Εθνικο Μετσοβιο Πολυτεχνειο
Σχολη Ηλεκτρολογων Μηχανικων και Μηχανικων Υπολογιστων

Τομεας Τεχνολογιας Πληροφορικης και Υπολογιστων

ΑΝΑΛΥΣΗ ΔΥΝΑΜΙΚΩΝ ΔΙΑΔΙΚΑΣΙΩΝ ΔΙΑΜΟΡΦΩΣΗΣ ΑΠΟΨΗΣ ΜΕ ΧΡΗΣΗ ΤΕΧΝΙΚΩΝ ΚΥΡΤΗΣ ΒΕΛΤΙΣΤΟΠΟΙΗΣΗΣ

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Αλέξανδρος
Τσιγώνιας-Δημητριάδης

Επιβλέπων: Δημήτρης Φωτάκης
Επίκουρος Καθηγητής ΕΜΠ

Αθήνα, Αύγουστος 2018

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ

ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ

ΑΝΑΛΥΣΗ ΔΥΝΑΜΙΚΩΝ ΔΙΑΔΙΚΑΣΙΩΝ ΔΙΑΜΟΡΦΩΣΗΣ ΑΠΟΨΗΣ ΜΕ ΧΡΗΣΗ ΤΕΧΝΙΚΩΝ ΚΥΡΤΗΣ ΒΕΛΤΙΣΤΟΠΟΙΗΣΗΣ

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Αλέξανδρος
Τσιγώνιας-Δημητριάδης

Επιβλέπων: Δημήτρης Φωτάκης
Επίκουρος Καθηγητής ΕΜΠ

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή στις 27 Αυγούστου 2018.

....	....	....
Δημήτρης Φωτάκης	Νικόλαος Παπασπύρου	Αριστείδης Παγουρτζής
Επίκουρος Καθηγητής ΕΜΠ	Αναπληρωτής Καθηγητής ΕΜΠ	Επίκουρος Καθηγητής ΕΜΠ

....

Αλέξανδρος Τσιγώνιας-Δημητριάδης
(Διπλωματούχος Ηλεκτρολόγος Μηχανικός & Μηχανικός Υπολογιστών Ε.Μ.Π.)

Περίληψη

Στη σύγχρονη εποχή, τα κοινωνικά δίκτυα παίζουν έναν πολύ σημαντικό ρόλο στην καθημερινή μας ζωή. Οι άνθρωποι συμμετέχουν σε διαφορετικά κοινωνικά δίκτυα για να ενημερωθούν, να ανταλάξουν απόψεις με τους άλλους και να γνωστοποιήσουν τις σκέψεις τους. Οι κοινωνικές επιστήμες, όπως η κοινωνιολογία και τα οικονομικά, μελετούν από πάντα τις σχέσεις μεταξύ των ανθρώπων και έχουν προσπαθήσει να εξηγήσουν πώς τα μέλη μιας ομάδας διαμορφώνουν τις απόψεις τους για κάποιο ζήτημα. Το γεγονός ότι τα κοινωνικά δίκτυα μεγαλώνουν με ταχύτατους ρυθμούς έχει ως αποτέλεσμα οι αλληλεπιδράσεις μεταξύ των χρηστών να πολλαπλασιάζονται και να γίνονται πιο πολύπλοκες. Ως εκ τούτου, δημιουργήθηκε η ανάγκη να μοντελοποιήσουμε με μαθηματικό τρόπο την ανταλλαγή απόψεων και να αναλύσουμε τα δυναμικά συστήματα που προκύπτουν. Ενδιαφερομάστε κυρίως να μάθουμε αν οι αλληλεπιδράσεις αυτές μπορούν να οδηγήσουν το σύστημα σε κατάσταση ισορροπίας και πόσο γρήγορα συμβαίνει αυτό, σε περίπτωση που επιτυγχάνεται. Σε αυτή τη διπλωματική εργασία, ξεκινάμε παρουσιάζοντας δύο από τα πιο σημαντικά μοντέλα διαμόρφωσης άποψης και σύγχρονα ερευνητικά αποτελέσματα που βασίζονται σε αυτά τα μοντέλα. Συνεχίζουμε, δείχνοντας πώς μπορούμε να χρησιμοποιήσουμε τεχνικές από την κυρτή βελτιστοποίηση ώστε να αναλύσουμε και να κατανοήσουμε καλύτερα τη δυναμική αυτών των μοντέλων. Πιο συγκεκριμένα, θα προσπαθήσουμε να συνδέσουμε τον αλγόριθμο gradient descent, την online κυρτή βελτιστοποίση καθώς και τη στοχαστική βελτιστοποίηση με το μοντέλο Friedkin-Johnsen. Τέλος, θα εισάγουμε ένα πιθανοτικό μοντέλο βασισμένο στο FJ και θα εξετάσουμε τις συνθήκες σύγκλισής του.

Λέξεις-κλειδιά: Δυναμική Διαμόρφωση Άποψης, Κυρτή Βελτιστοποίηση, Στοχαστική Βελτιστοποίηση, Κοινωνικά Δίκτυα, Αλγοριθμική Θεωρία Παιγνίων

Abstract

In today's world, social networks play a major role in our everyday lives. People participate in different social networks in order to get informed, exchange opinions with others and externalize their thoughts. Social sciences, such as sociology and economics, have long been studying the relationships between individuals and have tried to explain how groups of people form their opinions. As social networks grow rapidly, such interactions proliferate and become much more complex. This created the need to formulate mathematically the exchange of opinions and analyze the dynamics of these systems. We are mainly interested in finding whether an equilibrium can be reached as a result of such interactions and if the latter holds, how fast this happens. In this thesis, we start by presenting two of the most influential opinion formation models and contemporary research based on them. We continue by showing how convex optimization techniques can help us gain insight into the analysis of opinion formation models. More specifically, we link gradient descent, online convex optimization and stochastic optimization mainly with the FJ model. Finally, we introduce a randomized variant of the FJ model and examine its convergence properties.

Keywords: Opinion Dynamics, Convex Optimization, Stochastic Optimization, Social Networks, Algorithmic Game Theory

Ευχαριστίες

Στη διάρκεια αυτής της διπλωματικής είχα την τύχη να συνεργαστώ με δύο εξαιρετικούς ανθρώπους. Θα ήθελα, αρχικά, να ευχαριστήσω τον επιβλέποντα καθηγητή μου, κύριο Φωτάκη, για την καθοδήγησή του και τις πολύτιμες συμβουλές του τόσο εντός όσο και εκτός των πλαισίων της διπλωματικής. Η μεταδοτικοτητά του και ο τρόπος που αντιμετωπίζει τους φοιτητές τον καθιστούν πρότυπο δασκάλου. Θα ήθελα επίσης να ευχαριστήσω θερμά τον υποψήφιο διδάκτορα Στρατή Σκουλάκη για την αμέριστη βοήθειά του, την προσπάθειά του στο να κατανοήσω εις βάθος τις έννοιες που μελέτησα και την προθυμία του να απαντάει πάντα στις ερωτήσεις μου.

Με την παρούσα εργασία, κλείνει ένας σημαντικός κύκλος της ζωής μου και ανοίγει ένας νέος. Σε αυτό το σημείο, οφείλω ένα μεγάλο ευχαριστώ στους γονείς μου για την υποστήριξη που μου έδειξαν σε όλες μου τις αποφάσεις και για τις αρχές και την αγάπη για τη γνώση που μου εμφύσησαν από μικρή ηλικία. Βεβαίως, ευχαριστώ θερμά τους ανθρώπους του εργαστηρίου ΜΟΠ. Περνώντας ατελείωτες ώρες στο χώρο αυτό, ήταν ο λόγος που η καθημερινότητά μου γινόταν πιο ενδιαφέρουσα και αποτέλεσαν ένα σημαντικό στήριγμα τους τελευταίους μήνες. Ένα μεγάλο ευχαριστώ οφείλω και στους φίλους μου, που ως αναπόσπατο κομμάτι της ζωής μου υπήρξαν μαζί μου στα εύκολα και στα δύσκολα. Τέλος, ευχαριστώ όλους αυτούς τους ανθρώπους που ήρθαν στη ζωή μου και έμειναν ή έφυγαν, γεμίζοντας αυτά τα έξι φοιτητικά χρόνια με αξέχαστες εμπειρίες.

Περιεχόμενα

0 Εκτεταμένη ελληνική περίληψη **vi**

1 Introduction **1**

1.1 Brief overview of the chapters 3

2 Preliminaries **4**

2.1 Linear Algebra 4

2.2 Stochastic processes 6

2.3 Convex optimization 7

2.4 Probability 8

2.5 Algorithmic Game Theory 10

2.5.1 Games and Potential Functions 12

2.5.1.1 Definitions and properties 12

3 Opinion formation models and their convergence properties **14**

3.1 The De-Groot Model 14

3.2 The Friedkin-Johnsen Model 17

3.3 Variants of the FJ-model 21

3.3.1 A deterministic variant with outdated information 21

3.3.2 A randomized asynchronous variant 27

4 Optimization and opinion formation models **31**

4.1 Gradient Descent 32

4.1.1 Best-response dynamics and Gradient Descent 35

4.1.2 The FJ-model and Gradient Descent 36

4.2 Online convex optimization 37
4.2.1 Online Gradient Descent 39
4.2.2 The bandit version 41
4.3 Stochastic optimization 42
4.3.1 Motivation and history 42
4.3.2 A limited information protocol for the FJ model 44
4.3.3 Variance reduction and an analysis in epochs 47

5 Stochastic Gradient Descent and the FJ model 49

5.1 An opinion formation model with limited information 49
5.2 A variant of the FJ model with increased sampling 53
5.2.1 How close to the equilibrium point can we get? 54
5.2.2 Convergence with high probability 56

Βιβλιογραφία 59

0 Εκτεταμένη ελληνική περίληψη

Το διαδίκτυο, ίσως η σημαντικότερη εφεύρεση του 20ου αιώνα, μαζί με την ισχυρή παρουσία των κοινωνικών μέσων στην καθημερινή μας ζωή, οδήγησαν σε έναν νέο διασυνδεδεμένο κόσμο. Σήμερα, οι άνθρωποι μπορούν να επικοινωνούν και να ανταλλάσσουν πληροφορίες γρήγορα σε όλο τον κόσμο. Αυτό έχει ως αποτέλεσμα τη δημιουργία τεράστιων και σύνθετων κοινωνικών δικτύων που συνεχίζουν να αναπτύσσονται ραγδαία. Λόγω αυτής της επαναστατικής εποχής, η πληροφορική έχει επίσης γνωρίσει μια τεράστια ανάπτυξη κατά τη διάρκεια των τελευταίων δεκαετιών. Μεταξύ των ζητημάτων που ο τομέας μας προσπαθεί να αντιμετωπίσει, είναι επίσης η ανάλυση των πτυχών αυτών των δικτύων, τόσο από θεωρητική όσο και από πρακτική σκοπιά. Ένα από αυτά τα υποπεδία εξετάζει τον τρόπο ποθ οι χρήστες σε ένα κοινωνικό δίκτυο ανταλλάζουν απόψεις και πώς διαμορφώνουν τη συμπεριφορά τους σχετικά με ορισμένα θέματα, όταν επηρεάζονται από άλλους. Το πεδίο αυτό ονομάστηκε *Opinion Dynamics* και έχει μελετηθεί διεξοδικά τα τελευταία χρόνια. Φυσικά, οι επιστήμονες έχουν προσπαθήσει από καιρό να κατανοήσουν πώς οι άνθρωποι αλληλεπιδρούν μεταξύ τους και πώς το κοινωνικό περιβάλλον επηρεάζει τη διαμόρφωση των απόψεών τους.

Ψυχολόγοι, κοινωνιολόγοι, πολιτικοί επιστήμονες και οικονομολόγοι έχουν συμμετάσχει εδώ και καιρό σε αυτή τη γραμμή έρευνας. Ερωτήσεις όπως το τι είναι γενετικά προκαθορισμένο και πώς το κοινωνικοοικονομικό περιβάλλον επηρεάζει την ανάπτυξη των ανθρώπων, πώς η συλλογική σοφία λειτουργεί και πώς οι άνθρωποι αποφασίζουν για το τι να ψηφίσουν, για να τοποθετηθούν στο πολιτικό φάσμα και να εκφράσουν τη γνώμη τους για διάφορα πολιτικά ζητήματα, έχουν σταθεί στον πυρήνα των προαναφερθέντων επιστημών. Η ανάγκη να κατανοήσουμε καλύτερα και από μια διαφορετική οπτική γωνία πώς οι άνθρωποι ανταλλάσσουν γνώμες και επηρεάζουν τους άλλους και πώς συμβαίνει η διάχυση της πληροφορίας σε μεγάλα σύνθετα κοινωνικά δίκτυα, οδήγησε στατιστικούς, μαθηματικούς και επιστήμονες πληροφορικής να προσπαθήσει να συλλάβουν τη συμπεριφορά των πρακτόρων μέσω μαθηματικών μοντέλων. Το πρώτο σημαντικό μοντέλο, δημοσιεύθηκε από τον Αμερικανό στατιστικό Morris H. DeGroot το 1974. Το κοινωνικό δίκτυο μοντελοποιείται ως ένα γράφημα $G(V, E)$ και σύμφωνα με έναν απλό κανόνα ανανέωσης των απόψεων, μια ομάδα ανθρώπων μπορεί να συγκλίνει στην ίδια άποψη σχετικά με ένα συγκεκριμένο θέμα.

Τα αποτελέσματά του αποτέλεσαν την απαρχή μιας πολύ γόνιμης γραμμής έρευνας και σήμερα έχουμε στην κατοχή μας καλώς ορισμένα μαθηματικά πλαίσια με ενδιαφέρουσες ιδιότητες. Το μοντέλο του DeGroot είναι εκφραστικό, αλλά πολύ απλό και ο κύριος στόχος είναι να σχεδιάσουμε νέα πιο σύνθετα μοντέλα, τα οποία διατηρούν τις ωραίες ιδιότητες και προσωμοιάζουν καλύτερα τη συμπεριφορά των πρακτόρων στην πραγματικότητα. Η ίσως πιο ενδιαφέρουσα γενίκευση του μοντέλου DeGroot προήλθε από τη συνεργασία του κοινωνιολόγου Noah Friedkin και του μαθηματικού Eugene Johnsen. Το άρθρο αυτό δημοσιεύθηκε το 1990 και από τότε πολλή έρευνα, τόσο θεωρητικής όσο και πειραματικής φύσεως, έχει παραχθεί με το μοντέλο FJ ως βασική ιδέα του. Μερικά από αυτά τα μοντέλα πρόκειται να μελετηθούν διεξοδικά αργότερα σε αυτήν την διπλωματική εργασία. Θα ονομάσουμε αυτά τα μοντέλα γραμμικά, λόγω της μορφής των κανόνων ανανέωσης και του γεγονότος ότι χρησιμοποιούμε εργαλεία από τη γραμμική άλγεβρα,τις στοχαστικές διεργασίες και τη θεωρία γραφημάτων για την ανάλυση των ιδιοτήτων σύγκλισης αυτών των μοντέλων. Ωστόσο, αυτά τα μοντέλα υποθέτουν ότι το υποβόσκων κοινωνικό δίκτυο δεν αλλάζει με την πάροδο του χρόνου.

Φαίνεται πιο κοντά στην πραγματικότητα να υποθέσουμε ότι οι άνθρωποι μπορεί να αρχίσουν να λαμβάνουν υπόψη τη γνώμη των νέων πρακτόρων, να σταματήσουν την ανταλλαγή απόψεων με άλλους ή να αλλάξουν το βαθμό εμπιστοσύνης προς ορισμένους από τους γείτονές τους σε ένα κοινωνικό δίκτυο. Ως εκ τούτου, αναπτύχθηκε μια άλλη γραμμή έρευνας, όπου οι γνώμες των πρακτόρων και των υποκείμενων δικτύων συνεξελίσσονται. Σε αυτή την κατηγορία, τα πιο γνωστά μοντέλα είναι το Hegselmann-Krause, το οποίο είναι ντετερμινιστικό και το Deffuant-Weisbuch μοντέλο, το οποίο είναι τυχαιοκρατικό. Αυτά τα μοντέλα μπορεί να έχουν μεγαλύτερο ενδιαφέρον, αλλά το εξελισσόμενο δυνμικό δίκτυο καθιστά την ανάλυση τους πολύ δύσκολη. Απαιτούνται ριζικά διαφορετικές τεχνικές, όπως η ιδέα της s-ενέργειας ενός πολυπρακτορικού συστήματος. Άλλες ερευνητικές δουλειές αναλύουν αυτές τις προαναφερθείσες διαδικασίες σχηματισμού γνώμης από μία παιγνιοθεωρητική σκοπιά: Σχεδιάζουμε ένα παίγνιο στο οποίο οι παίχτες παίζουν εγωιστικά και προσπαθούν να ελαχιστοποιήσουν τη συναρτήση κόστους τους . Παρακινημένοι από το γεγονός ότι οι άνθρωποι σε μια κοινωνία σπάνια φθάνουν σε συμφωνία για ένα θέμα, προσπαθούν να ποσοτικοποιήσουν το κόστος διαφωνίας που προκύπτει και να αναλύσουν την κοινωνική έκβαση τέτοιων παιγνίων. Τέλος, η διαμόρφωση άποψης όπου οι γνώμες είναι 0 ή 1 (αυτή είναι η περίπτωση, για παράδειγμα, στην ψηφοφορία) έχει μελετηθεί και Bayesian προσεγγίσεις στην κοινωνική μάθηση έχουν εξεταστεί. Σε κάθε περίπτωση, τα κοινωνικά και οικονομικά δίκτυα έχουν γίνει μέρος της καθημερινής μας ζωής και η κατανόηση των πτυχών τους αποτελεί μείζον μέλημα της επιστημονικής κοινότητας.

Σε αυτή τη διπλωματική, δεν πρόκειται να επεκταθούμε περαιτέρω στα συνεξελικτικά μοντέλα διαμόρφωσης άποψης, αλλά πρόκειται να επικεντρωθούμε αντ' αυτού σε γραμμικά μοντέλα, που συνδέονται στενά με το μοντέλο FJ. Πολλά διαφορετικά μαθηματικά εργαλεία έχουν χρησιμοποιηθεί για την ανάλυση τέτοιων δυναμικών συστημάτων. Μία από τις προσπάθειές μας θα είναι να συσχετίσουμε τη δυναμική διμόρφωση άποψης με την κυρτή βελτιστοποίηση και να δείξουμε ότι αλγόριθμοι και ιδέες από την κυρτή βελτιστοποίηση μπορούν να μας βοηθήσουν να αποκτήσουμε μια βαθύτερη κατανόηση όσον αφορά το πώς η δυναμική των εν λόγω συστημάτων εξελίσσεται με την πάροδο του χρόνου και πώς και κάτω από ποιες συνθήκες μπορεί να επιτευχθεί σύγκλιση σε κάποιο σύνολο απόψεων. Για παράδειγμα, η ενσωμάτωση της ιδέας του αλγορίθμου gradient descent(αν και από διαφορετική σκοπιά) έχει ήδη γίνει στον κλάδο του market theory. Συμπερασματικά, οι γνωστοί αλγόριθμοι κυρτής βελτιστοποίησης καθώς και τα αποτελέσματα της στοχαστικής βελτιστοποίησης μπορούν να γίνουν ένα πολύ χρήσιμο εργαλείο, τόσο στη θεωρία όσο και στην πράξη. Η άλλη προσπάθειά μας, αν και συνδέεται άμεσα με την προηγούμενη, θα είναι να εξετάσουμε μοντέλα διαμόρφωσης άποψης με μερική πληροφορία και τυχαιότητα. Το κίνητρο για μια τέτοια κατεύθυνση έρευνας πηγάζει από το γεγονός ότι τα κοινωνικά δίκτυα γίνονται ολοένα και μεγαλύτερα. Ως εκ τούτου, ένας χρήστης πιθανότατα θα ανταλλάξει πληροφορίες σχετικά με ένα θέμα με ένα υποσύνολο της κοινωνικής γειτονιάς του, δεδομένου ότι το να ρωτήσει όλους τους γείτονές τους έχει υψηλό υπολογιστικό κόστος και είναι μη ρεαλιστικό. Στόχος μας είναι να σχεδιάσουμε πρωτόκολλα με περιορισμένη ανταλλαγή πληροφορίας, τα οποία έχουν μια επιθυμητή αναμενόμενη συμπεριφορά και των οποίων η ταχύτητα σύγκλισης σε ένα σταθερό σύνολο απόψεων δεν είναι πολύ χειρότερη σε σύγκριση με το ισοδύναμο ντετερμινιστικό μοντέλο. Συνεπώς, θα εστιάσουμε επίσης σε τυχαιοκρατικά πρωτόκολλα, θα τα διατυπώσουμε ως στοχαστικά προβλήματα βελτιστοποίησης και θα αναλύσουμε τις ιδιότητες σύγκλισης τους.

Μοντέλα Διαμόρφωσης Άποψης και οι ιδιότητες σύγκλισής τους

Σε αυτό το κεφάλαιο, πρόκειται να παρουσιάσουμε μερικά σημαντικά μοντέλα διαμόρφωσης άποψης, από παλαιότερη σημαντική αλλά και σύγχρονη έρευνα.

Το μοντέλο του DeGroot

Το κοινωνικό δίκτυο αντιπροσωπεύεται από ένα κατευθυνόμενο γράφημα με βάρη $G(V, E)$, κατά το οποίο κάθε κόμβος αντιπροσωπεύει έναν παίχτη στο κοινωνικό δίκτυο και το βάρος w_{ij} κάθε ακμής στο γράφημα αντιπροσωπεύει την εμπιστοσύνη που ο κόμβος i έχει για τον j. Μπορεί επίσης να υπάρχουν self loops w_{ii} στο γράφημα που αντιπροσωπεύουν το βαθμό πίστης του κάθε κόμβου στην άποψή του. Χωρίς απώλεια της γενικότητας θεωρούμε ότι $\forall i \in V : \sum_{(i,j)\in E} w_{ij} = 1$. Καταρχάς, όλοι οι κόμβοι έχουν μια αρχική άποψη που δηλώνεται από το διάνυσμα $x(0)$ και ο κανόνας ανανέωσης σε κάθε χρονικό βήμα παίρνει τη μορφή $x(t+1) = T \cdot x(t)$, όπου T είναι ο πίνακας ανανέσης του μοντέλου. Ο πίνακας T περιέχει τα βάρη των ακμών του γραφήματος μας $G(V, E)$ και ορίζεται ως:

$$T_{ij} = \begin{cases} w_{ij} & \text{if } (i,j) \in E \\ 0 & \text{if } (i,j) \notin E \end{cases}$$

Σημειώστε εδώ ότι το γράφημα μπορεί είτε να είναι κατευθυνόμενο είτε μη κατευθυνόμενο το οποίο οδηγεί σε μια ασύμμετρη και μια συμμετρική έκδοση του μοντέλου αντίστοιχα.

Theorem 0.1. Το μοντέλο του DeGroot με n πράκτορες συγκλίνει στο μοναδικό σημείο ισορροπίας

$$x^* = \lim_{t\to\infty} x(t) = \lim_{t\to\infty} T^t x(0) \quad (1)$$

για κάθε αρχικό διάνυσμα απόψεων $x(0) \in [0,1]^n$ αν και μόνο αν η αλυσίδα Markov με πίνακα μετάβασης T είναι *μη υποβιβάσιμη* και *απεριοδική*.

Το μοντέλο Friedkin-Johnsen

Η διαφορά στο μοντέλο FJ είναι ότι κάθε πράκτορας έχει επίσης μια προσωπική άποψη s_i, η οποία ορίζεται στην αρχή της διαδικασίας, παραμένει σταθερή και συμμετέχει με κάποιο βάρος στον κανόνα ανανέωσης του κάθε πράκτορα. Είναι φυσικό να υποθέσουμε ότι αυτή η εσωτερική γνώμη υπάρχει όταν οι άνθρωποι ανταλλάσσουν απόψεις σχετικά με ένα συγκεκριμένο θέμα.

Ο κανόνας ανανέωσης λαμβάνει τώρα την ακόλουθη μορφή:

$$x_i(t+1) = \sum_{\substack{j \in N_i \\ j \neq i}} w_{ij} x_j(t) + w_{ii} s_i$$

Μπορούμε επίσης να γράψουμε την παραπάνω εξίσωση με πίνακες:

$$x(t) = A^t s + \sum_{n=0}^{t-1} A^n B s \tag{2}$$

Στην περίπτωση των επίμονων πρακτόρων, απαιτούμε τουλάχιστον ένας πράκτορας να δίνει βάρος στην προσωπική του αμετάβλητη άποψη, έτσι ώστε ο πίνακας B να μην είναι πάλι 0. Τότε, υπάρχει σημείο ισορροπίας για το σύστημα:

$$x(\infty) = \lim_{t \to \infty} x(t) = \sum_{n=0}^{\infty} A^n B s = (I - A)^{-1} B s \tag{3}$$

Lemma 0.2. Η ταχύτητα σύγκλισης στην περίπτωση αυτή είναι γεωμετρική, σύμφωνα με τη μεγαλύτερη ιδιοτιμή του πίνακα A.

$$\|e(t)\|_\pi \leq (\lambda_A)^t \|e(0)\|_\pi \tag{4}$$

Μία ντετερμινιστική παραλλαγή με παρωχημένη πληροφορία

Σε αυτήν την περίπτωση, οι άνθρωποι λαμβάνουν υπόψη τις έρευνες και τις δημοσκοπήσεις, οι οποίες αποτυπώνουν συνολικά την άποψη της κοινωνίας και ενδέχεται να επηρεάσουν σημαντικά τη διαμόρφωση της τελικής άποψης. Κατά συνέπεια, είναι λογικό να εισαχθεί ένας όρος που περιγράφει την επιρροή που έχει η γνώμη της κοινωνίας, πέρα από τις μεμονωμένες γνώμες των πρακτόρων και να ποσοτικοποιηθεί επίσης αυτή η επιρροή. Χωρίς να μπούμε σε λεπτομέρειες σχετικά με τους πίνακες, ο κανόνας ανανέωσης γράφεται ως:

$$x(t) = Ax(t-1) + Bx(T_D) + Cs \tag{5}$$

Theorem 0.3 (Σύγκλιση με παρωχημένη πληροφορία)**.** Ορίζουμε μια διάρκεια L για κάθε εποχή και υποθέτουμε ότι στην αρχή κάθε εποχής η μέση άποψη της κοινωνίας αναγγέλλεται στους πράκτορες. Το μοντέλο μας περιγράφεται από τον κανόνα ανανέωσης $x(t) = Ax(t-1) + Bx(T_D) + Cs$ και υποθέτουμε ότι $\|A\|_\infty \leq 1 - \frac{1}{n}$ και $\|B\|_\infty \leq \frac{n-1}{n^2}$. Τότε για κάθε C, κάθε $\vec{s} \in [0,1]^n$ και $\varepsilon > 0$ το παραπάνω μοντέλο συγκλίνει *varepsilon*-κοντά στο μοναδικό σημείο ισορροπίας $x^* = (\mathbb{I} - (A + B))^{-1}Cs$ μετά από $\mathcal{O}\left(n^2 \ln \frac{n^2\|C\|_\infty}{\varepsilon}\right)$ εποχές.

Κυρτή βελτιστοποίηση και μοντέλα διαμόρφωσης άποψης

Σε αυτό το κεφάλαιο, θα θέλαμε να γεφυρώσουμε τους αλγόριθμους βελτιστοποίησης και τη δυναμική διαμόρφωση άποψης. Οι επαναληπτικές διαδικασίες των αλγορίθμων βελτιστοποίησης που πρόκειται να περιγράψουμε, θα μπορούσαν επίσης να εφαρμοστούν στα μοντέλα διαμόρφωσης γνώμης και να γίνουν ένα ισχυρό εργαλείο στην ανάλυση τέτοιων δυναμικών συστημάτων. Θα ξεκινήσουμε με τον gradient descent, στη συνέχεια με online κυρτή βελτιστοποίηση και, τέλος, θα δούμε πώς θα μπορούσαμε να δανειστούμε ιδέες από τη στοχαστική βελτιστοποίηση, προκειμένου να σχεδιάσουμε νέα τυχαιοκρατικά πρωτόκολλα και να αναλύσουμε τις ιδιότητες σύγκλισής τους.

Το μοντέλο FJ και ο αλγόριθμος Gradient Descent

Μπορούμε να γράψουμε τον κανόνα ανανέωσης του μοντέλου FJ ως μια gradient descent διαδικασία πάνω σε μια συνάρτηση $f(x)$. Αν ορίσουμε τη συνάρτηση $f(x) = \|x - A \cdot x - B \cdot s\|_2^2$ και τον πίνακα $\Gamma = \left(\nabla^2 f(x^t)\right)^{-1} = \frac{1}{2}(I-A)^{-1}$, τότε:

$$\begin{aligned} x^t &= x^{t-1} - \Gamma \cdot \nabla f(x^t) \\ &= x^{t-1} - \frac{(I-A)^{-1}}{2} \cdot 2\left((I-A)x^{t-1} - Bs\right)(I-A) \\ &= x^{t-1} - (I-A)x^{t-1} + Bs \\ &= Ax^{t-1} + Bs \end{aligned}$$

Στοχαστική βελτιστοποίηση

Γενικά, οι μέθοδοι στοχαστικής βελτιστοποίησης χρησιμοποιούν τον ακόλουθο κανόνα ανανέωσης:

$$x_{k+1} \leftarrow \underset{y \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ \frac{1}{2\eta} \|y - x_k\|_2^2 + \langle \widetilde{\nabla}_k, y \rangle + \psi(y) \right\}$$

όπου:

- η είναι το βήμα που κάνουμε
- $\widetilde{\nabla}_k$ είναι ένα τυχαίο διάνυσμα, το οποίο ικανοποιεί $\mathbf{E}\left[\widetilde{\nabla}_k\right] = \nabla f(x_k)$ και αποτελεί έναν αμερόληπτο εκτιμητή του gradient.

Πιο τυπικά, μπορούμε να πούμε ότι μας δίνεται πρόσβαση σε ένα noisy oracle που ορίζουμε ως

$$\mathcal{O}(x_k) = \widetilde{\nabla}_k \text{ s.t. } \mathbf{E}\left[\widetilde{\nabla}_k\right] = \nabla f(x_k), \mathbf{E}\left[\left\|\widetilde{\nabla}_k\right\|^2\right] \leq G^2$$

Ένα πρωτόκολλο με περιορισμένη ανταλλαγή πληροφορίας Το κίνητρο για το μοντέλο που πρόκειται να παρουσιάσουμε, είναι ότι στα κοινωνικά δίκτυα είναι πολύ σπάνιο για έναν πράκτορα να ρωτήσει όλους τους γείτονές του, προκειμένου να ενημερώσει τη γνώμη του. Το μοντέλο FJ υποθέτει ότι οι χρήστες επικοινωνούν με όλους τους γείτονές τους σε κάθε βήμα και λαμβάνουν υπόψη ένα σταθμισμένο μέσο όρο των απόψεων αυτών.

Λαμβάνοντας υπόψη ένα γράφημα $G(V,)$, το διάνυσμα των αρχικών απόψεων $vecs$ και έναν όρο $gamma(t) in (0, 1)$, ο κανόνας ανανέωσης λαμβάνει την ακόλουθη μορφή, η οποία αναφέρει ότι η γνώμη ενός πράκτορα στο επόμενο βήμα είναι ένας κυρτός συνδυασμός της προηγούμενης γνώμης του και της γνώμης που έκανε sampling κατά τη διάρκεια αυτού του γύρου:

$$x_i(t) = \gamma(t) x_i(t-1) + (1 - \gamma(t)) r_i(t) \tag{6}$$

όπου $r_i(t)$ είναι το τυχαίο μας δείγμα από κάποιο γείτονα, το οποίο το τραβάμε με πιθανότητα w_{ij}:

$$r_i(t) = \begin{cases} x_j(t-1), & \text{if } j \neq i \\ s_i, & \text{if } j = i \end{cases}$$

Theorem 0.4. Δεδομένου ενός μη κατευθυνόμενου γραφήματος $G(V, E)$ επιλέγουμε $\gamma(t) = \dfrac{t-1}{t}$ και ορίζουμε το διάνυσμα απόστασης απο το σημείο ισορροπίας $e(t) = y(t) - x^*$. Τότε ισχύει ότι

$$\lim_{t \to \infty} \|\mathbb{E}[\vec{x}(t)] - x^*\|_\infty = 0.$$

Στοχαστική βελτιστοποίηση και το μοντέλο FJ

Απλοποιώντας το μοντέλο, και θεωρώντας έναν μη κατευθυνόμενο d-κανονικό γράφο μπορούμε να ξαναγράψουμε τον κανόνα ανανέωσης ως:

$$x_i^t = \frac{t-1}{t} x_i^{t-1} + \frac{R_i^t}{t} \tag{7}$$

Μπορούμε σε αυτήν την περίπτωση να ορίσουμε και μια ισχυρά κυρτή συνάρτηση δυναμικού:

$$\Phi(x^t) = \sum_{i<j} \left(x_i^t - x_j^t \right)^2$$

Μετά από υπολογισμούς μπορούμε να το διατυπώσουμε σαν ένα πρόβλημα στοχαστικής βελτιστοποίησης:

$$\boxed{x^t = x_i^{t-1} - \frac{1}{t} S^t \ , \ \mathbb{E}\left[S^t\right] = \frac{1}{2d} \cdot \nabla\Phi\left(x^{t-1}\right)}$$

και τέλος να δείξουμε το ακόλουθο άνω φράγμα για την αναμενόμενη απόσταση από το σημείο ισορροπίας:

$$\boxed{\mathbb{E}\left[\left\|x^t - x^*\right\|_2^2\right] \leq \alpha \cdot \left\|x^{t-1} - x^*\right\|_2^2 + \frac{1}{t^2} \cdot \mathbb{E}\left[\left\|S^t - \frac{\nabla\Phi(x^{t-1})}{2d}\right\|_2^2\right]}$$

Μια παραλλαγή του FJ με αυξημένη δειγματοληψία

Αντί να ενημερώνουν τις απόψεις τους σε κάθε βήμα, περιμένουν έναν αριθμό γύρων και συνεχίζουν τη δειγματοληψία έως ότου έχουν φθάσει σε μια κοντινή εκτίμηση του πραγματικού μέσου όρου. Θέλουμε για κάθε πράκτορα να τραβήξουμε τόσο πολλά δείγματα, έτσι ώστε ο μέσος όρος που προκύπτει με τα δείγματα που παίρνει να είναι με υψηλή πιθανότητα ϵ-κοντά στον πραγματικό μέσο όρο του μοντέλου FJ. Αυτό μπορεί να γραφεί ως:

$$x^{t+1} = A \cdot y^t + B \cdot s$$
$$\left\|y^{t+1} - x^{t+1}\right\|_\infty \leq \epsilon$$

Υποθέτουμε ότι στο γύρο t είμαστε στο σημείο y^t. Ως εκ τούτου, x^{t+1} θα είναι το επόμενο σημείο μας στην ντετερμινιστική περίπτωση, και στη συνέχεια απαιτούμε εκτίμηση ϵ-κοντά.

Αν παίρνουμε έναν σταθερό μεγάλο αριθμό δειγμάτων σε κάθε γύρο, μπορούμε να φτάσουμε όσο κοντά θέλουμε στο σημείο ισορροπίας? Η απάντηση που δίνουμε είναι αρνητική.

Lemma 0.5. Στην παραλλαγή του FJ που ορίσαμε, αν υποθέσουμε σταθερό μέγεθος δείγματος σε κάθε επανάληψη, σε βάθος χρόνου μπορούμε να δείξουμε ότι για καθορισμένο $\epsilon > 0$, μπορούμε να συγκλινουμε σε μια $\frac{\epsilon}{1 - \|A\|_\infty}$ - μπάλα γύρω από το σημείο ισορροπίας.

Lemma 0.6. Μετά από έναν πεπερασμένο αριθμό γύρων t, μπορούμε να φτάσουμε σε μία $\frac{c \cdot \epsilon}{1 - \|A\|_\infty}$ - μπάλα γύρω από το σημείο ισορροπίας, όπου $c > 1$.

Lemma 0.7. Μετά από t γύρους(ή εποχές), με πιθανότητα τουλάχιστον $1 - \delta$ έχουμε ότι $\left\|y^t - x^*\right\|_\infty \leq \frac{c \cdot \varepsilon}{1 - \|A\|_\infty}$ εάν κατά τη διάρκεια κάθε γύρου τραβάμε τουλάχιστον $\frac{\ln(2nt/\delta)}{2\varepsilon^2}$ δείγματα.

Τέλος χρειάζεται να αυξάνουμε το μέγεθος του δείγματος μας από γύρο σε γύρο, καθώς έτσι μπορούμε και να αναπροσαρμόζουμε την παράμετρο ε. Με αυτόν τρόπο, μπορούμε να κάνουμε το ε όσο μικρό επιθυμούμε και έτσι να συνεχίσουμε να κινούμαστε προς το σημείο ισορροπίας.

1 Introduction

The Internet, probably the most important invention of the 20th century, along with the strong presence of the social media in our everyday lives, have led to a whole new interconnected world. Nowadays, people can communicate and exchange information rapidly across the globe. This has resulted in the creation of huge and complex social networks which continue to grow rapidly. Due to this revolutionary era, computer science has also exhibited a vast development over the recent decades. Among the issues that our field tries to address, is also the analysis of aspects of such networks , both from a theoretical and a practical perspective. One of these subfields examines how the agents in a social network exchange opinions and how they shape their behavior about certain topics when influenced by others. This field was given the name *opinion dynamics* and has been studied exhaustively over the last years. Of course, scientists have long attempted to understand how people interact with each other and how the social environment affects the formation of their opinions.

Psychologists, sociologists, political scientists and economists have long been involved in this line of research. Questions like what's genetically determined and how the socioeconomic environment influences the development of people, how collective wisdom works and how people decide on what to vote, place themselves in the political spectrum and express an opinion about various political issues have been standing in the core of the aforementioned sciences. The need to understand better and from a different perspective how people exchange opinions and influence others and how the information diffusion occurs in large complex social networks, led statisticians, mathematicians and computer scientists to try to capture the behavior of the agents through mathematical frameworks. The first influential paper was published by American statistician Morris H. DeGroot in 1974 [[1]]. DeGroot represented the social network as a graph $G(V, E)$ and provided a simple model according to which a group of people can reach a consensus about a certain topic.

His results pointed the direction to a very fruitful line of research and up until now we possess rigorous mathematical frameworks with interesting properties. DeGroot's model is expressive, yet very simple and the main goal was to design new more complex models, which maintain the nice properties and resemble more the behavior of the agents in reality. The perhaps most interesting generalization of DeGroot's model came from the cooperation of the sociologist Noah Friedkin and the mathematician Eugene Johnsen. The article [[2]] was published in 1990 and since then a lot of research, both theoretical and experimental in nature, has been produced based on the FJ model as its core idea i.e. [[42], [3], [4], [5], [6], [7], [8]]. Some of these papers are also going to be studied later in this thesis. We are going to call such models, like DeGroot and FJ, linear models, due to the form of the update rules and the fact that we use tools from linear algebra, stochastic processes and graph theory to analyze the convergence properits of these models. Though, these models assume that the underlying social network does not change over time.

It seems closer to reality to assume that people might start taking into account the opinion of new agents, stop exchanging opinions with others or change the level of trust they put on some of their neighbors in a social network. Therefore, another line of research was developed, where the opinions of the agents and the underlying networks coevolve. In this category, the most influential models have been the Hegselmann-Krause model [[9]] in the deterministic case and the Deffuant-Weisbuch model [[10]] in the randomized one. These models might be of greater interest, but the evolving network makes the convergence analysis much more analysis. Fundamentally different techniques are required, such as [[11]] or the idea of the s-energy of a multiagent system [[12]]. Other works, such as [[13]] and [[14]], analyze such aforementioned opinion formation processes from a game-theoretic perspective: There is an underlying cost-minimization game and the agents try to myopically minimize their cost functions. Motivated by the fact that people in a society rarely reach an agreement about a topic, they try to quantify the disagreement cost incurred and analyze the social outcome of such games. Finally, binary opinion dynamics(this is the case, for example, in voting) has been studied [[15]], Bayesian approaches in social learning have been considered [[16]] and very recent research [[17]] might pose new interesting questions in the field. Nevertheless, social and economic networks, not only restricted in opinion dynamics, have become part of our everyday lives and understanding aspects of them is a major concern of the scientific community [[18]].

In this thesis, we are not going to discuss more about coevolutionary opinion formation models, but we are going to focus instead on linear models, closely related to the FJ model. A variety of mathematical tools has been used for the analysis of such dynamical systems. One of our attempts will be to relate opinion dynamics with convex optimization and show that algorithms and ideas from the convex optimization literature can help us gain insight in how the dynamics of such systems evolves over time and how and under which circumstances convergence can be reached. For instance, incorporating the idea of gradient descent(albeit from a different perspective) has already been done in markets [[19], [20]]. In conclusion, well-known convex optimization algorithms [[21]] as well as results in stochastic optimization can become a very useful tool, both in theory and in practice. Our other attempt, though directly linked with the previous one, will be to examine an analyze opinion formation models with *partial* information and possibly randomization. The motivation for such a research direction stems from the fact that social networks are growing larger and larger. Therefore, an agent will most likely exchange information about a topic with a small subset of her social neighborhood, since asking all of them becomes costly and unrealistic. Our goal is to design protocols with limited information exchange, which have a desirable expected behavior and whose rate of convergence to a fixed set of opinions is not much worse compared to the equivalent deterministic model. Consequently, we are also going to focus on randomized protocols, formulate them as stochastic optimization problems and analyze their convergence properties. Next, we give an overview of the organization of this thesis.

1.1 Brief overview of the chapters

This diploma thesis contains four main chapters.

In chapter 2, we are going to present all the basic definitions, theorems and notations that we are going to use in the other three chapters. We are going to discuss a bit more extensively about the concept of the potential function, since it will often appear later as a central idea in our proofs.

In chapter 3, we are going to introduce the most influential linear models, namely the DeGroot and the Friedkin-Johnsen model and briefly state their important properties without detailed proofs. Our focus will be mainly on more contemporary research work. For this purpose, we are going to examine two opinion formation models based on the FJ: A deterministic one with *outdated* information and an *asynchronous* randomized.

In chapter 4 we are going to bridge convex optimization and opinion dynamics. First we are going to show how first-order algorithms, mainly gradient descent, can be applied in the convergence analysis, then view opinion formation models as online repeated games and finally examine the limited information exchange through the lens of stochastic optimization.

Finally, in chapter 5 we are going to continue on the connection between stochastic optimization and the FJ model and alongside we are going to present another randomized variant, based again on the FJ.

2 Preliminaries

In this section, we are going to provide basic definitions and theorems which we are going to use throughout this thesis. We gather them all here, in the beginning, so that the reader can easily refer to them, whenever wished.

2.1 Linear Algebra

In chapter 3 we focus on opinion formation models and their convergence properties. Matrix properties and matrix norm properties are very useful in the convergence analysis of these models. In parts of chapter 5, we also apply similar concepts.

Definition 2.1 (Vector Norm)**.** A vector norm $\|x\|$ is any mapping from $\mathcal{R}^n$ to $\mathcal{R}$ with the following three properties:

1. $\|x\| > 0$, if $x \neq 0$
2. $\|\alpha x\| = |a|\|x\|$, for any $\alpha \in \mathcal{R}$
3. $\|x + y\| \leq \|x\| + \|y\|$

for any vector $x, y \in \mathcal{R}^n$.

The norm we are going to use the most is the Euclidean norm(or l_2 norm) which is defined as

$$\|x\|_2 = \left(\sum_{i=1}^{n} x_i^2 \right)^{1/2} = \sqrt{x^T \cdot x}$$

and can be interpreted as the length of a vector $x \in \mathcal{R}^n$. We will also define the following:

- The l_1-norm: $\|x\|_1 = \sqrt{\sum_i |x_i|}$
- The l_∞-norm: $\|x\|_\infty = \max_i |x_i|$
- The l_p-norm: $\|x\|_p = (\sum_i |x_i|^p)^{1/p}$, $p \geq 1$

Definition 2.2 (Dual norm)**.** Let $\|\cdot\|$ be any norm. Its dual norm is defined as

$$\|x\|_* = \max x^T y$$
$$s.t. \ \|y\| \leq 1.$$

Definition 2.3 (Matrix Norm)**.** . A matrix norm $\|A\|$ is any mapping from $\mathcal{R}^{n\times n}$ to $\mathcal{R}$ with the following three properties:

1. $\|A\| > 0$, if $A \neq 0$
2. $\|\alpha A\| = |a|\|A\|$, for any $\alpha \in \mathcal{R}$.
3. $\|A + B\| \leq \|A\| + \|B\|$

for any matrix $A, B \in \mathcal{R}^{n\times n}$.

One example is again the l_∞ norm, which is defined as $\|A\|_\infty = \max\limits_{1\leq i\leq n} \sum\limits_{j=1}^{n} |a_{ij}|$.

Definition 2.4 (Inner product)**.** Let V be a vector space over a field($\mathbb{R}$ or $\mathcal{C}$). We are going to denote the *inner product* as a function $\langle \cdot, \cdot \rangle : V \times V \to \mathcal{K}$

Regarding the eigenvalues and eigenvectors, as we know, the basic equation is $Ax = \lambda x$. The number λ is an eigenvalue of A. If for each eigenvalue we solve the system $(A - \lambda_i I)x_i = 0$, we can find the corresponding eigenvectors.

Definition 2.5 (Spectral radius)**.** The spectral radius of matrix A is defined as $\rho(A) = \max\{|\lambda| : \lambda \text{ is an eigenvalue of } A\}$

Lemma 2.1. For any matrix norm $\|\cdot\|$ and matrix A it holds:

$$\rho(A) = \lim_{k\to\infty} \left\|A^k\right\|^{1/k}$$

Lemma 2.2. If $\rho(A) < 1$, then $\lim\limits_{n\to\infty} \|A^n\| = 0$. If $\rho(A) > 1$, then $\lim\limits_{n\to\infty} \|A^n\| = \infty$.

Definition 2.6 (Primitive Matrix)**.** A *non-negative* square matrix A is called primitive if there is a k such that all the entries of A^k are positive.

Definition 2.7 (Irreducible matrix)**.** A *non-negative* square matrix A is called irreducible, if for any i, j there is a $k = k(i, j)$ such that $(A^k)_{ij} > 0$.

Theorem 2.3 (Perron-Frobenius theorem)**.** Let A be a non-negative primitive $r \times r$ matrix. There exists a real eigenvalue λ_1 with algebraic as well as geometric multiplicity one such that $\lambda_1 > 0$ and $\lambda_1 > |\lambda_j|$ for any other eigenvalue λ_j. Moreover, the left eigenvector u_1 and the right eigenvector v_1 associated with λ_1 can be chosen positive and such that $u_1^T v_1 = 1$.
Therefore, we can put the rest of the eigenvalues $\lambda_2, \ldots, \lambda_r$ in an order such that

$$\lambda_1 > |\lambda_2| \geq \cdots \geq \lambda_r.$$

Additionally, if A is stochastic, then $\lambda_1 = 1$. If A is substochastic, then $\lambda_1 < 1$.
If A is stochastic and irreducible with period $d > 1$, then there are exactly d distinct eigenvalues of modulus 1 and all other eigenvalues have modulus strictly less than 1.

2.2 Stochastic processes

A few definitions here are necessary, especially for the analysis of the DeGroot model [3.1] later.

Definition 2.8. We are going to call a *random walk* on a weighted (directed) graph the following process: First let $\forall i \in V : \sum_{(i,j)\in E} w_{ij} = 1$. We start from an arbitrary vertex v_0 at $t = 1$ and with probability w_{ij} we select one of the edges adjacent to v_0 and we proceed to the next vertex. At the next time step, we repeat the same process.

Definition 2.9 (Markov chain on G)**.** A finite discrete time Markov chain is a random walk on a weighted directed graph $G(V, E)$, s.t. $\forall i \in V : \sum_{(i,j)\in E} w_{ij} = 1$. Therefore, we define a sequence of random variables $X_0, X_1, \dots, X_t, \dots \in V$, s.t.

$$\Pr[X_t = v_i] = \Pr[\text{the Markov chain is at vertex } v_i \text{ at time step } t]$$

Note that we choose our next move, based solely on the current state of the chain, without using its history.

Definition 2.10 (Transition matrix)**.** Suppose that we start with an initial distribution $\pi(0) \in \mathbb{R}_{n\times 1}$. At time step t, we have π^t s.t. $\pi_i^t = \Pr[X_t = v_i]$ and the matrix $A_{n\times n}$ s.t. $A_{ij} = w_{ij}$. Then, we have

$$\pi^T(t) = \pi^T(t-1) \cdot A$$

and A is called the transition matrix.

Definition 2.11 (Stationary distribution)**.** A vector π^* for which it holds that $(\pi^*)^T = (\pi^*)^T \cdot A$, is called a stationary distribution. Basically, if a Markov chain has a stationary distribution, then from a time step t and onwards, this distribution will hold forever.

Definition 2.12 (Substochastic matrix)**.** A substochastic matrix is a square matrix A with non-negative entries, so that every row adds up to *at most* 1.

For the next definitions, we are going to be based on the graph G that we are given(alternative definitions also exist).

Definition 2.13 (Irreducibility)**.** A Markov chain is irreducible if and only if the underlying graph is strongly connected.

Definition 2.14 (Aperiodicity)**.** Let $l_1, l_2, \ldots, l_k$ be the lengths of all directed cycles of G. Then *periodicity* is simply the greatest common divisor(gcd) of all these lengths. A Markov chain is aperiodic if and only if its periodicity is 1.

Finally, we state the main important theorem that we are going to use later.

Theorem 2.4. Any *finite, irreducible and aperiodic* Markov chain with transition matrix A:

1. has a **unique** stationary distribution π^*
2. for **any** initial distribution $\pi(0)$ it holds that: $\lim_{t\to\infty}(\pi^T(0) \cdot A^t) = \pi^*$

2.3 Convex optimization

An important part of this thesis is devoted to convex optimization and opinion dynamics and how the former can be applied to the analysis of the latter. Here, we give the basic notation and definitions that we are going to encounter mainly in chapter 4.

Definition 2.15 (Convex function)**.** A function $f : \mathbb{R}^n \to \mathbb{R}$ is *convex* if $\mathbf{dom} f$ is a convex set and if for all x, $y \in \mathbf{dom} f$ and θ with $0 \leq \theta \leq 1$ we have

$$f(\theta x + (1-\theta)y) \leq \theta f(x) + (1-\theta) f(y)$$

Definition 2.16 (Sublevel set)**.** The α-sublevel set of a function $f : \mathbb{R}^n \to \mathbb{R}$ is defined as

$$C_\alpha = \{x \in \mathbf{dom} f | f(x) \leq \alpha\}$$

Definition 2.17 (Strong convexity)**.** If $f : \mathbb{R}^n \to \mathbb{R}$ is strongly convex on S, then there exists a $m > 0$ s.t.

$$\nabla^2 f(x) \succeq mI$$

for all $x \in S$.

Strong convexity has some nice properties that we are going to use, such as the following: For $x, y \in S$ we have

$$f(y) = f(x) + \nabla f(x)^T(y-x) + \frac{1}{2}(y-x)^T\nabla^2 f(z)(y-x)$$

for some z on the line segment $[x, y]$.

Definition 2.18 (Lipschitz function)**.** A function $f : S \to \mathbb{R}^m$, where $S \subset \mathbb{R}^n$, is called a Lipschitz function if there is a constant C(necessarily non-negative) s.t.

$$\|f(y) - f(x)\| \leq C\|y - x\|$$

for all $x, y \in S$.

Definition 2.19 (Smoothness condition)**.** A convex function $f : \mathbb{R}^n \to \mathbb{R}$ is *L-smooth* if $\forall x, y \in \mathbb{R}^n$ it satisfies

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|$$

Definition 2.20 (Euclidean ball)**.** A (Euclidean) ball in $\mathbb{R}^n$ has the form

$$B(x_c, r) = \{x | \|x - x_c\|_2 \leq r\} = \{x | (x - x_c)^T(x - x_c) \leq r^2\}$$

2.4 Probability

In this thesis, we examine not only full information protocols, but also opinion formation models with limited information exchange and the latter implies also randomization in the update procedure. Therefore, mainly in chapter 5, probabilistic arguments are needed and at this point it is necessary to state some important concentration inequalities.

Definition 2.21 (Variance of a random variable)**.** Let X be a continuous random variable with mean μ. The variance of X is

$$Var(X) = E((X - \mu)^2)$$

The definition is the same for discrete random variables. Important properties of the variance are the following:

1. If X and Y are independent then $Var(X+Y) = Var(X) + Var(Y)$.

2. For constants a and b, $Var(aX+b) = a^2 Var(X)$.

3. $Var(X) = E(X^2) - E(X)^2 = E(X^2) - \mu^2$.

Definition 2.22 (Union bound)**.** For any events $A_1, A_2, \ldots, A_n$ we have

$$\Pr\left(\bigcup_{i=1}^{n} A_i\right) \leq \sum_{i=1}^{n} \Pr(A_i)$$

Next, we are going to state some fundamental inequalities.

Theorem 2.5 (Chebyshev's inequality)**.** For any $\alpha > 0$,

$$\Pr(|X - E[X]| \geq a) \leq \frac{\mathrm{Var}[X]}{\alpha^2}$$

Next we are going to introduce the Chernoff bounds. Chernoff bounds can give *exponentially* decreasing bounds on the tail distribution and thus are widely used. As there are many different forms, we are going to start with the case of a sum of independent Bernoulli trials. Then we are going to give a more general form and finally we are going to state the fundamental Hoeffding's inequality.

Theorem 2.6 (Chernoff bounds)**.** Let $X = \sum_{i=1}^{n} X_i$, where $X_i = 1$ with probability p_i and $X_i = 0$ with probability $1 - p_i$ and all X_i are independent. Let $\mu = \mathbb{E}(X) = \sum_{i=1}^{n} p_i$. Then

- Upper tail: $\Pr(X \geq (1+\delta)\mu) \leq e^{-\frac{\delta^2}{2+\delta}\mu}$ for all $\delta > 0$
- Lower Tail: $\Pr(X \leq (1-\delta)\mu) \leq e^{-\frac{\mu\delta^2}{2}}$ for all $0 < \delta < 1$.

If we want to obtain a bound for the absolute value we can combine the two above inequalities and end up with the following:

Corollary 2.6.1. With X and $X_1, \ldots, X_n$ defined as before and $\mu = \mathbb{E}(X)$ it holds that

$$\Pr(|X - \mu| \geq \delta\mu) \leq 2e^{-\mu\delta^2/3} \text{ for all } 0 < \delta < 1.$$

The following bound applies to (bounded) random variables, regardless of their distribution.

Theorem 2.7. Let $X_1, X_2, \ldots, X_n$ be random variables s.t. $a \leq X_i \leq b$ for all i. Let $X = \sum_{i=1}^{n} X_i$ and set $\mu = \mathbb{E}(X)$. Then, for all $\delta > 0$:

- Upper tail: $\Pr(X \geq (1+\delta)\mu) \leq e^{-\frac{2\delta^2\mu^2}{n(b-a)^2}}$
- Lower tail: $\Pr(X \leq (1-\delta)\mu) \leq e^{-\frac{\delta^2\mu^2}{n(b-a)^2}}$

Finally we are going to state a very important bound for the probability that sums of bounded random variables are too large or too small.

Theorem 2.8 (Hoeffding's inequality)**.** Let $Z_1, Z_2, \ldots, Z_n$ be independent bounded random variables with $Z_i \in [a, b]$ for all i, where $-\infty < a < b < \infty$. Then it holds that

1. $\Pr\left(\frac{1}{n}\sum_{i=1}^{n}(Z_i - \mathbb{E}[Z_i]) \geq t\right) \leq e^{-\frac{2nt^2}{(b-a)^2}}$
2. $\Pr\left(\frac{1}{n}\sum_{i=1}^{n}(Z_i - \mathbb{E}[Z_i]) \leq -t\right) \leq e^{-\frac{2nt^2}{(b-a)^2}}$

for all $t \geq 0$.

2.5 Algorithmic Game Theory

Finally, we are going to present some basic notions which play a central role in the algorithmic game theory literature. Usually, as mentioned, we examine opinion formation models from a game-theoretic perspective: Given a social network, there is an underlying game with payoffs and selfish agents try to minimize their individual cost functions. In this section, we are going to define basic terms in the context of *cost-minimization* games and then focus on the important properties of the potential function(since it is sometimes our goal to model opinion exchange as a potential game). A cost minimization game consists of the following:

- a finite number k of players
- a finite strategy set S_i for each player i
- a cost function $C_i(s)$ for each player i, where $s \in S_1 \times S_2 \times \cdots \times S_k$ denotes a strategy profile or outcome.

Definition 2.23 (Pure Nash equilibrium)**.** A strategy profile s of a cost-minimization game is a *pure Nash equilibrium(PNE)* if for every player $i \in \{1, 2, \ldots, k\}$ and every unilateral deviation $s'_i \in S_i$,

$$C_i(s) \leq C_i(s'_i, s_{-i})$$

As we can see, it is a very strong notion, so it does not exist in all the games we examine. Directly linked with PNE is the following definition:

Definition 2.24 (Best-response dynamics)**.** *Best-response dynamics* is a procedure by which players search for a pure Nash equilibrium of a game. The procedure is simply the following:

- While the current outcome s is not a PNE:
  - Pick an arbitrary player i and an arbitrary beneficial deviation s'_i for player i and move to the outcome (s'_i, s_{-i}).

Best response dynamics can only halt at a PNE and it cycles if the game does not admit one.

Best-response, as we will analyze later, is suitable in potential games.
Sometimes, we need to relax the notion of PNE. Aiming for an approximate Nash Equilibrium can lead to faster and easier computation and is also an acceptable solution concept in many different settings in game theory.

Definition 2.25 (ϵ-Pure Nash Equilibrium)**.** For $\epsilon \in [0, 1]$, an outcome s of a cost-minimization game is an ϵ-pure Nash equilibrium(ϵ-PNE) if, for every player i and deviation $s'_i \in S_i$

$$C_i(s'_i, s_{-i}) \geq (1 - \epsilon) \cdot C_i(s)$$

Therefore, the difference with computing an exact NE is that a player might have an ϵ-move available to further decrease her cost function.

Accordingly, we can also define the ϵ-best response dynamics.

Definition 2.26 (ϵ-best response dynamics)**.** Using this procedure, players search for an ϵ-pure Nash equilibrium:

- While the current outcome s is not a ϵ-PNE:
  - Pick an arbitrary player i that has an ϵ-move, meaning a deviation s'_i with $C_i(s'_i, s_{-i}) < (1-\epsilon) \cdot C_i(s)$ and an arbitrary such move for player i, and move to the outcome (s'_i, s_{-i}).

Similar to before, ϵ-best response dynamics can halt only at an ϵ-PNE.

2.5.1 Games and Potential Functions

In this section, we are going to introduce the concept of the potential function, which is a very useful mathematical tool in optimization problems. First, we are going to give the necessary definitions and state some nice properties of the potential functions and the games which admit one. In fact, the concept of the potential function was indeed first introduced for the analysis of congestion games, in the seminal paper of Rosenthal[[22]]. There is a very rich bibliography around congestion games and we are not going to expand further here. Finally, as we already mentioned, the power of the potential function is not limited there. Due to its simplicity and usefulness, it appears in other elegant proofs in significant papers outside the field of congestion games.

2.5.1.1 Definitions and properties First, we are going to present the basic types of potential functions in the context of games. That is, we suppose that we have a game/model with n agents, defined as $G = (S_1, S_2, \ldots, S_n, C_1, C_2, \ldots, C_n)$. Each one has a cost function $C_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$. As we will later notice many times in this thesis, this is the case in *opinion dynamics* too.

Definition 2.27 (Exact Potential Function)**.** In the game we described before, suppose there exists a function $\Phi : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$, such that $\forall i \in 1, 2, \ldots, n$ and $\forall s_{-i} \in S_{-i}$, $\forall s_i, s_i' \in S_i$:

$$\Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}) = C_i(s_i, s_{-i}) - C_i(s_i', s_{-i}) \tag{8}$$

then we call Φ an exact potential function. A game which admits a potential function is subsequently called *potential game*.

Intuitively, the potential function captures globally the changes in the game. When a player deviates from her current strategy and the value of her cost function changes, the value of the potential function decreases or increases by exactly the same amount, regardless of which player deviates.

For completeness we are going to give two other definitions which generalize the concept of the potential function.

Definition 2.28 (Weighted Potential Function)**.** In the game we described before, suppose there exists a function $\Phi : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ and a vector $w \in \mathbb{R}_+^n$, such that $\forall i \in 1, 2, \ldots, n$ and $\forall s_{-i} \in S_{-i}$, $\forall s_i, s_i' \in S_i$:

$$\Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}) = w_i \left(C_i(s_i, s_{-i}) - C_i(s_i', s_{-i}) \right) \tag{9}$$

then Φ is called a *weighted potential function.*

We can further relax the concept by defining the following:

Definition 2.29 (Ordinal Potential Function)**.** In the game we described before, suppose there exists a function $\Phi : S_1 \times S_2 \times \cdots \times S_n \rightarrow \mathbb{R}$, such that $\forall i \in 1, 2, \ldots, n$ and $\forall s_{-i} \in S_{-i}$, $\forall s_i, s'_i \in S_i$:

$$C_i(s_i, s_{-i}) - C_i(s'_i, s_{-i}) > 0 \Longleftrightarrow \Phi(s_i, s_{-i}) - \Phi(s'_i, s_{-i}) > 0 \tag{10}$$

then we call Φ an ordinal potential function.

In other words, we take into account only the sign, and the potential function increases or decreases when a player's cost function increases or decreases accordingly.

Potential games were introduced in the seminal work of Monderer et al. [[23]] In this paper, it is stated the following important theorem:

Theorem 2.9. Every potential game admits (at least) one pure Nash Equilibrium.

Proof. Let s be a pure profile minimizing Φ. We will show that s is a Nash Equilibrium. Suppose by contradiction that s is not a NE. Thus, player i can improve by deviating to a new profile s' and subsequently her cost function will decrease. Since we have a potential game, it also holds that $\Phi(s') - \Phi(s) < 0$. This means that $\Phi(s') < \Phi(s)$, contradicting the fact the initial assumption that s minimizes Φ.

Corollary 2.9.1. If Φ is the potential function of our game, then every Nash Equilibrium is a local optimum of Φ.

The above simple proof allows us also to state the following proposition:

Corollary 2.9.2. In a finite potential game, from an arbitrary initial outcome, best-response dynamics converges to a pure Nash Equilibrium.

Proof. We remember from the definition of the best-response [2.24], that the cost function of an agent i who deviates from her strategy strictly decreases. Hence, the potential function also strictly decreases. Since we stated that the game is finite, best-response dynamics eventually halts and this will be for sure at a pure Nash Equilibrium.

3 Opinion formation models and their convergence properties

Human interactions give rise to the formation of different kinds of opinion in a society. The study of formations and dynamics of opinions has been one of the most important priorities in various scientific fields and recently also in computer science. The process by which new ideas, innovations and behaviors spread through a large social network can be thought of as a network interaction game. We model mathematically the opinions of the agents and we try to understand how opinions form and spread in social networks. We are mainly interested in designing models which express how the agents iteratively update their opinions and answer two essential questions in relation to them: First, if an equilibrium can emerge as a result of such interactions between the agents and under which assumptions this can hold. If there exists an equilibrium point, can agents reach a consensus about a certain topic? If we can answer positively the first question, we also try to define the convergence speed to such equilibrium, which is also a crucial factor that characterizes the efficiency of our model.

In this chapter, we are going to present some important opinion formation models, from the very beginning till contemporary work.

3.1 The De-Groot Model

The seminal network interaction model of information transmission, opinion formation and consensus formation is due to DeGroot [[1]]. DeGroot proposed a simple and natural model of how to reach a consensus. Individuals in a society start with initial opinions on a subject. Let these be represented by an n-dimensional vector of probabilities $x(0) = (x_1(0), x_2(0), ..., x_n(0))$. Each $x_i(0)$ lies in the interval [0,1] and can be thought of as an opinion of an agent about a topic. For example how much we believe that a certain project will have a positive outcome or how much we liked a restaurant that we recently visited. The social network is represented by a weighted directed graph $G(V, E)$ at which each node represents a node in the social network and the weight w_{ij} of each edge in the graph represents the trust that node i has to node j. There can also exist self loops w_{ii} in the graph that represent the stubborness of each node. Without loss of generality we consider that $\forall i \in V : \sum\limits_{(i,j)\in E} w_{ij} = 1$. At first all the nodes have an initial opinion which is denoted by the vector $x(0)$ and the update rule at each time step takes the form $x(t+1) = T \cdot x(t)$, where T is the updating matrix of DeGroot model. Matrix T contains the weights of the edges of our graph $G(V, E)$ and it is defined as:

$$T_{ij} = \begin{cases} w_{ij} & \text{if } (i,j) \in E \\ 0 & \text{if } (i,j) \notin E \end{cases}$$

Note here that the graph can be either directed or undirected which yields an asymmetric and a symmetric version of the model respectively. Furthermore, let T be a (row) stochastic matrix so that its entries across each row sum to one. Now we are ready to ask the two questions that concern us in opinion formation models:

- Under what conditions will the updating process we described before converge to a well-defined limit and what limit does it converge to?
- How fast does DeGroot model converge to an equilibrium point?

Following we are going to give an answer to the first question.

Given the vector $x(0)$ of the initial opinions, if we apply iteratively the update rule, after t time steps we get:

$$x(t) = T^t x(0) \tag{11}$$

Since T is stochastic we can apply to the DeGroot model nice properties from Markov chain theory. We are going to give the central theorem(proven in [[18]]) which shows under which conditions the opinion formation in the long run converges to a stable state.

Theorem 3.1. The DeGroot model with n agents converges to the unique equilibrium point

$$x^* = \lim_{t\to\infty} x(t) = \lim_{t\to\infty} T^t x(0) \tag{12}$$

for any initial vector of opinions $x(0) \in [0,1]^n$, if and only if the Markov chain with the transition matrix T is irreducible and aperiodic.

The definitions of irreducibility and aperiodicity are given in the introduction [2.13] [2.14].

In our setting these two conditions imply that for every agent i, there exists a time t_0 such that for every time $t \geq t_0$, i is influenced by all other agents. This is equivalent to the matrix T^{t_0} having only positive elements. Since

$$x^* = Tx^*$$

we call the limit vector x^* the Nash equilibrium of the model.

To make our point about convergence more clear we are going to give two examples, one that converges in the long run and one that does not.

Example 1. Suppose that we have 3 agents with the following edges and the respective weights:

This creates the following transition matrix:

$$\mathrm{T} = \begin{bmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

By updating the beliefs according to the DeGroot model we get:

$$T^2 = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \end{bmatrix} \qquad T^3 = \begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \end{bmatrix} \qquad T^4 = \begin{bmatrix} 1/4 & 1/2 & 1/4 \\ 1/2 & 1/4 & 1/4 \\ 1/2 & 1/2 & 0 \end{bmatrix}$$

If we iterate many times we finally get: $T^t = \begin{bmatrix} 2/5 & 2/5 & 1/5 \\ 2/5 & 2/5 & 1/5 \\ 2/5 & 2/5 & 1/5 \end{bmatrix}$

Therefore we have $\lim_{t\to\propto} A^t x(0) = \begin{bmatrix} 2/5 & 2/5 & 1/5 \\ 2/5 & 2/5 & 1/5 \\ 2/5 & 2/5 & 1/5 \end{bmatrix} \begin{bmatrix} x_1(0) \\ x_2(0) \\ x_3(0) \end{bmatrix}$ and the system converges to the unique equilibrium point:$x^* = \lim_{t\to\propto} A^t x(0) = \frac{2}{5}x_1(0) + \frac{2}{5}x_2(0) + \frac{1}{5}x_3(0)$.

This is indeed the stationary distribution of the Markov chain $\pi = \begin{bmatrix} \frac{2}{5} & \frac{2}{5} & \frac{1}{5} \end{bmatrix}$ because as we can check it holds that $\pi T = \begin{bmatrix} \frac{2}{5} & \frac{2}{5} & \frac{1}{5} \end{bmatrix} \begin{bmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} \frac{2}{5} & \frac{2}{5} & \frac{1}{5} \end{bmatrix} = \pi$

As we can also notice, the agents finally reach a consensus and the first two agents have twice the infuence that agent 3 has as we reach consensus in the long run.

Following we are going to illustrate an example, through which it will become more clear that if certain properties do not hold, the agents are never going to converge to this unique equilibrium point.

Example 2. We take the above graph and we change the outgoing edge of the third agent. Now she does not communicate any more with agent 2 but she listens only to agent 1. The transition matrix now becomes:

$$T = \begin{bmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

If we start again implementing the update rule we get:

$$T^2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{bmatrix} \qquad T^3 = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \qquad T^4 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{bmatrix}$$

If we continue further, the update matrix follows the exact same pattern: It oscillates and there is no convergence. This makes sense, since the first agent updates her opinion based equally on the opinions of the other two agents and the other two agents are based solely on agent 1. Thus they just end up swapping their opinions in each time step. The property of *aperiodicity* that we mentioned before does not hold here. All the cycles that we detect in this directed graph have a length which is a multiple of 2. This makes our matrix T *periodic* and convergence cannot be reached. So, we are stating again the important conclusion:

Corollary 3.1.1. A consensus is reached in the DeGroot model if and only if the group is *strongly connected* and *aperiodic*.

Next, we are going to study a variation of the DeGroot model and demonstrate all the important properties which yield the convergence rate. We are also going to analyze the convergence time of the above model by viewing it as a special case of the FJ-model.

3.2 The Friedkin-Johnsen Model

Another very important *linear* opinion formation model was introduced in the seminal work of Friedkin and Johnsen [[2]]. It can be viewed as a variation of the DeGroot model, which captures in a more natural way an opinion formation process.

The difference in FJ is that each agent holds also an *intrinsic* opinion s_i, which is defined in the beginning of the process, remains the same and participates with a weight factor in the update rule of each agent. It is natural to assume that this internal opinion exists when people exchange opinions about a certain opinions about a certain topic: It can be the bias that an agent has towards an issue, or the fact that the background each agent has, has already made them partially shape a belief about a given situation and as a result less conciliatory when other agents express their opinions to her.

Together with the intrinsic opinion s_i, the update rule takes now the following form:

$$x_i(t+1) = \sum_{\substack{j \in N_i \\ j \neq i}} w_{ij} x_j(t) + w_{ii} s_i$$

We can also write the above equation in matrix form by defining:

- Matrix **A** which has as elements all the w_{ij}'s and its diagonal elements are 0.
- Matrix **B**, for which it holds that $B_{ii} = w_{ii}$. The $w_i i$'s can be described as the *level of stubbornness* of the agents towards their intrinsic opinions.
- Vector **s**, which is the vector of the agents' intrinsic opinions.

Now the update rule becomes:

$$x(t+1) = Ax(t) + Bs \tag{13}$$

If we recursively apply the update rule we get that the vector of opinions at each time $t \geq 0$ is

$$x(t) = A^t s + \sum_{n=0}^{t-1} A^n Bs \tag{14}$$

In their work Ghaderi and Srikant [[3]] give an answer to the two questions we posed also for the DeGroot model. Next we are going to examine the existence of equilibrium and if there exists one, also the convergence time of the dynamics, in two different cases.

First, we will assume that we have no stubborn agents. This means that B = 0 and A is a row-stochastic matrix. The update rule becomes $x(t+1) = A^t s$, so it resembles the DeGroot model. Therefore, we know that our system admits a unique equilibrium point and it is left to show the convergence rate. Before that, we are going to state a lemma (similar to lemma 1 in [[3]]) for a special case:

Lemma 3.2. Consider a social network with no stubborn agents and opinions vector **s**. If for each agent i, the weights of edges connecting i to her neighborhood are of the form $1/d_i$, where d_i is the cardinality of the neighborhood. then the dynamics will converge to the following unique equilibrium

$$x_i(\infty) = \frac{1}{2|E|}\sum_{j=1}^{n}|N_j|s_j \ , \forall i \in V$$

Therefore, all the agents will adopt a unique opinion which is a convex combination of their initial opinions. Each agents participates in this final opinion with a weight proportional to the cardinality of her neighborhood.

We return again to the general case and continue on the convergence rate of the dynamics we described. In order to prove the convergence rate we are going to introduce a couple of helpful symbols. If π is the stationary distribution of the random walk on the social network graph, we define a norm in the vector space $\mathbb{R}^n$ with respect to the vector π, given by the following inner product

$$\langle z, y\rangle_\pi = \sum_{i=1}^{r} z(i)y(i)\pi(i)$$

$$\|z\|_\pi = \sqrt{\sum_{i=1}^{r} z(i)^2\pi(i)}$$

Additionally, we define the error as a vector

$$e(t) = x(t) - x(\infty)$$

For the following lemma we are not going to give the proof, which can be found in [[3]].

Lemma 3.3. In the case of no stubborn agents, for the error it holds that

$$\|e(t)\|_\pi \leq \rho_2^t \|e(0)\|_\pi \tag{15}$$

where $\rho_2 = \max_{i\neq 1}|\lambda_i|$ is the second largest eigenvalue of A

Finally, we can state the convergence time. A natural way to define the convergence time is the first time step that our error becomes less or equal to a very small positive number that we have defined:

$$\tau(v) = \inf\{t \geq 0 : \|e(t)\|_\pi \leq v\}$$

It is shown again in [[3]] that we can bound the convergence time as follows:

$$\left(\frac{1}{1-\rho_2}-1\right)\log\left(\frac{\|e(0)\|_\pi}{v}\right) \leq \tau(v) \leq \frac{1}{1-\rho_2}\log\left(\frac{\|e(0)\|_\pi}{v}\right) \tag{16}$$

In other words, the convergence time is $\Theta\left(\frac{1}{1-\rho_2}\right)$ as the number of agents grows.

In case of the stubborn agents, we demand that at least one agent possesses a level of stubbornness, so that B won't be zero again. In a connected social network graph, where matrix A has at least one row with sum less than one, this also means that A is *sub-stochastic*. This is an important property because we now know that all the eigenvalues of the matrix are less than one and hence $\lim_{t\to\infty} A^t = 0$. Again, by Perron-Frobenius theorem the largest eigenvalue is positive and real i.e. $0 < \lambda_1 < 1$. We define also the *spectral radius* of A as $\rho_1(A) = \lambda_1$. As a result, the equilibrium exists and is given as following

$$x(\infty) = \lim_{t\to\infty} x(t) = \sum_{n=0}^{\infty} A^n B s = (I-A)^{-1} B s \tag{17}$$

It is also worth noticing that since $B_{ii} = 0$ for all non-stubborn agents i, the initial opinions of non-stubborn agents will vanish eventually and have no effect on the equilibrium.

Finally the matrix form does not help us understand how the graph structure and the stubborn agents. In Ghaderi et. al [[3]] one can find out that by using standard arguments from stochastic processes such as random walks and hitting times, it can be proved that the above model converges to a unique equilibrium where the opinion of each agent is a convex combination of the initial opinions of the stubborn agents.

First of all, in order to characterize the convergence time we demand again that there is at least one stubborn agent. We introduce again the same error vector as before $e(t)$. We can also consider two groups of stubborn agents, the ones that are *fully stubborn* and the ones that are *partially stubborn*. For the first group it trivially holds that $e_i(t) = 0$. For the partially stubborn and the non-stubborn agents, using similar arguments as before, we can prove that the error reduces geometrically in each time step.

Lemma 3.4. The convergence of [13] to the equilibrium is geometric with a rate equal to the largest eigenvalue of A.

$$\|e(t)\|_\pi \leq (\lambda_A)^t \|e(0)\|_\pi \tag{18}$$

In order to bound the convergence time(we define it in exactly the same way) we use similar techniques as before and we obtain

$$\left(\frac{1}{1-\lambda_A}-1\right)\log\left(\frac{\|e(0)\|_\pi}{v}\right) \leq \tau(v) \leq \frac{1}{1-\lambda_A}\log\left(\frac{\|e(0)\|_\pi}{v}\right)$$

So, again we have $\tau(v) = \Theta\left(\frac{1}{1-\lambda_A}\right)$ as the number n of agents grows.

This concludes our analysis for the FJ-model. Above, we examined the static version of the FJ-model. This means that we assumed that the matrices A and B do not change over time. Though, a natural interpretation of how opinions are formed, is that the underlying social network graph also changes. The agents might discuss over a sequence of issues, or change the degree to which they are influenced by some of their neighbors. The analysis of such cases becomes nevertheless significantly more difficult and sometimes it becomes obscure how we can derive provable guarantees for the equilibrium and the convergence speed. Some variations of the FJ-model have been studied in the recent years and we refer the reader to a variesty of papers, such as [[7]] to find out more about the time-varying variations of FJ.

3.3 Variants of the FJ-model

In this section, we are going to discuss how we can extend the core idea of the FJ model in order to create new models. For this purpose we are going to study two models which are inspired from the FJ model. One will be deterministic and the second randomized.

3.3.1 A deterministic variant with outdated information

In this section we are going to briefly examine the convergence properties of a variation of the FJ model. As we saw, the Friedkin-Johnsen model has been studied extensively and admits some very nice properties. The motivation of the following work by Fotakis et al. [[8]] is that the FJ model does not capture global properties of a society which in many cases influence an agent's opinion. People get exposed constantly to public opinions and trends of the society. Take as an example a case where people can choose over a set of alternatives(i.e. elections). In this case, people take into account surveys and polls, which are viewed as the consensus view of the society and might significantly influence the opinion that they finally form. As a result, it makes sense to introduce a term which describes the influence that the opinion of the society has over the agents' individual opinions and quantify also this influence.

We are now ready to describe our setting. Following the work of Kleinberg et. al [[13]], we strive also to design a similar repeated game. In these kind of games, the goal is for the agents to reach an equilibrium point which is also a consensus, in a sense that all the agents adopt the same opinion. Of course, as shown also in the FJ model, the factors that we model in order to capture a natural behavior which seems close to reality, make consensus very difficult. Though, we still want the agents to reach an equilibrium point after a finite number of rounds. For this purpose, we introduce a cost function that each agent holds and we describe the dynamical system as a cost minimization game. The goal of the cost function is to capture the disagreement in the society. The agents try to minimize this disagreement and therefore try myopically to optimize their cost functions. The result of this parallel best-response dynamics becomes the update rule of our model. In [[13]], the repeated averaging of the neighboring opinions which occurs in the FJ model is viewed as a best response to a quadratic cost function. We will do the same, by extending this formulation and also incorporating the combined view of the society. Following [[8]], we start by introducing the cost funtion, which for agent i takes the form:

$$C_i(\vec{x}) = \sum_{j \in N(i)} w_{ij}(x_j - x_i)^2 + w_{ii}(x_i - s_i)^2 + d_i \left(\frac{\sum_{j=1}^{n} z_j}{n} - s_i \right)^2 \tag{19}$$

where

- $\vec{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ is the vector of the expressed opinions.
- We assume that the consensus view of the society is the average of the opinions. This is a rule that we can handle and captures many cases in practice, at least the ones where each agent participates the same in shaping the final opinion. Of course, there exists alternatives which can also model the general view of the societ.
- the factor $d_i > 0$ shows how much weight an agent puts on the consensus view of the society. Furthermore, since an agent tries to minimize a "disagreemnt" cost, a big d_i shows an agent's intention to drag the average opinion closer to her intrinsic opinion.

If the agents play simultaneously according to the best-response dynamics, for each agent i it holds that

$$\frac{\partial C_i(\vec{x})}{\partial x_i} = 0 \Leftrightarrow 2 \sum_{j \in N(i)} w_{ij}(x_i - x_j) + 2w_{ii}(x_i - s_i) + \frac{2d_i}{n} \left(\frac{\sum_{j \neq i} x_j}{n} + \frac{x_i}{n} - s_i \right) = 0$$

which leads to the following update rule:

$$x_i = \frac{\sum_{j \neq i} \left(w_{ij} - \frac{d_i}{n^2} \right) x_j + \left(w_{ii} + \frac{d_i}{n} \right) s_i}{\sum_{j \neq i} w_{ij} + w_i + \frac{d_i}{n^2}} \tag{20}$$

We can notice in the equation above that i's influence from some opinions x_j can be negative. The results we are going to present hold also for the case of negative influence. We are going to examine the convergence properties in the case where the aforementioned term can be negative and assuming that outdated information might be present. Because of these negative values, the opinions might also take values outside the $[0, 1]$ interval. We therefore allow also the opinions of the agents to take any real value. Following the formulation of the FJ model we are going to write the update rule in matrix form and then take advantage mainly of matrix norm properties to provide the convergence guarantees.

First, let's make one natural assumption. It seems costly and impractical to announce the average opinion of the society to the agents at each time step. For example, a poll about the voting preferences in the upcoming elections will not be updated every day, but possibly every month and it might change a lot from time to time. Though, the agents have to decide based on possibly *outdated* information. Therefore, we assume that the average opinion will be announced to the agents every T rounds and we are going to refer to the time interval between two announcements as an *epoch*. The length of the epoch is fixed and epoch D lasts $L = T_{D+1} - T_D$. Let's define the following matrices:

- A where the diagonal elements will be zero and $A_{ij} = \frac{w_{ij}}{\sum_{j \neq i} w_{ij} + w_i + d_i/n^2}$.
- B is another square matrix, for which it holds that $B_{ij} = \frac{d_i/n^2}{\sum_{j \neq i} w_{ij} + w_i + d_i/n^2}$.
- Finally C will be a diagonal matrix with $C_{ii} = \frac{w_{ii} + d_i/n}{\sum_{j \neq i} w_{ij} + w_i + d_i/n^2}$

Then we can rewrite the update rule in matrix form as follows:

$$x(t) = Ax(t-1) + Bx(T_D) + Cs \tag{21}$$

Next, we are going to show that the above model can converge ϵ-close to the equilibrium. In order to state the convergence rate, we are going to examine the total decrease we achieve after each epoch. First, it is important to bound the l_∞ norms of A and B. We need to show that both can be less than 1, following the proof of the FJ model. For this purpose, we will have to make extra assumptions/add a couple constraints and we are going to give them a natural interpretation.

Remark 1. Assume that we have n agents and w_{ii} is the stubbornness of agent i. We do not want the agents to be overly stubborn, but we can also assume that they have at least some confidence on their intrinsic opinion s_i. Therefore, we can assume that $w_{ii} \geq \dfrac{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2}{n}$. Then $\|A\|_\infty \leq 1 - \dfrac{1}{n}$.

Proof. We know that $\|A\|_\infty = \max_i \left| \dfrac{\sum\limits_{i\neq j} w_{ij}}{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} \right|$. By adding and subtracting w_{ii} and d_i/n^2 we get that:

$$\begin{aligned}
\|A\|_\infty &= 1 - \frac{w_{ii}}{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} - \frac{d_i/n^2}{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} \\
&\leq 1 - \frac{w_{ii}}{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} \leq 1 - \frac{w_{ii}}{nw_{ii}} = 1 - \frac{1}{n}
\end{aligned}$$

Remark 2. As mentioned, the agents follow also the average opinion of the society which participates in the update rule. Though we do not wish the average opinion to be the dominant factor in forming their opinion in each time step. For this purpose, it is convenient to upper bound the weight factor d_i so that it holds: $d_i \leq \sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2$. In any case, different assumptions in the two cases might give us also desirable properties. The two that we posed seem also to have a natural interpretation.

Proof. We have that $\|B\|_\infty = \max_i \left| \dfrac{-\sum\limits_{j\neq i} d_i/n^2}{\sum\limits_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} \right|$. By using our last assumption we have for that agent i:

$$\|B\|_\infty = \frac{(n-1)d_i/n^2}{\sum_{j\neq i} w_{ij} + w_{ii} + d_i/n^2} \leq \frac{(n-1)d_i}{n^2 d_i} = \frac{n-1}{n^2}$$

Now we can proceed to the main theorem (variation of 3.1. in [8]).

Theorem 3.5 (Convergence with outdated information)**.** Fix a duration L for each epoch and assume that at the beginning of each epoch the average opinion of the society is announced to the agents. Our model is described by the update rule $x(t) = Ax(t-1) + Bx(T_D) + Cs$ and we assume, as shown, that $\|A\|_\infty \leq 1 - \frac{1}{n}$ and $\|B\|_\infty \leq \frac{n-1}{n^2}$. Then for any matrix C, any vector of initial opinions $\vec{s} \in [0,1]^n$ and any $\varepsilon > 0$ the above model converges ε-close to the unique equilibrium point $x^* = (\mathbb{I} - (A+B))^{-1}Cs$ after $\mathcal{O}\left(n^2 \ln \frac{n^2\|C\|_\infty}{\varepsilon}\right)$ epochs.

Proof. Suppose that we are at epoch D and we have proceeded r rounds in this epoch. We define also again the error vector as before, so that $e(t) = \|x(t) - x^*\|_\infty$. Then we have:

$$e(T_D + r) = \|Ax(T_D + r - 1) + Bx(T_D) + Cs - Ax^* - Bx^* - Cs\|_\infty$$

By the properties of the norm, we obtain:

$$\begin{aligned} e(T_D + r) &\leq \|A\|_\infty \|x(T_D + r - 1) - x^*\|_\infty + \|B\|_\infty \|x(T_D) - x^*\|_\infty \\ &\leq \left(1 - \frac{1}{n}\right) e(T_D + r - 1) + \frac{n-1}{n^2} e(T_D) \\ &\leq \left(1 - \frac{1}{n}\right)^2 e(T_D + r - 2) + \left(1 - \frac{1}{n}\right)\left(\frac{n-1}{n^2}\right) e(T_D) + \left(\frac{n-1}{n^2}\right) e(T_D) \end{aligned}$$

where in the second inequality we applied the update rule along with the matrix norm properties once more. If we apply iteratively the update rule we end up with:

$$e(T_D + r) \leq$$
$$\left(1 - \frac{1}{n}\right)^r e(T_D) + \left(\frac{n-1}{n^2}\right)\left(1 + \left(1 - \frac{1}{n}\right) + \left(1 - \frac{1}{n}\right)^2 + \cdots + \left(1 - \frac{1}{n}\right)^{r-1}\right) e(T_D)$$

Using the fact that for the first n terms of a geometric series it holds that

$$\sum_{k=0}^{r-1} \left(1 - \frac{1}{n}\right)^r = \frac{1 - (1 - 1/n)^r}{1/n}$$

we then obtain that

$$e(T_D + r) \leq \left(\left(1 - \frac{1}{n}\right)^r + \left(1 - \frac{1}{n}\right) \left(1 - \left(1 - \frac{1}{n}\right)^r\right)\right) e(T_D)$$
$$= \left(1 - \frac{1}{n} + \frac{1}{n}\left(1 - \frac{1}{n}\right)^r\right) e(T_D)$$

We want the quantity $\left(1 - \frac{1}{n} + \frac{1}{n}\left(1 - \frac{1}{n}\right)^r\right)$ to be less than 1, so that we have a decrease. Therefore:

$$\left(1 - \frac{1}{n} + \frac{1}{n}\left(1 - \frac{1}{n}\right)^r\right) < 1 \Leftrightarrow$$
$$\left(1 - \frac{1}{n}\right)^r < 1$$

which holds. Since we do an analysis in epochs, let's try to providee a bound for what happens from the last round of an epoch till the last round of the next epoch. We assume that each epoch lasts L rounds. Then we can give the following loose upper bound:

$$1 - \frac{1}{n} > \left(1 - \frac{1}{n}\right)^r \Leftrightarrow$$
$$\frac{1}{n} - \frac{1}{n^2} > \frac{1}{n}\left(1 - \frac{1}{n}\right)^r \Leftrightarrow$$
$$1 - \frac{1}{n^2} > 1 - \frac{1}{n} + \frac{1}{n}\left(1 - \frac{1}{n}\right)^r$$

Therefore, if we are at epoch D, we can guarantee that $e(T_{D+1}) \leq \left(1 - \frac{1}{n^2}\right) e(T_D)$. Finally, we will examine after how many *epochs* we can get ε-close to the equilibrium point. Following the proof for the FJ model [[3]], we know that in our model $x^* = (\mathbb{I} - (A + B))^{-1}Cs$ is the unique solution. Proceeding further, we need first to bound $e(0)$. First of all it holds that:

$$\|A + B\|_\infty \leq \|A\|_\infty + \|B\|_\infty \leq 1 - \frac{1}{n} + \frac{1}{n} - \frac{1}{n^2} = 1 - \frac{1}{n^2}$$

Using the fact that $\sum_{n=0}^{\infty} A = (\mathbb{I} - A)^{-1}$ we obtain:

$$\left\|(\mathbb{I} - (A+B))^{-1}\right\|_\infty \leq \sum_{k=0}^{\infty} \|A+B\|_\infty \leq \sum_{k=0}^{\infty} \left(1 - \frac{1}{n^2}\right)^k = n^2$$

Then due to the fact that $s \in [0,1]^n$ we have that $\|s\|_\infty \leq 1$ and

$$\|x^*\|_\infty = \left\|\mathbb{I} - (A+B)^{-1} C s\right\|_\infty \leq n^2 \|C\|_\infty$$

Finally, we can conclude that $e(0) = \|s - x^*\|_\infty \leq 1 + n^2\|C\|_\infty$. Using lemma [3.4], we need $\mathcal{O}\left(n^2 \ln \frac{n^2\|C\|_\infty}{\varepsilon}\right)$ epochs to get ε-away from the equilibrium point.

This concludes the proof.

3.3.2 A randomized asynchronous variant

Next, we are going to analyze a protocol proposed by Frasca et. al [[6]], which is also based on the seminal FJ model. There are two key differences: First, the update rule includes randomization, which means that in each round the agents which interact are chosen uniformly at random. Second the proposed dynamics is asynchronous, in contrast to the models we studied till now, which assume that all the agents apply the update rule simultaneously. This might be interpreted as an attempt to capture a more natural behavior in social networks, where the users are not likely to communicate with all of their neighbors and at the same time. Let's first describe the high level idea of this opinion formation model: At each time step a *randomly* chosen agent updates her opinion according to a convex combination of her own opinion(the one she holds right now), the opinion of one of her neighbors(also randomly chosen) and her intrinsic belief. The main conclusion of this approach and what we are going to prove is that the expected dynamics converges to a stable state and more specifically to the limit opinions of the synchronous dynamics, or in other words the FJ model. Therefore, we also conclude that the expected beliefs of the agents do not reach a consensus, which means that agreement in the society is not achieved.

We can formulate this asynchronous model as follows: Each agent starts as usual with her intrinsic opinion $x_i(0) = s_i \in [0,1]$. At each time step t, we sample uniformly at random an edge from the set of all the edges of our graph. Then, the two agents across the selected edge meet and i updates her opinion according to the rule we described before:

$$x_i(t+1) = \gamma_i((1-\lambda_{ij})x_i(k) + \lambda_{ij}x_j(k)) + (1-\gamma_i)u_i$$

$$x_j(k+1) = x_j(k) \quad \forall j \neq i \tag{22}$$

For the above model, we are going to study the expected dynamics of it. Before proceeding to the main theorem, we are going to make a strong but useful assumption in order to draw our conclusions: From now on, we assume that $\lambda_{ij} \in [0,1]$(which makes sense for the convex combination that we have) for all the edges and more importantly that Λ is a row-stochastic matrix.

Theorem 3.6 ([6])**.** The expected behavior of the above procedure can be viewed as a version of the FJ model. More specifically, if we examine the matrix form of the update rule it holds that

$$\mathbb{E}[x(t+1)] = \mathbb{E}[A(t)]\mathbb{E}[x(t)] + \mathbb{E}[B(t)]s$$

where:

- $\mathbb{E}[A(t)] = \mathbb{I} - \dfrac{1}{|E|}(D(I-\Gamma) + \Gamma(\mathbb{I}-\Lambda)$. Here we also define D, which is a diagonal matrix with each d_{ii} standing for the degree of agent i.
- $\mathbb{E}[B(t)] = \dfrac{1}{|E|}D(\mathbb{I}-\Gamma)$, where D is defined as before.

Proof. Let's begin by writing the update rule [22] in matrix form. We use the standard basis vectors to describe the fact that only agent i updates her opinion from this round to the next and we can easily notice that:

$$x(t+1) = (\mathbb{I} - e_ie_i^T(\mathbb{I}-\Gamma))(\mathbb{I} + \gamma_{ij}e_ie_j^T - \gamma_{ij}e_ie_i^T)x(t) + e_ie_i^T(\mathbb{I}-\Gamma)s$$

Now, as we see, suppose that at time step t the edge (i,j) was chosen. Since we choose uniformly at random one edge we can write the above in the following simplified form:

$$x(t+1) = A(t)x(t) + B(t)s$$

where at time step the probability of choosing A_{ij} and B_{ij} respectively is $\dfrac{1}{|E|}$. The expected value of the above rule is

$$\mathbb{E}[x(t+1)] = \mathbb{E}[A(t)]\mathbb{E}[x(t)] + \mathbb{E}[B(t)]s$$

and let's calculate the expected value of each matrix separately. For $A(t)$ we have that

$$\begin{aligned}\mathbb{E}[A(t)] &= \frac{1}{|E|}\sum_{(i,j)\in E} A_{ij} \\ &= \frac{1}{|E|}\sum_{i<j}\sum_{j\in N_i}(\mathbb{I}-e_ie_i^T(\mathbb{I}-\Gamma))(\mathbb{I}+\gamma_{ij}e_ie_j^T-\gamma_{ij}e_ie_i^T) \\ &= \frac{1}{|E|}\sum_{i<j}\sum_{j\in N_i}(\mathbb{I}+\gamma_{ij}(e_ie_j^T-e_ie_i^T)-e_ie_i^T(\mathbb{I}-\Gamma)-e_ie_i^T(\mathbb{I}-\Gamma)\gamma_{ij}(e_ie_j^T-e_ie_i^T)\end{aligned}$$

First of all $\sum\limits_{i<j}\sum\limits_{j\in N_i}\mathbb{I} = |E|\cdot\mathbb{I}$.

Next we have that $\sum\limits_{i}\sum\limits_{j\in N_i} e_ie_i^T(\mathbb{I}-\Gamma) = \sum\limits_{i}|d_i|e_ie_i^T(\mathbb{I}-\Gamma) = D(\mathbb{I}-\Gamma)$.

Similarly, $\sum\limits_{i\in V}\sum\limits_{j\in N_i} -e_ie_i^T\Gamma\lambda_{ij}e_ie_i^T = \sum\limits_{i\in V} -e_ie_i^T\Gamma e_ie_i^T = -\Gamma$, since we know that Γ is a row-stochastic matrix.

Finally, $\sum\limits_{i\in V}\sum\limits_{j\in N_i} e_ie_i^T\Gamma\lambda_{ij}e_ie_j^T = \Gamma\Lambda$.

Therefore, we end up with

$$\mathbb{E}[A(t)] = \mathbb{I} - \frac{1}{|E|}(D(\mathbb{I}-\Gamma)+\Gamma(\mathbb{I}-\Lambda)) \tag{23}$$

If we calculate $\mathbb{E}[B(t)]$ similarly: $\mathbb{E}[B(t)] = \frac{1}{|E|}\sum\limits_{(i,j)\in E} B_{ij} = \frac{1}{|E|}\sum\limits_{i\in V}\sum\limits_{j\in N_i} e_ie_i^T(\mathbb{I}-\Gamma)$ and we obtain

$$\mathbb{E}[B(t)] = \frac{1}{|E|}D(\mathbb{I}-\Gamma) \tag{24}$$

This concludes the proof.

Finally, we can show that the above system converges to an equilibrium point. Working in the same way as in the FJ-model, we have to prove that $\mathbb{E}[A(t)] = \mathbb{I} - \frac{1}{|E|}(D(I-\Gamma)+\Gamma(\mathbb{I}-\Lambda)$ is invertible. The short proof for this can be found in [[6]]. Since this holds, then again we have that the asymptotic behavior of the dynamical system leads to:

$$\begin{aligned}x^* &= \lim_{t\to\infty}\mathbb{E}[x(t)] \\ &= (\mathbb{I}-\mathbb{E}[A(t)])^{-1}\mathbb{E}[B(t)]s \\ &= (D(\mathbb{I}-\Gamma)+\Gamma(\mathbb{I}-\Lambda))^{-1}D(\mathbb{I}-\Gamma)s\end{aligned}$$

We conclude that the expected behavior of the dynamics resembles the behavior of the FJ model.

4 Optimization and opinion formation models

In this chapter, we are going to focus mainly on techniques to solve convex optimization problems, a special class of mathematical optimization problems. Formulating a problem as a convex optimization problem provides us with great advantages: It is a field which has been extensively studied for around a century, and by now we already possess several numerical techniques, which allow us to solve such problems very efficiently. Although at first algorithms for optimizing over convex functions have been primarily studied for theoretical purposes, throughout the last decades, a rich field of practical applications has been discovered. Some scientific fields, where convex optimization has been proven very useful, are, among others, automatic control systems, data analysis, statistics and finance.

Because of this huge practical interest of optimization algorithms, there has been a vast development, considering the design of new algorithms, over the last years. The main motivation of these studies, has been the rapidly developing field of machine learning. At this point, we should mention the two most fundamental first-order methods and the most well-known second order method in order to locate a local(or global) optimum of a function: namely gradient descent, mirror descent and Newton's method. Later on, a lot of attention of the scientific community switched to the field of online convex optimization, where, as the name implies, we are trying to optimize over data that we don't know. Our goal is to design efficient algorithms and compare their behavior with their deterministic equivalent, or in other words, how well our algorithm performs compare to an optimal one, which has the advantage of knowing all the input data in advance. Finally, the massive amount of data that we have to handle nowadays when trying to find the optimal solution in complex problems, led to a revolution in the field of continuous optimization. Based on the framework of *stochastic gradient descent*, algorithms which do not examine the whole input, but rather use randomization in order to avoid the computational burden of huge datasets, were developed. Surprisingly enough, we now possess techniques which give us much more computationally efficient algorithms and convergence guarantees which are comparable to the full information algorithms.

In this chapter, we would like to bridge optimization algorithms and opinion dynamics. The iterative processes of the seminal optimization algorithms that we are going to describe, could also apply to opinion formation models and become a powerful tool in the analysis of such dynamical systems. We are going to start with gradient descent, continue with online convex optimization and finally discuss how we could borrow ideas from stochastic optimization algorithms in order to design new randomized protocols in opinion dynamics and analyze their convergence properties.

Also keep in mind that, as we mentioned before, the equilibrium point of a potential game corresponds to a local minimum of the potential function [Theorem 2.9.1]. Moreover, if there exists a potential function and is also strongly convex [Def. 2.17], the equilibrium point is the (unique) global minimum of the potential function. Since in this chapter we are discussing about unconstrained minimization, techniques presented can be also helpful in case we attempt to minimize the potential function and therefore reach an equilibrium point.

Since we are not going very deep into technical details and analysis of all the useful properties, for a reader who is interested in studying extensively the field of convex optimization, there exist excellent textbooks [[21], [24], [25]]. Moreover, there is also a nice survey on optimization methods for large-scale machine learning [[26]].

4.1 Gradient Descent

Since we are referring to descent methods, this means that the algorithms satisfy the condition $f(x^{k+1}) < f(x^k)$. Furthermore, the search direction in a descent method must satisfy $\nabla f(x^k)^T \Delta x^k < 0$. Therefore, a selection for the search direction could be the negative gradient and the resulting method is called gradient descent. The update rule is the following:

$$x_{k+1} = x_k - \eta_k \nabla f(x_k) \tag{25}$$

The term η_k is the *step size* of our algorithm. If we have a convex function f, imagine the above iterative process as follows: We start from far away going down a hill, till we reach a "valley". The step size in each round indicates how fast we are trying to get there. Throughout the whole convergence analysis, we will assume that the step size is fixed for each round and also relatively "small", so that we do not take huge steps and end up oscillating, instead of reaching the valley. Now, we are ready to analyze the well-known convergence properties of gradient descent in the case that we have a convex function.

Theorem 4.1. Suppose we have a function $f : \mathbb{R}^n \to \mathbb{R}$ which is convex, differentiable and L-smooth. Then, after k gradient descent iterations with a fixed step size $\eta = \frac{1}{L}$ $\left(\text{also holds for } \eta < \frac{1}{L}\right)$, the following inequality will be satisfied:

$$f(x_k) - f(x^*) \leq \frac{L \cdot \|x_0 - x^*\|_2^2}{2k} \tag{26}$$

Proof. Assume the function $g : \mathbb{R}^n \to \mathbb{R}$, which we define as $g(x) = \frac{L}{2}x^T x - f(x)$. Because

$f(x)$ is convex, $g(x)$ will also be convex. Due to the convexity of $g(x)$ we have

$$
\begin{aligned}
& g(y) \geq g(x) + \nabla g(x)^T (y - x) \Leftrightarrow \\
& \frac{L}{2} y^T y - f(y) \geq \frac{L}{2} x^T x - f(x) + (Lx - \nabla f(x))^T (y - x) \Leftrightarrow \\
& f(y) \leq f(x) + \nabla f(x)^T (y - x) + \frac{L}{2} \|y - x\|_2^2.
\end{aligned}
$$

If we substitute $y = x - \eta \nabla f(x)$ in the above inequality and suppose we are at a time step t we obtain:

$$
\begin{aligned}
f(x_t - \eta \nabla f(x_t)) = f(x_{t+1}) &\leq f(x_t) + \nabla f(x_t)^T (x_t - \eta \nabla f(x_t) - x_t) + \frac{L}{2} \|x_t - \eta \nabla f(x_t) - x_t\|_2^2 \\
&= f(x_t) - \eta \|\nabla f(x_t)\|_2^2 + \frac{\eta^2 L}{2} \|\nabla f(x_t)\|_2^2 \\
&\leq f(x_t) - \eta \left(1 - \frac{\eta L}{2}\right) \|\nabla f(x_t)\|_2^2
\end{aligned}
$$

Since $\|\nabla f(x_t)\|_2^2$ is positive and becomes zero only at the local minima, we notice that if we set the fixed step size to be relatively small, so that $1 - \frac{\eta L}{2}$ does not become negative, then we are guaranteed decrease at each iteration.
Continuing, we wish to insert in our inequality the distance to the equilibrium point. For this purpose we now take advantage of the convexity of $f(x)$, which gives us

$$
\begin{aligned}
& f(x^*) \geq f(x_t) - \nabla f(x_t)^T (x_t - x^*) \Leftrightarrow \\
& f(x_t) \leq f(x^*) + \nabla f(x_t)^T (x_t - x^*)
\end{aligned}
$$

Combining the two inequalities yields:

$$
f(x_{t+1}) - f(x^*) \leq \nabla f(x_t)^T (x_t - x^*) - \eta \left(1 - \frac{\eta L}{2}\right) \|\nabla f(x_t)\|_2^2
$$

For the simplicity of our analysis, we are going to fix the step size to be $\eta = \frac{1}{L}$(but the proof naturally does not hold only for this value). Therefore the last inequality will take the following form:

$$
f(x_{t+1}) - f(x^*) \leq \nabla f(x_t)^T (x_t - x^*) - \frac{1}{2L} \|\nabla f(x_t)\|_2^2
$$

Adding and subtracting $\frac{L}{2} \cdot \|x_t - x^*\|_2^2$ results in

$$f(x_{t+1}) - f(x^*) \leq \frac{L}{2}\left(\frac{2}{L}\nabla f(x_t)^T(x_t - x^*) - \frac{1}{L^2}\|\nabla f(x_t)\|_2^2 + \|x_t - x^*\|_2^2 - \|x_t - x^*\|_2^2\right)$$
$$= \frac{L}{2}\left(\|x_t - x^*\|_2^2 - \left\|x_t - \frac{1}{L}\nabla f(x_t) - x^*\right\|_2^2\right)$$
$$= \frac{L}{2}\left(\|x_t - x^*\|_2^2 - \|x_{t+1} - x^*\|_2^2\right)$$

Summing for all the k iterations, we obtain:

$$\sum_{t=0}^{k-1}(f(x_{t+1}) - f(x^*)) \leq \frac{L}{2}\sum_{t=0}^{k-1}\left(\|x_t - x^*\|_2^2 - \|x_{t+1} - x^*\|_2^2\right) \Leftrightarrow$$
$$\sum_{t=0}^{k-1}(f(x_{t+1}) - f(x^*)) \leq \frac{L}{2}\left(\|x_0 - x^*\|_2^2 - \|x_k - x^*\|_2^2\right) \leq \frac{L}{2}\|x_0 - x^*\|_2^2$$

Since we proved in the beginning that $f(x)$ is guaranteed to decrease after each time step, based on the above inequality it holds that

$$f(x_k) - f(x^*) \leq \frac{L \cdot \|x_0 - x^*\|_2^2}{2k}$$

This concludes the proof. We managed to show that gradient descent for convex, differentiable functions with L-Lipschitz gradients is guaranteed to converge and attains a convergence rate of $\mathcal{O}(1/k)$.

4.1.1 Best-response dynamics and Gradient Descent

In this subsection, we are going to show how best-response dynamics and gradient descent can coincide.

Imagine that we have a cost minimization game. This means that each agent has her individual cost function and at each iteration she plays the best move in order to minimize her cost function immediately as much as possible. We suppose also, that all agents play simultaneously. Intuitively, since all players are searching a finite set of moves and try to reach the minima of their cost functions as fast as possible, they might also contribute to a "common good". In other words, the agents play their best response and also help the system as a whole to move towards a local minimum or Nash equilibrium. Therefore, it seems that the above procedure resembles gradient descent. Of course, the overall progress is still not guaranteed. In some cases, we can imagine that the selfish choice of move from an agent might increase the cost function of another agent. In any case, the fact that the agents try each one independently to minimize their cost functions does not imply an overall progress.

Let's examine a process inspired by the FJ model. Suppose we are given an undirected weighted graph $G(V, E)$, each agent at each time step expresses an opinion $x_i \in [0, 1]$ and also holds an intrinsic opinion $s_i \in [0, 1]$. In this model, each agents has a cost function of the form $C_i(x) = \sum\limits_{j \neq i} w_{ij}(x_i - x_j)^2 + w_{ii}(x_i - s_i)^2$ and the agents are trying in parallel to minimize their cost functions. In this case, we can show that the best-response dynamics is equivalent to performing gradient descent on the model's potential function.

Then the model admits an exact potential function

$$\Phi(x) = \sum_{i<j} w_{ij}(x_i - x_j)^2 + \sum_{i=1}^{n} w_{ii}(x_i - s_i)^2 \tag{27}$$

If each agent plays her best response, then for agent i it holds that

$$\frac{\partial C_i(x)}{\partial x_i} = 0 \Leftrightarrow 2\sum_{j \neq i} w_{ij}(x_i - x_j) + 2w_{ii}(x_i - s_i) = 0$$

$$\Leftrightarrow x_i = \frac{\sum\limits_{j \neq i} w_{ij}x_j + w_{ii}s_i}{\sum\limits_{j \neq i} w_{ij} + w_{ii}}$$

Thus, in our opinion formation model the update rule could take the form

$x_i^{t+1} = \frac{\sum\limits_{j \neq i} w_{ij}}{\sum\limits_{j \neq i} w_{ij} + w_{ii}} \cdot x_j^t + \frac{w_{ii}}{\sum\limits_{j \neq i} w_{ij} + w_{ii}} \cdot s_i$, $\forall t > 0$. Also note that for the potential function and for agent i it holds that:

$$\frac{\partial \Phi(x^t)}{\partial x_i} = 2 \left(\sum_{j \neq i} w_{ij} + w_{ii} \right) x_i^t - 2 \sum_{j \neq i} w_{ij} x_j^t - 2 w_{ii} s_i \Rightarrow$$
$$\frac{\partial^2 \Phi(x^t)}{(\partial x_i)^2} = 2 \left(\sum_{j \neq i} w_{ij} + w_{ii} \right)$$

Proceeding further, we can rewrite the update rule as follows:

$$x_i^{t+1} = x_i^t + \frac{2 \left(\sum\limits_{j \neq i} w_{ij} x_j^t + w_{ii} s_i - \sum\limits_{j \neq i} w_{ij} x_i^t - w_{ii} x_i \right)}{2 \left(\sum\limits_{j \neq i} w_{ij} + w_{ii} \right)}$$

Notice now that this is equivalent to writing $x_i^{t+1} = x_i^t - \frac{\partial \Phi(x^t)}{\partial x_i} \cdot \left(\frac{\partial^2 \Phi(x^t)}{(\partial x_i)^2} \right)^{-1}$.

Therefore, we can write the update rule for the vector to show that it follows a gradient descent procedure:

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \Gamma \cdot \nabla \Phi(\mathbf{x}^t) \tag{28}$$

where Γ is a diagonal matrix with $\Gamma_{ii} = \left(\frac{\partial^2 \Phi(x^t)}{(\partial x_i)^2} \right)^{-1}$ and 0 otherwise.

The benefit from showing this equivalence is that in order to prove that the game the agents play converges to an equilibrium point boils down to proving that performing gradient descent on the *strongly convex* potential function of our model, converges to the global minimum of the function.

4.1.2 The FJ-model and Gradient Descent

We return to the undirected FJ-model(the weights are symmetric). We are going to show that we can also write the update rule of FJ as a gradient descent iterative process. For this purpose, we define the function $f(x) = \|x - A \cdot x - B \cdot s\|_2^2$. We make two remarks about this function:

- It is easy to see that $f(x)$ is a strongly convex function.
- If x^* is the equilibrium point of the FJ model, then $f(x^*) = \|x^* - Ax^* - Bs\|_2^2 = 0$. Since $f(x)$ cannot become negative, we conclude that in any other case we have $f(x) \geq f(x^*)$ and $x^* = \arg\min_x f(x)$ is the unique minimizer of f.

Therefore, since $f(x) = \|(I-A)x - Bs\|_2^2$ we obtain:

$$\nabla f(x^t) = 2((I-A)x - Bs)(I-A) \Rightarrow$$
$$\nabla^2 f(x^t) = 2(I-A)$$

If we try to write the update rule of the undirected FJ model similar to performing gradient descent on $f(x)$, it will look again like: $x^t = x^{t-1} - \Gamma \cdot \nabla f(x^t)$. Setting $\Gamma = \left(\nabla^2 f(x^t)\right)^{-1} = \frac{1}{2}(I-A)^{-1}$:

$$\begin{aligned} x^t &= x^{t-1} - \frac{(I-A)^{-1}}{2} \cdot 2\left((I-A)x^{t-1} - Bs\right)(I-A) \\ &= x^{t-1} - (I-A)x^{t-1} + Bs \\ &= Ax^{t-1} + Bs \end{aligned}$$

This concludes the proof, as we managed to obtain the update rule of the FJ model in matrix form. Note that Γ is a matrix which does not change over time(since we assume that the structure of the underlying graph and the weights on the edges do not change over time), and the elements in each row result from the inverse of the Hessian Matrix of $f(x)$.

4.2 Online convex optimization

Definition 4.1. An *online convex programming problem* consists of a feasible set $F \subseteq \mathbb{R}^n$ and an infinite sequence $c^1, c^2, ...$, where each $c^t : F \rightarrow \mathbb{R}$ is a convex function.

At each time step t, an online convex programming algorithm selects a vector $x^t \in F$. After the vector is selected, it receives the cost function c^t.

Basically in the online case, all the information is **not** available before decisions are made. In other words, here we are not aware of the actual value of the convex cost function.

In order to analyze the performance of such algorithms, it makes sense to compare its performance with the best algorithm in hindsight that knows all of the cost functions and selects one fixed vector. For this purpose we introduce an important notion, called *regret*.

Definition 4.2. Given an algorithm A, and a convex programming problem $(F, \{f^1, f^2, ...\})$, if $\{x^1, x^2, ...\}$ are the vectors selected by A, then the cost of A until time T is $\sum_{t=1}^{T} f^t(x^t)$.

The cost of a static feasible solution $x \in F$ until time T is $\sum_{t=1}^{T} f^t(x)$.

The *regret* of algorithm A until time T is

$$R_A(T) = \sum_{t=1}^{T} f^t(x^t) - \min_{x \in F} \sum_{t=1}^{T} f^t(x)$$

In this section, we are going to link the online convex optimization frameworks with our opinion formation models. The results that we are going to present have been widely applied, when studying online decision-making games. From now on, we are going to refer to *repeated games.* These work in the following way:

Definition 4.3 (Repeated Games)**.** Suppose that we play repeatedly a game among n agents for T rounds.At each round t the game proceeds as follows:

- Each agent i picks a probability distribution p_i^t over her set of actions S
- An adversary picks a cost vector $f^t : S \to [0, 1]$ for agent i. It is important that the cost vectors are bounded by a value.
- According to the probability distribution chosen, the agent picks a strategy x^t. The agent suffers a "disagreement cost" equal to $f^t(x^t)$(with slight abuse of notation we dropped the indexes i).
- Afterwards, the agent usually receives some kind of feedback and re-adapts her strategy for the next round. We can divide the kind of feedback into two categories that we are going to examine:
  1. The "full information" feedback: Given the actions of the other agents, the whole cost vector is revealed to i after she suffers the loss. This means that she learns the values that she would have obtained, had she chosen a different strategy.
  2. The "bandit" feedback: Instead of being able to observe all about f^t, the player learns only the value of $f^t(x^t)$.

Now we have a rough description of the type of the problems we are studying. Again we have a cost-minimization game, where an adversary picks the cost vectors for the agents in each round. Therefore, the connection to opinion dynamics becomes straightforward. We have a model which works according to a fixed update rule. If we can define a convex cost function for each agent so that the cost-minimization game captures the behavior of the opinion formation model(we saw such examples in the previous chapters) then with the frameworks that we are going to describe, we can play a *regret minimization* game and analyze that way also the online version of the model.

Regret is used as a measure of performance and the common goal is to prove that the average regret of an algorithm approaches zero. As we said, we compare the cumulative cost we suffer after T iterations with the best fixed action in hindsight. If the average regret vanishes as the number of iterations grows, it means that the average cost per iteration of our randomized algorithm approaches the performance of the best fixed action. Thus, it is a desirable property to obtain. More formally:

Definition 4.4 (No-regret algorithm)**.** An algorithm is called no regret if for any adversary and number of rounds T it holds that $R_A(T) = o(T)$.

Now we are ready to proceed to the first general framework.

4.2.1 Online Gradient Descent

Let's examine what we can do in the first case, where we receive the "full information" feedback. We are going to imagine the algorithm as a repeated game between an agent and an adversary. In the multiagent setting, the agents who try individually and simultaneously to minimize their cost functions, should simply choose a strategy and use a no-regret algorithm. The "loss" each one will eventually suffer is an expected cost , dependent on the actions of the other players.

We are going to focus on *online gradient descent*, which was introduced by Zinkevich [[27]]. First, we are going to make some assumptions and then proceed to the algorithm and its convergence rate: Assume $S \subset \mathbb{R}^n$ is a closed convex set of diameter at most D. This simply means that for every $x, x' \in S, \|x - x'\|_2 \leq D$.

Algorithm 1: Online Gradient Descent Algorithm

given a convex set S, T
for $t = 1$ to T **do**
$y^{t+1} = x^t - \eta^t \nabla f^t(x^t)$
$x^{t+1} = \Pi_S(y^{t+1})$

Theorem 4.2 (Zinkevich '03)**.** Consider any sequence of differentiable functions chosen by the adversary $f^1, f^2, ..., f^T : S \to \mathbb{R}$ such that $\left\|\nabla f^t(x)\right\| \leq G$ for any $1 \leq t \leq T, x \in S$, i.e. G is an upper bound on the gradient magnitudes. Then for the sequence $x^1, x^2, ..., x^T$ where each x^t is the strategy chosen by the decision-maker, and by setting $\eta = \frac{D}{G\sqrt{T}}$, the regret is bounded as follows:

$$\sum_{t=1}^{T} f^t(x^t) - \min_{x \in S} \left(\sum_{t=1}^{T} f^t(x) \right) \leq DG\sqrt{T} \tag{29}$$

Proof. In order not to get into many technical details, we are going to prove the theorem for the case of linear functions, but essentially the same proof is extended to all convex functions. Therefore, for the proof consider that if the player chooses a strategy among d available strategies according to a probability distribution x and she receives a cost vector $w \in \mathbb{R}^d$, then the cost function becomes $w^T x$ and is linear and the update rule could be simply written as $x^{t+1} = x^t - \eta^t w^t$ and then project onto the convex set. Then:

$$\begin{aligned} \left\|x^{t+1} - x^*\right\|^2 &\leq \left\|x^t - \eta_t w^t - x^*\right\| \\ &= \left\|x^t - x^*\right\|^2 - 2\eta_t w^t (x^t - x^*)^2 + \eta_t^2 \left\|w^t\right\|^2 \end{aligned}$$

In our case, the best fixed action in hindsight would be $\min_{x^* \in S} \sum_{t=1}^{T} w^t x^*$. Therefore, rearranging the above inequality and summing over the iterations:

$$\begin{aligned} &2w^t \eta_t (x^t - x^*) \leq \left\|x^t - x^*\right\|^2 + \eta_t^2 \left\|w^t\right\|^2 - \left\|x^{t+1} - x^*\right\|^2 \Leftrightarrow \\ &\sum_{t=1}^{T} w^t x^t - \sum_{t=1}^{T} w^t x^* \leq \frac{1}{2\eta^t} \left\|x^1 - x^*\right\|^2 + \frac{\eta^t}{2} \sum_{t=1}^{T} \left\|w^t\right\|^2 \end{aligned}$$

If we substitute with a fixed step size $\eta = \dfrac{D}{G\sqrt{T}}$ then we get the regret bound.

Finally, we will explain which are the regret bounds that we get in case the cost functions of the agents are linear and in case they are strongly convex. We refer to all the cases, because each model needs an appropriate cost function for the agents, in case it can be described as a repeated game.

- *The linear cost function:* As we proved, we have a no regret algorithm, though we can do even better than online gradient descent. Consider the N different strategy choices, as "experts" that tell us their opinion about an issue. We consider the following process, which describes a pretty natural behavior. First of all, when choosing an expert, we should have a criterion to decide how good they are, otherwise the problem degrades to choosing one uniformly at random. We should therefore take into account the past actions of the experts, trust some more than others and make it more likely to choose the "good" ones. This is what the *multiplicative weights algorithm* does [[28], [29]]. We describe the process: Each expert has a weight factor w_i^t, which represents the "level of trust or preference" that we have for her. In the beginning, all experts have the same weight. As we proceed, we want to promote the reliable ones, and punish the bad ones. We normalize the sum of the weights to be 1 and at each of the T iterations the following happens:

1. Choose an expert according to the distribution $x^t = w^t / \sum_{i=1}^{N} w_i^t$.
2. Receive the cost vector $c^t \in [0,1]^N$ and suffer the total cost of $(c^t)^T x$.
3. Update the weights of the experts according to the following rule: $w_i^{t+1} = w_i^t (1-\epsilon)^{c_i^t}$.

Note that each expert is chosen with probability proportional to its current weight, that the probabilities can only decrease and that they decrease at an exponential rate. Note also that the multiplicative weights algorithm can be viewed as a special case of online gradient descent, or more precisely, as a mirror descent method [[30]]. Finally, it is shown that although the regret of this algorithm has the same asymptotic behavior with OGD in the number of rounds T we play, it behaves much better as the number of experts grows rapidly [[31]].

- *The convex cost function*: In any case, we can simply perform online gradient descent in the broader class of convex cost functions and obtain an average regret with a rate of convergence of $\mathcal{O}(1/\sqrt{T})$ and thus we are guaranteed a no-regret algorithm.
- *The strongly convex cost function*: In this special case, following a similar analysis, OGD gives us a logarithmic bound on the regret [[31]]. Thus the average regret has a rate of convergence of $\mathcal{O}(\log(T)/T)$

4.2.2 The bandit version

We can extend the idea of online gradient descent to the "bandit" setting: As we already mentioned, in this case, in each period only the cost $f^t(x_t)$ is revealed. In the previous setting, online gradient descent achieved regret bounds of $\mathcal{O}(\sqrt{T})$. This type of feedback has also been extensively studied. in different types of games and for different types of adversaries. We are not going to expand further in this thesis, though we are just going to describe a seminal algorithm which has served as a framework: The assumptions are that the cost functions are convex, bounded and L-Lipschitz, and that the adversary is oblivious: She chooses the functions *without* knowing any of the player's moves. In the bandit setting, for the algorithm of Flaxman et.al [[32]] the expected regret is bounded from $\mathcal{O}(T^{3/4})$. Notice that we are still able to achieve a regret of $o(T)$.

The core of this approach is that we can find a simple approximation of the gradient, just by evaluating our cost function $f^t(x^t)$ at a single random point. So, surprisingly, it is possible to use gradient descent without seeing anything more than the value of the function at a single point. For the analysis of the algorithm and more details, we refer the reader to [[32]].

4.3 Stochastic optimization

4.3.1 Motivation and history

Although gradient descent enjoys some nice properties, is very simple and admits good convergence rates it can become impractical in today's complex problems. The rapid development of machine learning together with the increasing amount of data created new needs: In large-scale machine learning we have to deal with huge datasets and using *full-gradient methods* in each iteration can be computationally very costly. A natural idea instead, is to try to reach the optimal solution by taking only one, or a small batch, of samples in each iteration. We want these random samples to be, of course, good estimators of the gradient. Therefore, new methods, called *stochastic gradient methods* have been developed and studied extensively over the last years. In this section, we are going to talk a bit about the stochastic optimization framework and try to link it again with opinion dynamics. In other words, discuss how we can borrow ideas from the stochastic optimization literature to incorporate in the design of our models, and how we can analyze the convergence properties is we assume stochastic variants of models that we have already studied.

Driven from the need to find better algorithms to solve optimization problems where the information is abundant, Bottou [[33]] and Zhang [[34]] published seminal works where they examine the performance of SGD methods in different settings.

In general, SGD methods use the following update rule

$$x_{k+1} \leftarrow \underset{y \in \mathbb{R}^d}{\operatorname{argmin}} \left\{ \tfrac{1}{2\eta} \|y - x_k\|_2^2 + \langle \widetilde{\nabla}_k, y \rangle + \psi(y) \right\}$$

where:

- η is the step length as before
- $\widetilde{\nabla}_k$ is a random vector satisfying $\mathbf{E}\left[\widetilde{\nabla}_k\right] = \nabla f(x_k)$ and is the term that me defined as *gradient estimator*.

In most cases the *proximal function* $\psi(y)$ equals zero and the update rule takes the more recognizable form $x_{k+1} \leftarrow x_k - \eta \widetilde{\nabla}_k$. Most of the algorithms in the literature choose one sample uniformly at random as the unbiased gradient estimator. More formally, we can say that we are given access to a noisy gradient oracle which we define as

$$\mathcal{O}(x_k) = \widetilde{\nabla}_k \text{ s.t. } \mathbf{E}\left[\widetilde{\nabla}_k\right] = \nabla f(x_k),\ \mathbf{E}\left[\left\|\widetilde{\nabla}_k\right\|^2\right] \leq G^2$$

The huge advantage of the above idea is that computing the gradient in each iteration is n times faster compared to the full gradient method, since it is only required to take one sample in each iteration. Although it is very efficient computational method, the estimator might be far away from the expected value in some iterations. Thus, the variance hurts the convergence rate and it is in fact proven that SGD cannot converge at a rate faster than $1/\varepsilon$ if we are trying to get ϵ-close to the minimum, even if the function we are trying to minimize it strongly convex and smooth.

Since the main problem in designing fast stochastic gradient descent algorithms is exactly the variance of the stochastic process, the scientific community focused in reducing the variance: If k is the number of iterations, it is desirable to choose the gradient estimator in such a way, that the variance approaches to zero as k grows. The well-known variance reduction technique was introduced by Schmidt et. al [[35]] and was incorporated in their well-known stochastic average gradient method(SAG). In their seminal work, by using also previous gradient values stored in a memory, they managed to improve the convergence rates of the stochastic gradient methods. We usually try to minimize convex minimization problems in the form of the following, which captures many practical cases in the field of machine learning:

$$\min_{x\in\mathbb{R}^d}\left\{F(x) = f(x) + \psi(y) = \frac{1}{n}\sum_{i=1}^{n} f_i(x) + \psi(x)\right\} \tag{30}$$

Usually, the assumptions are that $f(x)$ is a convex function and a finite average of n convex, smooth functions and $\psi(x)$ is convex as well. What we seek to find is an approximate minimizer satisfying $F(x) \leq F(x^*) + \varepsilon$.

In this setting(excluding possibly the term $\psi(x)$), Schimdt et. al showed in the general case that the convergence rate can be improved from $\mathcal{O}(1/\sqrt{k})$ to $\mathcal{O}(1/k)$ (where k is the number of iterations) and in the interesting case where the finite sum is strongly convex, they improved the convergence rate of $\mathcal{O}(1/k)$ to a linear convergence rate of the form $\mathcal{O}(\rho^k)$.

Around the field of variance-reduction techniques one can find excellent papers, such as [[36], [37], [38]]. Applying such techniques as in [[39]] we can reach an ε - minimizer for 30 by computing $\mathcal{O}\left((n+\kappa)\log\frac{1}{\varepsilon}\right)$ stochastic gradients, where with κ we denote the condition number of the problem.

The fact that the best iteration complexity known for variance reduction methods has a linear dependence on the term κ is a problem, since as we mentioned also in gradient descent, this exact number can lead our search direction and speed to go even terribly wrong. Catalyst [[40]] was a significant advance but it was not until recently that Allen-Zhu [[41]] presented an algorithm which enjoys an $\sqrt{\kappa}$ dependency. This factor is also optimal according to [[41]].

4.3.2 A limited information protocol for the FJ model

In this section we are going to view the Friedkin-Johnsen model as a stochastic optimization problem. In the next chapter, we are going to provide the general framework with more details. The motivation of the setting we are going to present, is that in social networks it is very rare for an agent to ask all of her neighbors in order to update her opinion. The FJ model, in order to establish its convergence properties, assumes that the agents ask all of their neighbors at each time step and take into account a weighted average of these opinions. In reality, though, it makes more sense to assume that the agents in each round ask a small subset of their neighbors about a certain topic and update their opinions. For this purpose, following the work of Fotakis et. al [[42]], we present a randomized protocol where each agent picks one of her neighbors and updates her opinion according to a convex combination of her current one and the opinion of the neighbor she sampled. The main result is that this simple protocol *converges in expectation* to the stable state of the FJ model.

First, we will describe the limited information protocol we are going to follow. Given an undirected graph $G(V,E)$, an initial vector of opinions $\vec{s}$ and a term $\gamma(t) \in (0,1)$, the update rule takes the following form, which states that the opinion of an agent at the next time step is a convex combination of her previous opinion and the opinion she sampled during this round:

$$x_i(t) = \gamma(t)x_i(t-1) + (1-\gamma(t))r_i(t) \tag{31}$$

Now, $r_i(t)$ is a discrete random variable which works as follows: Each agent samples simultaneously and randomly the opinion of one neighbor with probability equal to the weight of the respective edge w_{ij}(since for each agent we have normalized the weights to be $\sum_{j=1}^{n} = 1$) and sets $r_i(t)$ to be:

$$r_i(t) = \begin{cases} x_j(t-1), & \text{if } j \neq i \\ s_i, & \text{if } j = i \end{cases}$$

The expected value of $r_i(t)$ is the weighted sum, where each weight represents the possibility that a certain opinion will be selected:

$$\mathbb{E}[r_i(t)] = \sum_{j \neq i} w_{ij} x_j(t-1) + w_{ii} s_i$$

Therefore, taking expectation on both sides of the update rule, this implies for the vector of opinions that:

$$\mathbb{E}[\vec{x}(t)] = \gamma(t)\mathbb{E}[\vec{x}(t-1)] + (1-\gamma(t))(A\mathbb{E}[\vec{x}(t-1)] + B\vec{s}) \tag{32}$$

From now on, we are going to choose the value $\gamma(t) = \frac{t-1}{t}$, which yields an adaptive step size for the iterative process. For simplicity, we are also going to denote $y(t) = \mathbb{E}[\vec{x}(t)]$. Therefore, at round t+1 we have $y(t+1) = \frac{t-1}{t} \cdot y(t) + \frac{1}{t} \cdot (Ay(t) + Bs)$. Furthermore, for the equilibrium point it also holds that $x^* = \frac{t-1}{t} \cdot x^* + \frac{1}{t} \cdot (Ax^* + Bs)$. We notice about this particular value for $\gamma(t)$ the following:

- As t becomes larger, we put more weight in the previous opinion of the agent and less weight on the random sample we draw.
- Since, according to the update rule, the variance of the stochastic process does not decrease as the number of iterations grows, we are obligated to decay the step size as we get closer to the equilibrium point. This makes the convergence rate slower, as we take smaller steps as we proceed, but lowers the variance. Our goal is that the system will not oscillate close to the optimum.

Now, we can state the main theorem (similar as in [[42]]), which shows that in expectation the randomized protocol, as stated above, converges to the equilibrium point.

Theorem 4.3. Given an undirected graph $G(V, E)$, choose $\gamma(t) = \frac{t-1}{t}$ and denote $e(t) = y(t) - x^*$ as the error vector. Then it holds for the limited information FJ model that:

$$\lim_{t \to \infty} \|\mathbb{E}[\vec{x}(t)] - x^*\|_\infty = 0.$$

Proof.

$$\begin{aligned}
\|y(t+1)-x^*\|_\infty &= \left\|\frac{t-1}{t}\cdot y(t)+\frac{1}{t}A\cdot y(t)+\frac{1}{t}B\cdot s-\frac{t-1}{t}\cdot x^*-\frac{1}{t}A\cdot x^*-\frac{1}{t}B\cdot s\right\|_\infty \\
&= \left\|\frac{t-1}{t}\cdot y(t)+\frac{1}{t}A\cdot y(t)-\frac{t-1}{t}\cdot x^*-\frac{1}{t}A\cdot x^*\right\|_\infty \\
&= \left\|\left(\frac{t-1}{t}\cdot\mathbb{I}+\frac{1}{t}A\right)(y(t)-x^*)\right\|_\infty \\
&\leq \left\|\frac{t-1}{t}\cdot\mathbb{I}+\frac{1}{t}A\right\|_\infty\cdot\|y(t)-x^*\|_\infty \\
&\leq \|\epsilon(t)\|\cdot\left(\frac{t-1}{t}+\frac{\|A\|_\infty}{t}\right) \\
&\leq \epsilon(t)\left(1-\frac{1-\|A\|_\infty}{t}\right)\leq \epsilon(t)\cdot e^{-\frac{1-\|A\|_\infty}{t}}
\end{aligned}$$

We conclude that as we proceed in rounds it holds that:

$$\|y(t+1)-x^*\|_\infty \leq \|y(t)-x^*\|_\infty\cdot e^{-\frac{1-\|A\|_\infty}{t}}$$

The first inequality follows from the submultiplicative property, the second from the fact that $\|A+B\|\leq\|A\|+\|B\|$ and the fact that in this case the norm of the identity equals 1, and the third follows from the fact that $1-x<e^{-x}$, when $x<1$.

If we apply iteratively the inequality we just derived, we get that:

$$\begin{aligned}
\|y(t+1)-x^*\|_\infty &\leq e^{-\frac{1-\|A\|_\infty}{t}}\cdot e^{-\frac{1-\|A\|_\infty}{t-1}}\cdot e^{-\frac{1-\|A\|_\infty}{t-2}}\dots e^{-\frac{1-\|A\|_\infty}{1}}\cdot\epsilon(0) \\
&= e^{-(1-\|A\|_\infty)\sum_{i=1}^{t}\frac{1}{i}}\cdot\epsilon(0) \\
&\leq e^{-(1-\|A\|_\infty)\cdot logt}\cdot\epsilon(0) \\
&= \frac{1}{t^{1-\|A\|_\infty}}\cdot\epsilon(0)
\end{aligned}$$

Note that for the last inequality above we used a fact for the harmonic series, which states that $\log(N+1)<\sum_{n=1}^{N}\frac{1}{n}$. If we take the limit as $t\to\infty$, the proof is complete.

4.3.3 Variance reduction and an analysis in epochs

In this section, we are going to present two ideas which appear often in the recent stochastic optimization literature. These methods are mainly designed for machine learning problems. Although we cannot directly apply them in our settings, we will briefly present them for the sake of completeness and hope that their intuition and their high-level description could be potentially applied also in the analysis of the dynamical systems we study. Motivating these ideas a bit, we mentioned that we are forced to pick a step size in the update rule which decays over time, since as we get close to the optimal solution the variance of the random sampling, which remains the same over time, could prevent the iterative process from converging close to the point we are seeking. Following the work of Johnson and Zhang [[39]], we are going to present how we can potentially reduce the variance as the number of iterations grows and how analyzing the convergence of the proposed algorithm in epochs can be helpful. Here, with the term "epochs", we mean that we consider stages of fixed length. We let the algorithm run for a number of iterations and then we examine how the objective function we strive to minimize, is guaranteed to decrease from stage to stage. An analysis in epochs was included also in the opinion formation model we studied before [3.3.1].

Consider again the problem of minimizing an objective function as stated in 30. In this case, we have a separable function, which means that if we sample a random index, which corresponds to one of its components, then on expectation we have the gradient descent update rule. Consider now the following update rule:

$$x(t) = x(t-1) - \eta_t \left(\nabla f_i(x(t-1)) - \nabla f_i(x_s) + \nabla F(x_s)\right)$$

where:

- $i \in \{1, 2, \ldots, n\}$ is the random index that we draw
- x_s is a vector that we update every time after m iterations of the update rule. A choice for that could be the vector we compute in each m-th iteration.
- $F(x_s)$ is the full gradient that we compute only once every m iterations. Therefore, before we perform again the next m updates, we compute first $\nabla F(x_s) = \frac{1}{n}\sum_{i=1}^{n} \nabla f_i(x_s)$.

First we check that this is equivalent to performing stochastic gradient descent. Note that $\mathbb{E}\left[\nabla f_i(x_s)\right] = \nabla F(x_s)$ and then:

$$\mathbb{E}[x(t)|x(t-1)] = x(t-1) - \eta_t \mathbb{E}\left[\nabla f_i(x(t-1))\right] = x(t-1) - \eta_t F(x(t-1))$$

Now, we proceed to show the variance goes to zero as the number of iterations grows. As x_s and $x(t)$ get close to the equilibrium point, then $\nabla F(x_s)$ goes to zero. Furthermore, both the values of $\nabla f_i(x(t-1))$ and $\nabla f_i(x_s)$ are very close to $\nabla f_i(x^*)$ and thus under these conditions we obtain:

$$\nabla f_i(x(t-1)) - \nabla f_i(x_s) + \nabla F(x_s) \to \nabla f_i(x(t-1)) - \nabla f_i(x^*) \to 0$$

Although, as we mentioned, we cannot now apply this technique in our setting, this technique enables us to define a larger constant step size and obtain better convergence rates. Since we keep a vector x_s after every m iterations it is convenient for the analysis to bound the progress of the objective function towards the optimal solution after each epoch of fixed length m.

5 Stochastic Gradient Descent and the FJ model

In the previous chapter, we explained how algorithmic frameworks of convex optimization can relate to the opinion formation models and how the former can be applied to the latter in order to help us analyze how the iterative process evolves over time and state the convergence properties of these protocols. In this chapter, we continue with the FJ model and the stochastic optimization framework. First, we are going to describe how we can form as a stochastic optimization problem an opinion formation model with limited information(a simplified setting of the FJ model). Next, we are going to study a randomized variant of the FJ model with increased sampling and examine what guarantees we can provide.

5.1 An opinion formation model with limited information

For simplicity of our analysis, we are going to restrict the FJ-model to an *unweighted d-regular graph*. In the model presented, we assume that we are given a social graph and as the agents play a game in rounds, each agent communicates with her neighborhood. We denote N_i the set of agents who are neighbors of an agent i (i.e. $(i,j) \in E$). The opinions of the agents lie again in the interval $[0,1]$.

When the model is deterministic, which means that we know that in each round the agents ask all of their neighbors about their opinion, the update rule for round t is simply the average of the opinions in the neighborhood:

$$\forall x_i \in V : x_i^t = \frac{\sum\limits_{j \in N_i} x_j^{t-1}}{d} \tag{33}$$

Remark 3. We mentioned before that in the FJ-model there is also an *intrinsic opinion* s_i involved, which never changes. in the above rule, we did not take this intrinsic opinion into account, as it does not change the analysis that follows if included. One can imagine that there is a self loop for each agent which points to the intrinsic opinion or that there is a "dummy" vertex which represents the innate opinion, does not participate in the update rule and is linked to the corresponding agent. Then, the intrinsic opinion counts as much as each of the opinions of the neighbors in the update rule, since we are restricted to a d-regular unweighted graph. In this case, the update rule would take the form $x_i^t = \frac{\sum\limits_{j \in N_i} x_j^{t-1}}{d+1}$. So, for the rest of the analysis, we would just have to substitute the term d with $d+1$.

Now, let's model the randomized protocol with limited information. In each round, all the agents play simultaneously and each one chooses at round t uniformly at random one neighbor

$j \in N_i$ or a subset of the neighbors and learns her/their opinion(s) $x_j(t-1)$. Now the update rule becomes a *convex combination* of her previous opinion and the opinion(s) she sampled at round t:

$$x_i^t = \frac{t-1}{t} x_i^{t-1} + \frac{R_i^t}{t} \tag{34}$$

where with R_i^t we denote the opinion that the agent just learned.

For this game, we also introduce a **strongly convex** function:

$$\Phi(x^t) = \sum_{i<j} \left(x_i^t - x_j^t \right)^2$$

Remark 4. $\Phi(x)$ is a potential function for our game.

Proof. According to the definition of the potential function:

$$\begin{aligned}
&\Phi(x_i, x_{-i}) - \Phi(x_i', x_{-i}) \\
&= \sum_{i<j} (x_i - x_j)^2 - \sum_{i<j} (x_i' - x_j)^2 \\
&= \sum_{j \in N_i} (x_i - x_j)^2 - \sum_{j \in N_i} (x_i' - x_j)^2 \\
&= C_i(x_i, x_{-i}) - C_i(x_i', x_{-i})
\end{aligned}$$

where with $C_i(x)$ we denote the cost function of agent i. Therefore, $\Phi(x)$ captures the behavior of our game.

If we take the derivative at round t-1 w.r.t. agent i we obtain: $\dfrac{\partial \Phi(x^{t-1})}{\partial x_i} = 2dx_i^{t-1} - 2 \sum_{j \in N_i} x_j^{t-1}$.

Also, because of the random sampling, the expected value of the sample equals the average of the opinions of the neighborhood. By using also the above equation we notice that:

$$\begin{aligned}
\mathbb{E}\left[R_i^t\right] &= \frac{\sum_{j \in N_i} x_j^{t-1}}{d} \Rightarrow \\
\mathbb{E}\left[R_i^t\right] &= x_i^{t-1} - \frac{1}{2d} \cdot \frac{\partial \Phi(x^{t-1})}{\partial x_i}
\end{aligned}$$

We can also re-write the update rule as follows:

$$\begin{aligned}
x_i^t &= \frac{t-1}{t} x_i^{t-1} + \frac{R_i^t}{t} \\
&= x_i^{t-1} - \frac{1}{t} \left(x_i^{t-1} - R_i^t \right) \\
&= x_i^{t-1} - \frac{1}{t} S^t
\end{aligned}$$

where S^t is a random vector: $\forall i \in V$: $S_i^t = x_i^{t-1} - R_i^t$. Taking the expected value of S^t we obtain:

$$\begin{aligned}\mathbb{E}\left[S^t\right] &= x^{t-1} - \mathbb{E}\left[R^t\right] \\ &= x^{t-1} - x^{t-1} + \frac{1}{2d} \cdot \nabla\Phi\left(x^{t-1}\right) \\ &= \frac{1}{2d} \cdot \nabla\Phi\left(x^{t-1}\right)\end{aligned}$$

Therefore, we can finally write the update rule as follows:

$$\boxed{x^t = x_i^{t-1} - \frac{1}{t}S^t \ , \ \mathbb{E}\left[S^t\right] = \frac{1}{2d} \cdot \nabla\Phi\left(x^{t-1}\right)}$$

Therefore, the update rule now is equivalent to performing the stochastic gradient descent framework on the potential function of our model, since on expectation we obtain gradient descent. The step size is $\eta^t = \dfrac{1}{2dt}$ and as we notice, it has again to decay as the number of iterations grows and we get closer to the equilibrium point.

Before proceeding further we would like to show that while in the full-information FJ model we know its convergence properties and we know that in each round we get closer to the equilibrium, here the convergence depends on the variance of the stochastic process.

Let's bound the expected distance to the equilibrium point. Assume that we are at time step t and we want to relate the expected distance with the previous time step.

By applying the update rule and adding and subtracting the term $\dfrac{1}{2td}\nabla\Phi(x^{t-1})$ we get:

$$\begin{aligned}\mathbb{E}\left[\left\|x^t - x^*\right\|_2^2\right] &= \mathbb{E}\left[\left\|x^{t-1} - \frac{1}{t}S^t - x^*\right\|_2^2\right] \\ &= \mathbb{E}\left[\left\|x^{t-1} - \frac{1}{2td}\nabla\Phi(x^{t-1}) + \frac{1}{2td}\nabla\Phi(x^{t-1}) - \frac{1}{t}S^t - x^*\right\|_2^2\right] \\ &= \mathbb{E}\left[\left\|\left(x^{t-1} - \frac{1}{2td}\nabla\Phi(x^{t-1}) - x^*\right) - \frac{1}{2dt}\left(2d \cdot S^t - \nabla\Phi(x^{t-1})\right)\right\|_2^2\right]\end{aligned}$$

By expanding on the identity and the linearity of expectation:

$$\mathbb{E}\left[\left\|x^t - x^*\right\|_2^2\right] = \left\|x^{t-1} - \frac{1}{2td}\nabla\Phi(x^{t-1}) - x^*\right\|_2^2$$

$$- 2\cdot\frac{1}{2dt}\left(x^{t-1} - \frac{1}{2td}\nabla\Phi(x^{t-1}) - x^*\right)^T \left(2d\cdot\mathbb{E}\left[S^t\right] - \nabla\Phi(x^{t-1})\right) + \mathbb{E}\left[\left\|\frac{1}{t}\left(S^t - \frac{\nabla\Phi(x^{t-1})}{2d}\right)\right\|_2^2\right]$$

Notice now that $2d\cdot\mathbb{E}\left[S^t\right] - \nabla\Phi(x^{t-1}) = 0$ and we can also a constant $\alpha < 1$ to derive the inequality

$$\left\|x^{t-1} - \frac{1}{2td}\nabla\Phi(x^{t-1}) - x^*\right\|_2^2 \leq \alpha\cdot\left\|x^{t-1} - x^*\right\|_2^2$$

This leads to the following upper bound on the progress we make towards the equilibrium in each round:

$$\boxed{\mathbb{E}\left[\left\|x^t - x^*\right\|_2^2\right] \leq \alpha\cdot\left\|x^{t-1} - x^*\right\|_2^2 + \frac{1}{t^2}\cdot\mathbb{E}\left[\left\|S^t - \frac{\nabla\Phi(x^{t-1})}{2d}\right\|_2^2\right]}$$

As we notice now, the second term is exactly the variance of our random variable S^t. If we had no random sampling, but instead a full-information model, the second term would be zero and we could derive the convergence rate. Note that in the above inequality, if the variance was zero, we could derive a limear convergence rate.

From the above analysis two questions arise:

- Can we prove that the potential function decreases in expectation round after round? In other words, can we guarantee something for $\mathbb{E}\left[\Phi(x^t) \mid \Phi(x^{t-1}\right]$? Can we reach with provable guarantees an ε-minimizer of $\Phi(x^*)$?
- Note that if we increase the number of the samples in each round nothing changes in the above analysis. What happens though to the convergence rate?

The convergence rate of the potential function is of course affected by the number of samples we take. If we try to follow this kind of analysis, regarding the convergence on expectation of the potential function, we can establish a relation between the two. This approach does not yield any useful results and therefore, we omit the proof. Though, just for the sake of completeness, if we follow similar steps to the proofs in the descent methods and bound the variance inserted by each agent in terms of the number of samples, then we can mention that we can satisfy a linear convergence rate in each round of the form:

$$\mathbb{E}\left[\Phi(x^{t+1}) \mid \Phi(x^t)\right] - \Phi(x^*) \leq c(\varepsilon) \cdot (\Phi(x^t) - \Phi(x^*))$$

where there must be $c(\varepsilon) < 1$ and we can make the parameter ε arbitrarily small. This happens if at round $t+1$, we take at least $d_s \geq \dfrac{d \cdot n}{(\Phi(x^t) - \Phi(x^*)) \cdot \varepsilon}$ samples. This establishes a tradeoff between the number of samples in each round and how fast we are going to reach the global minimum. First, the larger the sample size, the faster we move in each round. Moreover, in each step we are going to have to sample a larger number in order to continue moving towards the global minimum. Though, the sample size is very large. Sampling in this case might be useful if we start far away from the equilibrium point. Though, after the value of the potential function decreases under a threshold value, then we need to take more samples than considering the deterministic update rule. This makes the random sampling method useless and thus the method not effective.

5.2 A variant of the FJ model with increased sampling

Consider the following protocol: We have the FJ model with limited information and the agents sample again the opinion of one neighbor in each round. Instead of updating their opinions at each time step, they wait a number of rounds and continue sampling until they have reached a close estimation of the actual averaging. In other words, assume that R_i^t is a random variable, for which it holds that $R_i^t = x_j^{t-1}$ with probability $1/d$, since we are working on an undirected d-regular graph. After r rounds we simply take the average of these samples:

$$\forall \text{ agent } i : x_i^t = \frac{1}{r} \sum_{k=1}^{r} R_{ik}^t.$$

We are going to study the convergence properties of this model. Before we state these results, let's model the above simple process. We want for each agent to draw that many samples, so that the averaging she performs with the samples she draws is *with high probability* ϵ-close to the deterministic averaging of the FJ model. This can be written as the FJ model in matrix form with the additional constraint of keeping our estimation no more than ϵ-far for each agent:

$$x^{t+1} = A \cdot y^t + B \cdot s$$
$$\left\|y^{t+1} - x^{t+1}\right\|_\infty \leq \epsilon$$

We suppose that at round t we are at the point y^t. Therefore, x^{t+1} would be our next point in the deterministic case, and then we demand that all of the agents estimate their coordinate ϵ-close. Let us assume too, that the additional constraint is satisfied in each round (exclude for the moment that this happens in each round with high probability). Also notice that if we have a large but constant sample size, this implies a fixed ϵ in our system.

5.2.1 How close to the equilibrium point can we get?

Ideally, we would like to show that we can get arbitrarily close to the equilibrium point and prove that once we reach such a state we remain there with high probability. This means that we want to obtain a characterization of the following form:

$$t \geq t_0 \ , \ \Pr[\left\|x^t - x^*\right\|_2 > \epsilon] < \eta$$

where η is an arbitrarily small positive number.

First, we should ask if such a characterization is even possible to obtain. We give a negative answer to that.

Lemma 5.1. In our randomized FJ model variant with constant sample size, in the long run we are guaranteed to converge inside an $\dfrac{\epsilon}{1 - \|A\|_\infty}$ - ball, for fixed ϵ.

Proof. We are going again to bound the distance to the equilibrium point. Assume without loss of generality that we are at time step $t+1$. We apply iteratively the update rule along with the constraint we introduced, in order to relate the distance at this time step or as t grows to infinity with the initial distance from the equilibrium point, so that we can notice the cumulative progress we make.

First, by adding and subtracting x^{t+1} and the subadditivity of the l_∞ norm we get that:

$$\left\|y^{t+1} - x^*\right\|_\infty = \left\|y^{t+1} - x^{t+1} + x^{t+1} - x^*\right\|_\infty \leq \left\|y^{t+1} - x^{t+1}\right\|_\infty + \left\|x^{t+1} - x^*\right\|_\infty$$

Applying the update rule and the constraint that we want to hold with high probability, using the fact that $x^* = A \cdot x^* + B \cdot s$ and the submultiplicative property of the norm:

$$\left\|y^{t+1} - x^*\right\|_\infty \leq \epsilon + \left\|A \cdot y^t + B \cdot s - A \cdot x^* - B \cdot s\right\|_\infty \leq \epsilon + \|A\|_\infty \cdot \left\|y^t - x^*\right\|_\infty$$

If we follow the same process a second time it yields:

$$\begin{aligned}\left\|y^{t+1}-x^*\right\|_\infty &\leq \epsilon + \|A\|_\infty \cdot \left\|y^t - x^t\right\|_\infty + \|A\|_\infty \cdot \left\|x^t - x^*\right\|_\infty \\ &\leq \epsilon \cdot (1+\|A\|_\infty) + \|A\|_\infty \cdot \left\|A \cdot y^{t-1} + B \cdot s - A \cdot x^* - B \cdot s\right\|_\infty \\ &\leq \epsilon \cdot (1+\|A\|_\infty) + \|A\|_\infty^2 \cdot \left\|y^{t-1} - x^*\right\|_\infty\end{aligned}$$

Therefore, we can apply iteratively the process $t-1$ times and we end up with:

$$\begin{aligned}\left\|y^{t+1}-x^*\right\|_\infty &\leq \epsilon \cdot (\|A\|_\infty + \|A\|_\infty^2 + \cdots + \|A\|_\infty^{t-1} + \|A\|_\infty^t) + \|A\|_\infty^{t+1} \cdot \left\|y^0 - x^*\right\|_\infty \\ &\leq \epsilon \cdot (\|A\|_\infty + \|A\|_\infty^2 + \cdots + \|A\|_\infty^{t-1} + \|A\|_\infty^t) + \|A\|_\infty^{t+1}\end{aligned}$$

The last inequality follows from the fact that $\left\|y^0 - x^*\right\|_\infty \leq 1$ since all the opinions belong in the interval $[0,1]$. Note that since we have assumed that all the agents give some weight to the intrinsic opinion they hold about a topic(i.e. the elements in the diagonal of B are strictly positive) then $\|A\|_\infty < 1$. Thus, as t grows the geometric series will eventually converge and we obtain

$$\lim_{t\to\infty} \left\|y^{t+1} - x^*\right\|_\infty \leq \frac{\epsilon}{1-\|A\|_\infty} \tag{35}$$

This concludes the proof.

Our final goal is to prove convergence with high probability. We know by now that we cannot assume that we can get arbitrarily close to the equilibrium point and therefore we cannot use probabilistic arguments for this case. Instead, the guarantee we have is that our system can reach an $\frac{\epsilon}{1-\|A\|_\infty}$ - ball. Next, we ask a simple question: How many rounds are there needed to reach this ball around convergence?

Lemma 5.2. After a finite number of rounds t we are guaranteed to reach an $\frac{c\cdot\epsilon}{1-\|A\|_\infty}$ - ball, where $c > 1$.

Proof. We will follow the previous proof. We bound again the distance $\left\|y^{t+1} - x^*\right\|_\infty$. By the end of the previous proof at some point we obtained the following inequality:

$$\left\|y^t - x^*\right\|_\infty \leq \epsilon \cdot (\|A\|_\infty + \|A\|_\infty^2 + \cdots + \|A\|_\infty^{t-2} + \|A\|_\infty^{t-1}) + \|A\|_\infty^t \cdot v(0)$$

where $v(0) = \left\|y^0 - x^*\right\|_\infty$ by the start of the iterative process. Here, we are not going to examine the asymptotic behavior. Instead, using the geometric series we obtain:

$$\left\|y^t - x^*\right\|_\infty \leq \epsilon \cdot \frac{1 - \|A\|_\infty^t}{1 - \|A\|_\infty} + \|A\|_\infty^t \cdot v(0)$$

Setting for example our constant $c = 2$, we want to know when we enter the corresponding ball:

$$\epsilon \cdot \frac{1 - \|A\|_\infty^t}{1 - \|A\|_\infty} + \|A\|_\infty^t \cdot v(0) \leq \frac{2\epsilon}{1 - \|A\|_\infty} \Leftrightarrow$$
$$\|A\|_\infty^t \left(v(0) - \frac{\epsilon}{1 - \|A\|_\infty} \right) \leq \frac{\epsilon}{1 - \|A\|_\infty}$$

and finally we get the lower bound:

$$t \geq \frac{\ln \left(\frac{\frac{\epsilon}{1 - \|A\|_\infty}}{v(0) - \frac{\epsilon}{1 - \|A\|_\infty}} \right)}{\ln \|A\|_\infty} \tag{36}$$

Of course, the initial point has to be outside the ball around x^*, otherwise the bound does not make any sense. Since $\ln\|A\|_\infty < 0$, then also it has to hold that:

$$\frac{\frac{\epsilon}{1 - \|A\|_\infty}}{v(0) - \frac{\epsilon}{1 - \|A\|_\infty}} < 1 \Leftrightarrow v(0) > \frac{2\epsilon}{1 - \|A\|_\infty}$$

5.2.2 Convergence with high probability

So now we are ready to prove that starting with an arbitrary vector y^0 not close to the equilibrium point we can converge inside the ball that we are guaranteed.

Lemma 5.3. After t rounds, with probability at least $1 - \delta$ we have $\left\|y^t - x^*\right\|_\infty \leq \frac{c \cdot \varepsilon}{1 - \|A\|_\infty}$ if we take at least $\frac{\ln(2nt/\delta)}{2\varepsilon^2}$ samples per round .

Proof. For this purpose, we are going to use Hoeffding's inequality [2.8]. The protocol indicates that all the agents take one sample at each round simultaneously and they wait r time steps(or take r samples) till they have an ϵ-close estimation of the deterministic update rule. Therefore in our case we have:

$$\Pr\left[\left| \frac{\sum_{j=1}^{r} X(i)}{r} - \frac{\sum_{j \in N_i} X(i)}{d} \right| > \varepsilon \right] \leq 2e^{-2r\varepsilon^2}$$

The above bounds the probability of the bad event for an agent, so that the constraint $\left|y_i^t - x_i^t\right| \leq \epsilon$ holds for agent i. In this section, we return to the analysis of our arguments on an unweighted d-regular graph. Therefore, the update rule becomes a simple averaging. Since each agent also holds an intrinsic opinion, imagine, for example that matrix A is a $d \times d$ matrix, which takes the values $A_{ij} = \frac{1}{d+1}$, if $i \neq j$ and all the diagonal elements are zero. This way also the sum of each row of the matrix is less than 1. Then we also put equal weight on the intrinsic opinion of an agent+ so matrix B takes the values $B_{ii} = \frac{1}{d+1}$ for all the agents and 0 otherwise.

Of course, we want for all the agents to avoid the bad outcome, so that with high probability $\left\|y^t - x^t\right\|_\infty$ holds. In this case we need a union bound on the n agents:

$$\begin{aligned}\Pr\left[\exists \text{ agent } i : \left|\frac{\sum_{j=1}^{r} X(i)}{r} - \frac{\sum_{j\in N_i} X(i)}{d}\right| > \varepsilon\right] &\leq \sum_{i=1}^{n} \Pr\left[\left|\frac{\sum_{j=1}^{r} X(i)}{r} - \frac{\sum_{j\in N_i} X(i)}{d}\right| > \varepsilon\right] \\ &\leq 2ne^{-2r\varepsilon^2}\end{aligned}$$

We need though the constraint to hold for each round in case we want with high probability to move towards the equilibrium point after each round. As we proved, we are guaranteed to get inside an $\dfrac{\epsilon}{1 - \|A\|_\infty}$ - ball after t or more rounds. using again the Hoeffding and Union bound:

$$\begin{aligned}\Pr\left[\exists t, i : \left|\frac{\sum_{j=1}^{r} X(i)}{r} - \frac{\sum_{j\in N_i} X(i)}{d}\right| > \varepsilon\right] &\leq \sum_{i=1}^{t}\sum_{j=1}^{n} \Pr\left[\left|\frac{\sum_{j=1}^{r} X(i)}{r} - \frac{\sum_{j\in N_i} X(i)}{d}\right| > \varepsilon\right] \\ &\leq 2nte^{-2r\varepsilon^2}\end{aligned}$$

Now we want to make that probability arbitrarily small, so assume :

$$2nte^{-2r\varepsilon^2} \leq \delta \Leftrightarrow -2r\varepsilon^2 \leq \frac{\delta}{2nt} \Leftrightarrow 2r\varepsilon^2 \geq \ln\left(\frac{2nt}{\delta}\right) \Leftrightarrow$$

$$r \geq \frac{\ln\left(\frac{2nt}{\delta}\right)}{2\varepsilon^2} \tag{37}$$

This gives a lower bound on the number of samples that we have to take in each round in order for our analysis to hold with high probability. This concludes the proof.

It is also common to relate the parameter δ with the number n of the agents, so we usually want that $\delta \geq \frac{1}{n^c}$, where c is an arbitrary constant, greater than 1. If we also make an assumption about the number of rounds and give a lower bound which is related to the number of agents, then by substituting in the above inequality we get an expression which is of the form $\Omega\left(\ln n/\varepsilon^2\right)$.

From the above lemma we can draw another conclusion. An approach with constant sample size does not work, since we described this case as a linear system, and even if the estimation error remains smaller than ε deterministically, we cannot hope for much. If we follow the analysis of the previous lemma and we increase the sample size in each round, since it is dependent on t and by readapting ε, we can actually get arbitrarily close to the equilibrium point. If we sample more and more, the variance will keep reducing and our system can converge in an $\frac{\epsilon}{1-\|A\|_\infty}$-ball around the equilibrium point, for any $\varepsilon > 0$.

Βιβλιογραφία

[1] M. H. Degroot, ``Reaching a consensus,'' *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118--121, 1974.

[2] N. E. Friedkin and E. C. Johnsen, ``Social influence and opinions,'' *The Journal of Mathematical Sociology*, vol. 15, no. 3-4, pp. 193--206, 1990.

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

[4] A. Das, S. Gollapudi, R. Panigrahy, and M. Salek, ``Debiasing social wisdom,'' in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '13*, ACM Press, 2013.

[5] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin, ``Novel multidimensional models of opinion dynamics in social networks,'' *IEEE Transactions on Automatic Control*, vol. 62, pp. 2270--2285, may 2017.

[6] P. Frasca, C. Ravazzi, R. Tempo, and H. Ishii, ``Gossips and prejudices: Ergodic randomized dynamics in social networks,'' *IFAC Proceedings Volumes*, vol. 46, no. 27, pp. 212 -- 219, 2013. 4th IFAC Workshop on Distributed Estimation and Control in Networked Systems (2013).

[7] A. V. Proskurnikov, R. Tempo, M. Cao, and N. E. Friedkin, ``Opinion evolution in time-varying social influence networks with prejudiced agents,'' *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11896 -- 11901, 2017. 20th IFAC World Congress.

[8] M. Epitropou, D. Fotakis, M. Hoefer, and S. Skoulakis, ``Opinion formation games with aggregation and negative influence,'' in *Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12-14, 2017, Proceedings*, pp. 173--185, 2017.

[9] R. Hegselmann and U. Krause, ``Opinion dynamics and bounded confidence: Models, analysis and simulation,'' *Journal of Artificial Societies and Social Simulation*, vol. 5, pp. 1--24, 2002.

[10] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, ``Mixing beliefs among interacting agents,'' *Advances in Complex Systems*, vol. 03, pp. 87--98, jan 2000.

[11] A. Bhattacharyya, M. Braverman, B. Chazelle, and H. L. Nguyen, ``On the convergence of the hegselmann-krause system,'' in *Proceedings of the 4th conference on Innovations in Theoretical Computer Science - ITCS '13*, ACM Press, 2013.

[12] B. Chazelle, ``The total s-energy of a multiagent system,'' *SIAM Journal on Control and Optimization*, vol. 49, no. 4, pp. 1680--1706, 2011.

[13] D. Bindel, J. Kleinberg, and S. Oren, ``How bad is forming your own opinion?,'' in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pp. 57--66, Oct 2011.

[14] K. Bhawalkar, S. Gollapudi, and K. Munagala, ``Coevolutionary opinion formation games,'' in *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, (New York, NY, USA), pp. 41--50, ACM, 2013.

[15] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, ``Binary opinion dynamics with stubborn agents,'' *ACM Trans. Econ. Comput.*, vol. 1, pp. 19:1--19:30, Dec. 2013.

[16] D. Acemoglu and A. Ozdaglar, ``Opinion dynamics and learning in social networks,'' *Dynamic Games and Applications*, vol. 1, pp. 3--49, Mar 2011.

[17] R. Abebe, J. Kleinberg, D. Parkes, and C. E. Tsourakakis, ``Opinion dynamics with varying susceptibility to persuasion,'' in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, (New York, NY, USA), pp. 1089--1098, ACM, 2018.

[18] M. O. Jackson, *Social and Economic Networks.* Princeton, NJ, USA: Princeton University Press, 2008.

[19] Y. K. Cheung, R. Cole, and N. Devanur, ``Tatonnement beyond gross substitutes?,'' in *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing - STOC '13*, ACM Press, 2013.

[20] B. Birnbaum, N. Devanur, and L. Xiao, ``Distributed algorithms via gradient descent for fisher markets,'' in *Proceedings of the 12th ACM Conference on Electronic Commerce*, ACM, June 2011.

[21] S. Boyd and L. Vandenberghe, *Convex Optimization.* New York, NY, USA: Cambridge University Press, 2004.

[22] R. W. Rosenthal, ``A class of games possessing pure-strategy nash equilibria,'' *International Journal of Game Theory*, vol. 2, pp. 65--67, Dec 1973.

[23] D. Monderer and L. S. Shapley, ``Potential games,'' *Games and Economic Behavior*, vol. 14, no. 1, pp. 124 -- 143, 1996.

[24] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course.* Springer Publishing Company, Incorporated, 1 ed., 2014.

[25] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization.* Society for Industrial and Applied Mathematics, 2001.

[26] L. Bottou, F. E. Curtis, and J. Nocedal, ``Optimization methods for large-scale machine learning,'' *SIAM Review*, vol. 60, pp. 223--311, jan 2018.

[27] M. Zinkevich, ``Online convex programming and generalized infinitesimal gradient ascent,'' in *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, pp. 928--935, AAAI Press, 2003.

[28] Y. Freund and R. E. Schapire, ``A decision-theoretic generalization of on-line learning and an application to boosting,'' *Journal of Computer and System Sciences*, vol. 55, no. 1, pp. 119 -- 139, 1997.

[29] S. Arora, E. Hazan, and S. Kale, ``The multiplicative weights update method: a meta-algorithm and applications,'' *Theory of Computing*, vol. 8, no. 6, pp. 121--164, 2012.

[30] N. Vishnoi, *A Mini-Course on Convex Optimization with a view towards designing fast algorithms.* October 2015.

[31] E. Hazan, ``Introduction to online convex optimization,'' *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157--325, 2016.

[32] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, ``Online convex optimization in the bandit setting: Gradient descent without a gradient,'' in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, (Philadelphia, PA, USA), pp. 385--394, Society for Industrial and Applied Mathematics, 2005.

[33] L. Bottou, ``Large-scale machine learning with stochastic gradient descent,'' in *Proceedings of COMPSTAT'2010* (Y. Lechevallier and G. Saporta, eds.), (Heidelberg), pp. 177--186, Physica-Verlag HD, 2010.

[34] T. Zhang, ``Solving large scale linear prediction problems using stochastic gradient descent algorithms,'' in *Proceedings of the Twenty-first International Conference on Machine Learning*, ICML '04, (New York, NY, USA), pp. 116--, ACM, 2004.

[35] M. Schmidt, N. Le Roux, and F. Bach, ``Minimizing finite sums with the stochastic average gradient,'' *Mathematical Programming*, vol. 162, pp. 83--112, Mar 2017.

[36] L. Xiao and T. Zhang, ``A proximal stochastic gradient method with progressive variance reduction,'' *SIAM Journal on Optimization*, vol. 24, pp. 2057--2075, jan 2014.

[37] S. Shalev-Shwartz and T. Zhang, ``Stochastic dual coordinate ascent methods for regularized loss,'' *Journal of Machine Learning Research*, vol. 14, no. 1, pp. 567--599, 2013.

[38] S. Shalev-Shwartz and T. Zhang, ``Accelerated mini-batch stochastic dual coordinate ascent,'' in *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'13, (USA), pp. 378--385, Curran Associates Inc., 2013.

[39] R. Johnson and T. Zhang, ``Accelerating stochastic gradient descent using predictive variance reduction,'' in *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'13, (USA), pp. 315--323, Curran Associates Inc., 2013.

[40] H. Lin, J. Mairal, and Z. Harchaoui, ``A universal catalyst for first-order optimization,'' in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, (Cambridge, MA, USA), pp. 3384--3392, MIT Press, 2015.

[41] Z. Allen-Zhu, ``Katyusha: The first direct acceleration of stochastic gradient methods,'' in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, (New York, NY, USA), pp. 1200--1205, ACM, 2017.

[42] D. Fotakis, D. Palyvos-Giannas, and S. Skoulakis, ``Opinion dynamics with local interactions,'' in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pp. 279--285, AAAI Press, 2016.

[43] T. Roughgarden, *Twenty Lectures on Algorithmic Game Theory.* New York, NY, USA: Cambridge University Press, 1st ed., 2016.

[44] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* New York, NY, USA: Cambridge University Press, 2005.

[45] G. Strang, *Introduction to Linear Algebra.* Wellesley, MA: Wellesley-Cambridge Press, fourth ed., 2009.

[46] S. Skoulakis, *Παίγνια διαμόρφωσης άποψης σε κοινωνικά δίκτυα.* Diploma Thesis, October 2014.

[47] C. D. Meyer, ed., *Matrix Analysis and Applied Linear Algebra.* Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2000.

[48] P. Brémaud, *Markov Chains.* Springer New York, 1999.

[49] 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, (New York, NY, USA), pp. 604--612, ACM, 2004.

[50] D. J. Foster, Z. Li, T. Lykouris, K. Sridharan, and É. Tardos, ``Fast convergence of common learning algorithms in games,'' *CoRR*, vol. abs/1606.06244, 2016.

[51] W. Hoeffding, ``Probability inequalities for sums of bounded random variables,'' *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13--30, 1963.

[52] P. G. Spirakis and P. N. Panagopoulou, ``Potential functions in strategic games,'' in *Computer Science – Theory and Applications*, pp. 283--297, Springer Berlin Heidelberg, 2013.