Εθνικο Μετσοβιο Πολυτεχνειο
Σχολη Ηλεκτρολογων Μηχανικων Και Μηχανικων Υπολογιστων

ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΜΩΝ (Co.Re.Lab.)

Δυναμικές Διαμόρφωσης Άποψης με Περιορισμένη Πληροφορία

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

Άνθιμου-Βαρδή Α. Κανδήρου

Αθήνα, Οκτώβριος 2018

Εθνικο Μετσοβιο Πολυτεχνειο
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ
ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΜΩΝ

Δυναμικές Διαμόρφωσης Άποψης με Περιορισμένη Πληροφορία

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

Άνθιμου Βαρδή Α. Κανδήρου

Επιβλέπων Καθηγητής: Δημήτριος Φωτάκης
Επίκουρος Καθηγητής ΕΜΠ

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 17^{η} Οκτωβρίου 2018.

....	....	....
Δημήτριος Φωτάκης Επ. Καθηγητής ΕΜΠ	Αριστείδης Παγουρτζής Αναπ. Καθηγητής ΕΜΠ	Ευάγγελος Μαρκάκης Επ. Καθηγητής ΟΠΑ

Αθήνα, Οκτώβριος 2018.

....

Άνθιμος -Βαρδής Α. Κανδήρος
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Περίληψη

Στη μελέτη των κοινωνικών δικτύων, ένα βασικό ερώτημα αφορά τον τρόπο με τον οποίο διαμορφώνονται οι απόψεις των παικτών. Ένα από τα πιο σημαντικά μοντέλα για την περιγραφή των αλληλεπιδράσεων μεταξύ των ατόμων του δικτύου είναι το μοντέλο Friedkin-Johnsen.

Ένα μειονέκτημα αυτού του μοντέλου είναι ότι υποθέτει πως κάθε παίκτης λαμβάνει υπ'οψιν του τις απόψεις *όλων* των φίλων του για να διαμορφώσει τη δική του. Στα σύγχρονα τεράστια κοινωνικά δίκτυα, αυτή η θεώρηση είναι μη ρεαλιστική. Στην εργασία αυτή προτείνουμε μια παραλλαγή του FJ μοντέλου, όπου κάθε παίκτης λαμβάνει περιορισμένη πληροφορία σε κάθε γύρο για τις απόψεις των γειτόνων του. Δείχνουμε ότι ο αλγόριθμος του νέου μοντέλου έχει την ιδιότητα no regret, το οποίο σημαίνει ότι οι παίκτες έχουν κίνητρο να συμπεριφερθούν με αυτό τον τρόπο. Επίσης, αποδεικνύουμε τη σύγκλιση του νέου μοντέλου στο ίδιο σημείο ισορροπίας με το αρχικό FJ. Επιπλέον, δείχνουμε ότι το μοντέλο μας έχει το γρηγορότερο ρυθμό σύγκλισης μεταξύ όλων των πρωτοκόλλων με την ιδιότητα no regret. Τέλος, δίνουμε παραδείγματα αλγορίθμων που δεν είναι no regret και επιτυγχάνουν γρηγορότερο ρυθμό σύγκλισης.

Λέξεις -Κλειδιά: Friedkin-Johnsen, Περιορισμένη Πληροφορία, Κάτω Φράγμα, Εκτίμηση Bernoulli

Abstract

In the study of social networks, a fundamental question concerns the way players in a network form their opinions. One of the most important models that describes the interactions between the players is the Friedkin-Johnsen model.

A drawback of this model is the assumption that each player considers the opinions of *all* of his friends before forming his own. In today's huge social networks, this assumption is unrealistic. In this thesis, we propose a variant of the FJ model, where each player receives limited information about his friends' opinions in each round. We show that the algorithm of the new model has the *no regret* property, which means that it is a natural choice for the players to adopt. Moreover, we prove the convergence of the new model to the same stable point as that of the original FJ model. Also, we show that our model is among the ones with the fastest convergence rate in the family of no regret protocols. Finally, we provide examples of algorithms that are not no regret and achieve faster convergence rates.

Key words: Friedkin-Johnsen, limited information, lower bound, Bernoulli estimation

Ευχαριστίες

Θα ήθελα να ευχαριστήσω θερμά τους καθηγητές του Εργαστηρίου Λογικής και Επιστήμης Υπολογιστών, κ. Ζάχο, κ. Παγουρτζή και κ. Φωτάκη για τις συμβουλές και την καθοδήγησή τους κατά τη διάρκεια των προπτυχιακών μου σπουδών.

Ιδιαίτερα, θα ήθελα να ευχαριστήσω τον κύριο Φωτάκη για τη βοήθεια και την υποστήριξη που μου παρείχε κατά τη διάρκεια της εκπόνησης της παρούσας διπλωματικής εργασίας. Η ενθάρρυνση και η εμπιστοσύνη που μου έδειξε ,τόσο σε ακαδημαϊκό όσο και σε προσωπικό επίπεδο, ήταν καθοριστικές για την εισαγωγή μου στην ερευνητική διαδικασία και για τα επόμενα βήματά μου στον ακαδημαϊκό χώρο.

Θα ήθελα επίσης να ευχαριστήσω όλα τα μέλη του Εργαστηρίου Λογικής και Υπολογιστών για το ευχάριστο περιβάλλον συνεργασίας που έχουν δημιουργήσει. Ιδιαίτερα θα ήθελα να ευχαριστήσω τους Στρατή Σκουλάκη και Βασίλη Κοντονή για την τεράστιας σημασίας συνεργασία τους στο ερευνητικό κομμάτι της παρούσας διπλωματικής, της οποίας τα αποτελέσματα δημοσιεύσαμε στο Conference on Web and Internet Economics [ΔΦΣ18] που θα γίνει το Δεκέμβριο του 2018. Χωρίς τη συμβολή τους η προσπάθεια αυτή δε θα είχε ολοκληρωθεί.

Θα ήθελα τέλος να ευχαριστήσω τους ανθρώπους που ήταν κοντά μου όλα αυτά τα χρόνια για την συμπαράσταση που μου προσέφεραν. Ιδιαίτερα θα ήθελα να ευχαριστήσω τους γονείς μου, που στάθηκαν στο πλευρό μου σε όλα τα στάδια της εκπαίδευσής μου και υποστήριξαν τις επιλογές μου.

Βαρδής Άνθιμος Κανδήρος
Αθήνα, Οκτώβριος 2018

Εκτεταμένη περίληψη στα Ελληνικά

Σε αυτό το Κεφάλαιο παρουσιάζουμε μια περίληψη των αποτελεσμάτων αυτής της διπλωματικής. Αρχικά, γίνεται μια ανοφορά στον τρόπο με τον οποίο οι άνθρωποι αλληλεπιδρούν σε ένα κοινωνικό δίκτυο. Ένα γνωστό μαθηματικό μοντέλο που περιγράφει αυτή την αλληλεπίδραση είναι το μοντέλο Friedkin-Johnsen , όπου η διαμόρφωση των απόψεων προσομοιάζει έναν επαναλαμβανόμενο αλγόριθμο που εκτελείται από τους ανθρώπους στο δίκτυο. Στη συνέχεια, εξετάζουμε διάφορους περιορισμούς του FJ μοντέλου όσον αφορά την περιγραφή πραγματικών κοινωνικών δικτύων. Συγκεκριμένα, υποστηρίζουμε ότι δεν είναι ρεαλιστικό ένα άτομο να ρωτά τις απόψεις όλων των φίλων του κάθε φορά που αλλάζει τη δική του. Στη συνέχεια, προτείνουμε μια παραλλαγή του μοντέλου με περιορισμένη πληροροφία, ώστε να αναπαριστά καλύτερα την πραγματικότητα και παρουσιάζουμε αποτελέσματα σύγκλισης για αυτή. Τέλος, στην τελευταία Ενότητα επιχειρηματολογούμε για το ότι δεν υπάρχει γρηγορότερος αλγόριθμος που να λύνει το πρόβλημα της διαμόρφωσης απόψεων.

Διαμόρφωση Απόψεων και το μοντέλο FJ

Η μελέτη της *Διαμόρφωσης Απόψεων* έχει πλούσια ιστορία(π.χ. [Θας08]) . Η Διαμόρφωση Απόψεων είναι μια *δυναμική διαδιδασία*, με την έννοια ότι κοινωνικά συνδεδεμένοι άνθρωποι(π.χ. οικογένεια, φίλοι, συνάδελφοι) ανταλλάσσουν πληροφορία και αυτό οδηγεί σε αλλαγές στις απόψεις που διατυπώνουν. Στις μέρες μας, η έλευση του Διαδικτύου και των Κοινωνικών μέσων καθιστά τη μελέτη της διαμόρφωσης απόψενω σε μεγάλα κοινωνικά δίκτυα ακόμη πιο σημαντική. Ρεαλιστικά μοντέλα της συμπεριφοράς των ανθρώπων έχουν μεγάλη πρακτική χρησιμότητα για πρόβλεψη συμπεριφορών, στοχευμένη διαφήμιση κ.λ.π. Στην προσπάθειά τους να τυποποιήσουν τη διαδικασία της διαμόρφωσης απόψεων, οι κοινωνιολόγοι έχουν προτείνει διάφορα μοντέλα. Σε αυτά τα μοντέλα, τα άτομα του δικτύου ονομάζονται *παίκτες*.

Η κοινή βάση όλων αυτών των μοντέλων, η οποία προτάθηκε από τον De Groot [DeG74] , είναι ότι οι απόψεις εξελίσσονται μέσω μιας διαδικασίας επαναλαμβανόμενων μέσων όρων των απόψεων που οι παίκτες συλλέγουν από τη ¨γειτονιά' τους. Στο μοντέλο του De Groot , η άποψη ενός παίκτη αναπαρίσταται ως ένας πραγματικός αριθμός στο διάστημα $[0, 1]$. Οι παίκτες επικοινωνούν με τους φίλους τους, το οποίο οδηγεί στην αλλαγή των απόψεών τους. Σε κάθε γύρο, κάθε παίκτης αλλάζει την άποψή τους παίρνοντας ένα σταθμισμένο μέσο όρο των απόψεων των φίλων του. Τυπικά, τα βάρη που τοποθετούν οι παίκτες στις απόψεις των γειτόνων τους αναπαρίστανται με ένα *πίνακα εμπιστοσύνης* T, όπου T_{ij} είναι το βάρος που ο παίκτης i βάζει στην άποψη του παίκτη j. Φυσικά, ο T είναι στοχαστικός πίνακας, το οποίο σημαίνει ότι το άθροισμα των στοιχείων κάθε γραμμής είναι 1. Αν συμβολίσουμε με $z(t)$ το διάνυσμα των απόψεων την περίοδο t, τότε την περίοδο $t+1$ το νέο διάνυσμα απόψεων που διαμορφώνεται αφού έχουν γίνει όλες οι αλληλεπιδράσεις είναι $z(t + 1) = Tz(t)$. Παρόλο που αυτό το μοντέλο δεν περιγράφει με ακρίβεια τις πραγματικές κοινωνικές αλληλεπιδράσεις, η απλότητα και φυσικότητά του παρέχουν μια διαίσθηση για τους παράγοντες που επηρεάζουν τη διαμόρφωση απόψεων.

Μια κάπως διαφορετική προσέγγιση δίνεται σε ένα μοντέλο που προτάθηκε από τους Hegselmann και Krause(HK model)στο [HK02]. Σε αυτό το μοντέλο, οι απόψεις των παικτών είναι πάλι πραγματικοί αριθμοί στο $[0, 1]$. Επιπλέον, κάθε παίκτης i έχει ένα *παράγοντα εμπιστοσύνης* ϵ_i. Σε κάθε γύρο, ο παίκτης i ανανεώνει την άποψή του ως το μέσο όρο των απόψεων των παικτών που βρίσκονται σε ϵ_i απόσταση από τη δική του άποψη. Αυτό σημαίνει ότι οι κοινωνικές σχέσεις ενός παίκτη μπορεί να αλλάζουν καθώς περνά ο χρόνος. Αυτό είναι προφανώς μια πιο ρεαλιστική αναπαράσταση των κοινωνικών σχέσεων, καθώς οι άνθρωποι τείνουν να επηρεάζονται περισσότερο από τις απόψεις των ομοϊδεατών τους.

Ένα άλλο ενδιαφέρον μοντέλο προτάθηκε στο [ΔΝΑΩ00]. Σε αυτό το μοντέλο, σε κάθε γύρο δύο παίκτες επιλέγονται τυχαία με ομοιόμορφο τρόπο από το σύνολο όλων των ζευγών στον πληθυσμό. Αυτοί οι δύο παίκτες αλλάζουν τις απόψεις τους, υπό την προϋπόθεση ότι η διαφορά στις απόψεις τους δεν είναι μεγαλύτερη από ένα κατώφλι d. Η λογική πίσω από τη συνθήκη κατωφλίου είναι ότι οι παίκτες επηρεάζουν ο ένας τον άλλο μόνο όταν οι απόψεις τους είναι ήδη κάπως παρόμοιες, διαφορετικά δεν μπαίνουν καν στον κόπο να συζητήσουν. Αυτή η συμπεριφορά μπορεί να οφείλεται σε έλλειψη αμοιβαίας κατανόησης , σύγκρουση συμφερόντων ή κοινωνική πίεση.

Οι εργασίες μετά τον DeGroot εστίασαν στη διατύπωση γενικών συνθηκών, κάτω από τις οποίες τέτοιες διαδικασίες διαμόρφωσης απόψεων θα συγκλίνουν σε μια τελική κατάσταση συμφωνίας, στην οποία όλοι οι παίκτες έχουν την ίδια άποψη. Ωστόσο, αυτή η έμφαση στην επίτευξη συμφωνίας περιορίζει το φάσμα την διαδικασίας μοντελοποίησης σε ένα συγκεκριμένο τύπο δυναμικών διαμόρφωσης άποψης, όπου όλες οι απόψεις του συνόλου συγκλίνουν. Όπως παρατήρησε ο κοινωνιολόγος David Krackhardt

> Δε θα πρέπει να αγνοήσουμε το γεγονός ότι στον πραγματικό κόσμο η συμφωνία συνήθως δεν επιτυγχάνεται. Οι περισσότεροι επιστήμονες κοινωνικών δικτύων αναγνωρίζουν αυτό το γεγονός και δεν επικεντρώνονται σε μια κατάσταση ισορροπίας όπου επικρατεί συμφωνία. Αντ'αυτού, είναι πιο πιθανό να προσπαθήσουν να εξηγήσουν την έλλειψη συμφωνίας (τη μεταβλητότητα) σε απόψεις και στάσεις που εμφανίζονται σε πραγματικά κοινωνικά περιβάλλοντα. [Κρα09]

Στο μοντέλο που μας ενδιαφέρει, το οποίο προτάθηκε από τους Friedkin καιJohnsen στο [ΦΘ90], η συμφωνία των απόψεων συνήθως δεν επιτυγχάνεται. Στο μοντέλο FJ, έχουμε ένα μη κατευθυνόμενο γράφημα χωρίς βάρη $G(V, E)$, το οποίο αναπαριστά ένα κοινωνικό δίκτυο. Το γράφημα έχει n κόμβους. Κάθε κόμβος στο δίκτυο αναπαριστά έναν παίκτη. Η ύπαρξη ακμής μεταξύ δύο παικτών συμβολίζει ότι σχετίζονται με κάποιο τρόπο. Κάθε παίκτης έχει μια άποψη, η οποία είναι ένας πραγματικός αριθμός στο $[0, 1]$. Το μοντέλο καθορίζει μια συγκεκριμένη διαδικασία, μέσω της οποίας κάθε παίκτης ανανεώνει την άποψή του. Αυτή η διαδιδασία εκτελείται σε γύρους. Κάθε παίκτης $i \in V$ έχει μια τιμή $s_i \in [0, 1]$, η οποία αναπαριστά την εσωτερική του άποψη και μένει σταθερή με το πέρασμα του χρόνου. Επιπλέον, στο γύρο t κάθε παίκτης i παράγει σαν έξοδο έναν αριθμό $x_i(t) \in [0, 1]$, που είναι η άποψή του στον τρέχον γύρο. Η διαδικασία τρέχει ως εξής. Αρχικά, το $x_i(0)$ μπορεί να έχει οποιαδήποτε τιμή στο $[0, 1]$. Στο γύρο t, κάθε παίκτης i ανανεώνει την άποψή του $x_i(t)$ χρησιμοποιώντας τον ακόλουθο κανόνα:

$$x_i(t) = \frac{\sum_{j \in N_i} w_{ij} x_j(t-1) + w_{ii} s_i}{\sum_{j \in N_i} w_{ij} + w_{ii}}, \tag{1}$$

, όπου N_i είναι το σύνολο των γειτόνων του παίκτη i και το διάνυσμα $x(t)$ των απόψεων των παικτών στο γύρο t. Επίσης, το βάρος w_{ij} που αντιστοιχεί στην ακμή $(i, j) \in E$ μετράει πόση επίδραση έχει ο j στην άποψη του i και το βάρος w_{ii} ποσοτικοποιεί πόσο επιρρεπής είναι ο i στην υιοθέτηση απόψεων που διαφέρουν από την εσωτερική του άποψη.

Συνεπώς, κάθε παίκτης υπολογίζει το σταθμισμένο μέσο όρο των απόψεων των γειτόνων του στον προηγούμενο γύρο και σχηματίζει τη νέα του άποψη. Αυτό το είδος διαδικασίας ανανέωσης καλείται μια *δυναμική*, μια απλή διαδικασία που δεν αλλάζει από γύρο σε γύρο και εκτελείται με τον ίδιο τρόπο από όλους τους παίκτες. Άλλα παραδείγματα δυναμικών δίνονται στα [ΔεΓ74, ΦΘ90, HK02, ΔΝΑΩ00].

Αξίζει να παρατηρήσουμε ότι λόγω της σταθεράς s_i σε κάθε επανάληψη, ο επαναλαμβανόμενος υπολογισμός του μέσου όρου δε θα οδηγήσει απαραίτητα σε συμφωνία τους παίκτες. Με αυτό τον τρόπο, το μοντέλο διαχωρίζει την εγγενή προσωπική άποψη s_i και την συνολική άποψη x_i ενός παίκτη. Η τελευταία εκφράζει ένα συμβιβασμό μεταξύ της επίμονης τιμής και των απόψεων που εκφράζουν οι φίλοι του παίκτη i. Αυτή η διαφοροποίηση μεταξύ και συναντάται και σε εμπειρικές μελέτες που προσπαθούν να συνδέσουν ισχυρές απόψεις όπως ο πολιτικός προσανατολισμός με διαφορές στην εκπαίδευση και στο υπόβαθρο των ανθρώπων. Μάλιστα, ορισμένες μελέτες εξετάζουν και γενετικούς παράγοντες που επηρεάζουν τέτοιες απόψεις. [ΑΦΗ+05]

Το μοντέλο FJ είναι ένα από τα πιο επιδραστικά στη μελέτη της Διαμόρφωση Απόψεων. Έχει ένα πολύ απλό κανόνα ανανέωσης, που το καθιστά κατάλληλο για την περιγραφή φυσικής συμπεριφοράς. Επίσης, οι βασικές υποθέσεις του συμφωνούν με τα εμπειρικά αποτελέσματα σχετικά με τον τρόπο με τον οποίο διαμορφώνονται οι απόψεις [ΑΦΗ+05, Κρα09]. Έχει μελετηθεί επίσης υπό το πρίσμα της Θεωρίας Παιγνίων. Στο [BKO11] θέτουν το ερώτημα του πως θα ποσοτικοποιηθεί το κόστος της διάσταση απόψεων, σε περίπτωση που οι παίκτες δεν καταλήξουν τελικά σε συμφωνία. Για να απαντήσουν αυτή την ερώτηση, θεώρησαν τον κανόνα ανανέωσης ενός παίκτη σαν αλγόριθμο ελαχιστοποίησης μιας τετραγωνικής συνάρτησης που εκφράζει το κόστος της ασυμφωνίας. Βασισμένοι σε αυτό το κόστος όρισαν το ακόλουθο *παίγνιο διαμόρφωσης απόψεων*. Κάθε κόμβος i είναι ένας ωφελιμιστής παίκτης του οποίου η στρατηγική είναι η δημόσια άποψη x_i που εκφράζει. Για κάθε παίκτη ορίζεται μια συνάρτηση κόστους, η οποία έχει ως ορίσματα τις απόψεις όλων των παικτών στο δίκτυο. Για τον παίκτη i, η συνάρτηση κόστους είναι

$$C_i(x_i, x_{-i}) = \sum_{j \in N_i} w_{ij}(x_i - x_j)^2 + w_{ii}(x_i - s_i)^2 \tag{2}$$

όπου με x_{-i} συμβολίζουμε το διάνυσμα των απόψεων όλων των άλλων παικτών εκτός του i. Μας ενδιαφέρει να μελετήσουμε ποια είναι η καλύτερη επιλογή που μπορεί να κάνει ένας παίκτης για την άποψή του, δεδομένου ότι οι απόψεις των άλλων παικτών είναι σταθερές. Σε αυτή την περίπτωση, έχουμε συνάρτηση μίας μεταβλητής, η οποία είναι μια γνησίως κυρτή συνάρτηση, το οποίο συνεπάγεται ότι υπάρχει μοναδικό σημείο ελαχίστου το οποίο μηδενίζει την παράγωγο. Έστω λοιπόν ότι ο παίκτης γνωρίζει τις απόψεις των παικτών της χρονική στιγμή. Τότε, μηδενίζοντας την παράγωγο και κάνοντας απλές πράξεις, προκύπτει ότι ο παίκτης για να ελαχιστοποιήσει το κόστος του στον επόμενο γύρο, πρέπει να επιλέξει

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

Παρατηρούμε ότι αυτός είναι ο κανόνας ανανέωσης του μοντέλου 1. Αυτό σημαίνει ότι το μοντέλο FJ ανήκει σε μια κατηγορία αλγορίθμων που ονομάζονται δυναμικές *best response* , όπου κάθε παίκτης επιλέγει ως στρατηγική του στο γύρο t αυτή που ελαχιστοποιεί το κόστος του, υπολογισμένο με βάση τις στρατηγικές των άλλων παικτών στον προηγούμενο γύρο $t-1$. Ο αλγόριθμος αυτός στο δικό μας παίγνιο είναι ο Αλγόριθμος 1, όπου $x_{-i}(t-1)$ προφανώς συμβολίζει το διάνυσμα των απόψεων όλων των παικτών εκτός του i τη χρονική στιγμή $t-1$. Αυτή η οπτική γωνία είναι πολύ

Algorithm 1 Δυναμική Best response

Αρχικά, ο παίκτης i έχει άποψη $x_i(0)$.

1: Σε κάθε γύρο t

2: Ο παίκτης i ανανεώνει ως εξής

$$x_i(t) = \operatorname*{argmin}_{x \in [0,1]} C_i\left(x, x_{-i}(t-1)\right)$$

πιο πλήρης, καθώς μπορούν να μελετηθούν πολλές διαφορές πτυχές της διαδικασίας διαμόρφωσης απόψεων ορίζοντας κατάλληλα παίγνια. Περαιτέρω εργασίες θεωρούν παραλλαγές του παραπάνω παιγνίου και μελετούν ιδιότητες σύγκλισης διαφόρων δυναμικών που μπορεί να ακολουθήσουν οι παίκτες.

Για να κατανοήσουμε καλύτερα γιατί οι παίκτες μπορεί να θέλουν να ακολουθήσουν μία δυναμική, πρέπει πρώτα να ορίσουμε μια πολύ σημαντική έννοια της Θεωρίας Παιγνίων, την *Ισορροπία Nash* ενός παιγνίου.

Definition 1 ([N+50]). *Ας υποθέσουμε ότι $x \in \mathbf{R}^n$ είναι το διάνυσμα των στρατηγικών όλων των παικτών. Τότε, το είναι μια* Ισορροπία Nash *αν για κάθε παίκτη και για κάθε στρατηγική του παίκτη, ισχύει*

$$C_i(y, x_{-i}) \geq C_i(x_i, x_{-i})$$

Ουσιαστικά, μια ισορροπία Nash είναι ένα σύνολο από στρατηγικές τέτοιες ώστε κανένας παίκτης να μη μπορεί να μειώσει το κόστος του αλλάζοντας τη στρατηγική του, αν οι άλλοι παίκτης κρατήσουν τις στρατηγικές τους σταθερές. Γι αυτό το λόγο, αν οι παίκτες του παιγνίου φτάσουν κάποια στιγμή σε αυτή την κατάσταση, είναι απίθανο κάποιος παίκτης να θέλει να αλλάξει τη στρατηγική του, γι αυτό καλείται *ισορροπία*. Η έννοια της ισορροπίας Nash είναι θεμελιώδης στις Οικονομικές επιστήμες, καθώς παρέχει έναν τρόπο πρόβλεψης της τελικής συμπεριφοράς των παικτών σε ένα παίγνιο. Θα ασχοληθούμε με το πρόβλημα του υπολογισμού αυτής της ισορροπίας Nash από τους παίκτες.

Η έννοια της *δυναμικής* λαμβάνει πλέον άλλο ένα νόημα: είναι ένας κανόνας ανανέωσης που μπορεί να οδηγήσει τους παίκτες στον υπολογισμό της Ισορροπίας Nash του παιγνίου. Αποδεικνύεται ότι το παίγνιο έχει μοναδική Ισορροπία. Θα δείξουμε ότι αν όλοι οι παίκτες εκτελούν τον Αλγόριθμο 1 , οι απόψεις των παικτών συγκλίνουν στο x^*.

Theorem 1. *Αν x^* είναι το μοναδικό σημείο Ισορροπίας του παιγνίου και $a = \min \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}}$, τότε:*

$$\|x(t) - x^*\|_\infty \leq (1-a)^t \|x(0) - x^*\|_\infty$$

Αυτό φυσικά αποτελεί πλεονέκτημα για το μοντέλο FJ , καθώς δείχνει ότι οι παίκτες οδηγούνται σε μια κατάσταση ισορροπίας όσον αφορά το κόστος τους. Το ερώτημα που τίθεται τώρα είναι κατά πόσο αυτό το μοντέλο είναι ρεαλιστικό. Για να υλοποιηθεί ο κανόνας ανανέωσης 1, πρέπει κάθε παίκτης να μαθαίνει τις απόψεις όλων των φίλων του σε κάθε γύρο. Όμως, σε πραγματικά κοινωνικά δίκτυα, ένα άτομο είναι πιο πιθανό να επηρεαστεί στο σχηματισμό άποψης από λίγους ανθρώπους, με τους οποίους μιλάει συχνά. Αυτό μας οδηγεί στο να αλλάξουμε το μοντέλο, ώστε να αντανακλά αυτή την πραγματικότητα.

Περιορισμένη Πληροφορία

Για να καλύψουμε την αδυναμία του μοντέλου να περιγράψει την πραγματική διαδικασία διαμόρφωσης πληροφορίας, παραλλάσσουμε το μοντέλο στο [BKO11]. Ερμηνεύουμε το βάρος w_{ij} ως μέτρο του πόσο συχνά συναντάει ο i τον j. Παρατηρούμε ότι το C_i δεν είναι πλέον ντετερμινιστική συνάρτηση των προηγούμενων απόψεων.

Definition 2. *Για ένα διάνυσμα απόψεων $x \in [0,1]^n$, το κόστος ασυμφωνίας του παίκτη i είναι μια τυχαία μεταβλητή $C_i(x_i, x_{-i})$ που ορίζεται ως εξής:*

- *Ο παίκτης i συναντά έναν από τους γείτονες j με πιθανότητα $p_{ij} = w_{ij} / \sum_{j \in N_i} w_{ij}$.*
- *Ο παίκτης δέχεται κόστος $C_i(x_i, x_{-i}) = (1-a_i)(x_i - x_j)^2 + a_i(x_i - s_i)^2$, όπου $\alpha_i = w_{ii}/(\sum_{j \in N_i} w_{ij} + w_{ii})$.*

Definition 3. *Συμβολίζουμε με $I = (P, s, \alpha)$ ένα παίγνιο διαμόρφωσης απόψεων, όπου P είναι ένας $n \times n$ πίνακας με μη αρνητικά στοιχεία p_{ij}, με $p_{ii} = 0$ και $\sum_{j=1}^n p_{ij}$ είναι είτε 0 είτε 1, $s \in [0,1]^n$ είναι το διάνυσμα των εσωτερικών απόψεων, $\alpha \in (0,1]^n$ το διάνυσμα αυτο-εμπιστοσύνης.*

Παρατηρούμε ότι η μέση τιμή του κόστους σε κάθε γύρο για κάθε παίκτη είναι το κόστος του στο κανονικό μοντέλο. Γι αυτό, η Ισορροπία ως προς το αναμενόμενο κόστος ασυμφωνίας είναι ξανά x^*.

Αυτό το παίγνιο μας παρέχει ένα αρχέτυπο για όλες τις δυναμικές που θα μελετήσουμε στην παρούσα εργασία. Στο γύρο t, κάθε παίκτης i επιλέγει μια άποψη $x_i(t)$ και υφίσταται ένα κόστος ασυμφωνίας με βάση το γείτονα που συνάντησε τυχαία. Στο τέλος του γύρου t, πληροφορείται μόνο για την άποψη αυτού του γείτονα που συνάντησε και μπορεί να χρησιμοποιήσει αυτή την πληροφορία για να ανανεώσει την άποψή του στον επόμενο γύρο. Προφανώς, διαφορετικοί κανόνες ανανέωσης οδηγούν σε διαφορετικές *δυναμικές*, που όμως όλες έχουν τον κοινό περιορισμό της περιορισμένης πληροφορίας. Το ερώτημα τώρα είναι αν υπάρχει κάποιος κανόνας ανανέωσης που να επιτυγχάνει σύγκλιση των απόψεων στο , δεδομένης της περιορισμένης πληροφορίας. Επίσης, είναι ενδιαφέρον να εξετάσουμε πώς επηρεάζεται ο ρυθμός σύγκλισης από αυτή την απαίτηση.

Μία ιδέα για τον κανόνα ανανέωσης προέρχεται από τη στρατηγική *Follow the Leader*. Αυτή η στρατηγική, που είναι από τις πρώτες που μελετήθηκαν στη θεωρία παιγνίων, λέει με απλά λόγια «παίξε το βέλτιστο με βάση αυτά που έχεις παρατηρήσει». Η στρατηγική περιγράφεται αναλυτικά στον Αλγόριθμο 2. Όπως μπορούμε να δούμε, ο αλγόριθμος ουσιαστικά εκτελεί

Algorithm 2 Δυναμική Follow the Leader

1: Αρχικά $x_i(0) = s_i$ για όλους τους παίκτες i.
2: Στο γύρο $t \geq 0$ κάθε παίκτης i:
 3: Συναντά το γείτονα με δείκτη W_i^t, $\mathbf{P}\left[W_i^t = j\right] = p_{ij}$.
 4: Υφίσταται κόστος $(1-\alpha_i)(x_i(t) - x_{W_i^t}(t))^2 + a_i(x_i(t) - s_i)^2$ και μαθαίνει την άποψη $x_{W_i^t}(t)$.
 5: Ανανεώνει την άποψή του $x_i(t+1) = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t}(1-\alpha_i)(x - x_{W_i^\tau}(\tau))^2 + \alpha_i(x-s_i)^2$ (3)

επαναλαμβανόμενη λήψη μέσων όρων με τις τιμές που λαμβάνει. Η ιδέα είναι ότι μετά από μερικές επαναλήψεις, οι απόψεις των γειτόνων έχουν φτάσει αρκετά κοντά στο x^*, οπότε ο μέσος όρος των απόψεων σε κάθε γύρο είναι κοντά στο μέσο των γειτόνων εκείνη τη χρονική στιγμή, το οποίο κάνει το αρχικό FJ. Αυτό μας οδηγεί τελικά στο ακόλουθο αποτέλεσμα για τη σύγκλιση του αλγορίθμου στο σημείο ισορροπίας.

Theorem 2. *Έστω $I = (P, s, \alpha)$ ένα παίγνιο διαμόρφωσης απόψεων με σημείο ισορροπίας $x^* \in [0,1]^n$. Το διάνυσμα απόψεων $x(t) \in [0,1]^n$ που παράγεται από τον κανόνα ανανέωσης (3) μετά από t γύρους ικανοποιεί*

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq C\sqrt{\log n}\frac{(\log t)^{3/2}}{t^{\min(1/2,\rho)}},$$

όπου $\rho = \min_{i\in V} a_i$ και C είναι μια σταθερά.

Η απόδειξη του Θεωρήματος 2 μπορεί να χωριστεί σε δύο μέρη. Στο πρώτο μέρος προσπαθούμε να βρούμε μια αναδρομική σχέση για το σφάλμα $\|x(t) - x^*\|_\infty$. Φυσικά, αυτή η ποσότητα είναι μια τυχαία μεταβλητή, συνεπώς η αναδρομική σχέση θα ικανοποιείται με μεγάλη πιθανότητα. Η ακριβής διατύπωση δίνεται στο ακόλουθο θεώρημα.

Theorem 3. *Έστω $e(t)$ η λύση της ακόλουθης αναδρομικής σχέσης,*

$$e(t) = \delta(t) + (1-\rho)\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t}$$

όπου $e(0) = \|x(0) - x^*\|_\infty$ και

$$\delta(t) = \sqrt{\frac{\ln(\pi^2 n t^2/6p)}{t}}$$

. Τότε,

$$\mathbf{P}\left[\text{για όλα τα } t \geq 1,\ \|x(t) - x^*\|_\infty \leq e(t)\right] \geq 1 - p$$

Το δεύτερο μέρος της απόδειξης συνίσταται στο να φράξουμε την ακολουθία $e(t)$, βασισμένοι στην παραπάνω αναδρομική. Αυτό εμπεριέχει κυρίως τεχνική δουλειά. Ειδικότερα, θα αποδείξουμε το ακόλουθο.

Theorem 4. *Έστω $e(t)$ μια ακολουθία που ικανοποιεί την αναδρομική σχέση*

$$e(t) = \delta(t) + (1-\alpha)\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \text{ και } e(0) = \|x(0) - x^*\|_\infty,$$

όπου $\delta(t) = \sqrt{\frac{\ln(Dt^{2.5})}{t}}$, $\delta(0) = 0$, *και* $D > \mathrm{e}^{2.5}$ *είναι μια θετική σταθερά. Τότε*

$$e(t) \leq \sqrt{2.5\ln(D)}\frac{(\ln t)^{3/2}}{t^{\min(\rho,\,1/2)}}.$$

Το Θεώρημα 4 δίνει κατέυθείαν ένα άνω φράγμα για το χρόνο σύγκλισης της δυναμικής, λόγω του Θεωρήματος 3.

Στην επόμενη ενότητα, θα δούμε ότι εκτός από την απλότητά του, ο Αλγόριθμος 2 είναι μια «λογική επιλογή» για τους παίκτες ενός παιγνίου. Αυτό σημαίνει ότι είναι στο συμφέρον ενός παίκτη να χρησιμοποιήσει αυτό τον αλγόριθμο, διότι του εξασφαλίζει χαμηλό συνολικό κόστος. Με άλλα λόγια, όταν έχουν περάσει πολλοί γύροι, οι παίκτες δε θα έχουν *μετανιώσει* για τις επιλογές τους. Δηλαδή, το συνολικό κόστος ασυμφωνίας ενός παίκτη που ακολουθεί έναν τέτοιο κανόνα είναι κοντά στο συνολικό κόστος ασυμφωνίας που θα είχε υποστεί αν έπαιζε μια σταθερή στρατηγική. Αυτή η ιδιότητα πρέπει να ισχύει ανεξάρτητα από το πώς ανανεώνουν τις απόψεις τους οι άλλοι παίκτες και ποιους γείτονες επιλέγει να συναντήσει ένας παίκτης σε κάθε γύρο. Αυτή η ισχυρή ιδιότητα, που καλείται ιδιότητα *no regret*, καθιστά μια δυναμική διαδικασία φυσική επιλογή για την περιγραφή της συμπεριφοράς των παικτών και περιγράφεται στην επόμενη ενότητα.

Το μοντέλο OCO

Για να ορίσουμε τι σημαίνει να είναι ένας αλγόριθμος *no regret*, θα πρέπει πρώτα να περιγράψουμε το περιβάλλον μέσα στο οποίο ο αλγόριθμος λειτουργεί. Καθώς οι παίκτες λαμβάνουν σε κάθε γύρο σαν είσοδο μια γειτονική άποψη, φαίνεται ότι το πιο κατάλληλο μοντέλο για την περιγραφή της κατάστασης είναι αυτό της *Άμεσης Κυρτής Βελτιστοποίησης*(Online Convex Optimization ή OCO).

Στο μοντέλο OCO , ένας αλγόριθμος είναι ένας παίκτης που λαμβάνει επαναλαμβανόμενες αποφάσεις. Αφού πάρει μια απόφαση, ο παίκτης υφίσταται κόστος, ανάλογα με την απόφασή του. Η συνάρτηση κόστους δεν είναι γνωστή στον παίκτη και μπορεί ακόμη και να εξαρτάται από την απόφαση που παίρνει ο παίκτης. Αυτό είναι ένα *άμεσο* περιβάλλον, διότι ο παίκτης δεν ξέρει πώς θα εξελιχθούν οι συναρτήσεις κόστους. Φυσικά, για να έχει νόημα να μελετηθεί ένα τέτοιο πρόβλημα, πρέπει να επιβάλλουμε ορισμένους περιορισμούς στο μοντέλο. Κατάρχήν, οι συναρτήσεις κόστους πρέπει να είναι φραγμένες, αλλιώς σε κάθε γύρο θα μπορούσε το κόστος να αυξάνεται, με αποτέλεσμα η απόδοση του αλγορίθμου να μην είναι ποτέ καλή. Επίσης, το σύνολο αποφάσεων πρέπει να έχει κάποιου είδους δομή, όπως θα δούμε στη συνέχεια.

Τώρα θα ορίσουμε τυπικά όλες αυτές τις έννοιες. Το σύνολο αποφάσεων είναι ένα κυρτό υποσύνολο K του $\mathbf{R}^n$ και οι συναρτήσεις κόστους είναι κυρτές συναρτήσεις στο K. Έστω ότι $\mathcal{F}$ είναι μια οικογένεια από κυρτές φραγμένες συναρτήσεις που είναι διαθέσιμες στον αντίπαλο. Αυτές είναι οι πιθανές συναρτήσεις κόστους. Ο αλγόριθμος τρέχει για έναν αριθμό επαναλήψεων T. Στην επανάληψη t, ο παίκτης επιλέγει $x_t \in K$. Αφού ο παίκτης δεσμευτεί σε αυτή την επιλογή, η συνάρτηση $f_t \in \mathcal{F}$ αποκαλύπτεται και υφίσταται κόστος $f_t(x_t)$. Έχουμε τον ακόλουθο ορισμό.

Definition 4. *Ένας αλγόριθμος A για το πρόβλημα OCO μϵ σύνολο συναρτήσϵων $\mathcal{F}_{s,\alpha}$και σύνολο αποφάσϵων $\mathcal{K} = [0, 1]$ ϵίναι μια ακολουθία από αποφάσϵις $(A_t)_{t=0}^{\infty}$ όπου $A_t : \mathcal{F}^t \mapsto \mathcal{K}$.*

Ο χρόνος εκτέλεσης ενός OCO αλγορίθμου ορίζεται ως ο συνολικός χρόνος που χρειάζεται για να παραχθούν οι αποφάσεις για όλους τους γύρους. Αυτό συνήθως εξαρτάται από τη διάσταση του K, τον αριθμό γύρων T και τις παραμέτρους των συναρτήσεων κόστους. Τώρα θα προσπαθήσουμε να βρούμε ένα σωστό ορισμό για το τι είναι ένας καλός αλγόριθμος σε αυτό το μοντέλο. Αφού το μοντέλο OCO έχει την αφετηρία του στη θεωρία παιγνίων, η έννοια της αποδοτικότητας θα προκύψει από αυτό τον κλάδο. Ορίζουμε το *regret* ενός αλγορίθμου ως τη διαφορά του συνολικού κόστους σε όλους τους γύρους και τους κόστους της καλύτερης σταθερής στρατηγικής.

Τυπικά, έστω ένας αλγόριθμος $\mathcal{A}$ που δεδομένων των προηγούμενων αποφάσεων του παίκτη παίρνει μια απόφαση για τον τρέχοντα γύρο. Τότε, ορίζουμε το regretτου $\mathcal{A}$ μετά από T γύρους ως:

$$\text{regret}_T(\mathcal{A}) = \sup_{\{f_1,\dots,f_T\}\subseteq\mathcal{F}} \left\{ \sum_{t=1}^{T} f_t(x_t) - \min_{x\in K} \sum_{t=1}^{T} f_t(x) \right\} \tag{4}$$

Διαισθητικά, ένας αλγόριθμος θα είναι αποδοτικός αν το regretτου είναι υπογραμμικό ως προς τον αριθμό των γύρων, δηλαδή αν $\text{regret}_T(\mathcal{A}) = o(T)$. Αυτό σημαίνει ότι ο αλγόριθμος επιτυγχάνει κατά μέσο όρο εξίσου καλό κόστος όσο η καλύτερη δυνατή σταθερή στρατηγική, που θα ήταν το καλύτερο που θα μπορούσε να πετύχει ένας συμβατικός αλγόριθμος. Με άλλα λόγια, ακόμα και αν ο αντίπαλος επιλέγει τις συναρτήσεις κόστους κακοπροαίρετα, το συνολικό κόστος του παίκτης θα παραμένει κατά μία έννοια ελάχιστο. Αυτό καθιστά τους νο ρεγρετ αλγορίθμους συμφέρουσες επιλογές για τους παίκτες σε ένα παίγνιο.

Τώρα που ορίσαμε τι σημαίνει για έναν αλγόριθμο να είναι regret, θα προσπαθήσουμε να διατυπώσουμε το πρόβλημα της διαμόρφωσης απόψεων υπό περιορισμένη πληροφορία σαν πρόβλημα OCO. Αυτό το πρόβλημα θα μοιάζει με ένα «παιχνίδι», το οποίο θα παίζεται μεταξύ του παίκτη και ενός αντιπάλου. Στο γύρο $t \geq 0$:

1. *Ο παίκτης διαλέγει μια τιμή $x_t \in [0, 1]$.*

2. *Ο αντίπαλος παρατηρεί το x_t ανδ σϵλϵςτς α $b_t \in [0, 1]$*

3. *Ο παίκτης υφίσταται κόστος $f(x_t, b_t) = (1 - \alpha)(x_t - b_t)^2 + \alpha(x_t - s)^2$.*

όπου τα s, a είναι σταθερές. Ο σκοπός του παίκτη είναι να διαλέξει x_t βασισμένος στις προηγούμενες επιλογές $(b_0, \dots, b_{t-1})$ ώστε να ελαχιστοποιήσει το συνολικό κόστος του. Στην περίπτωσή μας, το σύνολο στρατηγικών του προβλήματος OCO είναι το $K = [0, 1]$ και το σύνολο των συναρτήσεων είναι $\mathcal{F}_{s,\alpha} = \{x \mapsto (1 - \alpha)(x - b)^2 + \alpha(x - s)^2 : b \in [0, 1]\}$. Ως αποτέλεσμα, κάθε επιλογή από σταθερές s, a οδηγεί σε ένα διαφορετικό πρόβλημα OCO . Αφού οι συναρτήσεις στο $\mathcal{F}_{s,\alpha}$ ορίζονται μοναδικά από τον αριθμό $b \in [0, 1]$, ο ορισμός παίρνει την ακόλουθη μορφή.

Definition 5. *Ένας αλγόριθμος A για το OCO πρόβλημα μϵ $\mathcal{F}_{s,\alpha}$ και $\mathcal{K} = [0, 1]$ ϵίναι μια ακολουθία συναρτήσϵων $(A_t)_{t=0}^{\infty}$ όπου $A_t : [0, 1]^t \mapsto [0, 1]$.*

Παρατηρούμε ότι ενώ στο περιβάλλον OCO ο αντίπαλος επιλέγει το b_t για να προκαλέσει υψηλό κόστος στον παίκτη, στο μοντέλο FJ το b_t είναι η άποψη του παίκτη που επιλέχθηκε τυχαία. Συνεπώς, αν ο αλγόριθμος εγγυάται no regretστον παίκτη, τότε ο παίκτης θα επιτύχει χαμηλό κόστος ακόμα και στην ακραία περίπτωση που οι απόψεις των γειτόνων τύχει να είναι οι ίδιες με αυτές που θα επέλεγε ένας αντίπαλος. Συνεπώς, είναι λογικό να υποθέσουμε ότι ένας παίκτης επιλέγει το σύμφωνα με το αλγόριθμο $\mathcal{A}_i$ για το πρόβλημα OCO με $\mathcal{F}_{s_i,\alpha_i}$. Υπάυτό το πρίσμα, αν μπορούσαμε να αποδείξουμε ότι ο Αλγόριθμος 2 είναι no regret, τότε θα ήταν φυσιολογική συμπεριφορά για τους παίκτες και άρα θα είχε νόημα σαν επέκταση του μοντέλου FJ. Πράγματι, θα αποδείξουμε το ακόλουθο Θεώρημα.

Theorem 5. *Έστω η συνάρτηση* $f : [0,1]^2 \mapsto [0,1]$ *με* $f(x,b) = (1-\alpha)(x-b)^2 + \alpha(x-s)^2$ *για κάποιες σταθερές* $s, \alpha \in [0,1]$. *Έστω* $(b_t)_{t=0}^{\infty}$ *μια τυχαία ακολουθία με* $b_t \in [0,1]$. *Ιφ* $x_t = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$ *τότε για όλα τα* t, $\sum_{\tau=0}^{t} f(x_\tau, b_\tau) \le \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau) + O(\log t)$.

Μέχρι στιγμής έχουμε δείξει ότι ο Αλγόριθμος 2 συγκλίνει με υπογραμμικό ρυθμό στην ισορροπία και ότι είναι λογική επιλογή για παίκτες που προσπαθούν να ελαχιστοποιήσουν το κόστος τους υπό συνθήκες αβεβαιότητας. Τι θα συνέβαινε όμως αν υπήρχε ένας γρηγορότερος αλγόριθμος που επιτυγχάνει αυτούς τους στόχους; Όπως προκύπτει, αυτό είναι αδύνατο να πραγματοποιηθεί.

Ένα κάτω φράγμα

Θα θέλαμε να δείξουμε ότι δεν υπάρχει πρωτόκολλο ή αλγόριθμος που να συγκλίνει με εκθετικό ρυθμό στην ισορροπία του παιγνίου και επιτυγχάνει no regret. Καθώς στο μοντέλο η μόνη είσοδος του αλγορίθμου είναι οι συναρτήσεις που παρέχονται από τον αντίπαλο, ένα τέτοιο πρωτόκολλο δε θα έπρεπε να δέχεται καμία άλλη είσοδο πέρα από την άποψη ενός γείτονα σε κάθε γύρο. Αυτό σημαίνει ότι καμία επιπλέον πληροφορία για τη δομή του δικτύου δε θα είναι γνωστή στον αλγόριθμο. Αυτό μας κινητοποιεί να ορίσουμε τέτοια πρωτόκολλα ως πρωτόκολλα που έχουν *άγνοια* για το γράφημα στο οποίο λειτουργούν.

Definition 6. *Ένα* Πρωτόκολλο Άγνοιας *A είναι μια ακολουθία από συναρτήσεις* $(A_t)_{t=0}^{\infty}$ *όπου* $A_t : [0,1]^{t+2} \mapsto [0,1]$.

Ένας πρωτόκολλο άγνοιας μαζί με ένα συγκεκριμένο παίγνιο διαμόρφωσης απόψεων μας δίνουν μια *δυναμική άγνοιας*.

Definition 7. *Έστω ένα πρωτόκολλο άγνοιας A. Για ένα δεδομένο παίγνιο διαμόρφωσης απόψεων* $I = (P, s, \alpha)$ *ο αλγόριθμος A παράγει μια δυναμική άγνοιας* $x_A(t)$ *που ορίζεται ως εξής:*

- *Αρχικά κάθε παίκτης i επιλέγει την άποψη* $x_i^A(0) = A_0(s_i, \alpha_i)$
- *Στο γύρο* $t \ge 1$, *κάθε παίκτης i διαλέγει την άποψη* $x_i^A(t) = A_t(x_{W_i^0}(0), \ldots, x_{W_i^{t-1}}(t-1), s_i, \alpha_i)$, *όπου* W_i^t *είναι ο γείτονας που ο i συναντά στο γύρο t.*

Από τον παραπάνω ορισμό παρατηρούμε ότι η άποψη του κάθε παίκτη εξαρτάται μόνο από τις προηγούμενες απόψεις που έχει δει και από τις σταθερές a_i, s_i. Θα θέλαμε να δείξουμε ότι δεν υπάρχει πρωτόκολλο άγνοιας με εκθετικό ρυθμό σύγκλισης για το πρόβλημά μας, που θα σήμαινε το ίδιο και για δυναμικές. Για να το πετύχουμε αυτό, χρησιμοποιούμε μια κλασική τεχνική που λέγεται *αναγωγή*. Δηλαδή, θα δείξουμε ότι αν ένα τέτοιο πρωτόκολλο υπήρχε, τότε θα μπορούσε να χρησιμοποιηθεί για να λύσει ένα πρόβλημα, για το οποίο ξέρουμε πως δεν υπάρχουν αρκετά «γρήγορες » λύσεις. Μένει να βρούμε ένα κατάλληλο τέτοιο πρόβλημα.

Η ιδέα έρχεται από τον κλάδο της Στατιστικής Εκτίμησης. Σε ένα πρόβλημα στατιστικής εκτίμησης, έχουμε πρόσβαση σε δείγματα που ακολουθούν μια συγκεκριμένη κατανομή και σκοπός μας είναι να υπολογίσουμε μια συγκεκριμένη ποσότητα που αφορά την κατανομή, για παράδειγμα τη μέση τιμή. Ειδικότερα, θα μας φανεί χρήσιμο το πρόβλημα της *Εκτίμησης Bernoulli*, όπου πρέπει να εκτιμήσουμε την τιμή της πιθανότητας p μια τυχαίας μεταβλητής βασισμένοι σε έναν αριθμό δειγμάτων που ακολουθούν αυτή την κατανομή.

Για να λύσουμε το πρόβλημα της εκτίμησης, κατασκευάζουμε μια συνάρτηση που ονομάζεται *εκτιμήτρια*. Αφού η τυχαία μεταβλητή παίρνει τιμές στο $\{0,1\}$, προκύπτει ο ακόλουθος ορισμός μιας εκτιμήτριας.

Definition 8. *Μια εκτιμήτρια* $\theta = (\theta_t)_{t=1}^{\infty}$ *είναι μια ακολουθία συναρτήσεων,* $\theta_t : \{0,1\}^t \mapsto [0,1]$.

Συνήθως, το t είναι ο αριθμός δειγμάτων της κατανομής που λαμβάνουμε. Τότε, η εκτιμήτρια δίνει σας έξοδο ένα πραγματικό αριθμό(στην περίπτωσή μας στο $[0,1]$), ο οποίος αναπαριστά την εκτιμώμενη ποσότητα της κατανομής, δηλαδή τη μέση τιμή. Για να μετρήσουμε την αποτελεσματικότητα μιας εκτιμήτριας, ορίζουμε το *ρίσκο* , που είναι ουσιαστικά το αναμενόμενο σφάλμα μιας εκτιμήτριας από την πραγματική τιμή που προσπαθούμε να υπολογίσουμε.

Definition 9. *Έστω P μια κατανομή Bernoulli με μέση τιμή p και P^t η αντίστοιχη t-διάστατη κατανομή γινόμενο. Το* ρίσκο *μιας εκτιμήτριας $\theta = (\theta_t)_{t=1}^{\infty}$ είναι $\mathbf{E}_{(X_1,\ldots,X_t)\sim P^t}\left[|\theta_t(X_1,\ldots,X_t) - p|\right]$, το οποίο θα συμβολίσουμε με $\mathbf{E}_p\left[|\theta_t(X_1,\ldots,X_t) - p|\right]$ ή $\mathbf{E}_p\left[|\theta_t - p|\right]$ για συντομία.*

Το risk ποσοτικοποιεί το σφάλμα της εκτιμώμενης τιμής $\hat{p} = \theta_t(Y_1,\ldots,Y_t)$ ως προς την πραγματική παράμετρο p καθώς αυξάνεται ο αριθμός δειγμάτων. Καθώς το p είναι άγνωστο, κάθε λογική εκτιμήτρια $\theta = (\theta_t)_{t=1}^{\infty}$ πρέπει να εγγυάται ότι $\lim_{t\to\infty} \mathbf{E}_p\left[|\theta_t - p|\right] = 0$ για όλα τα p. Όσο πιο γρήγορα συγκλίνει αυτή η ποσότητα στο 0, τόσο καλύτερη θεωρείται η εκτιμήτρια.

Τώρα θα προσπαθήσουμε να καταλάβουμε γιατί το πρόβλημα της εκτίμησης Bernoulli είναι τόσο σημαντικό για εμάς. Όπως θα δούμε, το πρόβλημα του υπολογισμού του σημείου ισορροπίας στο μοντέλο της περιορισμένης πληροφορίας είναι πολύ παρόμοιο με αυτό του υπολογισμού της πιθανότητας p μια τυχαίας μεταβλητής Bernoulli. Διαισθητικά, αυτό συμβαίνει διότι η τιμή της ισορροπίας συνδέεται στενά με τις πιθανότητες επιλογής p_{ij}. Συνεπώς, ο υπολογισμός του x^* οδηγεί στον υπολογισμό αυτών των πιθανοτήτων, το οποίο μοιάζει με το πρόβλημα της εκτίμησης Bernoulli.

Πιο συγκεκριμένα, έστω ότι έχουμε μια τοπολογία με ένα κόμβο στο κέντρο και άλλους δύο κόμβους συνδεδεμένους με αυτόν. Έστω ότι p είναι η πιθανότητα ο κεντρικός κόμβος να επικοινωνήσει με τον αριστερό. Ουσιαστικά, η επιλογή γείτονα καθορίζεται από μια τυχαία μεταβλητή Bernoulli , η οποία παίρνει τιμή 1 με πιθανότητα p. Με απλούς υπολογισμούς προκύπτει ότι το σημείο ισορροπίας x^* αυτού του δικτύου είναι γραμμική συνάρτηση του p, που σημαίνει ότι αν μπορούμε να υπολογίσουμε γρήγορα το x^*, θα μπορούμε να βρούμε γρήγορα και το p. Το ακόλουθο Θεώρημα περιγράφει ακριβώς το αποτέλεσμα.

Theorem 6. *Έστω A ένας κανόνας ανανέωσης άγνοιας τέτοιος ώστε για όλα τα παίγνια $I = (P,s,\alpha)$,*

$$\lim_{t\to\infty} t^{1+c}\mathbf{E}\left[\|x_A(t) - x^*\|_\infty\right] = 0.$$

Τότε υπάρχει μια εκτιμήτρια $\theta_A = (\theta_t^A)_{t=1}^{\infty}$ έτσι ώστε για κάθε $p \in [0,1]$, $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p\left[|\theta_t^A - p|\right] = 0$.

Άρα, το μόνο πρόβλημα είναι να δείξουμε ότι δεν υπάρχει εκτιμήτρια Bernoulli που να επιτυγχάνει «γρήγορη» σύγκλιση για όλα τα $p \in [0,1]$. Θα περίμενε κανείς ότι αυτό το πρόβλημα θα είχε λυθεί προ πολλού από την κοινότητα της Στατιστικής. Ωστόσο, για αυτό το ερώτημα, δεν υπάρχει ικανοποιητική απάντηση. Στη συνέχεια της εργασίας θα εξηγήσουμε τα πιο σημαντικά εργαλεία της Στατιστικής για την απόδειξη κάτω φραγμάτων και γιατί δεν ταιριάζουν στην περίπτωσή μας. Ωστόσο, με ένα απλό και κομψό επιχείρημα μπορούμε να αποδείξουμε το ακόλουθο Θεώρημα.

Theorem 7. *Έστω $\theta = (\theta_t)_{t=1}^{\infty}$ μια εκτιμήτρια Bernoulli με ρυθμό σφάλματος $\mathbf{E}_p\left[|\theta_t - p|\right]$. Τότε, για κάθε $c > 0$, το σύνολο όλων των $p \in [0,1]$ που είναι τέτοια ώστε $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p\left[|\theta_t - p|\right] = 0$ έχει μέτρο Lebesgue 0.*

Το Θεώρημα μας λέει ουσιαστικά ότι «σχεδόν για όλα» τα $p \in [0,1]$, η εκτιμήτρια δεν επιτυγχάνει γρήγορη σύγκλιση. Αυτό είναι αρκετό για να αποδείξουμε ότι καμία δυναμική διαδικασία με άγνοια του γραφήματος δεν συγκλίνει γρηγορότερα από $O\left(1/t^{1+c}\right)$.

Περιεχόμενα

Chapter 1

Introduction

This Chapter serves as an overview to the various results presented in this thesis. It begins with a discussion of the way people interact in a social network. A popular model for this behavior is the Friedkin-Johnsen model, where opinion formation takes the form of a repeated algorithm that run by the people in the network. We then proceed to discuss various limitations of the FJ-model in describing real social networks. In particular, we argue that a person cannot possibly ask all his friends about their opinions every time he changes his. We propose a limited information variant of the model that renders it more realistic and present some convergence results for it. Finally, in the last Section we argue that there is no faster algorithm that could solve the problem of opinion formation.

1.1 Opinion Dynamics and the FJ-model

The study of *Opinion Formation* has a long history (see e.g. [Jac08]). Opinion Formation is a *dynamic process* in the sense that socially connected people (e.g. family, friends, colleagues) exchange information and this leads to changes in their expressed opinions over time. Today, the advent of the internet and social media makes the study of opinion formation in large social networks even more important; realistic models of how people form their opinions by interacting with each other are of great practical interest for prediction, advertisement etc. In an attempt to formalize the process of opinion formation, several models have been proposed over the years. The people in the network are called *agents*.

The common assumption underlying all these models, which dates back to DeGroot [DeG74], is that opinions evolve through a form of repeated averaging of information collected from the agents' social neighborhoods. In DeGroot's model, the opinion of each agent is represented as a real number is $[0, 1]$. The agents are able to communicate with their friends, which results in a change in their opinion. In each round, an agent changes his opinion by taking a weighted average of the opinions of their friends. Formally, the weight they put on each other's opinions is represented by a trust matrix T, where T_{ij} is the weight that agent i puts in j's opinion. Of course, T is stochastic, meaning that the sum of the elements in each row equals 1. If $z(t)$ is the vector of opinions during period t, then in period $t+1$ the new opinion vector formed after all the interactions take place is $z(t+1) = Tz(t)$. Despite the fact that this model doesn't accurately describe real social interactions, it is simple and natural enough to provide us with an intuition about the various factors that influence opinion formation.

A quite different was provided in a model proposed by Hegselmann and Krause(HK) in [HK02]. In this model, the opinions of the agents are again real numbers in $[0, 1]$. Additionally, each agent i has a *confidence factor* ϵ_i. In each round, agent i updates his opinion as the average of the opinions of the other agents that are within ϵ_i distance from his opinion. This means that the social connections of an agent might change as time passes. This is clearly a more realistic representation of social interactions, as people tend to be influence the most by others with similar opinions.

Another interesting model was proposed in [DNAW00]. In this model, in each round two agents are selected uniformly at random from the population. These two agents readjust their opinions provided that their current difference in opinion is not greater than a threshold d. The rationale for the threshold condition is that agents only interact when their opinion are already close enough; otherwise they do not even bother to discuss. The reason for such behaviour might be for instance lack of understanding, conflicts of interest or social pressure.

The body of work following DeGroot has developed a set of general conditions under which such processes will converge to a state of consensus, in which all nodes hold the same opinion. This emphasis on consensus, however, restricts the focus of the modeling activity to a specific type of opinion dynamics, where the opinions of the group all come together. As the sociologist David Krackhardt has observed

> We should not ignore the fact that in the real world consensus is usually not reached. Recognizing this, most traditional social network scientists do not focus on an equilibrium of consensus. They are instead more likely to be concerned with explaining the lack of consensus (the variance) in beliefs and attitudes that appears in actual social influence contexts.[Kra09]

In the model we are interested in, which was proposed by Friedkin and Johnsen in[FJ90], consensus is not generally reached. In the *FJ* model, we have an unweighted undirected graph $G(V, E)$ representing a social network. The network has n nodes. Each node in the network is called an *agent*. An edge between two agents exists if and ony if they know each other in some way. Each agent has an opinion, which is a real number in $[0, 1]$. The model specifies a certain procedure, by which each agent changes his opinion. This procedure runs in rounds. Each agent $i \in V$ has a value $s_i \in [0, 1]$, which represents his internal opinion and remains constant over time. Furthermore, in round t each agent i outputs a number $x_i(t) \in [0, 1]$, which is his opinion at the current round. The procedure runs as follows. Initially , $x_i(0)$ can have an arbitrary value in $[0, 1]$. At round t, each agent i updates his opinion $x_i(t)$ using the following rule

$$x_i(t) = \frac{\sum_{j \in N_i} w_{ij} x_j(t-1) + w_{ii} s_i}{\sum_{j \in N_i} w_{ij} + w_{ii}}, \tag{1.1}$$

where N_i is the set of neighbors of agent i and $x(t)$ the vector of the opinions of the agents at round t. Also, the weight w_{ij} associated with the edge $(i, j) \in E$ measures the extent of the influence that j poses on i and the weight $w_{ii} > 0$ quantifies how susceptible i is in adopting opinions that differ from her internal opinion s_i.

Hence, each agent takes a weighted sum of the mean of his neighbors values at the previous round and of his internal opinion and forms the new opinion. This kind of update rule is called a *dynamic*, a simple procedure that doesn't change in each round and is executed in the same way by all the agents. Examples of other dynamic procedures are given in [DeG74, FJ90, HK02, DNAW00].

Note that because of the presence of s_i as a constant in each iteration, repeated averaging will not in general bring all nodes to the same opinion. In this way, the model distinguishes between an individual's intrinsic belief s_i and her overall opinion x_i ; the latter represents a compromise between the persistent value of s_i and the expressed opinions of others to whom i is connected. This distinction between s_i and x_i also has parallels in empirical work that seeks to trace deeply held opinions such as political orientations back to differences in education and background, and even to explore genetic bases for such patterns of variation [AFH$^+$05].

The FJ model is one of most influential models for opinion formation. It has a very simple update rule, making it plausible for modeling natural behavior and its basic assumptions are aligned with empirical findings on the way opinions are formed [AFH$^+$05, Kra09]. It has also been studied under a game theoretic viewpoint. Bindel et al. posed the question of how one would quantify the cost of the lack of consensus, in case consensus is not reached. To answer the question, they considered a player's update rule as the minimizer of a quadratic disagreement cost function and based on it they defined the following opinion formation game [BKO11]. Each node i is a selfish agent whose strategy is the public opinion x_i that she expresses. For each agent, there is a cost function, which takes as arguments all the opinions in the network. Intuitively, this function models the disagreement cost between an agent and her neighbors. For agent i, this cost function is

$$C_i(x_i, x_{-i}) = \sum_{j \in N_i} w_{ij}(x_i - x_j)^2 + w_{ii}(x_i - s_i)^2 \tag{1.2}$$

where x_{-i} denotes the vector of the opinions of all the other agents except i. We notice that C_i is a strongly convex function on x, so it has a unique minimum point which is also a zero of the derivative of the function. After some straightforward calculations, we get that

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

which is the same update rule that was discribed by the FJ model 1.1. Thus, the best response dynamics in game 1.2 is the update rule of the FJ model. This means that the FJ model belongs to the class of algorithms called *best response dynamics*, where each player chooses as his strategy at round t the one that minimizes his cost, evaluated using the strategies of other players at round $t-1$. In our game, the best response dynamics is the Algorithm 3 , where $x_{-i}(t-1)$

Algorithm 3 Best response dynamics

Initially, agent i has opinion $x_i(0)$.

1: In each round t
2: Agent i updates as follows

$$x_i(t) = \operatorname*{argmin}_{x \in [0,1]} C_i\left(x, x_{-i}(t-1)\right)$$

obviously denotes the vector of opinions of all the other agents except i at time $t-1$. This framework is much more comprehensive since different aspects of the opinion formation process can be easily captured by defining suitable games. Subsequent works considered variants of the above game and studied the convergence properties of various dynamics that the players could follow. In [BGM13] they study convergence of dynamics in coevolutionary opinion formation games, where nodes form their opinions by maximizing agreements with friends weighted by the strength of the relationships, which in turn depend on difference in opinion with the respective friends. In [BFM16] they study opinion formation games with dynamic social influences and provide algorithms for computing all Pure Nash Equilibria of them. In [EFHS17] they define a similar game to ours, where players are influenced not only by neighboring opinions, but also by "global trends" of the opinions in the network. They prove the existence of a unique Nash equilibrium in this setting and study best response dynamics. In order to understand why players would want to follow a best response dynamic, we first have to define a very important concept in game theory, the *Nash Equilibrium* of a game.

Definition 10 ([N$^+$50])**.** *Suppose $x \in \mathbf{R}^n$ is the vector of strategies of all the players. Then, x is a* Nash Equilibrium *if for every player i and for every strategy y of player i, it holds*

$$C_i(y, x_{-i}) \geq C_i(x_i, x_{-i})$$

Essentially, a Nash Equilibrium is a set of strategies such that no player can reduce his cost by deviating from his strategy, if all the other players maintain their strategy. Thus, if the game ever reaches this state x, it is unlikely that any of the players will want to change their strategy, hence the word *equilibrium.* Therefore, the computation of all the possible Nash equilibria(there can be more than one) is arguably a useful task in order to understand the behavior of the players. The Nash equilibrium is a fundamental concept in the study of games and has found numerous applications in Economic Theory, which are beyond the scope of this thesis. Its wide use and applicability stems from the fact that it is considered a realistic modelling of the limiting behavior of the agents in a game. In this thesis, we will be interested in how fast the agents can compute the Nash equilibrium of the game we just defined.

Now, the concept of dynamics has an additional meaning: it is an update rule that the agents follow in order to compute the Nash Equilibrium of the game.

As a result, the process of opinion formation of the agents can also be viewed as a search for the equilibrium point of the game. An interesting question is to study the convergence of such dynamics. For example, in [YOA+13] they study the convergence in distribution of opinions in a binary opinion dynamics setting where some agents are "stubborn", meaning that they don't change their opinion but only influence others. Also, in [FGV16] they study the convergence properties of a "noisy" best response dynamics, where the uncertainty is due to the fact that people are not fully rational. In Section 2.1 we are going to prove that game 1.2 has a unique Nash equilibrium x^*. The proof is follows naturally from the definition.

Theorem 8. *The game where players have costs given by 1.2 has a unique Nash Equilibrium* $x^* \in [0,1]^n$ *, which satisfies*

$$x_i^* = \frac{\sum_{j \in N_i} w_{ij} x_j^* + w_{ii} s_i}{\sum_{j \in N_i} w_{ij} + w_{ii}} \tag{1.3}$$

for each agent $i \in V$*.*

We are also going to prove that if the agents run Algorithm 3 , their opinions converge to x^*. This is also proven in [GS14], where they bound the convergence time of the FJ model in various graph topologies. The proof consists of showing that the distance of $x(t)$ from the equilibrium x^* shrinks by at least $1 - a$ in each round.

Theorem 9. *If* x^* *is the unique convergence point of the game and* $a = \min \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}}$*, then:*

$$\|x(t) - x^*\|_\infty \leq (1-a)^t \|x(0) - x^*\|_\infty$$

A natural question that arises now is whether the dynamics of the FJ model is a realistic representation of how agents form their opinion in today's world. It is clear that in order to implement rule 1.1, each agent should learn all of his neighbors' opinions in each round. But, in real social networks, a person could have hundreds or thousands of friends and only be influence by a few close ones in order to form an opinion. This prompts us to modify the model in order to be consistent with real behavior. This will be done in the next section.

1.2 Limited Information

Many recent works study the Nash equilibrium x^* of the opinion formation game defined in [BKO11] under various perspectives. For example, in [GTT13], they consider the sum of all the opinions to be a sign of the positive opinion of the network on a particular product. For a fixed T, they consider the problem of finding T nodes, such that if these nodes have fixed opinion 1, the resulting Nash equilibrium maximizes the sum of opinions. Also, in [AKPT18] they again want to maximize the sum of opinions in the network. However, this time we are given a constant k and should find k nodes and parameters $[l_i, r_i]$ where the a_i will vary in these nodes. Finally ,in [MMT17] they define the Disagreement and Polirization quantities as a function of the Equilibrium x^* of a given network and try to minimize their sum over all graph topologies that have the same total sum of edge weights.

The reason for this scientific interest is evident: the equilibrium x^* is considered as an appropriate way to model the final opinions formed in a social network, since the *well established* FJ model converges to it.

Our work is motivated by the fact that there are notable cases in which the FJ model is not an appropriate model for the dynamic of the opinions, due to the large amount of information exchange that it implies. More precisely, at each round its update rule (1.1) requires that every agent learns all the opinions of her social neighbors. In today's large social networks where users usually have several hundreds of friends it is highly unlikely that, each day, they learn the opinions of all their social neighbors. In such environments it is far more reasonable to assume that individuals randomly meet a small subset of their acquaintances and these are the only opinions that they learn. Such information exchange constraints render the FJ model unsuitable for modeling the opinion formation process in such large networks and therefore, it is not clear whether x^* captures the limiting behavior of the opinions. Similar work is done in [FPS16], where they provide convergence results for limited information variants of the Heglesmann-Krause model and the FJ model. Although the considered limited information variant of the FJ model is very similar to ours, their convergence results are much weaker, since they concern the expected value of the opinion vector. Also, in [HCM17] they examine the equilibrium convergence properties of no-regret learning with exponential weights in potential games in a framework where players have access to a noisy estimate of their payoff vectors. Moreover, in [MS17] they study dynamics' long-run behavior when players have either noiseless or noisy information on their payoff gradients. In this work we ask:

Question. *Is the equilibrium x^* an adequate way to model the final formed opinions in large social networks? Namely, are there simple variants of the FJ model that require limited information exchange and converge fast to x^*? Can they be justified as natural behavior for selfish agents under a game-theoretic solution concept?*

To address these questions, one could define precise dynamical processes whose update rules require limited information exchange between the agents and study their convergence properties. Instead of doing so, we describe the opinion formation process in such large networks as *dynamics* of a suitable opinion formation game that captures these information exchange constraints. This way we can study general classes of *dynamics* (e.g. no regret dynamics) without explicitly defining their update rule. The opinion formation game that we consider is a variant of the game in [BKO11] based on interpreting the weight w_{ij} as a measure of how frequently i meets j. Notice how C_i is no longer a deterministic function of the previous opinions.

Definition 11. *For a given opinion vector $x \in [0,1]^n$, the disagreement cost of agent i is the random variable $C_i(x_i, x_{-i})$ defined as follows:*

- *Agent i meets one of her neighbors j with probability $p_{ij} = w_{ij}/\sum_{j\in N_i} w_{ij}$.*
- *Agent i suffers cost $C_i(x_i, x_{-i}) = (1-a_i)(x_i-x_j)^2 + a_i(x_i-s_i)^2$, where $\alpha_i = w_{ii}/(\sum_{j\in N_i} w_{ij} + w_{ii})$.*

Definition 12. *We denote an instance of the opinion formation game of Definition 11 as $I = (P, s, \alpha)$, where P is a $n \times n$ matrix with non-negative elements p_{ij}, with $p_{ii} = 0$ and $\sum_{j=1}^{n} p_{ij}$ is either 0 or 1, $s \in [0,1]^n$ is the internal opinion vector, $\alpha \in (0,1]^n$ the self confidence vector.*

If i, j are not neighbors, then the corresponding entry in the matrix P is 0. Note that the expected disagreement cost of each agent in the above game is the same as the disagreement cost in [BKO11] (scaled by $\sum_{j\in N_i} w_{ij} + w_{ii}$). Hence, its Nash equilibrium with respect to the expected disagreement cost is again x^*.

Definition 13. *For a given instance $I = (P, s, \alpha)$ the equilibrium $x^* \in [0,1]^n$ is the unique solution of the following linear system, for every $i \in V$, $x_i^* = (1 - \alpha_i) \sum_{j \in N_i} p_{ij} x_j^* + a_i s_i$.*

We denote as W_i^t the neighbor that agent i met at round t, which is a random variable whose probability distribution is determined by the instance $I = (P, s, \alpha)$ of the game, $\mathbf{P}\left[W_i^t = j\right] = p_{ij}$. Another parameter of an instance I that we often use is $\rho = \min_{i \in V} \alpha_i$.

This game provides us with a general template of all the *dynamics* examined in this paper. At round t, each agent i selects an opinion $x_i(t)$ and suffers a disagreement cost based on the opinion of the neighbor that she randomly met. At the end of round t, she is informed only about the opinion and the index of this neighbor and may use this information to update her opinion in the next round. Obviously different update rules lead to different *dynamics*, however all of these respect the information exchange constraints: at every round each agent learns the opinion of *just one* of her neighbors. Question 1.2 now takes the following more concrete form.

Question. *Can the agents update their opinions according to the limited information that they receive such that the produced opinion vector $x(t)$ converges to the equilibrium x^*? How is the convergence rate affected by the limited information exchange? Are there dynamics that ensure that the cost that the agents experience is minimal?*

To answer Question 1.2, we examine an algorithm in which all agents update their opinions according to the *Follow the Leader* principle. Since each agent i must select $x_i(t)$, before knowing which of her neighbors she will meet and what opinion her neighbor will express, this update rule says «*play the best according to what you have observed*». For a given instance (P, s, a) of the game the Follow the Leader dynamics $x(t)$ is defined in Dynamics 4.

Algorithm 4 Follow the Leader dynamics

1: Initially $x_i(0) = s_i$ for all agents i.
2: At round $t \geq 0$ each agent i:
  3: Meets neighbor with index W_i^t, $\mathbf{P}\left[W_i^t = j\right] = p_{ij}$.
  4: Suffers cost $(1 - \alpha_i)(x_i(t) - x_{W_i^t}(t))^2 + a_i(x_i(t) - s_i)^2$ and learns the opinion $x_{W_i^t}(t)$.
  5: Updates her opinion $x_i(t+1) = \mathrm{argmin}_{x \in [0,1]} \sum_{\tau=0}^{t} (1 - \alpha_i)(x - x_{W_i^\tau}(\tau))^2 + \alpha_i (x - s_i)^2$
(1.4)

As we can see, the algorithm essentially performs a repeated averaging over the values it receives. The idea is that after some iterations, the opinions of the neighbors will have converged close enough to x^*, so the averaging of the opinions in each round is close to the average of the neighbors, which is what the original FJ model did. Of course, this is just an intuition, which we will make precise in Chapter 2. More specifically, in continuous optimization, one usually measures the performance of an algorithm by establishing some upper bound on the quantity $\|x(t) - x^*\|$, where $x(t)$ is the output of the algorithm and x^* is the point we would like to compute. Obviously, this bound will depend on the number of rounds t. If the bound is of the form ρ^t, where $\rho \in (0, 1)$, then the convergence rate is said to be *linear*. Linear convergence is usually the best rate one could hope for in an optimization algorithm. On the other hand, a bound of the form $1/t^c$ for some $c > 0$ is called *sublinear*. The main theorem of this section states that the dynamic procedure 4 converges to equilibrium with sublinear rate. Since we are dealing with uncertainty, the convergence metric used in Theorem 10 is $\mathbf{E}\left[\|x(t) - x^*\|_\infty\right]$ where the expectation is taken over the random meeting of the agents.

Theorem 10. *Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 11 with equilibrium $x^* \in [0,1]^n$. The opinion vector $x(t) \in [0,1]^n$ produced by update rule (1.4) after t rounds satisfies*

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq C \sqrt{\log n} \frac{(\log t)^{3/2}}{t^{\min(1/2, \rho)}},$$

where $\rho = \min_{i \in V} a_i$ and C is a universal constant.

The proof of Theorem 10 can be divided into two parts. The first part consists of finding a suitable recurrence relation for $\|x(t) - x^*\|_\infty$. Of course, this quantity is a random variable, hence the recurrence will be satisfied with high probability. The precise formulation is given in the following theorem.

Theorem 11. *Let $e(t)$ the solution of the following recursion,*

$$e(t) = \delta(t) + (1-\rho)\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t}$$

where $e(0) = \|x(0) - x^*\|_\infty$ and

$$\delta(t) = \sqrt{\frac{\ln(\pi^2 n t^2 / 6p)}{t}}$$

. Then,

$$\mathbf{P}\left[\text{for all } t \geq 1,\ \|x(t) - x^*\|_\infty \leq e(t)\right] \geq 1 - p$$

The second part of the proof consists of bounding the sequence $e(t)$, based on this recurrence, which is mainly a technical task. Specifically, we are going to prove the following theorem.

Theorem 12. *Let $e(t)$ be a function satisfying the recursion*

$$e(t) = \delta(t) + (1-\alpha)\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \quad \text{and } e(0) = \|x(0) - x^*\|_\infty,$$

where $\delta(t) = \sqrt{\frac{\ln(Dt^{2.5})}{t}}$, $\delta(0) = 0$, and $D > \mathrm{e}^{2.5}$ is a positive constant. Then

$$e(t) \leq \sqrt{2.5 \ln(D) \frac{(\ln t)^{3/2}}{t^{\min(\rho,\, 1/2)}}}.$$

Theorem 12 automatically gives an upper bound on the convergence time of the dynamic, because of Theorem 11. As a final step, we translate this result to one involving the expectation of the error, so Theorem 10 follows.

In the next section we argue that, apart from its simplicity, update rule 4 is also a "natural" choice of algorithm for the players to follow. This means that the selfish players actually have an incentive to use Algorithm 4 because it ensures low total cost for them. In other words, after many rounds have passed, the agents will have not *regretted* their choices so far. Namely, the total disagreement cost of an agent that follows such a rule is close to the total disagreement cost that she would experience by selecting the best fixed opinion in hindsight. The latter must hold regardless of the way the other agents update their opinions and of the neighbors that the agent gets to meet. This powerful property, called the *no regret property*, renders no regret dynamics natural dynamics for describing the behavior of agents and is described in the next Section.

1.3 The OCO framework

In order to define what it means for an algorithm to be no-regret, we first have to describe the setting in which this algorithm works. Since the agent receives in each round as input a neighboring opinion, it seems that the most appropriate framework to model this situation is that of *Online Convex Optimization*(OCO).

In the OCO model, an algorithm is identified as a player who iteratively makes decisions. After making a decision, the player suffers a loss, depending on his decision. The loss is unknown to the decision maker beforehand and can even be chosen by an adversary depending on the decision just made. This is an *online* setting, because the player doens't know how the loss functions are going to go. Some restrictions are necessary in order for this model to be meaningful. Firstly, the loss functions in each round should be bounded, otherwise an adversary could always increase the loss so that the performance of the algorithm is never good. Also, the decision set of the player should be "structured" in some way.

We now formally define all these concepts. The presentation is similar to the one given in [Haz16].The decision set is a convex subset K of $\mathbf{R}^n$ and the loss functions are convex functions over K. Suppose that $\mathcal{F}$ is a family of convex bounded functions available to the adversary. These are the potential cost functions. The algorithm runs for a number of iterations. At iteration t, the player chooses $x_t \in K$. After the player commits to this choice, a function $f_t \in \mathcal{F}$ is revealed and he receives cost $f_t(x_t)$. Let T be the total number of iterations of the algorithm. The following definition clarifies the concepts.

Definition 14. *An algorithm A for the OCO problem with function space $\mathcal{F}_{s,\alpha}$ and decision space $\mathcal{K} = [0,1]$ is a sequence of functions $(A_t)_{t=0}^{\infty}$ where $A_t : \mathcal{F}^t \mapsto \mathcal{K}$.*

The running time of an OCO algorithm is defined as the total time to produce the decisions for all rounds. This time typically depends of the dimensionality of K, on the number of rounds T and on the parameters of the cost functions. We now try to find a suitable definition for a "good" OCO algorithm. Since the model is game-theoretic in nature, our notion of efficiency comes from game theory. We define the *regret* of an algorithm as the difference between the total cost suffered in all rounds and the cost of the best fixed decision in hindsight.

Formally, let $\mathcal{A}$ be an algorithm that given a certain decision history of the player makes a decision for the current round. Then, we define the regret of $\mathcal{A}$ after T rounds as:

$$\text{regret}_T(\mathcal{A}) = \sup_{\{f_1,\ldots,f_T\} \subseteq \mathcal{F}} \left\{ \sum_{t=1}^{T} f_t(x_t) - \min_{x \in K} \sum_{t=1}^{T} f_t(x) \right\} \tag{1.5}$$

Intuitively, an algorithm performs well if its regret is sublinear in the number of rounds T, i.e. $\text{regret}_T(\mathcal{A}) = o(T)$. This means that the algorithm performs on average as well as the best fixed decision in hindsight, which would be the performance of an ordinary algorithm. In other words, even if the other agents selected their opinions maliciously, the player's total experienced cost would still be in a sense minimal. This makes a no-regret algorithm a sound and logical choice for selfish agents in a game.

This is highlighted in many recent works. In [CBL03], the authors offer a unified way to prove that a variety of game theoretic algorithms have low regret, which means that they consider the regret to be the crucial property that such algorithms should satisfy. Moreover, in [EMN09] they consider a specific sub-class of socially conave games and prove that any no-regret algorithm converges to the Nash Equilibrium of such games. Furthermore, in [KPT11] they study the performance of protocols for load balancing in distributed computing. In their analysis, they assume that all players follow no regret strategies. Finally, in [SALS15] they state that no-regret algorithms are a strong match for playing games because their regret bounds hold even in adversarial environments. These are just a few examples of the growing interest of the community about no-regret algorithms as a model for natural behaviour.

Perhaps one might ponder about how easy it is for an algorithm to satisfy the no regret property, if it is at all possible. *Online Gradient Descent* is an influential no regret algorithm proposed by Zinkevic in [Zin03] for the general OCO problem, where the adversary can select any convex function with bounded gradient. It is essentially a form of gradient descent, adapted in order to function in the online setting. Thus, in most OCO problems at least one no regret algorithm exists.

We now present some examples of problems where the OCO framework could be a useful model.

Prediction from expert advice. One of the most important problems in prediction theory is the so-called *Expert problem.* The decision maker has to choose advice among one of n experts. This procedure is executed for several rounds, and each time the cost of choosing the opinion of an expert is arbitrary. The goal of the decision maker is to make decisions that have similar cost to the best expert in hindsight.

This problem has an elegant formulation as an OCO problem. One can think of a decision in each round as a probability distribution on the n experts. Therefore, the set of decisions K is the n-dimensional simplex, that is

$$K = \left\{ x \in \mathbf{R}^n : \sum_{i=1}^{n} x_i = 1, x_i \geq 0 \text{ for all i} \right\}$$

which is a convex set. Suppose that g_t is the cost vector at round t, which means that $g_t(i)$ is the cost of expert i in round t. Then, the cost function of round t is the expected cost of choosing an expert according to the distribution of x, so $f_t(x) = g_t^\intercal x$, which is linear and thus convex.

Online spam filtering. A spam filtering system takes emails as input and should decide whether each one is spam or valid. If d is the number of words in the dictionary, then we represent an email by a vector $x \in \mathbf{R}^d$, where x has value 1 in the entries that correspond to words that appear in the email and 0 to the rest. To decide whether an email is spam or valid, we learn a filter $x \in \mathbf{R}^d$. Specifically, an email $a \in \mathbf{R}^d$ is classified by the filter $x \in \mathbf{R}^d$ as $\hat{y} = \operatorname{sgn}\langle x, a\rangle$(1 is for spam, -1 is for valid).

In the OCO model, the set of decisions is the set of all filters in $\mathbf{R}^d$ that have bounded Euclidean norm, i.e. that lie in some ball of fixed radius. The cost functions are determined according to a stream of emails coming to the system. At round t the player chooses a filter x_t. Then a pair (y_t, a_t) arrives. The cost incurred is then $l(y_t, a_t)$, where $\hat{y_t}$ is the classification of a_t according to the filter x_t and l is a convex loss function.

Now that we have defined what it means for an algorithm to be *no regret*, we shall try to formulate our problem of limited information exchange in the language of the OCO framework. Based on the cost 11 that the agents experience in the new setting, we consider an appropriate *Online Convex Optimization* problem. This problem can be viewed as a «game» played between an adversary and a player. *At round* $t \geq 0$,

1. *the player selects a value* $x_t \in [0, 1]$.
2. *the adversary observes the* x_t *and selects a* $b_t \in [0, 1]$
3. *the player receives cost* $f(x_t, b_t) = (1-\alpha)(x_t - b_t)^2 + \alpha(x_t - s)^2$.

where s, α are constants in $[0, 1]$. The goal of the player is to pick x_t based on the history $(b_0, \ldots, b_{t-1})$ in a way that minimizes her total cost. In our case the feasibility set of the OCO problem is $\mathcal{K} = [0, 1]$ and the set of functions is $\mathcal{F}_{s,\alpha} = \{x \mapsto (1-\alpha)(x-b)^2 + \alpha(x-s)^2 : b \in [0, 1]\}$. As a result, each selection of the constants s, α leads to a different OCO problem. Since the functions in $\mathcal{F}_{s,\alpha}$ are uniquely determined by the number $b \in [0, 1]$, Definition 15 takes the following form

Definition 15. *An algorithm A for the OCO problem with $\mathcal{F}_{s,\alpha}$ and $\mathcal{K} = [0, 1]$ is a sequence of functions $(A_t)_{t=0}^{\infty}$ where $A_t : [0, 1]^t \mapsto [0, 1]$.*

Notice that while in the OCO setting the adversary chooses b_t in order to incur high cost to the player, in the FJ model b_t is the opinion of the player that was selected randomly. Thus, if the algorithm ensures no regret to the player, then the player achieves small cost even in the extreme case that the neighbors' opinions happen to be the same as the "adversarially" picked ones. Hence, it is reasonable to assume that each agent i selects $x_i(t)$ according to no regret algorithm A_i for the OCO problem with $\mathcal{F}_{s_i,\alpha_i}$ Under this perspective, if we could prove that Algorithm 4 is a no regret algorithm, then it follows that it is a rational choice for selfish agents and thus a natural limited information variant of the FJ model.

Other works concern the convergence properties of dynamics based on no regret learning algorithms. In [FV97] they use the term *calibrated* learning rule, which is essentially a no regret rule. They prove that in two player games such rules lead to coarse correlated equilibrium in the limit. In [FS99] they present a simple algorithm that uses the multiplicative-weight methods of Littlestone and Warmuth[LW94] and prove that it is no-regret. Also, in [SA00] they propose a new and simple adaptive procedure for playing a game: *regret-matching*. In this procedure, players may depart from their current play with probabilities that are proportional to measures of regret for not having used other strategies in the past. It is shown that this adaptive procedure guarantees that, with probability one, the empirical distributions of play converge to the set of correlated equilibria of the game. In addition, in [CHM17] they show that, in n-person finite generic games that admit unique Nash equilibrium, the strategy vector converges locally and fast to it. They also provide conditions for global convergence. Our results fit in this line of research since we show that for a game with infinite strategy space, the strategy vector (and not the time-averaged) converges to the Nash equilibrium x^*.

We now present the key steps for proving the no regret property for our algorithm. The full proof is given in Section 2.2. This can be derived by the more general results in [HAK07]. Specifically, Hazan et al. present algorithms that that achieve regret $O(\ln T)$ for an arbitrary sequence of strictly convex functions with bounded first and second derivatives. However, we give a short and simple proof that may be of interest. We first prove that a similar strategy that also takes into account the value b_t admits no regret (Lemma 1). Obviously, knowing the value b_t before selecting x_t is in direct contrast with the OCO framework, however proving the no regret property for this algorithm easily extends to establishing the no regret property of *Follow the Leader*.

Lemma 1. *Let $(b_t)_{t=0}^{\infty}$ be an arbitrary sequence with $b_t \in [0,1]$. Let $y_t = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau)$ then for all t,*

$$\sum_{\tau=0}^{t} f(y_\tau, b_\tau) \leq \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau).$$

Now we can understand why *Follow the Leader* admits no regret. Since the cost incurred by the sequence y_t is at most that of the best fixed value, we can compare the cost incurred by x_t with that of y_t. Since the functions in $\mathcal{F}_{s,\alpha}$ are quadratic, the extra term $f(x, b_t)$ that y_t takes into account doesn't change dramatically the minimum of the total sum. Namely, x_t, y_t are relatively close.

Lemma 2. *For all $t \geq 0$, $f(x_t, b_t) \leq f(y_t, b_t) + 2\frac{1-\alpha}{t+1} + \frac{(1-\alpha)^2}{(t+1)^2}$.*

These lemmas are enough to prove that our algorithm is no-regret.

Theorem 13. *Consider the function $f : [0,1]^2 \mapsto [0,1]$ with $f(x,b) = (1-\alpha)(x-b)^2 + \alpha(x-s)^2$ for some constants $s, \alpha \in [0,1]$. Let $(b_t)_{t=0}^{\infty}$ be an arbitrary sequence with $b_t \in [0,1]$. If $x_t = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t-1} f(x, b_\tau)$ then for all t, $\sum_{\tau=0}^{t} f(x_\tau, b_\tau) \leq \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau) + O(\log t)$.*

So far, we have established that Algorithm 4 converges with sublinear rate to the equilibrium and that it is a natural choice for selfish players. What if there was a faster algorithm that achieved all these things? As it turns out, this is impossible to accomplish. The next section discusses this issue.

1.4 A lower bound

We would like to show that there exists no protocol that converges exponentially fast to the equilibrium of the game and also achieves no regret. Since in the OCO framework the only input of an algorithm are the functions provided by the opponent, such a protocol would have no other input apart from the neighboring opinions in each round. This means that no additional information about the structure of the graph is known to the algorithm. This motivates us to define these protocols that are *oblivious* to the structure in which they operate.

Definition 16. *A graph oblivious update rule A is a sequence of functions $(A_t)_{t=0}^{\infty}$ where $A_t : [0,1]^{t+2} \mapsto [0,1]$.*

A graph-oblivious update rule together with an instance of the opinion formation game define an graph-oblivious *dynamic*.

Definition 17. *Let a graph oblivious update rule A. For a given instance $I = (P, s, \alpha)$ the rule A produces a graph oblivious dynamics $x_A(t)$ defined as follows:*

- *Initially each agent i selects her opinion $x_i^A(0) = A_0(s_i, \alpha_i)$*
- *At round $t \geq 1$, each agent i selects her opinion $x_i^A(t) = A_t(x_{W_i^0}(0), \dots, x_{W_i^{t-1}}(t-1), s_i, \alpha_i)$, where W_i^t is the neighbors that i meets at round t.*

As can be seen from definition 17, the opinion of each agent only depends on the previous opinions he has received and on internal constants a_i, s_i.

We would like to show that no graph oblivious dynamics with linear convergence rate exist for our problem, which would also imply the same for no regret dynamics. To do this, we use a standard techinique called *reduction*. Thus, we will show that if such a protocol existed, then it would also solve a problem, for which we know that no efficient solution exists. It remains to find such a suitable problem.

The idea comes from Statistical Estimation. The field of Statistics and, in particular, *Inferential Statistics*, deals with sample data drawn from a population. The goal is to use patterns in the sample data to draw inferences about the population represented, accounting for randomness. A very important subfield of Inferential Statistics is *Estimation*, where the inferences take the form of estimating numerical characteristics of the data, such as the mean or the variance of the distribution. An interesting example that is relevant to our problem of computing the equilibrium is that of *Bernoulli Estimation*, where we have to estimate the value of the probability p of a Bernoulli random variable based on a number of samples that follow the same distribution.

In order to solve the problem of estimation, we construct a function called an *estimator*. Since the random variable takes values in $\{0, 1\}$, we get the following definition of an estimator.

Definition 18. *An estimator $\theta = (\theta_t)_{t=1}^{\infty}$ is a sequence of functions, $\theta_t : \{0,1\}^t \mapsto [0,1]$.*

Usually, t is the number of samples taken according to the distribution. Then, the estimator outputs a real value(in our case in $[0, 1]$), which represents an estimated quantity of the distribution, i.e. its expectation. To measure the efficiency of an estimator, we define the *risk*, which corresponds to the expected error of an estimator from the real value of the quantity to be computed.

Definition 19. *Let P be a Bernoulli distribution with mean p and P^t be the corresponding t-fold product distribution. The* risk *of an estimator $\theta = (\theta_t)_{t=1}^{\infty}$ is $\mathbf{E}_{(X_1,\dots,X_t)\sim P^t}[|\theta_t(X_1,\dots,X_t) - p|]$, which we will denote by $\mathbf{E}_p[|\theta_t(X_1,\dots,X_t) - p|]$ or $\mathbf{E}_p[|\theta_t - p|]$ for brevity.*

The risk $\mathbf{E}_p [|\theta_t - p|]$ quantifies the error rate of the estimated value $\hat{p} = \theta_t(Y_1, \dots, Y_t)$ to the real parameter p as the number of samples t grows. Since p is unknown, any meaningful estimator $\theta = (\theta_t)_{t=1}^{\infty}$ must guarantee that $\lim_{t\to\infty} \mathbf{E}_p [|\theta_t - p|] = 0$ for all p. The faster the quantity $\mathbf{E}_p [|\theta_t - p|]$ converges to 0, the better the estimator is considered to be. For example, *sample mean* has error rate $\mathbf{E}_p [|\theta_t - p|] \le \frac{1}{2\sqrt{t}}$ and this rate is typical for a lot of estimation problems in Statistics. Our goal will be to prove that the best rate one can achieve for the Bernoulli Estimation problem is $O\left(1/\sqrt{t}\right)$.

Question. *But why is the Bernoulli Estimation problem so important to us?*

As it turns out, the problem of estimating the equilibrium point in this limited information exchange model is very similar to that of computing the probability of a Bernoulli random variable. Intuitively, this is because the value of the equilibrium is closely related with the selection probabilities p_{ij}. Thus, computing x^* amounts to computing these probabilities, which is very similar to Bernoulli estimation.

More specifically, suppose that we have a graph topology with a node in the center and two other nodes connected to it. Suppose p is the probability that the central node talks with the left node. This is similar to having a Bernoulli random variable that takes value 1 with probability p, which means it is 1 when the center picks the left node and 0 if it picks the right. By straightforward calculations, the equilibrium point x^* of this network is a linear function of p, which means that if we could find x^* fast, we could also find p fast. The following Theorem summarises the reduction and is proven in Chapter 4.

Theorem 14. *Let A a graph oblivious update rule such that for all instances $I = (P, s, \alpha)$,*

$$\lim_{t\to\infty} t^{1+c}\mathbf{E}\left[\|x_A(t) - x^*\|_\infty\right] = 0.$$

Then there exists an estimator $\theta_A = (\theta_t^A)_{t=1}^{\infty}$ such that for all $p \in [0,1]$, $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p\left[|\theta_t^A - p|\right] = 0$.

Hence, it now remains to prove that there is no fast Bernoulli estimator for all $p \in [0,1]$. It appears that this problem would long have been resolved by the Statistics community. However, for this particular question, no definitive answer seems to have been given. In Chapter 4 we review important tools in Statistics for proving lower bounds on the risk that an estimator can achieve with t samples. We then explain why these approaches don't work for our problem. In Chapter 5 we prove by a simple argument a slightly stronger claim: for every Bernoulli estimator, for "almost all" $p \in [0,1]$ the estimator does not achieve fast estimation.

Theorem 15. *Let $\theta = (\theta_t)_{t=1}^{\infty}$ be a Bernoulli estimator with error rate $\mathbf{E}_p [|\theta_t - p|]$. For any $c > 0$, the set of all $p \in [0,1]$ such that $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p [|\theta_t - p|] = 0$ has Lebesgue measure 0.*

This implies that no graph oblivious dynamic procedure converges faster that $O\left(1/t^{1+c}\right)$.

To illustrate the necessity of the graph oblivious property in order for the lower bound to hold, in Chapter 3 we present two algorithms that achieve linear convergence rate by utilizing some additional information about the graph or the probability matrix P. The *row dependent dynamics* assumes that each player i knows the probabilities p_{ij} of selecting each of his neighbors. The protocol then uses them to simulate the original FJ model, by using "outdated" versions of the neighbors opinions. This strategy of course is not a natural behavior for the players, which is why it does not have the no regret property, as we show with a simple example. In addition, we present the FREQS algorithm, which works under the assumption that all neighbors are picked uniformly at random. The protocol then dictates that all players stop updating their opinions for some time and instead learn the values of all their neighbors enough times to reconstruct the true frequencies with good accuracy after a rounding step. The rounding is a crucial part in the proof of convergence, since it ensures that with nonzero probability the weighted average is known *exactly*.

Chapter 2

Convergence in the FJ model

In this Chapter, we provide proofs for some convergence Theorems mentioned in Chapter 1 about the FJ model. We begin by proving the uniqueness of Nash Equilibrium in the game with payoffs 1.2. and the convergence result for the simple FJ model. We then proceed to provide full proof of the no regret property of 4 discussed in Section 1.3, which renders it natural for players to adopt. Finally, we are going to prove that the dynamic procedure 4 converges to the Nash equilibrium of Game 11. Our general strategy will be to express the error from equilibrium in the form of a recurrence relation and then solve it. We are also interested in the rate of convergence of this procedure. As we will see, the convergence rate is not as fast as that of the original FJ model of Section 1.1. Chapters 4 and 5 offer some explanations for this behavior.

2.1 Convergence of the classical FJ model

We remind the reader that in the classical FJ model, the costs of the game are defined as follows:

$$C_i(x_i, x_{-i}) = \sum_{j \in N_i} w_{ij}(x_i - x_j)^2 + w_{ii}(x_i - s_i)^2 \tag{2.1}$$

We will first prove that this game has a unique Nash Equilibrium, which is denoted by x^*. The high level approach is the following: we will prove that every x^* that is a Nash Equilibrium must be the solution of a linear system of the form $Ax^* = b$. Then, using a lemma from linear algebra, we prove that the matrix A is invertible and thus there is a unique solution to the system and consequently a unique equilibrium.

We start by proving a sufficient condition for invertibility of a matrix, which is a special case of Gershgorin's Theorem[Var10].

Lemma 3. *Let $A \in \mathbf{R}^{n \times n}$ be a matrix such that*

$$a_{ii} = 1$$

,for all $i \in \{1, \dots, n\}$ and

$$\sum_{\substack{j=1 \\ j \neq i}}^{n} |a_{ij}| < 1$$

for all $i \in \{1, \dots, n\}$. Then A is invertible.

Proof. Suppose that A is not invertible. Then, A has an eigenvalue equal to 0, which means there exists a nonzero $z \in \mathbf{R}^n$ such that

$$Az = 0$$

We choose i such that $|z_i| = \|z\|_\infty$. By expanding the above equality at the i-th line, we obtain

$$z_i = -\sum_{\substack{j=1 \\ j\neq i}}^{n} a_{ij} z_j$$

Taking absolute values we obtain

$$|z_i| = \left|\sum_{\substack{j=1 \\ j\neq i}}^{n} a_{ij} z_j\right| \leq \sum_{\substack{j=1 \\ j\neq i}}^{n} |a_{ij} z_j| \leq \sum_{\substack{j=1 \\ j\neq i}}^{n} |a_{ij}| \|z\|_\infty = \|z\|_\infty \sum_{\substack{j=1 \\ j\neq i}}^{n} |a_{ij}| < \|z\|_\infty$$

which is a contradiction due to the definition of z_i. In the derivation we have used the hypothesis about the sum of elements in each row. Hence, A doesn't have a zero eigenvalue and thus is invertible. □

Now we prove the existence and uniqueness of Nash Equilibrium in our setting.

Theorem 8. *The game where players have costs given by 1.2 has a unique Nash Equilibrium $x^* \in [0,1]^n$, which satisfies*

$$x_i^* = \frac{\sum_{j\in N_i} w_{ij} x_j^* + w_{ii} s_i}{\sum_{j\in N_i} w_{ij} + w_{ii}} \tag{1.3}$$

for each agent $i \in V$.

Proof. We are first going to prove that a Nash Equilibrium of the game should necassarily satisfy equation 1.3. If x^* is a Nash Equilibrium, then by definition 10

$$C_i(x_i^*, x_{-i}^*) \leq C_i(y, x_{-i}^*)$$

for every player i and $y \in [0,1]$. This means that x_i^* is a minimum point of the function $f_i(x) = C(x, x_{-i}^*)$. Since the function f_i is continuous in $[0,1]$, for each i we get

$$f_i'(x_i^*) = 0$$

which is equation 1.3.
We are now going to prove that there is a unique vector $x^* \in \mathbf{R}^n$ that satisfies equation 1.3 for each i. We define the $n \times n$ matrix $A = \{a_{ij}\}$ as follows

$$a_{ij} = \begin{cases} 1 & \text{,if i =j} \\ 0 & \text{,if i and j aren't connected} \\ -\frac{w_{ij}}{\sum_{j\in N_i} w_{ij} + w_{ii}} & \text{if i and j are connected} \end{cases}$$

We also define vector $b \in \mathbf{R}^n$ where $b_i = \frac{w_{ii} s_i}{\sum_{j\in N_i} w_{ij} + w_{ii}}$. Clearly, equation 1.3 can be rewritten as

$$Ax^* = b \tag{2.2}$$

It suffices to show that equation 2.2 has a unique vector solution x^*. This is equivalent to showing that matrix A is invertible. It is easy to observe that for every row i

$$\sum_{\substack{j=1 \\ j\neq i}}^{n} |a_{ij}| = \sum_{j\in N_i} \frac{w_{ij}}{\sum_{j\in N_i} w_{ij} + w_{ii}} = \frac{\sum_{j\in N_i} w_{ij}}{\sum_{j\in N_i} w_{ij} + w_{ii}} < 1$$

,since $w_{ii} > 0$ for all i. Thus, Lemma 3 implies that A is invertible and the linear system has a unique solution. □

We now prove that the best response dynamics converge to the Nash equilibrium x^*. In the proof, we take advantage of Equation 2.2, which is satisfied by the x^*, in order to prove a simple recurrence property.

Theorem 9. *If x^* is the unique convergence point of the game and $a = \min \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}}$, then:*

$$\|x(t) - x^*\|_\infty \leq (1-a)^t \|x(0) - x^*\|_\infty$$

Proof. Let i be an agent and $t \geq 1$ a round. Then by Equations 1.1 and 1.3 we get

$$\begin{aligned}
|x_i(t) - x_i^*| &= \left| \frac{\sum_{j \in N_i} w_{ij} x_j(t-1)}{\sum_{j \in N_i} w_{ij} + w_{ii}} + \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}} s_i - \left(\frac{\sum_{j \in N_i} w_{ij} x_j^*}{\sum_{j \in N_i} w_{ij} + w_{ii}} + \frac{w_{ii}}{\sum_{j \in N_i} w_{ij} + w_{ii}} s_i \right) \right| \\
&= \left| \frac{\sum_{j \in N_i} w_{ij} (x_j(t-1) - x_j^*)}{\sum_{j \in N_i} w_{ij} + w_{ii}} \right| \\
&\leq \frac{\sum_{j \in N_i} w_{ij} \left| x_j(t-1) - x_j^* \right|}{\sum_{j \in N_i} w_{ij} + w_{ii}} \\
&\leq \frac{\sum_{j \in N_i} w_{ij}}{\sum_{j \in N_i} w_{ij} + w_{ii}} \|x(t-1) - x^*\|_\infty \\
&\leq (1-a) \|x(t-1) - x^*\|_\infty
\end{aligned}$$

Thus

$$\|x(t) - x^*\|_\infty \leq (1-a) \|x(t-1) - x^*\|_\infty$$

Continuing in the same way we get

$$\|x(t) - x^*\|_\infty \leq (1-a) \|x(t-1) - x^*\|_\infty \leq \ldots \leq (1-a)^t \|x(0) - x^*\|_\infty$$

□

This proof highlights the importance of the interal opinions of the agents in this model. If the agents didn't have internal opinions, then $a_i = 0$ and we wouldn't have convergence to equilibrium.

Thus far, we have shown that the FJ model is a rational dynamics procedure for the agents to adopt and leads to exponential convergence to the equilibrium x^* of the game. There is yet another interpretation of this result. That is, it can be viewed as a gradient descent algorithm performed simultaneously on all agents. Specifically, each coordinate of the vector of gradient descent corresponds to the opinion of one of the agents. Hence, at each time step, each agent updates the corresponding coordinate using the rule of gradient descent and the previous opinions. The cost function that this rule tries to minimize is a quadratic form. The details of how to write the FJ model as a gradient descent step are given in [Ep14].

2.2 Follow the Leader is no-regret

In this Section we will provide the full proof that Algorithm 4 is no regret. We first prove that a similar strategy that also takes into account the value b_t admits no regret (Lemma 1). Obviously, knowing the value b_t before selecting x_t is in direct contrast with the OCO framework, however proving the no regret property for this algorithm easily extends to establishing the no regret property of *Follow the Leader*.

Lemma 1. *Let* $(b_t)_{t=0}^{\infty}$ *be an arbitrary sequence with* $b_t \in [0,1]$. *Let* $y_t = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t} f(x,b_\tau)$ *then for all* t,

$$\sum_{\tau=0}^{t} f(y_\tau, b_\tau) \le \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau).$$

Proof. By definition of y_t, $\sum_{\tau=0}^{t} f(y_t, b_\tau) = \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau)$, so

$$\begin{aligned}
\sum_{\tau=0}^{t} f(y_\tau, b_\tau) - \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau) &= \sum_{\tau=0}^{t} f(y_\tau, b_\tau) - \sum_{\tau=0}^{t} f(y_t, b_\tau) \\
&= \sum_{\tau=0}^{t-1} f(y_\tau, b_\tau) - \sum_{\tau=0}^{t-1} f(y_t, b_\tau) \\
&\le \sum_{\tau=0}^{t-1} f(y_\tau, b_\tau) - \sum_{\tau=0}^{t-1} f(y_{t-1}, b_\tau)
\end{aligned}$$

The last inequality follows by the fact that $y_{t-1} = \operatorname{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t-1} f(x,b_\tau)$. Inductively, we prove that $\sum_{\tau=0}^{t} f(y_\tau, b_\tau) \le \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, b_\tau)$. □

Now we can understand why *Follow the Leader* admits no regret. Since the cost incurred by the sequence y_t is at most that of the best fixed value, we can compare the cost incurred by x_t with that of y_t. Since the functions in $\mathcal{F}_{s,\alpha}$ are quadratic, the extra term $f(x, b_t)$ that y_t takes into account doesn't change dramatically the minimum of the total sum. Namely, x_t, y_t are relatively close.

Lemma 2. *For all* $t \ge 0$, $f(x_t, b_t) \le f(y_t, b_t) + 2\frac{1-\alpha}{t+1} + \frac{(1-\alpha)^2}{(t+1)^2}$.

Proof. We first prove that for all t,

$$|x_t - y_t| \le \frac{1-\alpha}{t+1}. \tag{2.3}$$

By definition $x_t = \alpha s + (1-\alpha)\frac{\sum_{\tau=0}^{t-1} b_\tau}{t}$ and $y_t = \alpha s + (1-\alpha)\frac{\sum_{\tau=0}^{t} b_\tau}{t+1}$.

$$\begin{aligned}
|x_t - y_t| &= (1-\alpha)\left|\frac{\sum_{\tau=0}^{t-1} b_\tau}{t} - \frac{\sum_{\tau=0}^{t} b_\tau}{t+1}\right| \\
&= (1-\alpha)\left|\frac{\sum_{\tau=0}^{t-1} b_\tau - t b_t}{t(t+1)}\right| \\
&\le \frac{1-\alpha}{t+1}
\end{aligned}$$

The last inequality follows from the fact that $b_\tau \in [0,1]$. We now use inequality (2.3) to bound the difference $f(x_t, b_t) - f(y_t, b_t)$.

$$\begin{aligned}
f(x_t, b_t) &= \alpha(x_t - s)^2 + (1-\alpha)(x_t - y_t)^2 \\
&\le \alpha(y_t - s)^2 + 2\alpha\,|y_t - s|\,|x_t - y_t| + \alpha\,|x_t - y_t|^2 \\
&\quad + (1-\alpha)(y_t - y_t)^2 + 2(1-\alpha)\,|y_t - y_t|\,|x_t - y_t| + (1-\alpha)\,|x_t - y_t|^2 \\
&\le f(y_t, b_t) + 2\,|x_t - y_t| + |y_t - x_t|^2 \\
&\le f(y_t, b_t) + 2\frac{1-\alpha}{t+1} + \frac{(1-\alpha)^2}{(t+1)^2}
\end{aligned}$$

□

Now we are ready to prove that the algorithm is no regret.

Theorem 13. *Consider the function* $f : [0,1]^2 \mapsto [0,1]$ *with* $f(x,b) = (1-\alpha)(x-b)^2 + \alpha(x-s)^2$ *for some constants* $s, \alpha \in [0,1]$. *Let* $(b_t)_{t=0}^{\infty}$ *be an arbitrary sequence with* $b_t \in [0,1]$. *If* $x_t = \text{argmin}_{x\in[0,1]} \sum_{\tau=0}^{t-1} f(x,b_\tau)$ *then for all* t, $\sum_{\tau=0}^{t} f(x_\tau, b_\tau) \leq \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x,b_\tau) + O(\log t)$.

Proof. Theorem 13 easily follows by Lemma 1

$$\begin{aligned}\sum_{\tau=0}^{t} f(x_\tau, b_\tau) &\leq \sum_{\tau=0}^{t} f(y_\tau, b_\tau) + \sum_{\tau=0}^{T} 2\frac{1-\alpha}{\tau+1} + \sum_{\tau=0}^{t} \frac{(1-\alpha)^2}{(\tau+1)^2} \\ &\leq \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, y_\tau) + 2(1-\alpha)(\log t + 1) + (1-\alpha)\frac{\pi^2}{6} \\ &\leq \min_{x\in[0,1]} \sum_{\tau=0}^{t} f(x, y_\tau) + O(\log t)\end{aligned}$$

□

Hence, the *Follow the Leader* dynamics is a natural behavior for the agents in this limited information setting. In the next sections, we will show that it also converges to x^* quite fast. This reveals that the equilibrium x^* is a robust choice for modeling the limiting behavior of the opinions of agents since, even in our limited information setting, there exist simple and natural dynamics that converge to it.

2.3 Convergence of Follow the Leader

2.3.1 Results

We will first state formally what we would like to prove.

Theorem 10. *Let* $I = (P, s, \alpha)$ *be an instance of the opinion formation game of Definition 11 with equilibrium* $x^* \in [0,1]^n$. *The opinion vector* $x(t) \in [0,1]^n$ *produced by update rule (1.4) after* t *rounds satisfies*

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq C\sqrt{\log n}\frac{(\log t)^{3/2}}{t^{\min(1/2,\rho)}},$$

where $\rho = \min_{i\in V} a_i$ *and* C *is a universal constant.*

The proof of Theorem 10 can be divided into two parts. The first part consists of finding a suitable recurrence relation for $\|x(t) - x^*\|_\infty$. Of course, this quantity is a random variable, hence the recurrence will be satisfied with high probability. The idea is the following. Suppose W_i^t is the random choice that agent i makes at time t. Then, as we have shown, the update rule for agent i can be written in the form:

$$x_i(t) = (1-\alpha_i)\frac{\sum_{\tau=1}^{t} x_{W_i^\tau(\tau-1)}}{t} + \alpha_i s_i$$

We will rewrite the above rule using indicator functions. The quantity $\mathbf{1}[W_i^\tau = j]$ is equal to 1 if agent i selects neighbor j at time τ and 0 otherwise. We now deduce the following:

$$\begin{aligned}x_i(t) &= (1-\alpha_i)\frac{\sum_{\tau=1}^{t} x_{W_i^\tau}(\tau-1)}{t} + \alpha_i s_i \\ &= (1-\alpha_i)\sum_{\tau=1}^{t}\frac{\sum_{j\neq i}\mathbf{1}[W_i^\tau = j]x_j(\tau-1)}{t} + \alpha_i s_i \\ &= (1-\alpha_i)\sum_{j\neq i}\frac{\sum_{\tau=1}^{t}\mathbf{1}[W_i^\tau = j]x_j(\tau-1)}{t} + \alpha_i s_i\end{aligned}$$

For a fixed $j \neq i$, the term $\sum_{\tau=1}^{t} \mathbf{1}[W_i^\tau = j]/t$ is the empirical frequency with which i selected j until time t. We now proceed to bound the difference $|x_i(t) - x_i^*|$. Using Theorem 8 we obtain:

$$\begin{aligned}
|x_i(t) - x_i^*| &= \left| (1-\alpha_i) \sum_{j \neq i} \frac{\sum_{\tau=1}^{t} \mathbf{1}[W_i^\tau = j] x_j(\tau - 1)}{t} + \alpha_i s_i - (1-\alpha_i) \sum_{j \neq i} p_{ij} x_j^* - \alpha_i s_i \right| \\
&= \left| (1-\alpha_i) \sum_{j \neq i} \left(\frac{\sum_{\tau=1}^{t} \mathbf{1}[W_i^\tau = j] x_j(\tau - 1)}{t} - p_{ij} x_j^* \right) \right| \\
&\leq |(1-\alpha_i) \sum_{j \neq i} \left| \frac{\sum_{\tau=1}^{t} \mathbf{1}[W_i^\tau = j] x_j(\tau - 1)}{t} - p_{ij} x_j^* \right|
\end{aligned}$$

Ideally, the empirical frequencies of selecting the neighbors would be the probabilities p_{ij}. In other words, ideally we would have

$$\frac{\sum_{\tau=1}^{t} \mathbf{1}[W_i^\tau = j]}{t} = p_{ij}$$

and the error $e(t)$ would satisfy a recurrence of the form

$$e(t) = (1-\rho) \frac{\sum_{i=1}^{t} e(t-1)}{t} \tag{2.4}$$

which would give a rate $O\left(\frac{1}{t^\rho}\right)$. But this situation is not likely to happen, because there is always some difference between the empirical frequencies and the actual probabilities. This difference can easily be quantified by the Hoeffding inequality(see Appendix A), which shows that with high probability the empirical frequencies converge to the actual probabilities relatively fast. This will result in a modified recurrence relation. The precise formulation is given in the following theorem.

Theorem 11. *Let $e(t)$ the solution of the following recursion,*

$$e(t) = \delta(t) + (1-\rho) \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t}$$

where $e(0) = \|x(0) - x^*\|_\infty$ and

$$\delta(t) = \sqrt{\frac{\ln(\pi^2 n t^2 / 6p)}{t}}$$

. Then,

$$\mathbf{P}\left[\text{for all } t \geq 1,\ \|x(t) - x^*\|_\infty \leq e(t)\right] \geq 1 - p$$

The second part of the proof consists of bounding the sequence $e(t)$, base on this recurrence, which is mainly a technical task. Specifically, we are going to prove the following theorem.

Theorem 12. *Let $e(t)$ be a function satisfying the recursion*

$$e(t) = \delta(t) + (1-\alpha) \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \text{ and } e(0) = \|x(0) - x^*\|_\infty,$$

where $\delta(t) = \sqrt{\frac{\ln(Dt^{2.5})}{t}}$, $\delta(0) = 0$, and $D > \mathrm{e}^{2.5}$ is a positive constant. Then

$$e(t) \leq \sqrt{2.5 \ln(D)} \frac{(\ln t)^{3/2}}{t^{\min(\rho,\, 1/2)}}.$$

Lemma 12 automatically gives an upper bound on the convergence time of the dynamic, because of Theorem 11. As a final step, we translate this result to one involving the expectation of the error, so Theorem 10 follows. In the remainder of this chapter, we focus on the proofs of these two Theorems.

2.3.2 Formulation of Recurrence

We first prove Theorem 11. We remind that W_i^τ denotes the agent j that i met at round τ and that this happens with probability p_{ij} and x^* is the unique equilibrium point of the instance $I = (P, s, \alpha)$. At first we prove that with probability at least $1 - p$, for all $t \geq 1$ and all agents i:

$$\left| \frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau}}{t} - \sum_{j \neq i} p_{ij} x_j^* \right| \leq \delta(t) \tag{2.5}$$

where $\delta(t) = \sqrt{\frac{\log(\pi^2 n t^2/(6p))}{t}}$. We notice that the quantity $\sum_{\tau=1}^{t} x^*_{W_i^\tau}$ depends on the random choices of the neighbors. Thus, this result means that the true frequencies of the neighbors are not very far from the probabilities of selection. The reason we chose this particular expression for $\delta(t)$ is to fit with the Hoeffding inequality and the union bound that we are going to use in the following.

Since W_i^τ are independent random variables with $\mathbf{P}[W_i^\tau = j] = p_{ij}$ and $\mathbf{E}\left[x^*_{W_i^\tau}\right] = \sum_{j \neq i} p_{ij} x_j^*$. We now apply Hoeffding's inequality 14 to the independent variables W_i^τ and get

$$\mathbf{P}\left[\left| \frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau}}{t} - \sum_{j \neq i} p_{ij} x_j^* \right| > \delta(t)\right] < 6p/(\pi^2 n t^2).$$

To bound the probability of error for all rounds $t = 1$ to ∞ and all agents i, we apply the union bound

$$\sum_{t=1}^{\infty} \mathbf{P}\left[\max_i \left| \frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau}}{t} - \sum_{j \neq i} p_{ij} x_j^* \right| > \delta(t)\right] \leq \sum_{t=1}^{\infty} \frac{6}{\pi^2} \frac{1}{t^2} \sum_{i=1}^{n} \frac{p}{n} = p$$

As a result with probability $1 - p$ we have that for all $t \geq 1$ and all agents i,

$$\left| \frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau}}{t} - \sum_{j \neq i} p_{ij} x_j^* \right| \leq \delta(t) \tag{2.6}$$

We have proven that with the outcome will be close to the ideal, where the frequencies of the neighbors are exactly the p_{ij}. Thus, we can expect the recurrence relation that will arise to be $\delta(t)$ "close" to the one we formulated in 2.4. As it turns out, this is exactly the case. Assume that $\|x^\tau - x^*\|_\infty \leq e(\tau)$ for all $\tau \leq t - 1$. Then

$$\begin{aligned}
x_i(t) &= (1 - \alpha_i) \frac{\sum_{\tau=1}^{t} x_{W_i^\tau}(\tau - 1)}{t} + \alpha_i s_i \\
&\leq (1 - \alpha_i) \frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau} + \sum_{\tau=1}^{t} e(\tau - 1)}{t} + \alpha_i s_i \qquad (2.7) \\
&\leq (1 - \alpha_i) \left(\frac{\sum_{\tau=1}^{t} x^*_{W_i^\tau}}{t} + \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right) + \alpha_i s_i \\
&\leq (1 - \alpha_i) \left(\sum_{j \in N_i} p_{ij} x_j^* + \delta(t) + \frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right) + \alpha_i s_i \qquad (2.8) \\
&\leq x_i^* + \delta(t) + (1 - \alpha) \left(\frac{\sum_{\tau=0}^{t-1} e(\tau)}{t} \right)
\end{aligned}$$

We get (2.7) from the induction step and (2.8) from inequality (2.6). Similarly, we can prove that $x_i(t) \geq x_i^* - \delta(t) - (1 - \alpha) \frac{\sum_{\tau=1}^{t} e(\tau)}{t}$. As a result $\|x(t) - x^*\|_\infty \leq e(t)$ and the induction is complete. Therefore, we have that with probability at least $1 - p$, $\|x(t) - x^*\|_\infty \leq e(t)$ for all $t \geq 1$. Now we just have to solve this recurrence relation of $e(t)$, which is done in the next Section.

2.3.3 Solution of Recurrence

We now present the proof of Lemma 12, which provide a bound for the solution of the recursion of the previous Section. We first try to express the term $e(t+1)$ as a function of just the previous value $e(t)$. To do this, observe that for all $t \geq 0$ the function $e(t)$ satisfies the following recursive relation

$$e(t+1) = e(t)\left(1 - \frac{\rho}{t+1}\right) + \delta(t+1) - \delta(t) + \frac{\delta(t)}{t+1} \tag{2.9}$$

For $t = 0$ we have that

$$e(1) = (1-\rho)e(0) + \delta(1) = (1-\rho)e(0) + \sqrt{\ln D} \tag{2.10}$$

Observe that for $D > \mathrm{e}^{2.5}$, $\delta(t)$ is decreasing for all $t \geq 1$. Therefore, $\delta(t+1) - \delta(t) + \frac{\delta(t)}{t+1} \leq \frac{\delta(t)}{t+1}$ and from equations (2.9) and (2.10) we get that for all $t \geq 0$

$$e(t+1) \leq e(t)\left(1 - \frac{\rho}{t+1}\right) + \frac{\sqrt{\ln(D(t+1)^2)}}{(t+1)^{3/2}} \leq e(t)\left(1 - \frac{\rho}{t+1}\right) + \frac{\sqrt{2\ln(D(t+1))}}{(t+1)^{3/2}}$$

Now let $g(t) = \frac{\sqrt{2\ln(Dt)}}{t^{3/2}}$. By a standard inductive argument we obtain for all $t \geq 1$

$$\begin{aligned}
e(t) &\leq (1 - \frac{\rho}{t})e(t-1) + g(t) \\
&\leq (1 - \frac{\rho}{t})(1 - \frac{\rho}{t-1})e(t-2) + (1 - \frac{\rho}{t})g(t-1) + g(t) \\
&\leq (1 - \frac{\rho}{t}) \cdots (1-\rho)e(0) + \sum_{\tau=1}^{t} g(\tau) \prod_{i=\tau+1}^{t} (1 - \frac{\rho}{i}) \\
&\leq \frac{e(0)}{t^\rho} + \sum_{\tau=1}^{t} g(\tau) e^{-\rho \sum_{i=\tau+1}^{t} \frac{1}{i}} \\
&\leq \frac{e(0)}{t^\rho} + \sum_{\tau=1}^{t} g(\tau) e^{-\rho(H_t - H_\tau)} \\
&\leq \frac{e(0)}{t^\rho} + e^{-\rho H_t} \sum_{\tau=1}^{t} g(\tau) e^{\rho H_\tau} \\
&\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{2}}{t^\rho} \sum_{\tau=1}^{t} \tau^\rho \frac{\sqrt{\ln(D\tau)}}{\tau^{3/2}} \\
&\leq \frac{e(0)}{t^\rho} + \frac{\sqrt{2\ln D}}{t^\rho} \sum_{\tau=1}^{t} \frac{\sqrt{\ln \tau}}{\tau^{3/2-\rho}}
\end{aligned}$$

It remains to bound the sum on the right side by some logarithmic function. We observe that

$$\sum_{\tau=1}^{t} \frac{\sqrt{\ln \tau}}{\tau^{3/2-\rho}} \leq \int_{\tau=1}^{t} \frac{\sqrt{\ln \tau}}{\tau^{3/2-\rho}} \mathrm{d}\tau \tag{2.11}$$

since, $\tau \mapsto \frac{\sqrt{\ln \tau}}{\tau^{3/2-\rho}}$ is a decreasing function of τ for all $\rho \in [0,1]$.

- If $\rho \leq 1/2$ then

$$\int_{\tau=1}^{t} \tau^\rho \frac{\sqrt{\ln \tau}}{\tau^{3/2}} \mathrm{d}\tau \leq \sqrt{\ln t} \int_{\tau=1}^{t} \frac{1}{\tau} \mathrm{d}\tau = (\ln t)^{3/2}$$

- If $\rho > 1/2$ then

$$\begin{aligned}\int_{\tau=1}^{t} \tau^{\rho} \frac{\sqrt{\ln \tau}}{\tau^{3/2}} \mathrm{d}\tau &= \int_{\tau=1}^{t} \tau^{\rho-1/2} \frac{\sqrt{\ln \tau}}{\tau} d\tau \\ &= \frac{2}{3} \int_{\tau=1}^{t} \tau^{\rho-1/2} ((\ln \tau)^{3/2})' \mathrm{d}\tau \\ &= \frac{2}{3} t^{\rho-1/2} (\ln t)^{3/2} - (\rho - 1/2) \frac{2}{3} \int_{\tau=1}^{t} \tau^{\rho-3/2} (\ln \tau)^{3/2} \mathrm{d}\tau \\ &\leq \frac{2}{3} t^{\rho-1/2} (\ln t)^{3/2}\end{aligned}$$

Hence, in any case,

$$e(t) \leq \sqrt{2.5 \ln(D)} \frac{(\ln t)^{3/2}}{t^{\min(\rho,\, 1/2)}}$$

Finally, we would like to translate this result into one that involves the quantity $\mathbf{E}\left[\|x(t) - x^*\|_\infty\right]$. Hence, we now present a proof of Theorem 10.

Theorem 10. *Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 11 with equilibrium $x^* \in [0,1]^n$. The opinion vector $x(t) \in [0,1]^n$ produced by update rule (1.4) after t rounds satisfies*

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq C\sqrt{\log n} \frac{(\log t)^{3/2}}{t^{\min(1/2, \rho)}},$$

where $\rho = \min_{i \in V} a_i$ and C is a universal constant.

Proof. By Theorem 11 setting $p = \frac{1}{12\sqrt{t}}$ we have that

$$\mathbf{P}\left[\|x(t) - x^*\|_\infty \leq e(t)\right] \geq 1 - \frac{1}{12\sqrt{t}}$$

where $e(t)$ is the solution of the recursion $e(t) = \delta(t) + (1-\rho)\frac{\sum_{\tau=0}^{t-1} e_p(\tau)}{t}$ with $\delta(t) = \sqrt{\frac{\log(2\pi^2 n t^{2.5})}{t}}$. Since $2\pi^2 \geq e^{2.5}$, Lemma 12 applies and $e(t) \leq C\sqrt{\log n}\frac{\log t^{3/2}}{t^{\min(\rho,1/2)}}$ for some universal constant C. Finally,

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq \frac{1}{12\sqrt{t}} + (1 - \frac{1}{12\sqrt{t}}) C\sqrt{\log n}\frac{(\log t)^{3/2}}{t^{\min(\rho,1/2)}} \leq (C + \frac{1}{12})\sqrt{\log n}\frac{(\log t)^{3/2}}{t^{\min(\rho,1/2)}}$$

□

Although the protocol converges to the equilibrium x^*, the convergence rate is not so fast as that of original FJ model. In the next Chapter, we present other algorithms that take advantage of some additional knowledge about the graph and converge exponentially fast. This doesn't mean that our protocol is suboptimal. In Chapters 4 and 5 we will show that without additional information no protocol can do better than *Follow the Leader*.

Chapter 3

Graph awareness

In this chapter, we are going to investigate possible ways to achieve exponential convergence, if we allow the protocol to know some more details about the graph. In doing so, we underline some of the limitations of the model and so we discover the fundamental difficulties of computing the equilibrium. The first protocol is a *Row dependent dynamics*, meaning that each agent knows the probabilities of choice of his neighbors. The second protocol assumes rational probabilities of choice, which enables rounding schemes that yield faster algorithms.

3.1 Row dependent dynamics

3.1.1 Algorithm

As we will see in the following Chapters, the lower bound crucially depends on the fact that $\mathcal{A}$ is no-regret. At this point, a natural question is whether this exponential gap is a generic restriction of our imperfect information model. We answer this question in the negative. More precisely the reason that no-regret dynamics suffer slow convergence is that the update rule depends only on the expressed opinions. In this section we exhibit an update rule that depends on the graph G and achieves exponentially fast convergence.

In [BT97] they study the convergence of various optimization algorithms in a distributed setting, where each node at a specific time does not have completely updated information about the values of the other nodes. One of these algorithms involves finding the fixed point of a contraction map when every node has some outdated information. Also, in [CC16] they study the convergence properties of gradient descent under the assumption that every node has "bounded outdateness", which means that no node has information that is very old. Based on these works for asynchronous distributed minimization algorithms , we provide an update rule showing that information about the graph G combined with agents that do not act selfishly can restore the fast convergence rate. Our update rule depends not only on the expressed opinions of the neighbors that an agent i meets, but also on the i-th row of matrix P. The latter is interesting because information about the graph G combined with agents that do not act selfishly, can restore the exponential convergence rate. However, agents that choose to update their opinions according to Dynamics 5 may experience regret if some other agents play adversarially.

In update rule (3.1), each agent stores the *most recent* opinions of the random neighbors that she meets in an array and then updates her opinion according to their weighted sum (each agent knows row i of P). For a given instance $I = (P, s, \alpha)$ we call the produced dynamics *Row Dependent dynamics* (Dynamics 5).

The problem with this approach is that the opinions of the neighbors that she keeps in her array are *outdated*, i.e. the opinion of a neighbor of agent i has changed since their last meeting. The good news is that as long as this outdatedness is bounded we can still achieve fast convergence to the equilibrium. By bounded outdatedness we mean that there exists a number of rounds B such that all agents have met all their neighbors at least once from $t-B$ to t.

Remark 1. *Update rule (3.1), apart from the opinions and the indices of the neighbors that an agent meets, also depends on the the exact values of the weights p_{ij} and that is why* Row Dependent dynamics *converge fast. We mention that the reduction of Section 4.1 still holds even if the agents also use the indices of the neighbors that they meet to update their opinion. The latter implies that any update rule that ensures fast convergence would require from each agent i to be* aware *of the i-th row of matrix P. This also showcases the fundamental difficulty of our setting, which is to compute the weights p_{ij} and more generally the frequency of an event. In chapter 4 we use this observation to construct lower bounds by reducing the equilibrium computation to bernoulli estimation.*

Algorithm 5 Row Dependent dynamics

1: Initially $x_i(0) = s_i$ for all agent i.
2: Each agent i keeps an array M_i of length $|N_i|$, randomly initialized.
3: At round $t \geq 0$ each agent i:
  4: Meets neighbor with index W_i^t, $\mathbf{P}\left[W_i^t = j\right] = p_{ij}$.
  5: Suffers cost $(1-\alpha_i)(x_i(t) - x_{W_i^t}(t))^2 + a_i(x_i(t) - s_i)^2$ and learns $(x_{W_i^t}(t), W_i^t)$.
  6: Updates her array M_i and opinion: $M_i[W_i^t] \leftarrow x_{W_i^t}(t)$, $x_i(t+1) = (1-\alpha_i)\sum_{j \in N_i} p_{ij} M_i[j] + \alpha_i s_i \qquad (3.1)$

Notice that the protocol essentially consists of the simple FJ-model calculation, the only difference being that in each round the agent learns only one opinion instead of all of them.

We now provide an example that shows that Dynamics 5 doesn't have the no regret property.

Example 1. *The purpose of this example is to illustrate that the update rule (3.1) does not ensure the no regret property. If some agents for various reasons exhibit irrational or adversarial behavior, agents that adopt update rule (3.1) may experience regret. That is the reason that* Row Dependent dynamics *converge exponetially faster that any* no regret dynamics *including the* FTL dynamics.

Consider the instance of the game of Definition 11 consisting of two agents. Agent 1 adopts update rule (3.1) and has $s_1 = 0, \alpha_1 = 1/2, p_{12} = 1$ and agent 2 plays adversarially. Thus, s_2, α_2, p_{21} don't need to be specified. By update rule (3.1), $x_1(t) = x_2(t-1)/2$ and thus total disagreement cost that agent 1 experiences until round t is

$$\sum_{\tau=0}^{t} \frac{1}{2} x_1(t)^2 + \frac{1}{2}(x_1(t) - x_2(t))^2 = \sum_{\tau=0}^{t} \frac{1}{8} x_2(t-1)^2 + \frac{1}{2}(\frac{1}{2} x_2(t-1) - x_2(t))^2.$$

Since agent 2 plays adversarially, she selects $x_2(t) = 0$ if t is even and 1 otherwise. As a result, the total cost that agent 1 experiences is $\sum_{\tau=0}^{t} \frac{1}{2} x_1(t)^2 + \frac{1}{2}(x_1(t) - x_2(t))^2 \simeq 3t/8$. Now agent 1 regrets for no adopting the fixed opinion 1/3 during the whole game play. Selecting $x_1(t) = 1/3$ for all t, would incur him total disagreement cost $\sum_{\tau=0}^{t} \frac{1}{2}(1/3)^2 + \frac{1}{2}(1/3 - x_2(t))^2 \simeq 7t/36$ which is less than $3t/8$.

3.1.2 Convergence Analysis

First of all, we shall investigate what would happen if there was a window length B such that all agents have met all their neighbors at least once during that time. Intuitively, this means that it takes B rounds to complete a single iteration of the FJ-model. This is exactly what the next lemma says.

Lemma 4. *Let $\rho = \min_i a_i$, and $\pi_{ij}(t)$ be the most recent round before round t, that agent i met her neighbor j. If there exists $B > 0$ such that for all $t \geq B$, $t - B \leq \pi_{ij}(t)$ then, for all $t \geq kB$, $\|x(t) - x^*\|_\infty \leq (1-\rho)^k$.*

Proof. To prove our claim we use induction on k. For the induction base $k = 1$,

$$|x_i(t) - x_i^*| = |(1-\alpha_i) \sum_{j \in N_i} p_{ij}(x_j(\pi_{ij}(t)) - x_j^*)| \leq (1-\alpha_i) \sum_{j \in N_i} p_{ij}|x_j(\pi_{ij}(t)) - x_j^*| \leq (1-\rho)$$

since all the opinions lie in $[0,1]$. Assume that for all $t \geq (k-1)B$ we have that $\|x(t) - x^*\|_\infty \leq (1-\rho)^{k-1}$. For $k \geq 2$, we again have that

$$|x_i(t) - x_i^*| \leq (1-\rho) \sum_{j \in N_i} p_{ij}|x_j(\pi_{ij}(t)) - x_j^*|$$

Since $t - B \leq \pi_{ij}(t)$ and $t \geq kB$ we obtain that $\pi_{ij}(t) \geq (k-1)B$. As a result, the inductive hypothesis applies,

$$|x_j(\pi_{ij}(t)) - x_j^*| \leq (1-\rho)^{k-1}$$

for all $j \in N_i$ and $|x_i(t) - x_i^*| \leq (1-\rho)^k$. □

In Row Dependent dynamics there does not exist a fixed length window B that satisfies the requirements of Lemma 4. However we can select a length value such that the requirements hold with high probability. To do this observe that agent i simply needs to wait to meet the neighbor j with the smallest weight p_{ij}. Therefore, after $\log(1/\delta)/\min_j p_{ij}$ rounds we have that with probability at least $1-\delta$ agent i met all her neighbors at least once. Since we want this to be true for all agents, we shall roughly take $B = 1/\min_{p_{ij}>0} p_{ij}$. In the following we give the detailed argument that leads to Theorem 16, showing that the convergence rate of update rule (3.1) is fast. First, we turn our attention to the problem of calculating the size of window B, such that with high probability all agents have outdatedness at most B. We begin by stating a useful fact concerning the coupons collector problem.

Lemma 5. *Suppose that a collector picks coupons with different probabilities, where n is the number of distinct coupons. Let w be the minimum of these probabilities. If he selects $\ln n/w + c/w$ coupons, then:*

$$\mathbf{P}\left[\textit{collector hasn't seen all coupons}\right] \leq \frac{1}{\mathrm{e}^c}$$

Proof. Suppose p_i is the probability of picking coupon i in a round. We are going to bound the probability that the collector hasn't seen all the distinct coupons after picking t coupons. For a single coupon i, the probability of not picking that coupon after t rounds is

$$(1-p_i)^t \leq (1-w)^t$$

By a simple union bound argument, we get

$$\begin{aligned}\mathbf{P}\left[\text{collector hasn't seen all coupons}\right] &\leq \sum_{i=1}^{n} \mathbf{P}\left[\text{collector hasn't seen coupon } i\right] \\ &\leq n(1-w)^t \\ &\leq n\mathrm{e}^{-wt}\end{aligned}$$

Thus, by picking $t = \ln n/w + c/w$ coupons, the probability of not seeing all the distinct ones is:

$$\mathbf{P}\left[\text{collector hasn't seen all coupons}\right] \leq n\mathrm{e}^{-w(\ln n/w + c/w)} = \frac{1}{\mathrm{e}^c}$$

□

We will use Lemma 5 with the agents as collectors to show that with high probability there is a window B that satisfies the requirements we want.

Lemma 6. *Let $\pi_{ij}(t)$ be the most recent round before round t that agent i met agent j and $B = 2\ln(\frac{nt}{\delta})/\min_{ij} p_{ij}$. Then with probability at least $1-\delta$, for all $\tau \geq B$ and for all $i, j \in N_i$*

$$\tau - B \leq \pi_{ij}(\tau) \leq \tau - 1.$$

Proof. For simplicity we denote $w = \min_{ij} p_{ij}$. Consider an agent i at round $\tau \geq B$ where $B = 2\ln(\frac{nt}{\delta})/w$ and assume that there exists an agent $j \in N_i$ such that $\pi_{ij}(\tau) < \tau - B$. Agent i can be viewed as a coupon collector that has buyed B coupons but has not found the coupon corresponding to agent j. Since $N_i < n$ and $\min_{j \in N_i} p_{ij} \geq w$ by Lemma 5 we have that

$$\mathbf{P}\left[\text{there exists } j \in N_i \text{ s.t. } \pi_{ij}(\tau) < \tau - B\right] \leq \frac{\delta}{nt}$$

The proof follows by a union bound for all agent i and all round $B \leq \tau \leq t$.

□

By direct application of Lemma 4 and Lemma 6, we obtain the following corollary, which implies the exponential convergence of our update rule with high probability.

Corollary 1. *Let $x(t)$ the opinion vector produced by update rule (3.1) for the instance $I = (P, s, \alpha)$, then with probability at least $1-\delta$*

$$\|x(t) - x^*\|_\infty \leq \exp\left(-\frac{\rho t \min_{ij} p_{ij}}{2\ln(\frac{nt}{\delta})}\right)$$

where $\rho = \min_{i \in V} \alpha_i$.

Proof. Let $B = 2\ln(\frac{nt}{\delta})/\min_{ij} p_{ij}$. By Lemma 6 we have that with probability at least $1-\delta$, for all $i, j \in N_i$ and for all $\tau \geq B$,

$$\tau - B \leq \pi_{ij}(\tau)$$

As a result, with probability at least $1-\delta$ the requirements of Lemma 4 are satisfied, meaning that

$$\|x(t) - x^*\|_\infty \leq (1-\rho)^{\frac{t}{B}} \leq \exp\left(-\frac{\rho t \min_{ij} p_{ij}}{2\ln(\frac{nt}{\delta})}\right)$$

□

We can now state and prove the exponential convergence result, in a form involving only expectations.

Theorem 16. *Let $I = (P, s, \alpha)$ be an instance of the opinion formation game of Definition 11 with equilibrium $x^* \in [0,1]^n$ and let $\rho = \min_{i \in V} a_i$. The opinion vector $x(t) \in [0,1]^n$ produced by update rule (3.1) after t rounds satisfies*

$$\mathbf{E}\left[\|x(t) - x^*\|_\infty\right] \leq 2\exp(-\rho \min_{ij} p_{ij} \sqrt{t}/(4\ln(nt))).$$

Proof. Let $u(t) = \|x(t) - x^*\|_\infty$ and $w = \min_{ij} p_{ij}$. From Corollary 1 we obtain:

$$\mathbf{P}\left[u(t) > \exp\left(-\frac{\rho w t}{2\ln(\frac{nt}{\delta})}\right)\right] \leq \delta$$

for every probability $\delta \in [0,1]$. Also, since all the parameters of the problem lie in $[0,1]$, we have

$$\mathbf{E}\left[u(t) | u(t) > r\right] \leq 1$$

We use a standard method for converting results involving high probability to ones stated with expectation. By the conditional expectations identity, we get:

$$\begin{aligned}\mathbf{E}\left[u(t)\right] &= \mathbf{E}\left[u(t)|u(t) > r\right]\mathbf{P}\left[u(t) > r\right] + \mathbf{E}\left[u(t)|u(t) \leq r\right]\mathbf{P}\left[u(t) \leq r\right] \\ &\leq \delta + r\end{aligned}$$

where $r = \exp\left(-\frac{\rho w t}{2\ln(\frac{nt}{\delta})}\right)$. If we set $\delta = \exp\left(-\frac{\rho w\sqrt{t}}{2\ln nt}\right)$, then:

$$\mathbf{E}\left[u(t)\right] \leq \exp\left(-\frac{\rho w\sqrt{t}}{2\ln nt}\right) + \exp\left(-\frac{\rho w t}{2\ln(\frac{nt}{\delta})}\right)$$

We now evaluate r for our choice of probability δ:

$$\begin{aligned}r &= \exp\left(-\frac{\rho w t}{2\ln\left(\frac{nt}{p}\right)}\right) \\ &= \exp\left(-\frac{\rho w t}{2\ln\left(\frac{nt}{\exp\left(-\frac{\rho w\sqrt{t}}{2\ln nt}\right)}\right)}\right) \\ &= \exp\left(-\frac{\rho w t}{2\ln nt + 2\frac{\rho w\sqrt{t}}{2\ln nt}}\right) \\ &\leq \exp\left(-\frac{\rho w t}{4\ln(nt)\sqrt{t}}\right) \\ &= \exp\left(-\frac{\rho w\sqrt{t}}{4\ln(nt)}\right)\end{aligned}$$

Using the previous calculation, we obtain:

$$\begin{aligned}\mathbf{E}\left[u(t)\right] &\leq \exp\left(-\frac{\rho w\sqrt{t}}{2\ln(nt)}\right) + \exp\left(-\frac{\rho w\sqrt{t}}{4\ln(nt)}\right) \\ &\leq 2\exp\left(-\frac{\rho w\sqrt{t}}{4\ln(nt)}\right) \\ &= 2\exp\left(-\rho\min_{ij} p_{ij}\frac{\sqrt{t}}{4\ln(nt)}\right)\end{aligned}$$

□

A small detail is that the convergence rate is not exactly linear, but $O\left(\mathrm{e}^{-\sqrt{t}}\right)$. However, this is a lot faster than what we achieved with protocol 4. This shows the power of knowing matrix P.

3.2 A protocol for rational weights

3.2.1 Algorithm

In this Section, we provide another protocol that uses additional information about the graph to achieve fast convergence. Precisely, our update rule depends on the expressed opinions, the number of neighbors of each agent, and the number of agents n. It also makes the assumption that the neighbors are picked uniformly at random from each agent. We could say that the algorithm knows the "structure" of matrix P but not the specific weights.

The main idea of the protocol is straightforward: to counterbalance the imperfect information, the agents can spend some rounds to simulate one round of the original FJ-model, much like the row dependent dynamics of Section 3.1. To do this, they agree to stop updating their expressed opinion for a large enough window of rounds so that everybody learns, with high probability, *exactly* the average of the opinions of their neighbors. Following the ideas of Section 1.2 an agent could just average the opinions that she gets in this window. Unfortunately this would again result in a poly$(1/\varepsilon)$-round protocol. However this can be fixed by using the additional knowledge of the number of agents and the number of neighbors. Precisely, each agent i keeps an array with the frequencies of the different opinions that she observes. The catch is that at the end of the window, she rounds each frequency to the closest multiple of $1/d_i$, where d_i is the number of neighbors of agent i. Obviously, since all the probabilities are multiples of $1/d_i$, the true frequencies are also multiples of $1/d_i$. This rounding step is crucial to ensure the exponential convergence rate. To see that this works, first notice that if all agents stop updating their opinions for a number of rounds, each agents just needs to specify exactly how many of her neighbors share a specific observed opinion. If the length of the window is large the frequency of a specific opinion at the end of the window will be sufficiently close to the true frequency. Since the frequencies of the opinions that agent i observes can only be multiples of $1/d_i$, we can round the estimated frequencies to the closest multiple of $1/d_i$ to recover the true frequencies and use them to get the exact average of the opinions of the neighbors. This means that there is positive probability that we have the exact average at the end of this process. We can then amplify this probability with standard methods. We now present the algorithm.

Algorithm 6 FREQS algorithm

$x_i(0) \leftarrow s_i$.
$M_1 = O(\ln(n/\varepsilon))$, $M_2 = O(d^2 \ln(d))$
for $l = 1, \ldots, \ln(1/\varepsilon)$ **do**
  Keep a set A of tuples $(x, \text{freq}(x))$ of and an array B of length M_1.
  for $j = 1, \ldots, M_1$ **do**
    for $k = 1, \ldots, M_2$ **do**
      Get the opinion X_k of a random neighbor.
      if X_j is not in A **then**
        Insert $(X_k, 1)$ to A.
      else
        $(X_k, \text{freq}(X_k)) \leftarrow (X_k, \text{freq}(X_k) + 1)$.
      end if
    end for
    Divide all frequencies of A by M_2.
    Round all frequencies of A to the closest multiple of $1/d_i$.
    $B(j) \leftarrow \alpha_i \sum_{x \in A} \text{freq}(x) + (1 - \alpha_i)s_i$.
  end for
  $x_i(t) \leftarrow \text{majority}_j B(j)$.
end for

Essentially, we compute the frequencies of the opinions of the neighbors. Notice that the algorithm consists of three nested for loops. The innermost loop is for sampling in order to learn the frequencies within error $1/d_i$. The middle loop is to amplify the probability of guessing the frequencies correctly after the rounding. The outermost loop is for performing enough steps to come ϵ-close to the equilibrium. There is a subtle difference between FREQS and the previous ones. This algorithm needs the desired error ϵ from the solution as input in order to run. This is because it is needed in the amplification phase of the procedure. We can of course choose gradually smaller ϵ in order to converge to the equilibrium.

3.2.2 Analysis

The first question we should answer is how many rounds should we wait in order to find the frequencies of the opinions precisely. Essentially what we want is to find the probability density function of the random variable that is the opinion that comes in each round. To bound the length of this window we use a standard argument from density estimation theory and show that with $n^2 \log n$ rounds each agent knows the frequencies within error smaller than $1/n$ with constant probability, which then can be trivially amplified by repeating the procedure.

We next state a version of the standard VC-Inequality that we will use in our argument. Let P be a discrete distribution over $[n]$, and let $S_1, \dots, S_t$ be t i.i.d samples drawn from P, i.e. $(S_1, \dots, S_t) \sim P^t$. The empirical distribution $\hat{P}_t$ is the following estimator of the density of P.

$$\hat{P}_t(A) = \frac{\sum_{i=1}^{t} \mathbf{1}[S_i \in A]}{t}, \tag{3.2}$$

where $A \subseteq [n]$. In words, $\hat{P}_t$ is simply a frequency diagram, since it counts how many times the value i appeared in the samples $S_1, \dots S_t$. The goal of $\hat{P}_t$ is to find the density of the distribution P. We will use the following version of the classical result of Vapnik and Chervonenkis.

Lemma 7 (VC-inequality [DL12, Vap13])**.** *Let $\mathcal{A}$ be a collection of subsets of $\{1, \dots, n\}$ and let $S_{\mathcal{A}}(t)$ be the Vapnik-Chervonenkis shatter coefficient, defined by*

$$S_{\mathcal{A}(t)} = \max_{x_1,\dots,x_t \in [n]} \left| \{\{x_1, \dots, x_t\} \cap A : A \in \mathcal{A}\} \right|.$$

Then

$$\mathbf{E}_{P^t}\left[\max_{A \in \mathcal{A}} \left|\hat{P}_t(A) - P(A)\right|\right] \leq 2\sqrt{\frac{\log 2S_{\mathcal{A}}(t)}{t}}$$

This result bounds the expected error of the estimator $\hat{P}_t$ using the so called *shatter coefficient*, which is a measure of how "many" subsets of $[n]$ we want to estimate the probability of. In our case, n is the number of different opinions of the neighbors, hence it is upper bounded by n. Also, we are only interested on the probability of the specific opinions, hence $\mathcal{A}$ consists only of singletons. These ideas are made precise in the following Lemma.

Lemma 8. *If $M_2 = 100n^2 \log(2n)$, then at the end of each rounding step, the frequencies are exactly correct with probability at least $3/5$.*

Proof. According to the update rule 6 all agents fix their opinions $x_i(t)$ for $M_1 \times M_2$ rounds. To estimate the sum of the opinions each agent estimates the frequencies k_j/d_i. Since the neighbors have at most d_i different opinions we can map the opinions to natural numbers in $[d_i]$. At each round the agent gets the opinion of a random neighbor and therefore the samples X_i that she observes are drawn from a discrete distribution P supported on $[d_i]$. If k_j be the (absolute) frequency of the opinion j namely the number of neighbors that express j as their opinion, then the probability $P(j)$ of opinion j is k_j/d_i. To learn the probabilities $P(j)$ using samples from P, we let $\mathcal{A} = \{\{1\}, \{2\}, \dots, \{d_i\}\}$ and use Lemma 7 to get that

$$\mathbf{E}_{P^m}\left[\max_{j \in [d_i]} \left|\hat{P}_m(j) - P(j)\right|\right] \leq 2\sqrt{\frac{\log 2d_i}{m}},$$

since $S_{\mathcal{A}} \leq n$. Therefore, an agent can draw $m = 100n^2 \log(2n)$ to learn the frequencies k_j/d_i within expected error $1/(5d)$. Notice now that the array A after line 14 corresponds to the empirical distribution of equation (3.2). Notice that if the agents have estimations of the frequencies k_j/d_i with error smaller than $1/d$, then by rounding them to the closest multiple of $1/2d_i$ they learn the frequencies exactly. Thus, we need to bound the probability that the agents learn the frequencies with error greater than $1/2d_i$. By Markov's inequality we have that

$$\mathbf{P}\left[\max_{j\in[d_i]}\left|\hat{P}_m(j)-P(j)\right|\geq\frac{1}{2d_i}\right]\leq 2d_i\mathbf{E}_{P^m}\left[\max_{j\in[d_i]}\left|\hat{P}_m(j)-P(j)\right|\right]<\frac{2}{5}$$

Hence, with probability at least 3/5 the rounded frequencies are exactly correct. □

We now present the probability amplification argument based on taking the majority. The technique is similar to proving the equivalence of various definitions of the complexity class BPP, see[Dou15].

Lemma 9. *If we run the rounding schema $O(\ln 1/\delta)$ times, then with probability at least $1-\delta$, the majority of the values we obtain will be the correct average.*

Proof. For each rounding that we perform we define a Bernoulli random variable X_i which is one if and only if the rounding yields the correct average. We know that all X_i have the same probability, which is greater than 3/5, as we showed in Lemma 8. By applying the Chernoff Bound 14 we get:

$$\mathbf{P}\left[\sum_{i=1}^{t}X_i<\frac{t}{2}\right]\leq\mathbf{P}\left[\left|\frac{\sum_{i=1}^{t}X_i}{t}-\mathbf{E}\left[\frac{\sum_{i=1}^{t}X_i}{t}\right]\right|>\frac{3}{5}-\frac{1}{2}\right]\leq 2\mathrm{e}^{-t/50}$$

Hence, if we run the rounding $t=50\ln\frac{2}{\delta}$ times, the majority will be correct with probability at least $1-\delta$. □

We now state the final Theorem about the convergence of Freqs.

Theorem 17. *Let $I=(G(V,E),s,a)$ be an instance of the opinion formation game of Definition 11 with $a>1/2$. Let d be the maximum degree of the graph G and $n=|V|$. There exists an update rule that after $O(d^2\log^2 n\log^2(1/\varepsilon))$ rounds achieves expected error $\mathbf{E}\left[\|x_t-x^*\|_\infty\right]\leq\varepsilon$.*

Proof. We know that, having computed the *exact* average of the opinions of the neighbors, $O(\log(1/\varepsilon))$ rounds are enough to achieve error ε. Since we need all nodes to succeed at computing the exact averages for $O(\log(1/\varepsilon))$ rounds, from the union bound we get that for $\delta<\frac{\varepsilon}{n\ln(1/\varepsilon)}$, with probability at least $1-\varepsilon$ the error is at most ε. This means that $M_2=O\left(\ln\frac{n}{\epsilon}+\ln\ln\frac{1}{\epsilon}\right)$. Finally, from the law of total expectation, after $T=O(d^2\log d\log(\varepsilon)(\log(n/\varepsilon)+\log\log(1/\varepsilon))$ rounds the expected error is $\mathbf{E}\left[\|x_T-x^*\|_\infty\right]=(1-\varepsilon)\varepsilon+\varepsilon\leq 2\varepsilon$. □

It is now evident that the main obstacle towards exponential convergence is the lack of knowledge of the p_{ij}. In the next Section, we will make this precise by introducing a related estimation problem in Statistics.

Chapter 4

A Statistical Lower bound for no regret processes

According to the convergence rate established in Section 2.3, we need $O\left(1/\varepsilon^2\right)$ rounds in order to be ε close to the equilibrium. This is a sublinear rate of convergence. A natural question is whether we could find an algorithm that achieves linear rate of convergence, which is the best one could hope for in an optimizaiton problem. In this Chapter, we prove that this is not possible, if the algorithm has no knowledge of the probabilities p_{ij} and only takes as input the previously seen opinions. In order to construct a lower bound on the running time of any such algorithm, we first have to identify the hardest task that this algorithm has to solve. In our case, the difficulty lies in estimating the probabilities p_{ij} of selecting the neighbours. Thus, by establishing some sort of *Statistical* lower bound about the estimation of a Bernoulli random variable, we have also bounded the running time of any such algorithm. In this Section, we first formulate precisely the reduction from distributed protocols to Bernoulli estimators. Then, we use a standard technique from Asymptotic Statistics to provide a lower bound for this estimation problem. Unfortunately, as we will discuss extensively, this lower bound is weak. For our purposes we developed a simple and novel lower bound for this estimation problem, which is presented in the next Section. Also, in Chapter 3 we show how we can obtain linear rate of convergence if the algorithm has additional knowledge about the graph of the matrix P.

4.1 Reducing FJ to Bernoulli estimation

In order to demonstrate the reduction form the computation of x^* to that of Bernoulli estimation, we must first precisely define the set of protocols that solve the distributed problem. These are the so-called *Graph Oblivious* rules, which receive no additional information for the graph apart from a neighboring opinion in each round. We repeat the definitions from Section 1.4.

Definition 16. *A graph oblivious update rule A is a sequence of functions $(A_t)_{t=0}^{\infty}$ where $A_t : [0,1]^{t+2} \mapsto [0,1]$.*

A graph-oblivious update rule together with an instance of the opinion formation game define an graph-oblivious *dynamic*.

Definition 17. *Let a graph oblivious update rule A. For a given instance $I = (P, s, \alpha)$ the rule A produces a graph oblivious dynamics $x_A(t)$ defined as follows:*

- *Initially each agent i selects her opinion $x_i^A(0) = A_0(s_i, \alpha_i)$*
- *At round $t \geq 1$, each agent i selects her opinion $x_i^A(t) = A_t(x_{W_i^0}(0), \ldots, x_{W_i^{t-1}}(t-1), s_i, \alpha_i)$, where W_i^t is the neighbors that i meets at round t.*

As can be seen from definition 17, the opinion of each agent only depends on the previous opinions he has received and on internal constants a_i, s_i. This is the minimum requirement that a protocol should satisfy, in order to prevent it from "cheating". In Chapter 3 we saw that even a small amount of additional information is enough to get past the "estimation difficulty" of the problem and achieve exponential rate. With these minimal assumptions about the protocols, we will show that $O\left(1/\sqrt{t}\right)$ is the best rate achievable.

To show the relationship between the limited information setting and Bernoulli estimation, we will show how we can use a Graph-oblivious protocol to solve the estimation problem. Essentially, we are reducing the estimation problem to the computation of the equilibrium point x^* of a certain instance of the opinion formation game. The following Theorem clarifies these thoughts.

Theorem 14. *Let A a graph oblivious update rule such that for all instances $I = (P, s, \alpha)$,*

$$\lim_{t\to\infty} t^{1+c}\mathbf{E}\left[\|x_A(t) - x^*\|_\infty\right] = 0.$$

Then there exists an estimator $\theta_A = (\theta_t^A)_{t=1}^{\infty}$ such that for all $p \in [0,1]$, $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p\left[|\theta_t^A - p|\right] = 0$.

Proof. We construct an estimator $\theta_A = (\theta_t^A)_{t=1}^{\infty}$ using the update rule A. Consider the instance I_p described in Figure 4.1. By straightforward computation, we get that the equilibrium point of the graph is $x_c^* = p/3, x_1^* = p/6 + 1/2, x_0^* = p/6$. Now consider the opinion vector $x_A(t)$ produced by the update rule A for the instance I_p. Note that for $t \geq 1$,

- $x_1^A(t) = A_t(x_c(0), \ldots, x_c(t-1), 1, 1/2)$
- $x_0^A(t) = A_t(x_c(0), \ldots, x_c(t-1), 0, 1/2)$
- $x_c^A(t) = A_t(x_{W_c^0}(0), \ldots, x_{W_c^{t-1}}(t-1), 0, 1/2)$

The key observation is that the opinion vector $x_A(t)$ is a deterministic function of the index sequence $W_c^0, \ldots, W_c^{t-1}$ and does not depend on p. Thus, we can construct the estimator θ_A with $\theta_t^A(W_c^0, \ldots, W_c^{t-1}) = 3x_c^A(t)$. Essentially, we are simulating the execution of the protocol by giving to the central node c the value of one of its two neighbors, depending on the value of the Bernoulli random variable at a given time. That is, for a given instance I_p the choice of neighbor W_c^t is given by the value of the Bernoulli random variable with parameter p ($\mathbf{P}\left[W_c^t = 1\right] = p$). Since the random variables W_c^t, which the estimator takes as input, follow $B(p)$, we know that the simulated protocol runs in the same way as it would in the instance I_p. As a result,

$$\mathbf{E}_p\left[|\theta_t^A - p|\right] = 3\mathbf{E}\left[|x_c^A(t) - p/3|\right] \leq 3\mathbf{E}\left[\|x_A(t) - x^*\|_\infty\right]$$

Since for any instance I_p, we have that

$$\lim_{t\to\infty} t^{1+c}\mathbf{E}\left[\|x_A(t) - x^*\|_\infty\right] = 0$$

, it follows that $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p\left[|\theta_t^A - p|\right] = 0$ for all $p \in [0,1]$. □

Theorem 14 says that the problem of computing the equilibrium point x^* is "harder" that estimating a Bernoulli probability p. In the next Sections, we will show some lower bounds about the estimation problem, which will also hold for our problem.

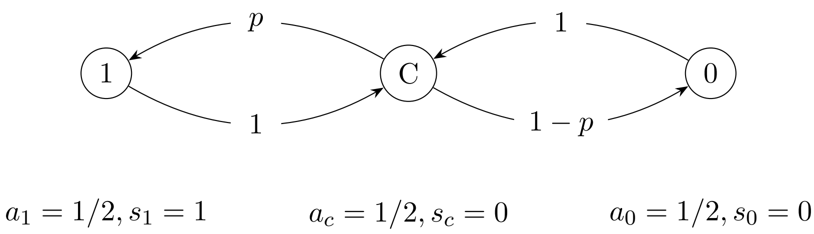

Figure 4.1

4.2 Statistical Lower Bounds

4.2.1 Minimax Risk

A very natural question in the field of estimation is how many samples one needs in order to estimate a quantity of a distribution. Using the definitions that we introduced in Section 1.4, the question can be restated as follows.

Question. *What is the smallest risk that can be obtained by an estimator of a quantity of a distribution, given t samples?*

The simplest quantity to compute is the expected value of the distribution. A good part of the Statistics literature during the 20th century is dedicated to answering this question by presenting various estimators that converge to the expected value as fast as possible. However, there seemed to be a barrier in how fast these estimators computed p. Specifically, in many cases no method could achieve risk lower that $O\left(1/\sqrt{t}\right)$. This led researchers to conclude that no faster reasonable estimator exists and to instead try and establish *lower bounds* in the risk of *any* estimator.

A special case of this computation is when the distribution is a Bernoulli with probability p. All the presented results will be stated for Bernoulli random variables, but can be easily generalized to more complex distributions.

The oldest sample complexity lower bound for estimation problems is the well-known Cramer-Rao inequality[Cra46, Rao92]. Let the function $\theta_t : \{0,1\}^t \mapsto [0,1]$ such that $\mathbf{E}_p\left[\theta_t\right] = p$ for all $p \in [0,1]$, then

$$\mathbf{E}_p\left[(\theta_t - p)^2\right] \geq \frac{p(1-p)}{t}. \tag{4.1}$$

Since $\mathbf{E}_p\left[|\theta_t - p|\right]$ can be lower bounded by $\mathbf{E}_p\left[(\theta_t - p)^2\right]$, we can apply the Cramer-Rao inequality and prove our claim in the case of *unbiased* estimators, $\mathbf{E}_p\left[\theta_t\right] = p$ for all t. However, this hypothesis is considered a rather restrictive one. Indeed, the proof of the inequality is merely a computation of the variance , which heavily relies on this property for differentiation with respect to p. A proof can be found in [MM08].

Sample complexity lower bounds without assumptions about the estimator are usually given as lower bounds for the *minimax risk*, which was defined [1] by Wald in [Wal39] as

$$\min_{\theta_t} \max_{p \in [0,1]} \mathbf{E}_p\left[|\theta_t - p|\right].$$

Minimax risk captures the idea that after we pick the best possible algorithm, an adversary inspects it and picks the worst possible $p \in [0,1]$ to generate the samples that our algorithm will

[1] Although the minimax risk is defined for any estimation problem and loss function, for simplicity, we write the minimax risk for estimating the mean of a Bernoulli random variable.

get as input. For simplicity, we denote the minimax risk for t samples as M_t. The methods of Le'Cam, Fano, and Assouad are well-known information-theoretic methods to establish lower bounds for the minimax risk. For more on these methods see [Yu97, Tsy08]. In the following sections, we will establish a lower bound of $\Omega(1/\sqrt{t})$ for the minimax risk for the case of estimating the mean of a Bernoulli. The exposition of LeCam's method in the following Sections is the same as in [Duc16].

4.2.2 From estimation to testing

In order to prove minimax lower bounds, a common method is to reduce the estimation problem to a testing problem. The idea is to show that the minimax risk can be lower bounded by the probability of failure of an estimator at a testing problem. Then, we can develop tools for lower bounding this probability. As a start, we give a precise definition of a *testing* problem. The discussion only involves Bernoulli distributions, but can easily be extended to arbitrary families of distributions.

Given an index set V of finite cardinality, consider a family of Bernoulli distributions $\{P_v\}_{v \in V}$. We denote by p_v the probability of distribution P_v. We conduct the following experiment. First, nature chooses uniformly an index $v \in V$. Then, t samples are drawn independently from distribution P_v, where v is the index that was chosen in the first step. The samples are given to an estimator, whose job is to decide what is the index $v \in V$ that was chosen is the first place. Since a Bernoulli random variable only takes the values 0 and 1, a *tester* is thus a function $\psi : [0,1]^t \mapsto V$ that takes the samples as input and outputs the "id" of the Bernoulli distribution. In the simplest case, V consists of only two elements, thus the problem is to *test*, or distinguish, which of the two distributions produces the samples.

Next, suppose that the family of Bernoulli distributions P_v is a 2δ *packing*, meaning that for all $v, v' \in V$ with $v \neq v'$, it holds:

$$|p_v - p_{v'}| \geq 2\delta$$

This hypothesis is a way to express that the distributions are "δ-far" from each other. The idea for the reduction from estimation to testing is simple: If we could estimate the probability of a Bernoulli distribution to within error δ, then we could also take the closest from the family P_v as the answer for our testing problem. Thus, the testing problem is in a sense "easier" than the estimation one. These ideas are made precise in the following lemma.

Lemma 10. *Suppose we have a family of Bernoulli distributions P_v that is a 2δ packing, for some $\delta > 0$. Then, the minimax risk M_t for Bernoulli distributions is lower bounded as follows:*

$$M_t \geq \delta \inf_{\psi} \mathbf{P}\left[\psi\left(X_1, \ldots, X_t\right) \neq V\right] = \delta \inf_{\psi} \mathbf{P}\left[\psi \textit{ fails the testing problem for the family } P_v\right]$$

Proof. First, we fix an estimator θ_t that takes t samples. Suppose that the samples come from a Bernoulli $B(p)$. Then:

$$\begin{aligned}\mathbf{E}_p\left[|\theta_t - p|\right] &\geq \delta \mathbf{E}_p\left[\mathbb{1}\left\{|\theta_t - p| > \delta\right\}\right] \\ &= \delta \mathbf{P}\left[|\theta_t - p| > \delta\right]\end{aligned}$$

In the latter probability is is implies that the samples come from $B(p)$. We now define a function ψ that attempts to solve the testing problem for this family of distributions. Specifically, ψ takes as input the output of estimator θ_t and outputs

$$\psi(\theta_t) = \operatorname*{argmin}_{v \in V} |\theta_t - p_v|$$

Thus, the function ψ tries to "guess" the value of p using estimator θ_t and then finds the Bernoulli with the closest p to the computed one. Now comes the important part of the reduction: if the samples come from one of the Bernoullis of family P_v and if estimator θ_t manages to compute the probability with error at most δ, then by the triangle inequality, all the other distributions will be at least δ far from the computed value. In other words, the closest Bernoulli to the computed one will be the correct one. Thus, if $|\theta_t - p| < \delta$, then ψ guesses the correct distribution. As a consequence of that, we get:

$$\mathbf{P}\left[|\theta_t - p_v| > \delta\right] \geq \mathbf{P}\left[\psi \text{ guesses wrongly}|V = v\right]$$

for all $v \in V$. Thus,

$$\mathbf{E}_{p_v}\left[|\theta_t - p_v|\right] \geq \delta \mathbf{P}\left[\psi \text{ guesses wrongly}|V = v\right]$$

for all $v \in V$. Hence, for a fixed estimator θ_t,

$$\begin{aligned}
\max_{p \in [0,1]} \mathbf{E}_p\left[|\theta_t - p|\right] &\geq \frac{1}{|V|} \sum_{v \in V} \mathbf{E}_{p_v}\left[|\theta_t - p_v|\right] \\
&\geq \delta \frac{1}{|V|} \sum_{v \in V} \mathbf{P}\left[\psi \text{ guesses wrongly}|V = v\right] \\
&= \delta \mathbf{P}\left[\psi \text{ fails the testing problem for the family } P_v\right]
\end{aligned}$$

□

The remaining challenge is to lower bound the probability of error in the underlying multi-way hypothesis testing problem by selecting a suitable class of Bernoulli distributions. To do this, there is a trade off in the choice of the separation δ. While a large δ increases the lower bound, as can be seen by Lemma 10, a smaller δ means the distributions are close to each other and hence it is harder to distinguish between them. Thus, smaller δ increases the probability of error of the tester. Usually, one attempts to choose the largest separation δ that guarantees a constant probability of error. In LeCam's method, the class consists of just two distributions, as we will see in the following.

4.2.3 LeCam's method

Le Cam's method, in its simplest form, provides lower bounds on the error in simple binary hypothesis testing testing problems. A very useful quantity when one studies these settings is the total variation distance between two distributions, which, roughly speaking, measures how "different" two distributions are.

Definition 20. *The* total variation distance *between probability distributions P and Q defined on a set X is defined as the maximum difference between probabilities they assign on subsets of X.*

$$\|P - Q\|_{TV} = \sup_{A \subseteq X} |P(A) - Q(A)|$$

The above definition is understood under the assumption that the set of values X is finite. If the set X is more complicated, the supremum may involve only sets whose preimage is a measurable set of our measurable space.

The total variation distance, as we shall see later in the course, is very important for verifying the optimality of different tests, and appears in the measurement of difficulty of solving hypothesis testing problems. This can already be seen by the way it is defined: it is the maximum difference between the two distributions. For example, if in a testing problem the total variation distance between all the distributions is small, then one expects that it will be hard to distinguish between the two. This is the idea on which LeCam's method is based.

LeCam's method uses a special case of the hypothesis testing problem, which is the *binary* testing problem. This means that there are two Bernoulli distributions P_1 and P_2 and our goal is to make a decision on whether P_1 or P_2 is the distribution generating the data we observe. Concretely, suppose that nature chooses one of the distributions P_1 or P_2 at random, and let $V \in \{1,2\}$ index this choice. Conditional on $V = v$, we then observe a sample X drawn from P_v. A *tester* is a function $\psi : X \mapsto \{1,2\}$. For every such function, by the law of conditional expectation, it holds:

$$\begin{aligned}\mathbf{P}[\psi(X) \neq V] &= \mathbf{P}[V=1]\,\mathbf{P}[\psi(X) \neq V | V=1] + \mathbf{P}[V=2]\,\mathbf{P}[\psi(X) \neq V | V=2] \\ &= \frac{1}{2}\mathbf{P}[\psi(X) \neq 1] + \frac{1}{2}\mathbf{P}[\psi(X) \neq 2]\end{aligned}$$

Thus, if we manage to lower bound the right side of the equation, we have a lower bound for the probability of error of the binary testing problem. This is done in the following lemma.

Lemma 11. *Let X be an arbitrary set. For any distributions P_1 and P_2 on X, we have*

$$\inf_{\psi}\left(\mathbf{P}[\psi(X) \neq 1] + \mathbf{P}[\psi(X) \neq 2]\right) \geq 1 - \|P_1 - P_2\|_{TV}$$

where the infimum is taken over all testers $\psi : X \mapsto \{1,2\}$.

Proof. Any test $\psi : X \mapsto \{1,2\}$ has an acceptance region, call it $A \subseteq X$, where it outputs 1 and a region A^c where it outputs 2. We have

$$\mathbf{P}[\psi(X) \neq 1] + \mathbf{P}[\psi(X) \neq 2] = P_1(A^c) + P_2(A) = 1 - P_1(A) + P_2(A)$$

Taking an infimum over such acceptance regions, we have

$$\inf_{\psi}\{\mathbf{P}[\psi(X) \neq 1] + \mathbf{P}[\psi(X) \neq 2]\} = \inf_{A \subseteq X}\{1 - P_1(A) + P_2(A)\} = 1 - \sup_{A \subseteq X}\{P_1(A) - P_2(A)\}$$

which yields the total variation distance as desired. □

Obviously, the smaller the total variation distance between P_1, P_2, the better the lower bound. Returning to the setting in which we receive t i.i.d. observations X_i , when $V = 1$ with probability $\frac{1}{2}$ and 2 with probability $\frac{1}{2}$, we have by Lemma 11

$$\inf_{\psi}\mathbf{P}[\psi(X_1, \ldots, X_t) \neq V] = \frac{1}{2} - \frac{1}{2}\|P_1^t - P_2^t\|_{TV}$$

where P_1^t, P_2^t are the t-fold product distributions of P_1, P_2 respectively. In this setting, X is equal to $\{0,1\}^t$. Hence, by applying Theorem 14 we get

$$M_t \geq \delta\left(\frac{1}{2} - \frac{1}{2}\|P_1^t - P_2^t\|_{TV}\right)$$

if $|p_1 - p_2| > 2\delta$. Now our only task is to select a suitable value for δ and compute the total variation distance of the two t-fold distributions. Obviously, the resulting lower bound will be a function of the number of samples t. Thus, it makes sense to set δ as some function of t. The following lemma provides the details.

Lemma 12. *For Bernoulli estimation, it holds*

$$M_t \geq \frac{1}{8\sqrt{2t}}$$

Proof. First, we define Bernoulli distributions P_1 and P_2 with $p_1 = \frac{1}{2} - \delta$, $p_2 = \frac{1}{2} + \delta$. The quantity δ will be defined later. We have that $|p_1 - p_2| = 2\delta$, thus

$$M_t \geq \delta\left(\frac{1}{2} - \frac{1}{2}\|P_1^t - P_2^t\|_{TV}\right) \tag{4.2}$$

By Pinsker's inequality we have

$$\|P_1^t - P_2^t\|_{TV}^2 \leq \frac{1}{2} D_{KL}(P_1^t, P_2^t) = \frac{t}{2} D_{KL}(P_1, P_2)$$

It remains to upper bound the KL divergence of two bernoullis. By definition, we have that

$$D_{KL}(P_1, P_2) = p_1 \log \frac{p_1}{p_2} + (1 - p_1) \log \frac{1 - p_1}{1 - p_2}$$

We now define the function

$$f(x) = x \log \frac{x}{p_2} + (1 - x) \log \frac{1 - x}{1 - p_2}$$

which is defined on $[p_1, p_2]$. This function is twice continuously differentiable, with derivatives:

$$f'(x) = \log \frac{x}{p_2} - \log \frac{1 - x}{1 - p_2}$$

$$f''(x) = \frac{1}{x} + \frac{1}{1 - x}$$

Thus, we can apply Taylor's Theorem, which states that there exists $\xi \in [p_1, p_2]$ such that:

$$f(p_1) = f_(p_2) + (p_1 - p_2) f'(p_2) + \frac{(p_1 - p_2)^2}{2} f''(\xi)$$

We now notice that

$$f(p_2) = f'(p_2) = 0$$

and $f(p_1) = D_{KL}(P_1, P_2)$. Thus

$$D_{KL}(P_1, P_2) = \frac{(p_1 - p_2)^2}{2} f''(\xi) = 2\delta^2 f''(\xi)$$

Without loss of generality, we can assume that $\delta < 1/4$, which means that

$$f''(x) \leq 8$$

for all $x \in [p_1, p_2]$. This means that

$$D_{KL} \leq 16\delta^2$$

We end up with the inequality

$$\|P_1^t - P_2^t\|_{TV} \leq \sqrt{\frac{t}{2} 16\delta^2} = 2\sqrt{2t}\delta$$

By setting $\delta = \frac{1}{4\sqrt{2t}}$ we have that

$$1 - \|P_1^t - P_2^t\|_{TV} \geq \frac{1}{2}$$

Hence

$$M_t \geq \frac{1}{8\sqrt{2t}}$$

which is what we wanted to prove. □

Obviously, this technique could be generalised for more complicated families of distributions and estimation problems. The simple application for Bernoulli estimation showcases the typical way to apply LeCam's method. One begins by bounding the total variation distance and then setting the value of δ in order to achieve constant probability of error, i.e. the quantity $1 - \|P_1^t - P_2^t\|_{TV}$ should be lower bounded by a constant. For more inequalities involving the total variation distance and the KL divergence, see [PDSS16]. Other methods, like Fano, generalize this technique by using a multiple hypothesis testing problem with more than two distributions[Yu97].

4.3 Limitations of LeCam

After proving this lower bound for the minimax risk, we should take a while and think about how it helps our problem. We remind the reader that the minimax risk was defined as

$$M_t = \min_{\theta_t} \max_{p\in[0,1]} \mathbf{E}_p \left[|\theta_t - p|\right]$$

By proving that $M_t \geq \frac{1}{8\sqrt{2t}}$, we show that for every estimator θ we can find a $p \in [0,1]$ such that the expected error after t rounds will be greater that $\frac{1}{8\sqrt{2t}}$. The key observation here is that for different values of t, the $p \in [0,1]$ that has large risk isn't necessarily the same. Ideally, we would like to prove a statement of the form

For every estimator θ there exists $p \in [0,1]$ such that for all times t it holds

$$\mathbf{E}_p \left[|\theta_t - p|\right] \geq \frac{1}{8\sqrt{2t}}$$

Instead, we have managed to prove the following

For every estimator θ and for every time t there exists $p \in [0,1]$ such that it holds

$$\mathbf{E}_p \left[|\theta_t - p|\right] \geq \frac{1}{8\sqrt{2t}}$$

The difference in the order of quantifiers might seem like a technical detail, but it has important consequences. The former statement is weaker than the latter. Notice how in the former claim, the probability p only depends on the estimator and not on the specific time t. If we had proved the former statement, then by the reduction of Theorem 14 it would follow that for every protocol we could define an instance where the risk would be at all times greater that $\frac{1}{8\sqrt{2t}}$. A sketch of the idea is the following: we define an estimator based on this protocol, as in Theorem 14. Then, we would find the $p \in [0,1]$ that has high risk for the specific estimator and we would run the protocol in the instance of Figure 4.1 with this p. Then, the error of the protocol would be high at all times. This is a convincing lower bound. On the contrary, if we were to apply the latter claim, then for every estimator and for every time t we would have to define a different instance of the opinion formation problem. This would result in a weird statement about the protocols, namely that for every protocol and for every time t we could find an instance where this protocol has high risk at that specific time t. This is not a satisfactory lower bound.

The bottom line of the conversation is that we want for every protocol(or estimator) to fix an instance(or a $p \in [0,1]$) and prove that the error(risk) for all times is high. This type of result seems almost impossible with the tools provided by LeCam or Fano, because the minimax risk is computed by "fixing" t. Surprisingly, there is a very simple and elegant argument not related to these methods that proves our claim. This will be presented in Section 5.2.

Chapter 5

Measure Theoretic Lower bounds

In Chapter 4 we investigated possible ways for obtaining lower bounds for estimation problems. An important observation was that the standard lower bounds on the number of samples could not be transferred to our problem, due to the dependence of the "difficult" instance of the problem from the number of samples. We are interested in showing stronger lower bounds, namely that for each estimator there is an instance where *at all times* it performs badly. In this Chapter we are going to prove something significantly stronger: for every estimator, *almost all* instances are hard, meaning that the convergence rate cannot be $o\left(1/t\right)$. The concept of *almost all* can be precisely defined be using tools from measure theory. Hence, we begin this chapter with a quick overview of the properties of measures and outer measures, which will be helpful in the proof.

5.1 Intro to Measure Theory

5.1.1 The problem of measure

For centuries, mathematicians were interested in the computation of the length, area and volume of various objects. All of these quantities are in fact "functions" defined on suitable spaces. For example, length is a function that assigns a number to each interval of the real line and area is a function defined on shapes of the 2d plane. Let's focus on the notion of length. For an interval of the form (a, b), it is intuitively obvious that the lenth is $b - a$.

An interesting question is whether we can generalize this notion of length to more complicated subsets of the real line. For example, the set

$$A = (a, b) \cup (c, d)$$

could be assigned quite intuitively a length of $b - a + d - c$, if $b < c$, since it is the union of two disjoint intervals. In a similar way, we could assign lenths to more complicated subsets of the real line. What is the most general class of subsets that could be assigned a length? What kind of properties should a function satisfy in order to be considered an appropriate generalization of length? The following sections answer these questions in detail.

5.1.2 σ-algebras

The first question we have to answer is what would be the domain of a function that assigns length. A possible solution is to consider all subsets of $\mathbf{R}$ as the domain. Unfortunately, there is no function that that is defined on all subsets of $\mathbf{R}$ and that can be considered a generalization of length. For details, see [Fol13]. However, it is possible to define length in very large collections of subsets, as we will see in the following Sections. Mathematicians came up with a very elegant way of defining collections of subsets where length can be defined. Such collections are called σ-algebras.

Definition 21. *A collection $\mathcal{F}$ of subsets of $\mathbf{R}$ is said to be a σ-albegra, if it satisfies the following properties:*

- $\mathbf{R} \in \mathcal{F}$
- *If $A \in \mathcal{F}$, then $A^c \in \mathcal{F}$.*
- *Let $\{A_i\}_{i=1}^{\infty}$ be a sequence of sets belonging to $\mathcal{F}$. Then*

$$\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}$$

all of these properties are quite intuitive for a collection of sets that have a length. For example, $\mathbf{R}$ should obviously be assigned a length of ∞. Also, if a set has a length, then it's complement could very easily be assigned a measure. Also, if we have a length for a family of sets, then obviously we can define the length of their union.

A direct consequence of the definition and DeMorgan's law is that σ-algebras are also closed under countable intersections. Simple examples of σ-albegras are $\{\emptyset, \mathbf{R}\}$ and $\mathcal{P}(\mathbf{R})$, the collection of all subsets of $\mathbf{R}$. We now present nother interesting way to construct σ-algebras from arbitrary collections of sets.

Suppose we have a collection $\mathcal{B}$ of subsets of $\mathbf{R}$, not necessarily a σ-albegra. Let $\mathcal{L}$ be the set of all σ-algebras that contain $\mathcal{B}$. We define

$$\mathcal{F} = \bigcap_{\mathcal{A} \in \mathcal{L}} \mathcal{A}$$

which is the intersection of all these σ-algebras. Note the σ- algebra $\mathcal{P}(\mathbf{R})$ contains $\mathcal{B}$, hence the set $\mathcal{L}$ is nonempty and the set $\mathcal{F}$ is well defined. Then, we have the following simple result.

Lemma 13. *The collection $\mathcal{F}$ is a σ-algebra. It is also the* smallest *σ-algebra that contains $\mathcal{B}$, in the sense that if a σ-algebra $\mathcal{F}'$ contains $\mathcal{B}$, then $\mathcal{F} \subseteq \mathcal{F}'$.*

Proof. The proof consists of checking the three properties that a σ-algebra should satisfy.

- $\mathbf{R}$ belongs in all $\mathcal{A} \in \mathcal{L}$ because they are σ-algebras, so it also belongs to their intersection.
- If $D \in \mathcal{F}$, then $D \in \mathcal{A}$ for all $\mathcal{A} \in \mathcal{L}$. Hence, $D^c \in \mathcal{A}$ for all $\mathcal{A} \in \mathcal{L}$, thus $D^c \in \mathcal{F}$.
- The closure under countable unions can be checked in the same way.

Thus, $\mathcal{F}$ is a σ-algebra. If $\mathcal{B} \subseteq \mathcal{F}'$ for some σ-algebra $\mathcal{F}'$, then $\mathcal{F}' \in \mathcal{L}$, thus $\mathcal{F} \subseteq \mathcal{F}'$. □

We say that $\mathcal{B}$ *generates* the σ-algebra $\mathcal{F}$. Lemma 13 gives us another way to construct a σ -algebra. That is, we begin with an arbitrary collection $\mathcal{B}$ of subsets of $\mathbf{R}$ and and take the σ-algebra generated by $\mathcal{B}$. A notable example is when we take $\mathcal{B}$ to be the set of all open subsets of $\mathbf{R}$. The generated σ-algebra is called the *Borel* σ-algebra of $\mathbf{R}$. It forms a rich class of subsets and plays an important role in the definition of the *Lebesgue integral*, which is beyond the scope of this thesis to discuss.

5.1.3 The Lebesgue Measure

As we previously discussed, it is important to understand what we aim at when we try to define a function generalizing the notion of length to arbitrary subsets. The length of a subset should be a *measure* of how "big" that subset is. Hence, from now on we will refer to this goal function as a *measure*. But what properties should a function satisfy in order to be sensibly considered a measure?

First of all, it makes sense that the domain of a measure is a σ-algebra, since it should be closed under some obvious operations like the union. It is also obvious that such a function should only take non-negative values, because it wouldn't make much sense to have a subset with "negative" length. Note that the function could also take the value $+\infty$. Here, we adopt the usual conventions about the extended real numbers for $+\infty$. Moreover, if we take the union of some sets that are all disjoint from one another, then the measure of the union should obviously be the sum of the measures of these sets. Hence, we give the following definition of a measure.

Definition 22. *Let $\mathcal{M}$ be a σ-algebra of $\mathbf{R}$. Then, a function $\mu : \mathcal{M} \mapsto [0, \infty]$ is called a* measure*, if it satisfies the following properties:*

- $\mu(\emptyset) = 0$
- *Let $\{A_i\}_{i=1}^{\infty}$ be a sequence of sets belonging to $\mathcal{M}$ that are pairwise disjoint. Then,*

$$\mu(\bigcup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \mu(A_i)$$

Now, our goal has been made precise. We want to find a σ-algebra $\mathcal{M}$ that contains all the intevals and a function $\mu : \mathcal{M} \mapsto [0, \infty]$ that satisfies the properties of measure. We also want the values of μ in simple intervals to coincide with the usual definition of length, that is $\mu((a, b)) = b - a$ for all $a < b$. We would obviously want $\mathcal{M}$ to be as large as possible.

Ideally, we would like the measure to be defined on $\mathcal{P}(\mathbf{R})$. However, it can be proved that no function satisfying Definition 22 can be defined on $\mathcal{P}(\mathbf{R})$. Some proofs can be found in [Fol13, Wag93, BT24]. The moral of these proofs is that $\mathbf{R}$ contains subsets which are so strangely put together that it is impossible to define a geometrically reasonable notion of measure for them, and the remedy for the situation is to discard the requirement that μ should be defined on all subsets of $\mathbf{R}$.

Indeed, it is possible to define a measure satisfying the preceding properties in the σ-algebra of *Lebesgue Measurable sets*. This measure is called the *Lebesgue* measure on the real line. The details of how this function is created are quite technical and thus cannot be presented here. Complete expositions about the Lebesgue measure are contained in [Fol13, RF88, Rud06]

We will instead list some important consequences of this construction. First, Borel sets are Lebesgue measurable. Also, it is possible for a set to have Lebesgue measure 0 but not be empty. For example, all countable subsets have Lebesgue measure 0. The point is that sets with measure 0 are "small" compared to other sets such as intervals. For example, if we prove that a property holds for all $x \in [0, 1]$ except for a set of measure 0, then we can be assured that the x that do not satisfy this property are very "rare" to find. In the following Section we are going to prove that every Bernoulli estimator will not be "fast" enough for *almost all* $p \in [0, 1]$. This means that the Bernoullis where it performs well are very few.

5.2 Counting the points of superefficiency

As we discussed in Section 4.3, the problem with the statistical lower bounds is that for every different time t we must find a different Bernoulli that has high risk for the specific estimator. We would like to show that for each estimator, there exists a Bernoulli, such that the estimator will have high risk at all times when it takes samples from this distribution. In the following Lemma, we will show something stronger. In particular, we prove that every estimator for almost all $p \in [0, 1]$ fails to achieve risk that is $o\left(1/t^{1+c}\right)$. The technique relies on a simple observation about where a p should lie in order for the estimator to converge fast. Essentially, we prove that the regions of superefficiency points shrink as the number of samples increases. Throughout the proof we use various properties of the Lebesgue measure. For proofs of these properties, see [RF88].

Theorem 15. *Let $\theta = (\theta_t)_{t=1}^{\infty}$ be a Bernoulli estimator with error rate $\mathbf{E}_p[|\theta_t - p|]$. For any $c > 0$, the set of all $p \in [0,1]$ such that $\lim_{t\to\infty} t^{1+c}\mathbf{E}_p[|\theta_t - p|] = 0$ has Lebesgue measure 0.*

Proof. Since θ_t is a function from $\{0,1\}^t$ to $[0,1]$, θ_t can have at most 2^t different values. Without loss of generality, we assume that θ_t takes the same value $\theta_t(x)$ for all $x \in \{0,1\}^t$ with the same number of 1's. For example, $\theta_3(\{1,0,0\}) = \theta_3(\{0,1,0\}) = \theta_3(\{0,0,1\})$. This is due to the fact that for any $p \in [0,1]$,

$$\sum_{0\le i\le t} \sum_{\|x\|_1 = i} |\theta_t(x) - p|\, p^i (1-p)^{t-i} \ge \sum_{0 \le i \le t} \binom{t}{i} \left| \frac{\sum_{\|x\|_1=i} \theta_t(x)}{\binom{t}{i}} - p \right| p^i(1-p)^{t-i}.$$

For any estimator θ with error rate $\mathbf{E}_p[|\theta_t - p|]$ there exists another estimator θ' that satisfies the above property and $\mathbf{E}_p[|\theta'_t - p|] \le \mathbf{E}_p[|\theta_t - p|]$ for all $p \in [0,1]$. Thus, we can assume that θ_t takes at most $t+1$ different values. Let

$$A = \{p \in [0,1] \ : \ \lim_{t\to\infty} t^{1+c}\ \mathbf{E}_{X\sim P^t}[|\theta_t(X) - p|] = 0\}.$$

We are going to prove that A has measure 0 and is thus measurable. Notice that,

$$A \subseteq \bigcup_{t=1}^{\infty} \bigcap_{k=t}^{\infty} A_k,$$

where $A_k = \{p \in [0,1] \ : \ R_k(p) < 1/2\}$, and $R_k(p) = k^{1+c}\ \mathbf{E}_{X\sim P^k}[|\theta_k(X) - p|]$. We have that $R_k : [0,1] \to [0, +\infty)$ is polynomial of degree k in p and therefore it is a measurable function. Thus, A_k is measurable. We now show that

$$A_k \subseteq B_k \coloneqq \{p \in [0,1] \ : \ k^{1+c} \min_{0\le i\le k} |\theta_k(i) - p| < 1\}.$$

We prove this by contradiction. Suppose that $p \in A_k$ but $p \notin B_k$. Since $p \in A_k$ we have that

$$R_k(p) = k^{1+c} \sum_{i=0}^{k} \binom{k}{i} |\theta_k(i) - p|\, p^i (1-p)^{k-i} \ge k^{1+c} \min_{0\le i\le k} |\theta_k(i) - p| \sum_{i=0}^{k} \binom{k}{i} p^i (1-p)^{k-i} \ge 1.$$

Since the functions $p \mapsto k^{1+c}|\theta_k(i) - p|$ are measurable, their pointwise minimum is measurable and therefore the sets B_k are also measurable. We next proceed to bound $\mu(B_k)$. Since θ_k can only take $k+1$ different values we have that there exist $k+1$ intervals (a_{k_i}, b_{k_i}) of length at most $2/k^{1+c}$ such that $B_k = \bigcup_{i=0}^{k} (a_{k_i}, b_{k_i})$. In particular:

$$a_{k_i} = \theta_k(i) - \frac{1}{k^{1+c}}$$

$$b_{k_i} = \theta_k(i) + \frac{1}{k^{1+c}}$$

as shown in Figure 5.1 Since μ is subadditive we have

$$\mu(B_k) \le \sum_{i=0}^{k} \frac{2}{k^{1+c}} = \frac{2(k+1)}{k^{1+c}}.$$

Now observe that

$$\mu(A) \le \mu\left(\bigcup_{t=1}^{\infty} \bigcap_{k=t}^{\infty} A_k\right) \le \mu\left(\bigcup_{t=1}^{\infty} \bigcap_{k=t}^{\infty} B_k\right) \le \sum_{t=1}^{\infty} \mu\left(\bigcap_{k=t}^{\infty} B_k\right) \le \sum_{t=1}^{\infty} \lim_{k\to\infty} \mu(B_k) = 0$$

Hence, $\mu(A) = 0$ and thus A is measurable. [1]

□

[1] A careful reader should notice that we wrote $\mu(A)$ without proving that A is Lebesgue measurable. A rigorous argument uses the *completeness* property of the Lebesgue measure, which says that if A is a subset of a set with Lebesgue measure 0, then A is Lebesgue measurable and has measure 0.

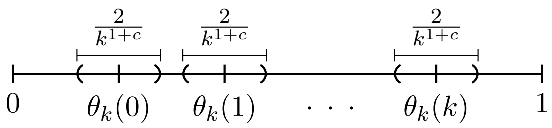

Figure 5.1: Estimator output at time k

Bibliography

[AFH$^+$05] John R. Alford, Carolyn L. Funk, John R. Hibbing, John R. Alford, and Carolyn L. Funk. Are political orientations genetically transmitted. *American Political Science Review*, pages 153–167, 2005.

[AKPT18] Rediet Abebe, Jon M. Kleinberg, David C. Parkes, and Charalampos E. Tsourakakis. Opinion dynamics with varying susceptibility to persuasion. *CoRR*, abs/1801.07863, 2018.

[BFM16] Vittorio Bilò, Angelo Fanelli, and Luca Moscardelli. Opinion formation games with dynamic social influences. In Yang Cai and Adrian Vetta, editors, *Web and Internet Economics*, pages 444–458, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[BGM13] Kshipra Bhawalkar, Sreenivas Gollapudi, and Kamesh Munagala. Coevolutionary opinion formation games. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 41–50, 2013.

[Bil08] Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.

[BKO11] David Bindel, Jon M. Kleinberg, and Sigal Oren. How bad is forming your own opinion? In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 57–66, 2011.

[BT24] Stefan Banach and Alfred Tarski. Sur la décomposition des ensembles de points en parties respectivement congruentes. *Fund. math*, 6(1):244–277, 1924.

[BT97] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.

[C$^+$52] Herman Chernoff et al. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.

[CBL03] Nicolò Cesa-Bianchi and Gábor Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51(3):239–261, Jun 2003.

[CC16] Yun Kuen Cheung and Richard Cole. A unified approach to analyzing asynchronous coordinate descent and tatonnement. *CoRR*, abs/1612.09171, 2016.

[CHM17] Johanne Cohen, Amélie Héliou, and Panayotis Mertikopoulos. Hedging under uncertainty: Regret minimization meets exponentially fast convergence. In *Algorithmic Game Theory - 10th International Symposium, SAGT 2017, L'Aquila, Italy, September 12-14, 2017, Proceedings*, pages 252–263, 2017.

[Çın11] Erhan Çınlar. *Probability and stochastics*, volume 261. Springer Science & Business Media, 2011.

[Cra46] Harald Cramer. *Mathematical Methods of Statistics*. Princeton Univ. Press, Princeton,NJ, 1946.

[DeG74] M.H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69:118–121, 1974.

[DFS18] Vasilis Kontonis Dimitris Fotakis, Vardis Kandiros and Stratis Skoulakis. Opinion dynamics with limited information. 2018.

[DL12] Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.

[DNAW00] Guillaume Deffuant, David Neau, Frédéric Amblard, and Gérard Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(1-4):87–98, 2000.

[Dou15] Mohammad Sadeq Dousti. Amplification lemma for bpp. 2015.

[Duc16] John Duchi. Lecture notes for statistics 311/electrical engineering 377. *URL: https://stanford.edu/class/stats311/Lectures/full_notes. pdf. Last visited on*, 2:23, 2016.

[EFHS17] Markos Epitropou, Dimitris Fotakis, Martin Hoefer, and Stratis 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*, pages 173–185, 2017.

[EMN09] Eyal Even-Dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 523–532, 2009.

[Ep14] *Σκουλάκης Ευστράτιος-παντελεήμων. παίγνια διαμόρφωσης ΄αποψης σε κοινωνικά δίκτυα. URL: http://artemis.cslab.ece.ntua.gr:8080/jspui/bitstream/123456789/12583/1/DT2014-0288.pdf*, 2014.

[FGV16] Diodato Ferraioli, Paul W. Goldberg, and Carmine Ventre. Decentralized dynamics for finite opinion games. *Theor. Comput. Sci.*, 648(C):96–115, October 2016.

[FJ90] Noah E. Friedkin and Eugene C. Johnsen. Social influence and opinions. *The Journal of Mathematical Sociology*, 15(3-4):193–206, 1990.

[Fol13] Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.

[FPS16] Dimitris Fotakis, Dimitris Palyvos-Giannas, and Stratis Skoulakis. Opinion dynamics with local interactions. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 279–285, 2016.

[FS99] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79 – 103, 1999.

[FV97] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1):40 – 55, 1997.

[GS14] Javad Ghaderi and R. Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Automatica*, 50(12):3209–3215, 2014.

[GTT13] Aristides Gionis, Evimaria Terzi, and Panayiotis Tsaparas. Opinion maximization in social networks. In *Proceedings of the 13th SIAM International Conference on Data Mining, May 2-4, 2013. Austin, Texas, USA.*, pages 387–395, 2013.

[HAK07] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192, Dec 2007.

[Haz16] Elad Hazan. Introduction to online convex optimization. *Found. Trends Optim.*, 2(3-4):157–325, August 2016.

[HCM17] Amélie Héliou, Johanne Cohen, and Panayotis Mertikopoulos. Learning with bandit feedback in potential games. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6372–6381, 2017.

[HK02] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal Artificial Societies and Social Simulation*, 5, 2002.

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.

[Jac08] M.O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.

[KPT11] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Load balancing without regret in the bulletin board model. *Distributed Computing*, 24(1):21–29, Sep 2011.

[Kra09] David Krackhardt. A plunge into networks. *Science*, 326(5949):47–48, 2009.

[LW94] Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.

[MM08] Adam Merberg and Steven J Miller. Course notes for math 162: Mathematical statistics the cramér-rao inequality. 2008.

[MMT17] Cameron Musco, Christopher Musco, and Charalampos E. Tsourakakis. Minimizing polarization and disagreement in social networks. *CoRR*, abs/1712.09948, 2017.

[MS17] Panayotis Mertikopoulos and Mathias Staudigl. Convergence to nash equilibrium in continuous games with noisy first-order feedback. In *56th IEEE Annual Conference on Decision and Control, CDC 2017, Melbourne, Australia, December 12-15, 2017*, pages 5609–5614, 2017.

[MU05] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge university press, 2005.

[N$^+$50] John F Nash et al. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.

[PDSS16] Pantelimon G Popescu, Silvestru Dragomir, EMIL I Sluşanschi, and OCTAVIAN N Stănăşilă. Bounds for kullback-leibler divergence. *Electronic Journal of Differential Equations*, 2016, 2016.

[Rao92] C Radhakrishna Rao. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in statistics*, pages 235–247. Springer, 1992.

[RF88] Halsey Lawrence Royden and Patrick Fitzpatrick. *Real analysis*, volume 32. Macmillan New York, 1988.

[Rud06] Walter Rudin. *Real and complex analysis.* Tata McGraw-Hill Education, 2006.

[SA00] Hart Sergiu and Mas-Colell Andreu. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[SALS15] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *NIPS*, pages 2989–2997, 2015.

[Shi96] Albert N Shiryaev. Probability, volume 95 of. *Graduate texts in mathematics*, page 81, 1996.

[Tsy08] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation.* Springer, New York ; London, 1 edition edition, November 2008.

[Vap13] Vladimir Vapnik. *The nature of statistical learning theory.* Springer science & business media, 2013.

[Var10] Richard S Varga. *Geršgorin and his circles*, volume 36. Springer Science & Business Media, 2010.

[Wag93] Stan Wagon. *The Banach-Tarski Paradox*, volume 24. Cambridge University Press, 1993.

[Wal39] Abraham Wald. Contributions to the Theory of Statistical Estimation and Testing Hypotheses. *The Annals of Mathematical Statistics*, 10(4):299–326, December 1939.

[YOA+13] Mehmet Ercan Yildiz, Asuman E. Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Trans. Economics and Comput.*, 1(4):19:1–19:30, 2013.

[Yu97] Bin Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, New York, NY, 1997.

[Zin03] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, pages 928–935. AAAI Press, 2003.

Appendix A

Hoeffding's Inequality

Suppose $\{X_i\}_{i=1}^{\infty}$ is a family of independent random variables with common mean value $\mu \in \mathbf{R}$. For each n, we define the random variable

$$S_n = \sum_{i=1}^{n} X_i$$

We notice that $\mathbf{E}[S_n/n] = \mu$. An early question posed in Probability Theory was to find the limiting behavior of the quantity S_n/n as n becomes large. Intuitively, one could guess that as n tends to infinity, this quantity would get closer and closer to the mean value μ. This observation is made mathematically precise in one of the fundamental results of modern Probability Theory, the *Law of large numbers*.

Theorem 18 (Kolmogorov's law of large numbers)**.** *Suppose $\{X_i\}_{i=1}^{\infty}$ is a family of independent random variables that have the same distribution as a generic random variable X. If $\mu = \mathbf{E}[X]$ exists(infinite values are allowed) then*

$$\mathbf{P}\left[\lim_{n\to\infty} \frac{S_n}{n} = \mu\right] = 1$$

The proof of this important theorem can be found in standard textbooks of measure theoretic probability[Bil08, Çın11, Shi96]. This theorem roughly tells us that as n becomes larger, the distribution of S_n/n becomes more and more concentrated around it's mean value μ. In computer science, sums of independent random variables frequently arise in the analysis of randomized algorithms. Sometimes, the performance of a randomized algorithm crucially depends on how fast the quantity S_n/n concentrates around it's mean. Thus, we need to quantify how fast the convergence is in the law of large numbers. The answer is provided by an elegant inequality, which also holds with minimal assumptions.

Lemma 14 (Hoeffding's Inequality[Hoe63, C$^+$52])**.** *Let $X_1, \ldots, X_t$ be independent random variables such that $0 \le X_i \le 1$ and let $X = (X_1 + \ldots + X_t)/t$. Then for all $t > 0$,*

$$\mathbf{P}[|X - \mathbf{E}[X]| \ge \lambda] \le 2e^{-2n\lambda^2}$$

The proof of this important result can be found in[MU05]. It says that the convergence for bounded random variables is fast. As we explained, this inequality is especially useful when we would like to know how much a sum of random variables deviates from the mean value. A particular setting where it can be used is if all the X_i follow the distribution $B(p)$(bernoulli with probability p). Then, Hoeffding inequality quantifies what is the probability that the mean of the values is close to p. In our problem, this will be useful with p replaced by the frequencies p_{ij} of meeting one's neighbors.