T$ for the weights w_i and Z be the normalizing constant such that $\sum_{i \in V} \tilde{\pi}_i = 1$. Then the following theorem exists.

Theorem 5.4.1. *For the error vector the following inequality exists:*

$$||\tilde{e}(t)||_{\tilde{\pi}} \leq (\lambda_A)^t ||\tilde{e}(0)||_{\tilde{\pi}}$$

where λ_A is the largest eigenvalue of A.

Proof. From the previous analysis, we know that for the error vector holds the following equation:

$$e(t) = Ae(t-1)$$

Consider the transition probability matrix of the random walk we defined previously on this chapter.

$$P = \begin{pmatrix} A_{n \times n} & B_{n \times n} \\ I_{n \times n} & 0 \end{pmatrix}$$

By definition of reversibility, one can easily check that the pair (P,π) is reversible, as:

$$\pi(i)P_{ij} = \pi(j)P_{ji}$$

This implies in particular, that $\pi(i)$, is a stationary distribution of P, the unique one as P is irreducible.
Furthermore, the following equation also holds $\forall i, j \in V$:

$$\pi(i)A_{ij} = \pi(j)A_{ji}$$

Then, by minor abuse of terminology, we would also call A reversible with respect to the distribution $\tilde{\pi} = [\frac{w_i}{\pi(A)} : i \in V]^T$, where $\pi(A)$ is the normalization constant. We can also express the reversibility of A, by saying that:

$$A^* = D^{\frac{1}{2}} A D^{-\frac{1}{2}}$$

is a symmetric matrix, where D is the diagonal matrix with elements of $\tilde{\pi} = [\frac{w_i}{\pi(A)} : i \in V]^T$ as diagonal entries, $D = diag\{\pi(1), \pi(2), ..\}$.

Since A^* is symmetric, its eigenvalues are real, it is diagonalizable, and the sets of right and left eigenvectors are the same.

We now introduce norms on $\mathbb{R}^n$ that are linked to the stationary distribution. Let $l^2(\pi)$ be the real vector space $\mathbb{R}^n$ endowed with the scalar product:

$$\langle x, y \rangle_\pi = \sum_i x(i) y(i) \pi(i)$$

and the corresponding norm

$$||x||_\pi = \sum_i x(i)^2 \pi(i)$$

Choosing a set of right (and therefore left) eigenvectors of A^*, $w_1, w_2, ..., w_n$, we have an orthonormal basis in $\mathbb{R}^r$. Defining u and v as follows:

$$w = D^{-\frac{1}{2}} u \text{ and } w = D^{\frac{1}{2}} v$$

we have

$$u = Dv$$

The matrices A and A^* have the same eigenvalues, and moreover, v is a right eigenvector of A and u a left eigenvector of A. Orthonormality of w_i's is with respect to the usual Euclidean Norm and is equivalent to orthonormality in $l^2(\pi)$ of $\{v_1, v_2, ..v_n\}$. That is:

$$\langle v_i, v_j \rangle_\pi = \delta_{ij}$$

Following the same argument:

$$\langle u_i, u_j \rangle_{\frac{1}{\pi}} = \delta_{ij}$$

Using $\{v_1, v_2, ...v_n\}$ as a basis for $\mathbb{R}^n$, every vector in $\mathbb{R}^n$ can be written as $\vec{x} = \sum_i a_i v_i$ and taking the scalar product in $l^2(\pi)$ with v_j gives $\langle x, v_j \rangle_\pi = a_j$.
Thus, we can write:

$$\vec{e}(t) = \sum_{i=1}^{n} \langle \vec{e}(t), v_i \rangle_{\tilde{\pi}} v_i$$

As $Av_i = \lambda_i v_i$,

$$A\vec{e}(t) = \sum_{i=1}^{n} \lambda_i \langle \vec{e}(t), v_i \rangle_{\tilde{\pi}} v_i$$

This rate will be governed by the largest eigenvalue, as the others will converge more quickly:

$$\begin{aligned} ||\vec{e}(t+1)||^2_{\tilde{\pi}} &= \sum_i \lambda_i^2 \langle \vec{e}(t), v_i \rangle^2_{\tilde{\pi}} ||v_i||^2_{\tilde{\pi}} \\ &= \sum_i \lambda_i^2 \langle \vec{e}(t), v_i \rangle^2_{\tilde{\pi}} \\ &\leq \lambda_A^2 \sum_i \langle \vec{e}(t), v_i \rangle^2_{\tilde{\pi}} \\ &= \lambda_A^2 ||\vec{e}(t)||^2_{\tilde{\pi}} \end{aligned}$$

Thus,

$$||\vec{e}(t+1)||_{\tilde{\pi}} \leq \lambda_A ||\vec{e}(t)||_{\tilde{\pi}}$$

and

$$||\vec{e}(t)||_{\tilde{\pi}} \leq \lambda_A^t ||\vec{e}(0)||_{\tilde{\pi}}$$

□

It is useful to remind at this point that as A is a substochastic matrix, $|\lambda_A| < 1$. In order to define the time that the opinion formation game converges, we fix a bound v arbitrarily close to 0. When the value of the error vector is at most equal to v, we say that the game is converged.

More formally, the convergence time is defined as $\tau(\mathrm{v}) = inf\{t \geq 0 : ||\vec{e}(t)||_{\tilde{\pi}} \leq \mathrm{v}\}$ for some fixed v. Thus, by the previous theorem we have:

$$\left(\frac{1}{1-\lambda_A} - 1\right) \leq \frac{\tau(\mathrm{v})}{log(||\vec{e}(t)||_{\tilde{\pi}}/\mathrm{v})} \leq \frac{1}{1-\lambda_A}$$

and as n grows $\tau(\mathrm{v}) = \Theta\left(\frac{1}{1-\lambda_A}\right)$.

Corollary 5.4.1. *The convergence time on the opinion formation game is $\tau(v) = \Theta\left(\frac{1}{1-\lambda_A}\right)$ where λ_A is the largest eigenvalue of A.*

5.5 Bounding the Price of Anarchy

The previous sections of the chapter wrap up certain aspects of the problem very neatly, although we did not estimate yet how "good" is the equilibrium that the game converges. As we have seen by now, in the opinion formation game a set of individuals explore the "landscape" of solutions in order to minimize their personal costs, until they reach the global minimum, the unique Nash equilibrium of the game.

However, the updates on individuals' opinions do not necessarily improve the social welfare. As example 6.1.1 illustrated, the Nash Equilibrium total cost can be worse than that of the social optimum. But how much worse?

In chapter 3, we defined the Price of Anarchy to be the ratio of the cost of the best Nash Equilibrium solution to the cost of the social optimum. This quantity reflects the blow-up in cost that we incur due to the requirement that our solution must be stable in the face of the individuals' self-interested behaviors. In the rest of the chapter we try to quantify the loss of efficiency resulting from selfish and uncoordinated behavior of the individuals.

Since we proved that the opinion formation game is potential, one could propose the comparison between the potential function of the game and the social optimum in order to measure the Price of Anarchy. More precisely, as the potential function differs from the social cost function, their minima will be different. Next, we present a simple upper bound of 2 on the Price of Anarchy based on these observations.

Theorem 5.5.1. *On the opinion formation game, $PoA \leq 2$.*

Proof. If we define $\vec{x}$ as the vector of opinions in the Nash Equilibrium and $\vec{y}$ as the vector of the opinions in the social optimum solution, we can write the following bound for PoA:

$$PoA = \frac{SC(\vec{x})}{SC(\vec{y})} \leq 2\frac{P(\vec{x})}{P(\vec{y})} \leq 2\frac{P(\vec{y})}{P(\vec{y})} = 2$$

as $P(\vec{x}) \leq SC(\vec{x}) \leq 2P(\vec{x}), \forall x$.

□

In order to find more robust bounds of the Price of Anarchy, next we use a novel "local smoothness" proof framework that we discussed in 3.4, as developed by Roughgarden and al. [79]. Roughly, smoothness controls the cost of a set of "one-dimensional perturbations" of an outcome, as a function of both the initial outcome and the perturbations. Intuitively, the opinion formation game is smooth if the objective function value of the pure Nash equilibrium $\vec{z}$ can be bounded using the following minimal recipe:

1. Let $\vec{o}$ be the social optimum outcome of the game
2. Invoke the Nash equilibrium hypothesis once per individual, to derive that each individual i's cost in the Nash equilibrium $\vec{z}$ is at most as high as if he played $\vec{o}_i$ instead.
3. Use the inequalities of the previous step, to prove that the objective function value of $\vec{z}$ is at least some fraction of that of $\vec{o}$.

The local smoothness technique requires the inequalities only for nearby pairs of outcomes, rather than for all pairs of outcomes. Thus, to show that the game is (λ, μ)-smooth, we need to prove that for a fixed profile $\vec{o}$, the following inequality holds :

$$\sum_{i \in V} c_i(z_i, \vec{z}_{-i}) + (o_i - z_i)\frac{\partial c_i(z_i, \vec{z}_{-i})}{\partial z_i} \leq \lambda C(\vec{o}) + \mu C(\vec{z}) \tag{5.4}$$

At this point we remind that if a cost minimization game is (λ, μ)-smooth with respect to an outcome $\vec{o}$, every pure-strategy Nash equilibrium $\vec{z}$ has objective function value at most $\frac{\lambda}{1+\mu}$ times that of $\vec{o}$.
We now instantiate the local smoothness framework, proving that the opinion formation game is $(\frac{3}{4}, \frac{1}{3})$-smooth and that the Price of Anarchy is at most $\frac{9}{8}$.

In order to use local smoothness technique, the following conditions must apply:

1. Every individual's opinion is a real number ($\forall i \in V, z_i \in \mathbb{R}$).
2. The individual cost functions must be continuously differentiable.

These two requirements are satisfied in our game. To show that the game is $(\frac{3}{4}, \frac{1}{3})$-smooth, we need to prove that for a fixed profile $\vec{o}$, the following inequality holds.

$$\sum_{i \in V} c_i(z_i, \vec{z}_{-i}) + (o_i - z_i)\frac{\partial c_i(z_i, \vec{z}_{-i})}{\partial z_i} \leq \frac{3}{4}C(\vec{o}) + \frac{1}{3}C(\vec{z}) \tag{5.5}$$

Proof. In order to prove the inequality 5.5, we have to show that for every $i \in V$ the following holds.

$$\begin{aligned}(z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (z_i - z_j)^2 + (o_i - z_i) \left((2z_i - 2s_i) + \sum_{j \in N(i)} 2w_{ij}(z_i - z_j) \right) \leq \\ \frac{3}{4} \left((o_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (o_i - o_j)^2 \right) + \frac{1}{3} \left((z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (z_i - z_j)^2 \right)\end{aligned} \tag{5.6}$$

Thus, we have to prove the following two inequalities:

$$(z_i - s_i)^2 + (o_i - z_i)(2z_i - 2s_i) \leq \frac{3}{4}(o_i - s_i)^2 + \frac{1}{3}(z_i - s_i)^2 \tag{5.7}$$

and

$$\begin{aligned}\sum_{j \in N(i)} w_{i,j}(z_i - z_j)^2 + (o_i - z_i) \sum_{j \in N(i)} 2w_{ij}(z_i - z_j) \leq \\ \frac{3}{4} \sum_{j \in N(i)} w_{i,j}(o_i - o_j)^2 + \frac{1}{3} \sum_{j \in N(i)} w_{i,j}(z_i - z_j)^2\end{aligned} \tag{5.8}$$

5.7 inequality

Let $x = (z_i - s_i)$ and $y = (o_i - z_i)$. Then we can write inequality 5.7 as follows:

$$\begin{aligned}&x^2 + (y - x) \times 2x \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\ &x^2 + 2xy - 2x^2 \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\ &2xy \leq \frac{3}{4}y^2 + \frac{4}{3}x^2 \implies \\ &(\frac{\sqrt{3}}{4}y + \frac{2}{\sqrt{3}}x)^2 \geq 0\end{aligned}$$

which is true for every $x = (z_i - s_i), y = (o_i - z_i) \in \mathbb{R}$

5.8 inequality

In this direction, for the 5.8 inequality, we have to prove that $\forall j \in N(i)$,

$$w_{ij}(z_i - z_j)^2 + (o_i - z_i)w_{ij}2(z_i - z_j) \leq \frac{3}{4}w_{ij}(o_i - o_j)^2 + \frac{1}{3}w_{ij}(z_i - z_j)^2$$

Let $x = (z_i - z_j)$ and $y = (o_i - o_j)$. Furthermore, observing the initial equation we want to prove, one can notice the following: for every pair of $(i, j) \in E$, there will be the following expressions:

$$2(o_i - z_i)w_{ij}(z_i - z_j)$$

and

$$2(o_j - z_j)w_{ij}(z_j - z_i)$$

because of the cost functions c_i and c_j. Hence, by transforming the above, we have the following expressions for every pair of $(i, j) \in E$: $(\frac{o_i}{2} - \frac{o_j}{2} - \frac{z_i}{2} + \frac{z_j}{2})2w_{ij}(z_i - z_j)$ and

$(\frac{o_j}{2} - \frac{o_i}{2} - \frac{z_j}{2} + \frac{z_i}{2})2w_{ij}(z_j - z_i)$.
Therefore, we have to prove the following:

$$x^2 + \frac{y-x}{2}2x \leq \frac{3}{4}y^2 + \frac{1}{3}x^2$$
$$x^2 + yx - x^2 \leq \frac{3}{4}y^2 + \frac{1}{3}x^2$$
$$(\frac{1}{\sqrt{3}}x + \frac{\sqrt{3}}{2}y)^2 \geq 0$$

which is true for every $x = (z_i - s_i), y = (o_i - s_i) \in \mathbb{R}$ □

Now that we have proven a bound of $\frac{9}{8}$, can we do better? The following instance gives a clear answer to that question.

Example 5.5.1. *Consider the instance of a three node path, as seen in the figure below, in which the nodes have internal opinions 0, $\frac{1}{2}$ and 1 respectively. All the weights are considered equal to 1.*

1. *Nash Equilibrium:*
   The Nash Equilibrium of the game is when the vector of opinions is $\vec{x} = (\frac{1}{4}, \frac{1}{2}, \frac{3}{4})$. Thus, the cost on the Nash Equilibrium solution is $SC(\vec{x}) = \frac{3}{8}$.

2. *Social Optimum:*
   On the other hand, the vector of opinions that minimize the social optimum function is $\vec{y} = (\frac{1}{3}, \frac{1}{2}, \frac{2}{3})$. Thus, the social optimum cost is $SC(\vec{y}) = \frac{2}{6}$.

Therefore, the Price of Anarchy in this instance of the opinion game is

$$PoA = \frac{SC(\vec{x})}{SC(\vec{y})} = \frac{9}{8}.$$

This proves that the Price of Anarchy bound of $\frac{9}{8}$ is tight for the opinion formation game.

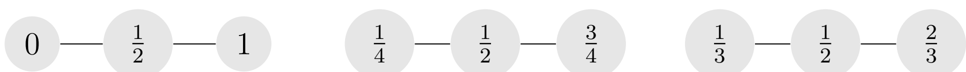

(a) *internal opinions* (b) *Nash Equilibrium* (c) *optimal solution*

Figure 5.2: Comparing Nash equilibrium and Social optimum situation

Corollary 5.5.1. *The opinion formation game has a tight bound of $\frac{9}{8}$ on the Price of Anarchy.*

In this chapter we discussed a variety of issues associated with the formation of opinions in Kleinberg's model. This model and its variants are tractable and powerful as tools for studying interactions between individuals and the form of opinions in many practical networks. By the above analysis, we have proven the unique existence of an equilibrium and have shed light on how the equilibrium behaves. Furthermore, in the last part of the chapter we constructed a tight bound of $\frac{9}{8}$ on the Price of Anarchy.

In closing this chapter, lets offer a couple of thoughts for further development of the theory. In many settings, it is clear that an individual's opinions on related issues are very similar. Thus, when an individual updates an opinion, he/she is influenced not only by his/her initial and his/her neighbors' opinions, but also by his/her opinions on related issues. This is the main setting of the next chapter, where we discuss a variation of Kleinberg's model which captures the above observations.

Chapter 6

Kleinberg-Bindel Model in Many Dimensions

So far, we discussed about Kleinberg's model [16] of opinion formation on social networks. For instance, consider a social network in which each individual has an initial opinion about the political situation of the country. Now, consider the same social network in which individuals hold opinions about a different issue, such as racism. As we analyzed in chapter 5, we can apply the Kleinberg's model in both cases. However, it can be clearly seen that an individual's opinions on the above issues are correlated and influence each other. Assume, for the shake of argument, an individual who's political positions are left-wing. In such cases we expect that his/her opinion about racism should be fairly negative.

The above example illustrates the expected nature of the social networks, which should reflect the interactions between an individual's opinions on different issues. Precisely, there are more than one dimensions that need to be considered, since the opinions on similar issues factor should not be neglected. Therefore, in this chapter we focus on developing a generic, multi-dimensional model that extends the Kleinberg's model, and enables to capture information about interactions between an individual's opinions on different issues.

The multi-dimensional model of the social network presented in this chapter is based on the basic profile of multi-dimensional and dynamic social networks. As we already discussed, the foundation of a social network is a structure made up of individuals, who are tied by a specific type of interdependency, such as friendship, kinship, common interest, financial exchange, and other. In order to represent such entities, the model assumes the representation of nodes and edges, where the nodes represent individuals and the edges represent the interconnections between them.

Furthermore, there are $K = \{1, 2, ...k\}$ issues and each individual i holds a vector of persistent internal opinions $\vec{s_i} = (s_i(1), s_i(2), ..., s_i(k))$, where each opinion $s_i(j)$ corresponds to the issue $j \in K$. In correspondence to the simple Kleinberg's model, each individual in addition to his/her vector of internal opinions, holds an external expressed opinion vector $\vec{z_i} = (z_i(1), z_i(2), ...z_i(k))$, where each opinion $z_i(j)$ corresponds to the issue $j \in K$. However, in contrast to what we discussed before, at each time step an individual i updates each opinion $z_i(j), j \in K$, influenced by his/her neighbor's external opinions on the issue j, his/her initial internal opinion on the issue j, and his/her external opinions on the other issues $j' \neq j$.

At this point, it is useful to remark some interesting observations. As in the case of the simple Kleinberg's model, an individual's neighbors are the persons he/she interacts with, where the extent of a neighbor's influence is reflected by the edge weight. However, such weights may different from one issue to another. For instance, individuals i, u may share a high weight in the issue j, whereas their weight in the issue j' could be fairly small. To illustrate this better, one can imagine two nodes of the social graph i, u and k connections between them, one for each issue. Notice that two individuals may be friends in one dimension, whereas in another they may not.

The aforementioned behavior can be expressed by a cost function that each individual holds and tries to minimize. In other words, an individual is continually prepared to improve his/her personal cost in response to changes made by his/her neighbors and to changes that are made to the other opinions that he/she holds. Quantifying the above, at each time step, an individual updates all his/her external opinions $z_i(j)$, in order to minimize the following quadratic function:

$$c(\vec{z_i}) = \sum_{j\in K} (z_i(j) - avg(z_i))^2 + \sum_{j\in K} (z_i(j) - s_i(j))^2 + \sum_{j\in K} \sum_{u\in N(i)} w^j_{i,u} (z_i(j) - z_u(j))^2 \quad (6.1)$$

where $avg(z_i) = \frac{\sum_{j\in K} z_i(j)}{k}$.

As a result, each individual updates each of his/her opinions to the average:

$$z_i(j) = \frac{\sum_{\substack{j'\neq j \\ j'\in K}} \frac{z_i(j')}{k} + s_i(j) + \sum_{u\in N(i)} w^j_{iu} z_u(j)}{2 + \frac{1}{k} + \sum_{u\in N(i)} w^j_{i,u}} \quad (6.2)$$

Again, we refer to this process as best response dynamics, as each individual's updates are based on his/her best response to the current situation.

In the next paragraphs of this chapter, we analyze in depth the properties of this model. In particular, in section 6.1 we examine whether an equilibrium exists for the opinion formation game. The positive results produce some further questions about the characterization, the convergence time and the optimality of the equilibrium which are discussed in the sections 6.2 and 6.4 respectively.

6.1 Existence of a Nash Equilibrium

As in the simple case of the one-dimensional model, we are interested in stable solutions where each player's best response is to stay put and not update any of his/her opinions. As we described in chapter 2, situations where no individual has an incentive to deviate are called Nash equilibrium situations[70]. However, since individuals act selfishly in order to minimize their personal cost, there is no evidence at this point that a Nash Equilibrium situation exists in the multi-dimensional opinion formation game. Therefore, in this section of the chapter we answer to the most fundamental question of the multi-dimensional game: Does the process of repeating updates on individuals' opinions ever stops?

Again, in this model the definition of the *social optimum* seems ideal to characterize the optimal solution of the game. More precisely, as we introduced in 2.5.2, a solution is social optimum if it minimizes the total cost to all individuals. One could understand this solution as if there was a supervisor who views all agents as equals and thus evaluates the quality of the solution by summing up their personal costs.

Clearly, it would be ideal if we could prove that the social optimum solution and the Nash equilibrium situation coincide. Thus, as a first step we should examine if the social optimum is a stable state for our game. Unfortunately, as the following example illustrates, social optimum could be an unstable solution for the opinion formation game.

Example 6.1.1. *Consider the simple instance of a two node path, as seen in the figure below, in which the individuals have vectors of internal opinions* $(0.5, 1)$ *and* $(1, 1.5)$ *respectively and the weight of their connecting edge is* 1. *One can easily prove that social optimum is* $SC(\vec{y}) = 0.4$, *which is minimized by the following vectors of opinions:* $(0.91, 0.99), (1.09, 1.01)$. *However, both individuals would deviate from these opinions, as they do not minimize individuals' personal costs.*

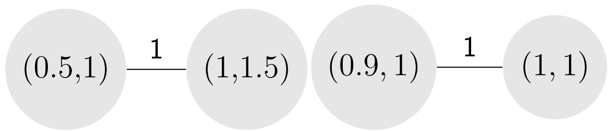

(a) *internal opinions* (b) *optimal solution*

Figure 6.1: A multi-dimensional two-node example

The previous example illustrates that the Nash Equilibrium, if exists, could be worse than the social optimum solution. Thus, we must not confuse the social optimum with the global minimum of the game, as social optimum does not necessarily decrease when an individual updates his/her opinion. The next step then, is to try to formulate a function that strictly decreases when an individual reduces his personal cost.

As we discussed in chapter 2, this is the case of a potential function. Precisely, there is a certain class of strategic games the so-called potential games [67], in which the incentive of all players to change their strategy can be expressed using a single global function called the potential function. Therefore, in potential games, in order to prove the existence of a Pure Nash Equilibrium, we have to prove the existence of a local minimum of the potential function.

According to the above analysis, our intuition suggests that the potential function should have the following property:

- If an individual i reduces his cost by $c_i - c_i'$, the same reduction should be made in the value of the potential function.

Next, with this property in mind, we construct formally a potential function for the game.

Theorem 6.1.1. *The multi-dimensional opinion formation game described above, is potential with the following potential function:*

$$\begin{aligned} P =& \sum_{i \in V} \sum_{j \in K} (z_i(j)^2 - z_i(j) avg(z_i)) + \sum_{i \in V} \sum_{j \in K} (z_i(j) - s_i(j))^2 + \\ & \sum_{i \in V} \sum_{j \in K} \sum_{u \in N(i)} \frac{1}{2} w_{i,u}^j (z_i(j) - z_u(j))^2 \end{aligned} \tag{6.3}$$

Proof. If P is the potential function, since the incentives of all players are mapped into one function the following must be applied:

$$\begin{aligned}
\frac{\partial P(\vec{z}_{-i}, z_i(j), \vec{z}_i(j^-))}{\partial z_i(j)} &= \frac{\partial c(\vec{z}_i)}{\partial z_i(j)} \\
&= (2 - \frac{2}{k} + 2 + \sum_{u \in N(i)} 2w_{i,u}^j) z_i(j) - 2s_i(j) \\
&- \sum_{j' \neq j} \frac{2z_i(j')}{k} - \sum_{u \in N(i)} 2w_{i,u}^j z_u(j), \forall i \in V, \forall j \in K
\end{aligned}$$

Thus,

$$\begin{aligned}
P(\vec{z_{-i}}, z_i(j), \vec{z}_i(j^-)) = \sum_{\substack{j' \neq j \\ j' \in K}} \left(-\frac{2}{k} z_i(j) z_i(j') \right) + (1 - \frac{1}{k}) z_i^2(j) + \left(z_i^2(j) - 2z_i(j)s_i(j) \right) + \\
\sum_{u \in N(i)} w_{i,u}^j \left(z_i^2(j) - 2z_i(j)z_u(j) \right) + C\left(z_i(j^-), z_{-i} \right)
\end{aligned}$$

Then, for $j = 1$,

$$\begin{aligned}
P(\vec{z_{-i}}, z_i(1), z_i(1^-)) = \sum_{\substack{j' \neq 1 \\ j' \in K}} \left(-\frac{2}{k} z_i(1) z_i(j') \right) + (1 - \frac{1}{k}) z_i^2(1) + \left(z_i^2(1) - 2z_i(j)s_i(j) \right) + \\
\sum_{u \in N(i)} w_{i,u}^1 \left(z_i^2(1) - 2z_i(1)z_u(j) \right) + C\left(z_i(1^-), z_{-i} \right)
\end{aligned} \tag{6.4}$$

and

$$\begin{aligned}
\frac{\partial P(\vec{z_{-i}}, z_i(1), \vec{z}_i(1^-))}{\partial z_i(2)} &= \frac{\partial C(z_{-i}, z_i(1^-))}{\partial z_i(2)} - \frac{2}{k} z_i(1) \\
&= (2 - \frac{2}{k} + 2 + \sum_{u \in N(i)} 2w_{i,u}^2) z_i(2) - 2s_i(2) - \\
&\sum_{u \in N(i)} 2w_{i,u}^2 z_u(2) - \sum_{j' \neq 2} \frac{2z_i(j')}{k} \\
&= \frac{\partial C\left(z_i(1^-), z_{-i} \right)}{\partial z_i(2)} - \frac{2}{k} z_i(1)
\end{aligned} \tag{6.5}$$

thus,

$$\begin{aligned}
C(z_{-i}, z_i(1^-)) = \sum_{\substack{j' \neq 1,2 \\ j' \in K}} \left(-\frac{2}{k} z_i(2) z_i(j') \right) + (1 - \frac{1}{k}) z_i^2(2) + \left(z_i^2(2) - 2z_i(2)s_i(2) \right) + \\
\sum_{u \in N(i)} w_{i,u}^2 \left(z_i^2(2) - 2z_i(2)z_u(2) \right) + C\left(z_i(1,2)^-, z_{-i} \right)
\end{aligned} \tag{6.6}$$

and by iterating this process $\forall j \in V$, we end up to the following.

$$\begin{aligned}
P(\vec{z_{-i}}, z_i(j), \vec{z}_i(j^-)) = \sum_{j \in K} (1 - \frac{1}{k}) z_i^2(j) + \sum_{j} \sum_{j' \neq j} \frac{z_i(j) z_i(j')}{k} + \sum_{j \in K} (z_i^2(j) - 2z_i(j)s_i(j)) + \\
\sum_{j \in K} \sum_{u \neq i} w_{i,u}^j (z_i^2(j) - 2z_i(j)z_u(j)) + C(z_{-i})
\end{aligned} \tag{6.7}$$

Now, for $i = 1$, we have

$$P(\vec{z_{-1}}, z_1(j), \vec{z}_1(j^-)) = \sum_{j \in K}(1 - \frac{1}{k})z_1^2(j) - \sum_{j}\sum_{j' \neq j}\frac{z_1(j)z_1(j')}{k} \\ \sum_{j \in K}(z_1^2(j) - 2z_1(j)s_1(j)) + \\ \sum_{j \in K}\sum_{u \neq 1} w_{1,u}^j(z_1^2(j) - 2z_1(j)z_u(j)) + C(z_{-1}) \tag{6.8}$$

and $\forall j \in K$,

$$\frac{\partial P(\vec{z_{-1}}, z_1(j), \vec{z}_1(j^-))}{\partial z_2(j)} = -2w_{1,2}^j z_1(j) + \frac{\partial C(z_{-1})}{z_2(j)} \tag{6.9}$$

thus,

$$(2 - \frac{2}{k} + 2 + \sum_{u \neq 2} 2w_{2,u}^j)z_2(j) - 2s_2(j) - \sum_{u \neq 2} 2w_{2,u}^j z_u(j) - \frac{2z_2(j')}{k} = \\ - 2w_{1,2}^j z_1(j) + \frac{\partial C(z_{-1})}{z_2(j)} \tag{6.10}$$

if we summarize the equation 6.10 $\forall j \in K$,

$$C(z_{-1}) = \sum_{j \in K}(1 - \frac{1}{k})z_2^2(j) - \sum_{j}\sum_{j' \neq j}\frac{z_2(j)z_2(j')}{k} + \sum_{j \in K}\left(z_2^2(j) - 2z_2(j)s_2(j)\right) + \\ \sum_{j \in K}\sum_{u \neq 2} w_{2,u}^j z_2^2(j) - \sum_{j \in K}\sum_{u \neq 1,2} w_{2,u}^j 2z_2(j)z_u(j) + C(z_{-(1,2)}) \tag{6.11}$$

in this direction, for $i = n$,

$$C(z_n) = \sum_{j \in K}(1 - \frac{1}{k})z_n^2(j) - \sum_{j}\sum_{j' \neq j}\frac{z_n(j)z_n(j')}{k} + \sum_{j \in K}\left(z_n^2(j) - 2z_n(j)s_n(j)\right) + \\ \sum_{j \in K}\sum_{u \neq n} w_{n,u}^j z_n^2(j) + c \tag{6.12}$$

and therefore we can write,

$$P(\vec{z_{-i}}, z_i(j), z_i(j^-)) = \sum_{i \in V}\sum_{j \in K}(z_i^2(j)(1 - \frac{1}{k} - \sum_{j' \neq j}\frac{z_i(j')z_i(j')}{k})) + \sum_{i \in V}\sum_{j \in K}(z_i^2(j) - 2z_i(j)s_i(j)) + \\ \sum_{i \in V}\sum_{j \in K}\sum_{u \neq i}\frac{1}{2}w_{i,u}^j(z_i(j) - z_u(j))^2 + c \tag{6.13}$$

for $c = \sum_{i \in V}\sum_{j \in K}(s_i(j))^2$,

$$P(\vec{z}_{-i}, z_i(j), z_i(j^-)) = \sum_{i \in V}\sum_{j \in K}(z_i(j)^2 - z_i(j)avg(z_i)) + \sum_{i \in V}\sum_{j \in K}(z_i(j) - s_i(j))^2 + \\ \sum_{i \in V}\sum_{j \in K}\sum_{u \in N(i)}\frac{1}{2}w_{i,u}^j(z_i(j) - z_u(j))^2 \tag{6.14}$$

Which is a potential function, since $\forall i \in V$,

$$P(z_i(j), z_i(j^-), z_{-i}) - P(z_i'(j), z_i(j^-), z_{-i}) = c_i(z_i(j), z_i(j^-), z_{-i}) - c_i(z_i'(j), z_i(j^-), z_{-i})$$

$\square$

At this point, we have proven that the multi-dimensional version of the opinion formation game is a potential game, which means that the set of minimizers of the potential function refines the set of Nash equilibria. Furthermore, the potential function we constructed, is a continuous and bounded function. Thus, a global minimum exists and our game has at least one Nash Equilibrium. Finally, one can easily prove that it is also a convex function and thus there are no local minima for the opinion formation game. These observations end up to the following corollary.

Corollary 6.1.1. *The multi-dimensional opinion formation game has a unique Pure Nash Equilibrium.*

6.2 Characterization of the Equilibrium

As proven in the previous section, the multi-dimensional opinion formation game has a unique pure Nash Equilibrium to which it converges. However, as in the one-dimensional version, we are interested in characterizing the stable solution of the game. In fact, we want to ascertain how each individual in the social network influences the equilibrium solution, as well as how his/her opinion in the stable state is affected by previous external opinions of his/her neighbors.

In this section, we characterize the equilibrium situation by extending the analysis of Ghaderi and al. [42] to more than one dimensions. Before examining the equilibrium state formally, it is useful to make some remarks. First, our intuition suggests that the equilibrium state is influenced by the weights that individuals put one another. Second, another factor that influences the equilibrium state is the variance of the individuals' opinion vectors. More precisely, if we examine an individual's opinion on an issue j, we can imagine his/her opinions on the other issues as opinions of neighbors on the same issue that share a weight of $1/K$ with him/her. To illustrate this, consider the trivial example where each individual puts zero weights on his/her neighbors' opinions and each individual's initial opinions are the same for every issue, meaning $\forall i \in V, s_i(1) = s_i(2) = s_i(3) = ... = s_i(k)$. In this case, the game is always in a stable state, as individuals have no incentive to deviate from their initial internal opinion vectors. By this, we understand that the more weight the individuals put each other and to their initial opinions as well as the bigger the variance of their opinion vectors, the more non-trivial is the convergence to a stable solution. Next we prove this intuition formally, by using some elements from Linear algebra.

At this point we remind that an individual i updates his/her opinion $z_i(j), j \in K$ to the average:

$$z_i(j) = \frac{\sum_{\substack{j' \neq j \\ j' \in K}} \frac{z_i(j')}{k}}{2 + \frac{1}{k} + \sum_{u \in N(i)} w_{i,u}^j} + \frac{s_i(j)}{2 + \frac{1}{k} + \sum_{u \in N(i)} w_{i,u}^j} + \frac{\sum_{u \in N(i)} w_{iu}^j z_u(j)}{2 + \frac{1}{k} + \sum_{u \in N(i)} w_{i,u}^j}$$

For the rest of the proof, in order to have a global aspect of the network, we will write the opinions in a matrix form. More specifically, in correspondence to the proof of the one-dimensional version of the opinion formation game, we define the vector $\vec{z}$ of length $n \times k$ that encompasses all the opinions of all individuals. In other words, the first k elements of the vector, correspond to the opinion vector of individual 1, $\vec{z}_1$, the following k elements correspond to the opinion vector of individual 2, and so on. Therefore, we can write $\vec{z} = [\vec{z}_1 \vec{z}_2 ... \vec{z}_n]^T$. Furthermore, we define a matrix $A_{nk \times nk}$, such that:

$$A_{iu} = \begin{cases} \frac{w_{iu}^j}{2+\frac{1}{k}+\sum_{u \in N(i)} w_{i,u}^j}, & \text{if } (i \text{ mod } k) = (u \text{ mod } k) \text{ and}(i \text{ div } k, u \text{ div } k) \in E \\ 0 & \text{otherwise} \end{cases}$$

In addition, we define a matrix $B_{nk \times nk}$ as follows.

$$B_{iu} = \begin{cases} \frac{1}{2+\frac{1}{k}+\sum_{u \in N(i)} w_{i,u}^j}, & \text{if } i = u \\ 0 & \text{if } i \neq u \end{cases}$$

Finally, we define the matrix $C_{nk \times nk}$:

$$C_{iu} = \begin{cases} \frac{\frac{1}{k}}{2+\frac{1}{k}+\sum_{u \in N(i)} w_{i,u}^j}, & \text{if } (i \text{ div } k) = (u \text{ div } k) \\ 0 & \text{otherwise} \end{cases}$$

Thus, we can write the individuals' opinions in a matrix form:

$$\begin{aligned} \vec{z}(t+1) &= A\vec{z}(t) + B\vec{s} + C\vec{z}(t) \\ &= (A+C)\vec{z}(t) + B\vec{s} \end{aligned} \tag{6.15}$$

One can easily verify that matrix $(A+C)_{nk \times nk}$, is an irreducible sub-stochastic matrix with the row-sum of the elements of each row being less than one. And by iterating the equation 6.15, arises the following formula:

$$\vec{z}(t+1) = (A+C)^t \vec{s} + \sum_{k=0}^{t-1} (A+C)^k B\vec{s}$$

One can keep track on the equilibrium in which the game converge, by increasing the above formula to its limit as $t \to \infty$. The results are described in the following theorem.

Theorem 6.2.1. *The pure Nash Equilibrium of the multi-dimensional opinion formation game is:*

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} (A+C)^k B\vec{s} = (I - A - C)^{-1} B\vec{s}$$

Proof. Let $\rho_1(A+C) = max_i|\lambda_i(A+C)|$ denote the spectral radius of $(A+C)$. By the Peron-Ferobenius theorem, the largest eigenvalue of $(A+C)$ should be real $1 > \lambda_1 > 0$ and $\rho_1(A+C) = \lambda_1$. Hence, $\lim_{t \to \infty}(A+C)^t = 0$ and

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} (A+C)^k B\vec{s} = (I + (A+C)^1 + ...)B\vec{s} = (I - A - C)^{-1} B\vec{s}$$

□

However, the above form does not give any insight on how the equilibrium depends on the graph structure. Therefore, next we describe the equilibrium in terms of explicit quantities that depend on the graph structure. For this, we define the graph $\hat{G}(\hat{V}, \hat{E})$ as follows:

- on the initial graph $G(V, E)$, we replace each node i with k new nodes $i(j)$, each for every opinion $z_i(j), j \in K$, and we connect each pair of these k new nodes by putting a weight of $1/k$ at each edge.
- furthermore, $\forall (i, u) \in E$, we connect all of their nodes $(i(j), u(j)), \forall j \in K$ and put the corresponding weights, $w^j_{i,u}$ on the edges.
- finally, we connect a new node $v_i(j)$ to each $i(j)$ node with a weight of 1, corresponding to the internal opinion of i on the issue j.

As a result, $\hat{G}(\hat{V}, \hat{E})$ is an undirected weighted graph with weights

$$
\begin{aligned}
&w^j_{i(j),u(j)}, \forall (i, u) \in E, \forall j \in K \\
&w_{i(j),v_i(j)} = 1, \forall (i(j), v_i(j)) \in \hat{E}, j \in K \\
&w_{i(j),i(j')} = 1/k, \forall (i(j), i(j')) \in \hat{E}, j, j' \in K
\end{aligned}
$$

Next, we define $w_i = \sum_{(i,u)\in\hat{E}} w_{i,j}$ as the weighted degree of node i. By this, we have

$$
w_i = \begin{cases} \sum_{(i,u)\in\hat{E}} (2 + \frac{1}{k} + \sum_{u\in N(i)} w^j_{i,u}) & \forall i(j) \text{ node} \\ 1 & \forall v_i(j) \text{ node} \end{cases}
$$

Lets consider now the random walk $Y(t)$ over $\hat{G}$ where the probability of transition from vertex $i(j)$ to vertex $u(j)$ is $P_{i(j)u(j)} = \frac{w^j_{iu}}{w_i}$. Assuming that the walk starts from some initial vertex $Y(0) = i(j), i \in V, j \in K$, we define τ_u as the first hitting time to vertex u, for any $u \in \hat{V}$. Next, follows a theorem that presents a different aspect from the one discussed above for the equilibrium state that the multi-dimensional opinion formation game converges.

Theorem 6.2.2. *Based on the random walk over the graph $\hat{G}$, the equilibrium state that the multi-dimensional opinion formation game ends is:*

$$
z_i(j)|_{t\to\infty} = \sum_{u\in V} P_{i(j)}(\tau = \tau_{v_u(j)}) \times \vec{s} + \sum_{j'\in K} P_{i(j)}(\tau = \tau_{v_i(j')}) \times \vec{s}
$$

for all $i \in V$, where $P_{i(j)}(\tau = \tau_r), r \in \hat{V} \setminus V$, is the probability that the random walk hits vertex r first among vertices in $\hat{V} \setminus V$, given that the random walk starts from $i(j)$.

This theorem offers a very simple perspective of the network. The reason is that if we examine myopically an individual's node $i(j)$, the nodes that correspond to his/her internal opinion $u_{i(j)}$ and to the other opinions that the individual holds, can be treated as neighbors of $i(j)$ node. If an individual i puts more weight to his/her internal opinion, the probability to hit first the corresponding $v_i(j)$ node is very high, while the probability to hit first his/her internal opinion on another issue is $1/k$. Furthermore, the more weight an individual i puts on a neighbor u regarding a certain issue j, the higher the probability hitting first the corresponding $v_{u(j)}$ node. Next, we construct a proof for the above.

Proof. The transition probability matrix of the random walk defined above is given by

$$P = \begin{pmatrix} (A+C)_{kn \times kn} & B_{kn \times kn} \\ I_{kn \times kn} & 0 \end{pmatrix}$$

Apparently, when the walk reaches $v_i(j)$, it returns to its corresponding individual $i(j)$ with probability 1. Nonzero elements of A correspond to transitions between individuals on a specific issue j. Nonzero elements of C correspond to transitions between an individual's opinions on different issues. Elements of B correspond to transitions from an individual to his/her corresponding $v_i(j)$. For each vertex $i \in V$, let $F_{i(j)u(j')} = P_{i(j)}(\tau = \tau_{u(j')})$ be the probability that the random walk hits $u(j')$ first, among vertices, given that the random walk starts from $i(j)$. Then, we can write the following recursive formula for $F_{i(j)u(j')}$ probabilities:

$$F_{i(j)u(j')} = B_{i(j)u(j')} + \sum_{k \in V} A_{i(j)k(j)} F_{ku(j')} + \sum C_{i(j)i(j')} F_{iu(j')}$$

Hence, writing the above in a matrix form,

$$\vec{F} = B + A\vec{F} + C\vec{F} \implies \vec{F} = (I - A - C)^{-1} B$$

and the equilibrium at each node $i \in V$ is a convex combination of initial opinions of individuals on the issues $j \in K$. $\square$

6.3 Convergence Time

Until now, we have proven that the opinion formation game has a unique Nash Equilibrium given by the following formula:

$$\vec{z}(\infty) = \sum_{k=0}^{\infty} (A+C)^k B\vec{s} = (I - A - C)^{-1} B\vec{s}$$

However, while eventual convergence to the equilibrium is important, it is also important to know how quickly individuals reach this stable solution. In other words, again we are interested in bounding the number of rounds of opinion updates that are needed in order to reach the equilibrium.

In fact, in the multi-dimensional version it is even less evident that the game converges in a reasonable time, since at each time step, there can be simultaneous updates from different individuals to the opinions they hold on more than one issue. Thus, although we know that the potential function decreases with a decrease in an individual's personal cost, on simultaneous updates the potential function does not necessarily decrease. For this, in the rest of this section, we will examine how quickly do the individuals reach the equilibrium, by generalizing the proof of the one dimensional opinion formation game to a multi-dimensional setting.

Again, we define the error vector, to know how the vector of opinions $\vec{z}(t)$, differ from the opinions in stable time $\vec{z}(\infty)$, at each time step t. For this let $\vec{e}(t) = \vec{z}(t) - \vec{z}(\infty)$ be the error vector. We can write the error vector as $e(t) = (A+C)e(t-1)$.

Proof. From the definition of e(t), we have:

$$
\begin{aligned}
e(t) &= (A+C)^t\vec{s} + \sum_{k=0}^{t-1}(A+C)^k B\vec{s} - \sum_{k=0}^{\infty}(A+C)^k B\vec{s} \\
&= (A+C)^k\vec{s} - \sum_{k=t}^{\infty}(A+C)^k B\vec{s} \\
&= (A+C)^t(\vec{s} - \sum_{k=0}^{\infty}(A+C)^k B\vec{s}) \\
&= (A+C)e(t-1)
\end{aligned}
$$

□

Let $\tilde{\pi} = [\frac{w_i}{Z} : i \in V]^T$ for the weights w_i and Z be the normalizing constant such that $\sum_{i\in V}\tilde{\pi}_i = 1$. Then the following theorem exists.

Theorem 6.3.1. *For the error vector the following inequality exists:*

$$||\tilde{e}(t)||_{\tilde{\pi}} \leq (\lambda_{A+C})^t||\tilde{e}(0)||_{\tilde{\pi}}$$

where λ_{A+C} is the largest eigenvalue of $(A+C)$.

Proof. From the previous analysis, we know that for the error vector holds the following equation:

$$e(t) = (A+C)e(t-1)$$

Consider the transition probability matrix of the random walk we defined previously on this chapter.

$$P = \begin{pmatrix} (A+C)_{kn\times kn} & B_{kn\times kn} \\ I_{kn\times kn} & 0 \end{pmatrix}$$

By definition of reversibility, one can easily check that the pair (P,π) is reversible, as:

$$\pi(i)P_{ij} = \pi(j)P_{ji}$$

This implies in particular, that $\pi(i)$, is a stationary distribution of P, the unique one as P is irreducible.
Furthermore, the following equation also holds $\forall i, j \in V$:

$$\pi(i)(A+C)_{ij} = \pi(j)(A+C)_{ji}$$

Then, by minor abuse of terminology, we would also call $A+C$ reversible with respect to the distribution $\tilde{\pi} = [\frac{w_i}{\pi(A+C)} : i \in V]^T$, where $\pi(A+C)$ is the normalization constant. We can also express the reversibility of $A+C$, by saying that:

$$(A+C)^* = D^{\frac{1}{2}}(A+C)D^{-\frac{1}{2}}$$

is a symmetric matrix, where D is the diagonal matrix with elements of $\tilde{\pi} = [\frac{w_i}{\pi(A+C)} : i \in V]^T$ as diagonal entries, $D = diag\{\pi(1), \pi(2), ..\}$.

Since $(A+C)^*$ is symmetric, its eigenvalues are real, it is diagonalizable, and the sets of right and left eigenvectors are the same.

Choosing a set of right (and therefore left) eigenvectors of $(A+C)^*$, $w_1, w_2, ..., w_n$, we have an orthonormal basis in $\mathbb{R}^r$. Defining u and v as follows:

$$w = D^{-\frac{1}{2}}u \text{ and } w = D^{\frac{1}{2}}v$$

we have

$$u = Dv$$

The matrices $(A+C)$ and $(A+C)^*$ have the same eigenvalues, and moreover, v is a right eigenvector of $(A+C)$ and u a left eigenvector of $(A+C)$. Orthonormality of w_i's is with respect to the usual Euclidean Norm and is equivalent to orthonormality in $l^2(\pi)$ of $\{v_1, v_2, ..v_n\}$.
Using $\{v_1, v_2, ...v_n\}$ as a basis for $\mathbb{R}^n$, every vector in $\mathbb{R}^n$ can be written as $\vec{x} = \sum_i a_i v_i$ and taking the scalar product in $l^2(\pi)$ with v_j gives $\langle x, v_j\rangle_\pi = a_j$.
Thus, we can write:

$$\vec{e}(t) = \sum_{i=1}^{n} \langle \vec{e}(t), v_i\rangle_{\tilde{\pi}} v_i$$

As $(A+C)v_i = \lambda_i v_i$,

$$(A+C)\vec{e}(t) = \sum_{i=1}^{n} \lambda_i \langle \vec{e}(t), v_i\rangle_{\tilde{\pi}} v_i$$

This rate will be governed by the largest eigenvalue, as the others will converge more quickly:

$$\begin{aligned}
||\vec{e}(t+1)||^2_{\tilde{\pi}} &= \sum_i \lambda_i^2 \langle \vec{e}(t), v_i\rangle^2_{\tilde{\pi}} ||v_i||^2_{\tilde{\pi}} \\
&= \sum_i \lambda_i^2 \langle \vec{e}(t), v_i\rangle^2_{\tilde{\pi}} \\
&\leq \lambda^2_{A+C} \sum_i \langle \vec{e}(t), v_i\rangle^2_{\tilde{\pi}} \\
&= \lambda^2_{A+C} ||\vec{e}(t)||^2_{\tilde{\pi}}
\end{aligned}$$

Thus,

$$||\vec{e}(t+1)||_{\tilde{\pi}} \leq \lambda_{A+C} ||\vec{e}(t)||_{\tilde{\pi}}$$

and

$$||\vec{e}(t)||_{\tilde{\pi}} \leq \lambda^t_{A+C} ||\vec{e}(0)||_{\tilde{\pi}}$$

□

It is useful to remind at this point that as $A+C$ is a substochastic matrix, $|\lambda_{A+C}| < 1$. In order to define the time that the opinion formation game converges, we fix a bound v arbitrarily close to 0. When the value of the error vector is at most equal to v, we say that the game is converged.

More formally, the convergence time is defined as $\tau(\text{v}) = inf\{t \geq 0 : ||\vec{e}(t)||_{\tilde{\pi}} \leq \text{v}\}$ for some fixed v. Thus, by the previous theorem we have:

$$\left(\frac{1}{1-\lambda_{A+C}} - 1\right) \leq \frac{\tau(\text{v})}{log(||\vec{e}(t)||_{\tilde{\pi}}/\text{v})} \leq \frac{1}{1-\lambda_{A+C}}$$

and as n grows $\tau(\text{v}) = \Theta\left(\frac{1}{1-\lambda_{A+C}}\right)$.

Corollary 6.3.1. *The convergence time on the opinion formation game is $\tau(v) = \Theta\left(\frac{1}{1-\lambda_{A+C}}\right)$ where λ_{A+C} is the largest eigenvalue of (A+C).*

6.4 Bounding the Price of Anarchy

To conclude our analysis on the multi-dimensional version of the opinion formation game, in this chapter we estimate how "good" is the equilibrium that the game converges. As we have seen by now, in the opinion formation game a set of individuals explore the "landscape" of solutions in order to minimize their personal costs, until they reach the global minimum, the unique Nash equilibrium of the game.

However, in previous chapters, we show that the updates on individuals' opinions do not necessarily improve the social welfare. Precisely, in chapter 3, we defined the Price of Anarchy to be the ratio of the cost of the best Nash Equilibrium solution to the cost of the social optimum. In the previous chapter, we proved that the one-dimensional version of the game has a tight bound of 9/8 on the price of Anarchy. Hence, we do not expect a lower bound for the multi-dimensional version, as it is a generalization of the simple case of the opinion formation game in one dimension. However, at this point there is no evidence that the price of Anarchy in this version has a bound of 9/8. Thus, in the rest of the chapter we try to quantify the loss of efficiency resulting from selfish and uncoordinated behavior of the individuals and calculate a tight bound for the price of Anarchy for our game.

Again, in order to find robust bounds for the Price of Anarchy, we use the "local smoothness" proof framework that we discussed in 3.4, as developed by Roughgarden and al. [79]. Roughly, smoothness controls the cost of a set of "one-dimensional perturbations" of an outcome, as a function of both the initial outcome and the perturbations. The local smoothness technique requires the inequalities only for nearby pairs of outcomes, rather than for all pairs of outcomes. Thus, to show that the game is (λ, μ)-smooth, we need to prove that for a fixed profile $\vec{o}$, the following inequality holds :

$$\begin{aligned}\sum_{i \in V} \sum_{j \in K} (c_i(z_i(j), \vec{z_i}(j-), \vec{z}_{-i}) + (o_i(j) - z_i(j)) \frac{\partial c_i(z_i(j), \vec{z_i}(j-), \vec{z}_{-i})}{\partial z_i}) \\ \leq \frac{3}{4} C(\vec{o}) + \frac{1}{3} C(\vec{z})\end{aligned} \tag{6.16}$$

Now, it is useful to remind that if a cost minimization game is (λ, μ)-smooth with respect to an outcome $\vec{o}$, every pure-strategy Nash equilibrium $\vec{z}$ has objective function value at most $\frac{\lambda}{1+\mu}$ times that of $\vec{o}$.
We now instantiate the local smoothness framework, proving that the multi-dimensional version of the opinion formation game is $(\frac{3}{4}, \frac{1}{3})$-smooth and that the Price of Anarchy is at most $\frac{9}{8}$.

In order to use local smoothness technique, the following conditions must apply:

1. Every individual's opinion is a real number ($\forall i \in V, \forall j \in K, z_i(j) \in \mathbb{R}$).
2. The individual cost functions must be continuously differentiable.

These two requirements are satisfied in our game. To show that the game is $(\frac{3}{4}, \frac{1}{3})$-smooth, we need to prove that for a fixed profile $\vec{o}$, the following inequality holds.

$$\begin{aligned}\sum_{i \in V} \sum_{j \in K} (c_i(z_i(j), \vec{z_i}(j-), \vec{z}_{-i}) + (o_i(j) - z_i(j)) \frac{\partial c_i(z_i(j), \vec{z_i}(j-), \vec{z}_{-i})}{\partial z_i}) \\ \leq \frac{3}{4} C(\vec{o}) + \frac{1}{3} C(\vec{z})\end{aligned} \tag{6.17}$$

Proof. If we replace the cost functions with their analytical expressions, the inequality 6.17 can be written as the summation of the following three inequalities:

1.

$$\begin{aligned}&\sum_{i\in V}\sum_{j\in K}\Big((z_i(j)-avg(z_i))^2+(o_i(j)-z_i(j))(2z_i(j)-2avg(z_i))\Big)\\&\leq\frac{3}{4}\sum_{i\in V}\sum_{j\in K}(o_i(j)-avg(o_i))^2+\frac{1}{3}\sum_{i\in V}\sum_{j\in K}(z_i(j)-avg(z_i))^2\end{aligned}\tag{6.18}$$

2.

$$\begin{aligned}&\sum_{i\in V}\sum_{j\in K}\big((z_i(j)-s_i(j))^2+(o_i(j)-z_i(j))(2z_i(j)-2s_i(j))\big)\leq\\&\frac{3}{4}\sum_{i\in V}\sum_{j\in K}(o_i(j)-s_i(j))^2+\frac{1}{3}\sum_{i\in V}\sum_{j\in K}(z_i(j)-s_i(j))^2\end{aligned}\tag{6.19}$$

3.

$$\begin{aligned}&\sum_{i\in V}\sum_{j\in K}\sum_{u\in N(i)}w^j_{i,u}(z_i(j)-z_u(j))^2+(o_i(j)-z_i(j))\sum_{u\in N(i)}2w^j_{i,u}(z_i(j)-z_u(j))\\&\leq\frac{3}{4}\sum_{i\in V}\sum_{j\in K}\sum_{u\in N(i)}w^j_{i,u}(o_i(j)-o_u(j))^2+\frac{1}{3}\sum_{i\in V}\sum_{j\in K}\sum_{u\in N(i)}w^j_{i,u}(z_i(j)-z_u(j))^2\end{aligned}\tag{6.20}$$

6.18 inequality

In order to prove the 6.18 inequality, we have to show that for every $i \in V$ the following holds.

$$\begin{aligned}&\sum_{j\in K}((z_i(j)-avg(z_i))^2+(o_i(j)-z_i(j))(2z_i(j)-2avg(z_i)))\\&\leq\frac{3}{4}\sum_{j\in K}(o_i(j)-avg(o_i))^2+\frac{1}{3}\sum_{j\in K}(z_i(j)-avg(z_i))^2\end{aligned}\tag{6.21}$$

At this point, we claim that we can write

$$\begin{aligned}&\sum_{j\in K}(o_i(j)-z_i(j))(2z_i(j)-2avg(z_i))=\\&\sum_{j\in K}(o_i(j)-avg(o_i)-z_i(j)+avg(z_i))(2z_i(j)-2avg(z_i))\end{aligned}\tag{6.22}$$

Precisely, in order to prove the equation 6.22, we can write the first part as follows.

$$\begin{aligned}&\sum_{j\in K}2(o_i(j)-z_i(j))(z_i(j)-avg(z_i))=\\&\sum_{j\in K}2(o_i(j)z_i(j)-o_i(j)avg(z_i)-z_i^2(j)+z_i(j)avg(z_i)=\\&2(K\times avg(o_iz_i)-K\times avg(o_i)avg(z_i)-K\times avg(z_i^2)+K\times avg^2(z_i))\end{aligned}\tag{6.23}$$

Regarding the second part of the 6.22 equation, we can write,

$$
\begin{aligned}
&\sum_{j\in K} 2(o_i(j) - avg(o_i) - z_i(j) + avg(z_i))(z_i(j) - avg(z_i)) = \\
&\sum_{j\in K} 2(o_i(j)z_i(j) - o_i(j)avg(z_i) - avg(o_i)z_i(j) + \\
&avg(o_i)avg(z_i) - z_i^2(j) + z_i(j)avg(z_i) + avg(z_i)z_i(j) - avg^2(z_i)) = \\
&2(K \times avg(o_iz_i) - K \times avg(o_i)avg(z_i) - K \times avg(o_i)avg(z_i) + \\
&K \times avg(o_i)avg(z_i) - K \times avg(z_i^2) + \\
&K \times avg^2(z_i) + K \times avg^2(z_i) - K \times avg^2(z_i)) = \\
&2(K \times avg(o_iz_i) - K \times avg(o_i)avg(z_i) - K \times avg(z_i^2) + K \times avg^2(z_i))
\end{aligned}
\tag{6.24}
$$

Clearly, by 6.23 and 6.24, the 6.22 holds. Therefore, if we set $x = (z_i(j) - avg(z_i))$ and $y = (o_i(j) - avg(o_i))$, we have to prove that $\forall i \in V, \forall j \in K$:

$$
\begin{aligned}
&x^2 + (y - x) \times 2x \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\
&x^2 + 2xy - 2x^2 \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\
&2xy \leq \frac{3}{4}y^2 + \frac{4}{3}x^2 \implies \\
&(\frac{\sqrt{3}}{4}y + \frac{2}{\sqrt{3}}x)^2 \geq 0
\end{aligned}
\tag{6.25}
$$

which is true for every $x = (z_i(j) - avg(z_i)), y = (o_i(j) - avg(o_i)) \in \mathbb{R}$

6.19 inequality

In this direction, for the 6.19 inequality, we set $x = (z_i(j) - s_i(j))$ and $y = (o_i(j) - s_i(j))$. Hence, we now have to prove that $\forall i \in V, \forall j \in K$

$$
\begin{aligned}
&x^2 + (y - x) \times 2x \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\
&x^2 + 2xy - 2x^2 \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \implies \\
&2xy \leq \frac{3}{4}y^2 + \frac{4}{3}x^2 \implies \\
&(\frac{\sqrt{3}}{4}y + \frac{2}{\sqrt{3}}x)^2 \geq 0
\end{aligned}
\tag{6.26}
$$

which is true for every $x = (z_i(j) - s_i(j)), y = (o_i(j) - s_i(j)) \in \mathbb{R}$

6.20 inequality

Finally, for the inequality 6.20, we have to prove that $\forall i \in V, \forall j \in K, \forall u \in N(i)$

$$
\begin{aligned}
&w_{i,u}^j(z_i(j) - z_u(j))^2 + (o_i(j) - z_i(j))2w_{i,u}^j(z_i(j) - z_u(j)) \\
&\leq \frac{3}{4}w_{i,u}^j(o_i(j) - o_u(j))^2 + \frac{1}{3}w_{i,u}^j(z_i(j) - z_u(j))^2
\end{aligned}
\tag{6.27}
$$

Now, we set $x = (z_i(j) - z_u(j))$ and $y = (o_i(j) - o_u(j))$. Furthermore, observing the initial equation we want to prove, one can notice the following. For every pair of $(i, u) \in E$,and $j \in K$ there will be the following expressions:

1. $$(o_i(j) - z_i(j))w_{iu}^j 2(z_i(j) - z_u(j))$$

2. $$(o_i(j) - z_i(j))w_{iu}^j 2(z_i(j) - z_u(j))$$

from the c_i and c_u, as one can clearly see from the equation 6.27. Thus, rearranging them, for every pair of $(i, u) \in E$ and for every $j \in K$ we have the expressions

$$(\frac{o_i(j)}{2} - \frac{o_u(j)}{2} - \frac{z_i(j)}{2} + \frac{z_u(j)}{2})2w_{iu}^j(z_i(j) - z_u(j))$$

and

$$(\frac{o_u(j)}{2} - \frac{o_i(j)}{2} - \frac{z_u(j)}{2} + \frac{z_i(j)}{2})2w_{iu}^j(z_i(j) - z_u(j))$$

Clearly, now we need to prove the following

$$\begin{aligned} & x^2 + \frac{y-x}{2}2x \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \\ & x^2 + yx - x^2 \leq \frac{3}{4}y^2 + \frac{1}{3}x^2 \\ & (\frac{1}{\sqrt{3}}x + \frac{\sqrt{3}}{2}y)^2 \geq 0 \end{aligned} \tag{6.28}$$

which is true for every $x = (z_i(j) - s_i(j)), y = (o_i(j) - s_i(j)) \in \mathbb{R}$ □

Thus, we proved an upper bound of $\frac{9}{8}$ for the multi-dimensional version of the opinion formation game. However, as we noticed before, this bound cannot be improved. This is because, the one-dimensional version of the opinion formation game has a tight bound of $\frac{9}{8}$ on the Price of Anarchy and since it is a special case of the game we examine in this chapter, the tightness must hold also for the multidimensional version.

Corollary 6.4.1. *The multi-dimensional opinion formation game has a tight bound of $\frac{9}{8}$ on the Price of Anarchy.*

Bibliography

[1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. Working Paper 14040, National Bureau of Economic Research, May 2008.

[2] D. Acemoglu and A. Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.

[3] D. Acemoglu, A. E. Ozdaglar, and A. ParandehGheibi. Spread of misinformation in social networks. *CoRR*, abs/0906.5007, 2009.

[4] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47–97, Jan 2002.

[5] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67 – 96, 1974.

[6] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 57–66, New York, NY, USA, 2005. ACM.

[7] E. Bakshy, D. Eckles, R. Yan, and I. Rosenn. Social influence in social advertising: Evidence from field experiments. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 146–161, New York, NY, USA, 2012. ACM.

[8] V. Bala and S. Goyal. Learning from neighbours. *The Review of Economic Studies*, 65(3):595–621, 1998.

[9] V. Bala and S. Goyal. Learning from neighbours. *The review of economic studies*, 65(3):595–621, 1998.

[10] V. Banerjee and A. V. Banerjee. A simple model of herd behavior. *Quart. J. Econom*, pages 797–818, 1992.

[11] A.-L. Barabasi. *Linked: How Everything Is Connected to Everything Else and What It Means*. Plume, 2003.

[12] J. A. Barnes. Class and Committees in a Norwegian Island Parish. *Human Relations*, 7:39–58, 1954.

[13] E. Ben-Naim. Opinion dynamics: Rise and fall of political parties. *EPL (Europhysics Letters)*, 69(5):671, 2005.

[14] E. Ben-Naim, P. Krapivsky, F. Vazquez, and S. Redner. Unity and discord in opinion dynamics. *Physica A: Statistical Mechanics and its Applications*, 330(1):99–106, 2003.

[15] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):pp. 992–1026, 1992.

[16] D. Bindel, J. Kleinberg, and S. Oren. How bad is forming your own opinion? In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 57–66, Oct 2011.

[17] M. Buchanan. *Nexus: Small Worlds and the Groundbreaking Science of Networks.* W.W. Norton, 2002.

[18] S. Chatterjee and E. Seneta. Towards Consensus: Some Convergence Theorems on Repeated Averaging. *J. Appl. Prob.*, 14:159–164, 1977.

[19] S. Choi, D. Gale, and S. Kariv. Behavioral aspects of learning in social networks: An experimental study. In *ADVANCES IN BEHAVIORAL AND EXPERIMENTAL ECONOMICS (IN THE ADVANCES IN APPLIED MICROECONOMICS SERIES), EDITED BY*, pages 25–61. John Morgan, JAI Press, 2005.

[20] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 67–73, New York, NY, USA, 2005. ACM.

[21] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[22] J. Coleman, E. Katz, H. Menzel, and C. U. B. of Applied Social Research. *Medical innovation: a diffusion study.* Advanced Study in Sociology. Bobbs-Merrill Co., 1966.

[23] A. M. Colman. *Game theory and its applications: In the social and biological sciences.* Psychology Press, 2013.

[24] F. Cucker and S. Smale. Emergent behavior in flocks. *Automatic Control, IEEE Transactions on*, 52(5):852–862, 2007.

[25] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pages 71–78, New York, NY, USA, 2006. ACM.

[26] M. de Oliveira. Isotropic majority-vote model on a square lattice. *Journal of Statistical Physics*, 66(1-2):273–281, 1992.

[27] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(1):87–98, 2000.

[28] M. H. Degroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[29] P. M. Demarzo, D. Vayanos, J. Zwiebel, P. Veronesi, and R. Zeckhauser. Persuasion bias, social influence, and unidimensional opinions. the quarterly. *Journal of Economics*, 2003.

[30] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW (Physics).* Oxford University Press, Inc., New York, NY, USA, 2003.

[31] K. Etessami and M. Yannakakis. On the complexity of nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, Apr. 2010.

[32] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 604–612, New York, NY, USA, 2004. ACM.

[33] F. Fagnani and S. Zampieri. Randomized consensus algorithms over large scale networks. *Selected Areas in Communications, IEEE Journal on*, 26(4):634–649, 2008.

[34] S. Fortunato, V. Latora, A. Pluchino, and A. Rapisarda. Vector opinion dynamics in a bounded confidence consensus model. *International Journal of Modern Physics C*, 16(10):1535–1551, 2005.

[35] J. French. *A Formal Theory of Social Power*. Reprint Series in Social Sciences. Irvington Publishers, 1993.

[36] N. E. Friedkin and E. C. Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–206, 1990.

[37] S. Galam. Local dynamics vs.B•social mechanisms: A unifying frame. *EPL (Europhysics Letters)*, 70(6):705, 2005.

[38] S. GALAM. Sociophysics: A review of galam models. *International Journal of Modern Physics C*, 19(03):409–440, 2008.

[39] S. Galam, Y. Gefen, and Y. Shapir. Sociophysics: A new approach of sociological collective behaviour. i. mean-behaviour description of a strike. *Journal of Mathematical Sociology*, 9(1):1–13, 1982.

[40] S. Galam and S. Moscovici. Towards a theory of collective phenomena: Consensus and attitude changes in groups. *European Journal of Social Psychology*, 21(1):49–74, 1991.

[41] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003.

[42] J. Ghaderi and R. Srikant. Opinion dynamics in social networks: A local interaction game with stubborn agents. *CoRR*, abs/1208.5076, 2012.

[43] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. Discussion Papers 777, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 1988.

[44] H. Gintis. *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton University Press, 2000.

[45] M. Gladwell. *The Tipping Point: How Little Things Can Make a Big Difference*. Little, Brown, 2006.

[46] B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, pages 112–149, 2010.

[47] F. Harary. A criterion for unanimity in french's theory of social power. *Studies in Social Power, Ann Arbor*, 1959.

[48] J. C. Harsanyi. Games with incomplete information played by b••bayesianb•• players, part iii. the basic probability distribution of the game. *Management Science*, 14(7):486–502, 1968.

[49] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5:1–24, 2002.

[50] R. A. Holley and T. M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann. Probab.*, 3(4):643–663, 08 1975.

[51] M. O. Jackson et al. *Social and economic networks*, volume 3. Princeton University Press Princeton, 2008.

[52] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules, 2002.

[53] E. Katz, P. Lazarsfeld, and C. U. B. of Applied Social Research. *Personal Influence; the Part Played by People in the Flow of Mass Communications, by Elihu Katz and Paul F. Lazarsfeld. With a Foreword by Elmo Roper*. Foundations of communications research. 1955.

[54] H. Kelman. Processes of opinion change. *Public Opinion Quarterly*, 25:57–78, 1961.

[55] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, STACS'99, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[56] U. Krause. A discrete nonlinear and non-autonomous model of consensus formation. In *Communications in Dierence Equations*, pages 227–236. Gordon and Breach Pub, 2000.

[57] Y.-M. Li and Y.-L. Shiu. A diffusion mechanism for social advertising over microblogs. *Decis. Support Syst.*, 54(1):9–22, Dec. 2012.

[58] Z. Li, W. Ren, X. Liu, and L. Xie. Distributed consensus of linear multi-agent systems with adaptive dynamic protocols. *Automatica*, 49(7):1986–1995, 2013.

[59] T. M. Liggett.

[60] F. Lima, A. Sousa, and M. Sumuor. Majority-vote on directed erdős–rényi random graphs. *Physica A: Statistical Mechanics and its Applications*, 387(14):3503–3510, 2008.

[61] J. Lorenz. Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, page 2007.

[62] J. Lorenz and D. A. Lorenz. On conditions for convergence to consensus. *Automatic Control, IEEE Transactions on*, 55(7):1651–1656, 2010.

[63] A. Madan, K. Farrahi, D. Gatica-Perez, and A. Pentland. Pervasive sensing to model political opinions in face-to-face networks. In K. Lyons, J. Hightower, and E. Huang, editors, *Pervasive Computing*, volume 6696 of *Lecture Notes in Computer Science*, pages 214–231. Springer Berlin Heidelberg, 2011.

[64] A. D. Mare, S. S. D. Catania, and V. Latora. Opinion formation models based on game theory, 2006.

[65] A. C. Martins. Continuous opinions and discrete actions in opinion dynamics problems. *International Journal of Modern Physics C*, 19(04):617–624, 2008.

[66] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application. In *Proceedings of the 6th ACM conference on Embedded network sensor systems*, pages 337–350. ACM, 2008.

[67] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.

[68] J. Moreno. *Who shall survive?: Foundations of sociometry, group psychotherapy and sociodrama*. Sociometry monographs. Beacon House, 1953.

[69] H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.

[70] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

[71] J. v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[72] M. E. J. Newman. Models of the small world. *J. Stat. Phys*, pages 819–841, 2000.

[73] M. E. J. Newman. The structure and function of complex networks. *SIAM REVIEW*, 45:167–256, 2003.

[74] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.

[75] A. Olshevsky and J. N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM J. Control Optim.*, 48(1):33–55, Feb. 2009.

[76] E. Rogers. *Diffusion of Innovations, 5th Edition*. Free Press, 2003.

[77] R. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[78] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 513–522, New York, NY, USA, 2009. ACM.

[79] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 255–267. SIAM, 2011.

[80] R. Schmitt-Beck. Paul f. lazarsfeld/bernard berelson/hazel gaudet, the people's choice. how the voter makes up his mind in a presidential campaign, new york/london 1944. In S. Kailitz, editor, *SchlG'Osselwerke der Politikwissenschaft*, pages 229–233. VS Verlag fG'Or Sozialwissenschaften, 2007.

[81] J. Scott. *Social Network Analysis: A Handbook.* SAGE Publications, 2000.

[82] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975.

[83] D. Strang and S. A. Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24(1):265–290, 1998.

[84] S. H. Strogatz. Exploring complex networks. *Nature*, 410(6825):268–276, 2001.

[85] Y. Su and J. Huang. Two consensus problems for discrete-time multi-agent systems with switching network topology. *Automatica*, 48(9):1988–1997, Sept. 2012.

[86] K. Sznajd-Weron. Sznajd model and its applications. *Acta Physica Polonica B*, 36(8):2537–2547, 2005.

[87] K. Sznajd-Weron and J. Sznajd. Opinion evolution in closed community. *International Journal of Modern Physics C*, 11(6):1157–1165, 2000.

[88] K. Sznajd-Weron and R. Weron. A simple model of price formation. *Int. J. Mod. Phys. C*, 13:115, 2002.

[89] J. Tavares, M. T. da Gama, and A. Nunes. Coherence thresholds in models of language change and evolution: The effects of noise, dynamics, and network of interactions. *Physical Review E*, 77(4):046108, 2008.

[90] J. N. Tsitsiklis. Problems in decentralized decision making and computation. Technical report, MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR INFORMATION AND DECISION SYSTEMS, 1984.

[91] J. N. Tsitsiklis, D. P. Bertsekas, M. Athans, et al. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE transactions on automatic control*, 31(9):803–812, 1986.

[92] T. W. Valente. Social network thresholds in the diffusion of innovations. *Social networks*, 18(1):69–89, 1996.

[93] J. Von Neumann and O. Morgenstern. Theory of games and economic behavior (1953 edition) princeton university press. *Princeton, NJ*, 1944.

[94] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.

[95] D. J. Watts. *Small Worlds: The Dynamics of Networks Between Order and Randomness.* Princeton University Press, Princeton, NJ, USA, 2003.

[96] D. J. Watts. *Six degrees: The science of a connected age.* WW Norton, 2004.

[97] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal. Meet, discuss, and segregate! *Complexity*, 7(3):55–63, 2002.

[98] I. WELCH. Sequential sales, learning, and cascades. *The Journal of Finance*, 47(2):695–732, 1992.