Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Υπολογισμός και Σύγκλιση σε Σημεία Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Κωνσταντίνα Π. Μπαϊρακτάρη

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

Εργαστήριο Λογικής και Επιστήμης Υπολογισμών
Αθήνα, Νοέμβριος 2019

Εθνικό Μετσόβιο Πολυτεχνείο
Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Υπολογισμός και Σύγκλιση σε Σημεία Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Κωνσταντίνα Π. Μπαϊρακτάρη

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

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 6η Νοεμβρίου 2019.

....	....	....
Δημήτριος Φωτάκης	Αριστείδης Παγουρτζής	Ευάγγελος Μαρκάκης
Αν. Καθηγητής Ε.Μ.Π.	Αν. Καθηγητής Ε.Μ.Π.	Επ. Καθηγητής Ο.Π.Α.

Εργαστήριο Λογικής και Επιστήμης Υπολογισμών
Αθήνα, Νοέμβριος 2019

....
Κωνσταντίνα Π. Μπαϊρακτάρη
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Περίληψη

Τα παίγνια διαμόρφωσης άποψης μοντελοποιούν την διαδικασία με την οποία οι άνθρωποι ανταλλάσσουν και διαμορφώνουν απόψεις. Αναλυτικότερα, περιγράφουν τα κοινωνικά σύνολα χρησιμοποιώντας δίκτυα, όπου οι κόμβοι είναι οι άνθρωποι και οι ακμές οι σχέσεις μεταξύ τους και υποθέτουν ότι οι άνθρωποι εκφράζουν ως άποψη ένα σταθμισμένο μέσο όρο των απόψεων των ατόμων με τα οποία έρχονται σε επαφή. Τα συνεξελικτικά παίγνια διαμόρφωσης άποψης συμπεριλαμβάνουν το γεγονός ότι οι σχέσεις μεταξύ των ανθρώπων διαμορφώνονται ταυτόχρονα με τις απόψεις τους.

Σε αυτή τη διπλωματική εργασία ασχολούμαστε με τον υπολογισμό σημείων ισορροπίας σε συνεξελικτικά παίγνια διαμόρφωσης άποψης. Πιο συγκεκριμένα, τα παίγνια που μελετάμε αποτελούν γενίκευση του μοντέλου Friedkin Johnsen στην οποία τα βάρη μεταξύ των παικτών είναι συνάρτηση των απόψεων που εκφράζουν και έχουν πάντα σημείο ισορροπίας γιατί είναι concave n-person games. Ως προς τη σύγκλιση μελετάμε δύο φυσικές δυναμικές, το best response και το follow the leader. Αρχικά, δείχνουμε ότι κάποια στιγμιότυπα τέτοιων παιγνίων μπορεί να έχουν πολλά σημεία ισορροπίας και σε τέτοιες περιπτώσεις το best response δεν συγκλίνει καθολικά. Επίσης, αποδεικνύουμε ότι για συναρτήσεις βαρών που δεν είναι διαφορίσιμες μπορεί το best response να μη συγκλίνει και τοπικά. Στη συνέχεια, χρησιμοποιώντας μία συνάρτηση δυναμικού, δείχνουμε ότι στα concave n-person games των οποίων οι συναρτήσεις κέρδους είναι diagonally strictly concave υπάρχει αλγόριθμος που συγκλίνει στο μοναδικό σημείο ισορροπίας. Στην περίπτωση των συνεξελικτικών παιγνίων διαμόρφωσης άποψης παρατηρούμε ότι αυτός ο αλγόριθμος είναι ισοδύναμος με το follow the leader, το οποίο επίσης δείχνουμε ότι εξασφαλίζει no-regret στους παίκτες που το ακολουθούν.

Λέξεις κλειδιά: Διαμόρφωση Άποψης, Κυρτή Βελτιστοποίηση, Συνεξελικτικά Μοντέλα

Abstract

The opinion formation games model the process according to which humans exchange and form their opinions. In particular, they represent social groups as networks, where the nodes correspond to people and the edges signify the relationships between them, assuming that people compute their opinions as a weighted average of the opinions of the people they encounter. The asymmetric coevolutionary opinion formation games are based on the fact that the relationships between humans evolve simultaneously with their opinions.

In this thesis we focus on the computation of equilibrium points in asymmetric coevolutionary opinion formation games. More specifically, the games we study are a generalization of the Friedkin Johnsen model in which the weights between players are a function of the opinions they express and they always have an equilibrium point since they are concave n-person games. We are particularly interested in the convergence of two natural dynamics, namely best response and follow the leader. Initially, we show that some instances of asymmetric coevolutionary opinions formation games might have multiple equilibria and in such cases the best response dynamic does not converge globally. Moreover, we prove that if the weight functions are not differentiable then best response might not even converge locally. Subsequently, using a potential function we show that there exists a dynamic that converges to the unique equilibrium point of a concave n-person game whose payoff functions are diagonally strictly concave. Particularly, in asymmetric coevolutionary opinion formation games we observe that this algorithm is equivalent to the follow the leader dynamic, which ensures no regret to the players that follow it.

Keywords: Opinion Formation, Convex Optimization, Coevolutionary Models

Ευχαριστίες

Η εκπόνηση της διπλωματικής μου εργασίας σηματοδοτεί την ολοκλήρωση των προπτυχιακών μου σπουδών στη σχολή των Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών. Σε αυτή μου την προσπάθεια υπήρξαν πολλοί που με βοήθησαν και για αυτό θα ήθελα να τους ευχαριστήσω. Ο επιβλέπων καθηγητής κ. Φωτάκης μέσω του μαθήματος των Διακριτών Μαθηματικών μου κίνησε το ενδιαφέρον να ασχοληθώ περαιτέρω με το αντικείμενο της Θεωρητικής Πληροφορικής και με στήριξε σε όλη τη διάρκεια της διπλωματικής με τη συνεχή του βοήθεια. Ιδιαίτερα σημαντική για την πραγματοποίηση της παρούσας εργασίας ήταν η συμβολή των μελών του Εργαστηρίου Λογικής και Επιστήμης Υπολογισμών Στρατή Σκουλάκη και Βαρδή Κανδήρου. Η φοιτητική μου εμπειρία ήταν μοναδική χάρη στους φίλους μου που γνώρισα στη σχολή και μαζί πορευτήκαμε αυτά τα πέντε χρόνια. Επιπρόσθετα θα ήθελα να ευχαριστήσω τους γονείς μου για τη συμπαράσταση και την υπομονή τους. Τέλος, αφιερώνω την διπλωματική μου στη γάτα μου Πολυάννα που μεγαλώσαμε μαζί αλλά δεν πρόλαβε να με δει διπλωματούχο.

Κωνσταντίνα

Contents

Chapter 1

Εκτεταμένη Ελληνική Περίληψη

Στο κεφάλαιο αυτό, συνοψίζουμε το περιεχόμενο της παρούσας διπλωματικής, δίνοντας βασικούς ορισμούς και θεωρήματα, χωρίς αποδείξεις.

1.1 Εισαγωγή

Οι άνθρωποι ως κοινωνικά όντα συζητάνε μεταξύ τους και αλληλεπιδρούν με αποτέλεσμα να διαμορφώνουν απόψεις, οι οποίες μεταβάλλονται καθώς έρχονται σε επαφή με άλλους ανθρώπους. Μια άποψη θα μπορούσε να θεωρηθεί ως ένας πραγματικός αριθμός σε ένα συνεχές διάστημα, όπως για παράδειγμα η βαθμολογία που βάζουν οι καταναλωτές σε ένα προϊόν. Για την μοντελοποίηση αυτής της δυναμικής διαδικασίας έχουν προταθεί μοντέλα που αντιστοιχούν το κοινωνικό σύνολο σε ένα δίκτυο, του οποίου κορυφές είναι οι άνθρωποι και οι ακμές συμβολίζουν τις σχέσεις μεταξύ τους. Τα βάρη των ακμών συμβολίζουν το πόσο εκτιμούν ο ένας τον άλλο και μπορεί οι ακμές να είναι κατυθυνόμενες ή όχι. Τα κύρια μοντέλα διαμόρφωσης άποψης υποθέτουν ότι οι άνθρωποι εκφράζουν ως γνώμη ένα μέσο όρο των απόψεων στις οποίες εκτίθενται.

Το 1990 οι Friedkin και Johnsen πρότειναν ένα μοντέλο [20] στο οποίο κάθε άτομο i έχει μια ιδιωτική άποψη s_i που παραμένει σταθερή και μία δημόσια άποψη x_i που μεταβάλλεται σε διακριτούς γύρους, και οι δύο στο διάστημα $[0, 1]$. Τα βάρη w_{ij} είναι μη αρνητικά και συμβολίζουν το πόσο συμπαθεί το άτομο i τον j. Παράλληλα κάθε άτομο έχει ένα μη μηδενικό βάρος w_i που αντιστοιχεί στο πόσο σίγουρο είναι για την ιδιωτική του άποψη. Επομένως σε κάθε γύρο t κάθε άτομο εκφράζει άποψη

$$x_i^t = \frac{\sum_{j \neq i} w_{ij} x_j^{t-1} + w_i s_i}{\sum_{j \neq i} w_{ij} + w_i}.$$

Αυτή η δυναμική διαδικασία θα καταλήξει σε μία σταθερή κατάσταση στην οποία κάθε άτομο θα έχει μία σταθερή άποψη που δε θα αλλάζει όταν παίρνει τον καινούριο μέσο όρο [21]. Γενικά οι απόψεις στις οποίες θα καταλήξουν τα άτομα μπορεί να είναι διαφορετικές μεταξύ τους λόγω των ιδιωτικών απόψεων που διατηρούν και λαμβάνουν υπόψη.

Το μοντέλο αυτό έχει και ένα ισοδύναμο παίγνιο στο οποίο κάθε άτομο έχει κόστος

$$c_i(\mathbf{x}^t) = \sum_{j \neq i} w_{ij}(x_i^t - x_j^t)^2 + w_i(x_i^t - s_i)^2,$$

το οποίο προκύπτει από το πόσο διαφωνεί με τους γείτονές του στο δίκτυο αλλά και το πόσο απέχει η άποψη που εκφράζει από αυτή που διατηρεί προσωπικά [5]. Αν κάθε παίκτης επιλέγει την άποψη που ελαχιστοποιεί το κόστος που είχε στον προηγούμενο γύρο τότε το παίγνιο θα εξελιχθεί με τον ίδιο τρόπο με τη δυναμική διαδικασία των μέσων όρων και θα καταλήξει σε ισορροπία, στην οποία κανένας δε θα έχει κίνητρο να αλλάξει άποψη δεδομένου ότι οι υπόλοιποι θα μείνουν σταθεροί στις απόψεις τους.

Μία παραλλαγή αυτού του μοντέλου είναι το μοντέλο Hegselmann Krause [23] του οποίου η βασική ιδεά είναι ότι οι άνθρωποι επηρρεάζονται όχι από όλους του άλλους αλλά μόνο από τους ανθρώπους που έχουν κοντινές απόψεις με αυτούς. Πιο συγκεκριμένα, στο γύρο t το άτομο i υπολογίζει τη γειτονιά του $N_i^t = \{j \neq i : |x_i^{t-1} - x_j^{t-1}| < \epsilon\}$ και σύμφωνα με αυτή αποκτά άποψη

$$x_i^t = \frac{\sum_{j \in N_i^t} x_j^{t-1} + s_i}{|N_i^t| + 1}.$$

Το ϵ είναι μία θετική σταθερά που καθορίζει το πόσο μπορεί να απέχουν οι απόψεις από τις οποίες επηρρεάζεται κάποιος. Αυτό το μοντέλο έχει πολλά σημεία ισορροπίας που αντιστοιχούν σε ένα διαχωρισμό των ατόμων σε διακριτές γειτονιές. Επιπρόσθετα η δυναμική διαδικασία σε αυτό το μόντελο πάντα θα συγκλίνει σε κάποιο σημείο ισορροπίας.

Συνήθως οι άνθρωποι γίνονται φίλοι με αυτούς που συμφωνούν και επομένως επηρρεάζονται περισσότερο από τις απόψεις τους. Βέβαια καθώς οι απόψεις τους αλλάζουν επηρρεάζονται και οι φιλίες τους. Τα συνεξελικτικά παίγνια διαμόρφωσης άποψης αποτελούν μία λογική γενίκευση του μοντέλου Friedkin Johnsen στην οποία τα βάρη μεταξύ των ατόμων δεν είναι σταθερά, αλλά εξαρτώνται από τις απόψεις που εκφράζουν οι υπόλοιποι [3]. Αναλυτικότερα, η συνάρτηση κόστους είναι

$$c_i(\mathbf{x}) = (1 - \alpha_i) \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i})(x_i - x_j)^2 + \alpha_i (x_i - s_i)^2$$

στην οποία το βάρος $q_{ij}(\mathbf{x}_{-i})$ μεταξύ του ατόμου i και του ατόμου j είναι συνεχής συνάρτηση όλων των δημόσιων απόψεων εκτός από αυτή του i και της σταθερής ιδιωτικής άποψης s_i. Τα βάρη είναι κανονικοποιημένα, δηλαδή για κάθε i ισχύει $\sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}) = 1$. Επίσης το α_i είναι μία σταθερά στο διάστημα $(0, 1]$ που συμβολίζει την αυτοπεποιήθηση του ατόμου i.

Ερωτήματα και Στόχοι

Το μοντέλο Friedkin Johnsen είναι στατικό με την έννοια ότι το δίκτυο που καθορίζει το πως αλληλεπιδρούν τα άτομα είναι σταθερό όσο εξελίσσεται η διαδικασία. Έχει αποδειχτεί ότι αυτό το μοντέλο έχει μοναδικό σημείο ισορροπίας στο οποίο συγκλίνει γραμμικά [5]. Αντίθετα το μοντέλο Hegselmann Krause και τα συνεξελικτικά παίγνια διαμόρφωσης άποψης είναι δυναμικά αφού κάθε άτομο μπορεί να επηρρεάζεται από διαφορετικούς ανθρώπους σε διαφορετικό βαθμό σε κάθε γύρο. Η δυναμική του μοντέλου Hegselmann Krause συγκλίνει σε σημείο ισορροπίας [2], ωστόσο δεν είναι γνωστό κατά πόσο τα συνεξελικτικά παίγνια υπάρχουν δυναμικές που να έχουν την ίδια ιδιότητα.

Στη διπλωματική αυτή εστιάζουμε στη μελέτη της συμπεριφοράς δύο διασθητικά φυσικών δυναμικών, του best response και του follow the leader, ως προς τη σύγκλιση σε σημεία ισορροπίας σε συνεξελικτικά παίγνια διαμόρφωσης άποψης. Γνωρίζουμε ότι και οι δύο αυτές δυναμικές διαδικασίες συγκλίνουν σε σημείο ισορροπίας στην ειδική περίπτωση που τα βάρη είναι σταθερά και το παίγνιο ανάγεται σε μοντέλο Friedkin Johnsen. Η ύπαρξη σημείου ισορροπίας εξασφαλίζεται από το θεώρημα του Rosen για την ύπαρξη ισορροπίας σε κοίλα παίγνια n-ατόμων [36] επειδή η συνάρτηση κόστους κάθε ατόμου είναι κυρτή ως προς την άποψή του. Αρχικά, μελετάμε τις διάφορες τεχνικές μεθόδους που μπορούν να χρησιμοποιηθούν για να αποδείξουν τη σύγκλιση μιας δυναμικής διαδικασίας. Αυτές συμπεριλαμβάνουν έννοιες από τη γραμμική άλγεβρα, την κυρτή βελτιστοποίηση και την άμεση κυρτή βελτιστοποίηση (online convex optimization). Στη συνέχεια, παραθέτουμε προηγούμενα αποτελέσματα σχετικά με την σύγκλιση δυναμικών σε ισορροπίες κοίλων παιγνίων n-ατόμων. Πιο συγκεκριμένα, παρουσιάζουμε την έννοια της διαγώνιας αυστηρής κοιλότητας που εξασφαλίζει ότι το σημείο ισορροπίας ενός κοίλου παιγνίου ν-ατόμων είναι μοναδικό και μελετάμε ένα δυναμικό σύστημα συνεχούς χρόνου που συγκλίνει στο σημείο ισορροπίας ενός κοίλου παιγνίου n-ατόμων αν οι συναρτήσεις κέρδους είναι διαγώνια αυστηρά κοίλες. Μία ερώτηση που προκύπτει είναι κατά πόσο υπάρχει ένα αντίστοιχο διακριτό δυναμικό σύστημα που να συγκλίνει στην ισορροπία. Δείχνουμε ότι όντως ένα τέτοιο δυναμικό σύστημα υπάρχει που μοιάζει με κάθοδο κλίσης με τη διαφορά ότι κάθε παίκτης βελτιστοποιεί μόνο τη δική του συνάρτηση κέρδους. Επιπρόσθετα, μελετάμε μία υποκλάση των κοίλων παιγνίων n-ατόμων που ονομάζονται κοινωνικά κοίλα παίγνια, για τα οποία αν όλοι οι παίκτες ακολουθούν στρατηγικές no-regret το παίγνιο συγκλίνει σε ισορροπία. Ωστόσο, αποδεικνύουμε ότι στη γενική περίπτωση τα συνεξελικτικά παίγνια διαμόρφωσης άποψης δεν ανήκουν σε αυτή την υποκλάση. Ακολούθως, δείχνουμε ότι αντίθετα με το μοντέλο Friedkin Johnsen στα συνεξελικτικά παίγνια διαμόρφωσης άποψης το best response μπορεί να μη συγκλίνει καθολικά σε στιγμιότυπα που έχουν πολλά σημεία ισορροπίας. Επιπρόσθετα, δείχνουμε ότι για συναρτήσεις βαρών που δεν είναι παραγωγίσιμες μπορεί το best response να μη συγκλίνει και τοπικά σε ένα σημείο ισορροπίας ανεξαρτήτως του πόσο κοντά σε αυτό βρίσκονται οι απόψεις. Παράλληλα, αποδεικνύουμε ότι ο αλγόριθμος Follow The Leader για τα συνεξελικτικά παίγνια διαμόρφωσης άποψης είναι no regret καθώς και ότι είναι ισοδύναμος με το διακριτό δυναμικό σύστημα που συγκλίνει για τα παίγνια των οποίων οι συναρτήσεις κέρδους είναι διαγώνια αυστηρά κοίλες.

1.2 Θεωρητικό Υπόβαθρο

Σε αυτή την παράγραφο παραθέτουμε τις βασικές έννοιες που είναι απαραίτητες για την κατανόηση των μεθόδων που χρησιμοποιούνται για να αποδείξουν τη σύγκλιση μιας δυναμικής διαδικασίας σε ισορροπία.

Θεωρούμε ότι μια δυναμική διαδικασία εξελίσσεται σε διακριτά βήματα κατά τα οποία κάθε πράκτορας επιλέγει κάποια τιμή σύμφωνα με κάποιο κανόνα. Σε αυτή την περίπτωση λέμε ότι το σύστημα είναι σε **ισορροπία** όταν η κατάσταση στην οποία βρίσκεται παραμένει σταθερή στο χρόνο. Από τη σκοπιά της θεωρίας παιγνίων έχουμε n παίκτες, κάθε ένας από τους οποίους επιλέγει μια στρατηγική $\mathbf{x}_i$ από ένα σύνολο S_i. Το διάνυσμα $\mathbf{x}$ περιλαμβάνει το σύνολο των στρατηγικών όλων των παικτών και ανήκει στο καρτεσιανό γινόμενο των επιμέρους συνόλων στρατηγικών. Σε κάθε γύρο του παιγνίου κάθε παίκτης παίζει τη στρατηγική του και έχει κόστος $c_i(\mathbf{x})$. Το σύνολο στρατηγικών $\mathbf{x}^*$ είναι σημείο **ισορ-**

ροπίας Nash για ένα παίγνιο αν για κάθε παίκτη i και κάθε στρατηγική $\mathbf{x}_i \in S_i$ ισχύει

$$c_i(\mathbf{x}^*) \leq c_i(\mathbf{x}_i, \mathbf{x}^*_{-i}).$$

Με άλλα λόγια σε αυτή την κατάσταση κανείς παίκτης δεν έχει κίνητρο να αλλάξει στρατηγική δεδομένου ότι οι υπόλοιποι θα παραμείνουν σταθεροί στις επιλογές τους. Μία στρατηγική λέγεται **best response** για κάποιο παίκτη αν ελαχιστοποιεί το κόστος του δεδομένου ότι οι υπόλοιποι παίκτες δεν αλλάζουν στρατηγική. Πιο συγκεκριμένα, η στρατηγική $\mathbf{x}_i$ είναι η ¨καλύτερη απάντηση' στη στρατηγική $\mathbf{x}_{-i}$ αν για κάθε $\mathbf{s}_i \in S_i$ έχουμε

$$c_i(\mathbf{x}) \leq c_i(\mathbf{s}_i, \mathbf{x}_{-i}).$$

Για να αποδείξουμε ότι μια δυναμική διαδικασία θα συγκλίνει σε ένα σημείο ισορροπίας μπορούμε να χρησιμοποιήσουμε μία **συνάρτηση δυναμικού** που αντιστοιχεί την κατάσταση του συστήματος σε πραγματικές τιμές και της οποίας τα τοπικά ελάχιστα θα αντιστοιχούν στα σημεία ισορροπίας του συστήματος. Επομένως για να δείξουμε ότι η δυναμική διαδικασία θα συγκλίνει σε σημείο ισορροπίας αρκεί να δείξουμε ότι η συνάρτηση δυναμικού θα συγκλίνει σε κάποιο τοπικό ελάχιστο.

Μία άλλη μέθοδος για την απόδειξη σύγκλισης μιας δυναμικής διαδικασίας εκμεταλλέυεται την ιδιότητα **no regret** που περιορίζει το συνολικό κόστος που έχει κάποιος που την ακολουθεί. Πιο αναλυτικά, ένας αλγόριθμος είναι no regret για τον παίκτη που τον ακολουθεί αν η διαφορά του συνολικού κόστους που θα έχει σε T γύρους από το ελάχιστο συνολικό κόστος που θα είχε αν πάντα επέλεγε την ίδια στρατηγική είναι της τάξης $o(T)$. Τυπικά το regret ποσοτικοποιείται ως

$$\text{regret}_T = \sup_{c_1,\ldots,c_T} \left(\sum_{t=1}^{T} c_t(\mathbf{x}_t) - \min_{\mathbf{x}\in K} \sum_{t=1}^{T} c_t(\mathbf{x}) \right),$$

όπου c_t είναι οι συναρτήσεις κόστους.

Ένας ευρέως διαδεδομένος αλγόριθμος που εξασφαλίζει no regret σε κάποιες περιπτώσεις είναι ο **Follow The Leader** σύμφωνα με τον οποίο ο παίκτης επιλέγει τη στρατηγική που θα του ελαχιστοποιούσε το συνολικό κόστος μέχρι το τελευταίο γύρο, δηλαδή

$$\mathbf{x}_T = \arg\min_{\mathbf{x}\in K} \sum_{t=1}^{T-1} c_t(\mathbf{x}).$$

1.3 Κοίλα Παίγνια n-ατόμων

Τα συνεξελικτικά παίγνια διαμόρφωσης άποψης έχουν πάντα σημείο ισορροπίας γιατί ανήκουν στην κατηγορία των κοίλων παιγνίων n ατόμων. Συμφωνα με τον ορισμό του [36] σε ένα κοίλο παίγνιο n ατόμων έχουμε n παίκτες, ο κάθε ένας με μία στρατηγική $\mathbf{x}_i$ που ανήκει σε ένα κυρτό και συμπαγές σύνολο S_i. Το σύνολο των στρατηγικών συμπεριέχεται στο διάνυσμα στρατηγικών $\mathbf{x}$ το οποίο ανήκει στο καρτεσιανό γινόμενο των επιμέρους συνόλων στρατηγικών των παικτών. Στο τέλος κάθε γύρου του παιγνίου κάθε παίκτης έχει κέρδος $u_i(\mathbf{x})$ το οποίο είναι συνεχής συνάρτηση του $\mathbf{x}$ και κοίλη στη στρατηγική $\mathbf{x}_i$ του ίδιου παίκτη.

Θεώρημα 1. *Κάθε κοίλο παίγνιο n ατόμων έχει σημείο ισορροπίας.*

Στα συνεξελικτικά παίγνια διαμόρφωσης άποψης η συνάρτηση κόστους ενός ατόμου είναι συνεχής στο διάνυσμα των απόψεων και κυρτή ως προς τη δική του άποψη. Επομένως το προηγούμενο θεώρημα ισχύει και για αυτά τα παίγνια που πάντα έχουν ισορροπία. Ωστόσο δε γνωρίζουμε πως μπορεί να υπολογιστεί αυτή η ισορροπία. Όπως αναφέραμε και προγουμένως μία τεχνική για να αποδειχθεί η σύγκλιση μιας δυναμικής διαδικασίας είναι με τη χρήση μιας συνάρτησης δυναμικού. Στα συνεξελικτικά παίγνια διαμόρφωσης άποψης είναι σχετικά δύσκολο να κατασκευαστεί μια τέτοια συνάρτηση λόγω της εγγενούς ασυμμετρίας που επιφέρουν οι συναρτήσεις βάρους. Μία υποκλάση κοίλων παιγνίων για τα οποία είναι γνωστή η ύπαρξη συνάρτησης δυναμικού είναι αυτά που έχουν συναρτήσεις κέρδους που είναι διαγώνια αυστηρά κοίλες.

Η έννοια της διαγώνιας αυστηρής κοιλότητας εξασφαλίζει ότι ένα κοίλο παίγνιο έχει μοναδικό σημείο ισορροπίας [36] και ορίζεται με τη χρήση της κλίσης (gradient) των συναρτήσεων κέρδους. Αξίζει να σημειωθεί ότι χρησιμοποιούμε την κλίση μιας συνάρτησης κέρδους ως προς τη στρατηγική του ίδιου πάικτη και τη συμβολίζουμε ως $\nabla_i u_i(\mathbf{x})$.

Ορισμός 1. *Οι συναρτήσεις κέρδους ενός κοίλου παιγνίου n ατόμων είναι διαγώνια αυστηρά κοίλες αν για κάθε δύο διαφορετικά διανύσματα στρατηγικών* $\mathbf{x}^*$ *και* $\bar{\mathbf{x}}$ *ισχύει*

$$\sum_{i=1}^{n}(\mathbf{x}_i^* - \bar{\mathbf{x}}_i)(\nabla_i u_i(\bar{\mathbf{x}}) - \nabla_i u_i(\mathbf{x}^*)) > 0.$$

Για μεγαλύτερη ευκολία μπορούμε για τον ορισμό της διαγώνιας αυστηρής κοιλότητας να πάρουμε το διάνυσμα

$$g(\mathbf{x}) = \begin{pmatrix} \nabla_1 u_1(\mathbf{x}) \\ \nabla_2 u_2(\mathbf{x}) \\ \vdots \\ \nabla_n u_n(\mathbf{x}) \end{pmatrix}$$

και να χρησιμοποιήσουμε τον πίνακα Τζακόμπι αυτού του διανύσματος που συμβολίζουμε με $G(\mathbf{x})$.

Ορισμός 2. *Οι συναρτήσεις κέρδους ενός κοίλου παιγνίου n ατόμων είναι διαγώνια αυστηρά κοίλες αν για κάθε στρατηγική* $\mathbf{x}$ *ο πίνακας* $(G(\mathbf{x}) + G(\mathbf{x})^T)$ *είναι αρνητικά ορισμένος.*

Στην περίπτωση αυτών των παιγνίων ορίζεται ένα δυναμικό συστημα σε συνεχή χρόνο το οποίο συγκλίνει στο μοναδικό σημείο ισορροπίας. Για κάθε πάικτη i

$$\dot{\mathbf{x}}_i(t) = \nabla_i u_i(\mathbf{x}).$$

Είναι σημαντικό να επισημάνουμε ότι αυτό το σύστημα μπορεί να γραφεί και με τη μορφή $\dot{\mathbf{x}}(t) = g(\mathbf{x}(t))$. Αν $g(\mathbf{x}(t)) = 0$, τότε το σύστημα έχει φτάσει σε σημείο ισορροπίας καθώς λόγω της μηδενικής παραγώγου οι στρατηγικές θα παραμείνουν σταθερές.

Θεώρημα 2. *Αν ο* $(G(\mathbf{x})+G(\mathbf{x})^T)$ *είναι αρνητικά ορισμένος για κάθε* $\mathbf{x} \in S$ *τότε η λύση του συστήματος* $\dot{\mathbf{x}}(t) = g(\mathbf{x}(t))$ *συγκλίνει στην ισορροπία Nash για οποιοδήποτε αρχικό σημείο στο* S.

Για να αποδειχτεί το προηγούμενο θεώρημα χρησιμοποιείται ως συνάρτηση δυναμικού η νόρμα του διανύσματος $g(\mathbf{x})$, η οποία όταν μηδενίζεται το σύστημα βρίσκεται σε ισορροπία.

Μία υποκλάση των κοίλων παιγνίων n ατόμων που έχουν σχετικό ενδιαφέρον είναι τα κοινωνικά κοίλα παίγνια. Για να ανήκει σε αυτή την κατηγορία ένα παίγνιο πρέπει να ικανοποιεί δύο συνθήκες. Πρώτα πρέπει να υπάρχει αυστηρός κυρτός συνδυασμός των συναρτήσεων κέρδους όλων των παικτών που να είναι κοίλος στο διάνυσμα $\mathbf{x}$ όλων των στρατηγικών. Επιπλέον πρέπει για κάθε παίκτη η συνάρτηση κέρδους του να είναι κυρτή ως προς τις στρατηγικές των υπόλοιπων παικτών. Σε αυτά τα παίγνια αν κάθε παίκτης ακολουθεί no regret δυναμική τότε ο μέσος όρος των στρατηγικών ως προς το χρόνο θα συγκλίνει στο σημείο ισορροπίας.

Θεώρημα 3. *Αν σε ένα κοινωνικά κοίλο παίγνιο όλοι οι παίκτες ακολουθούν κάποιο no regret αλγόριθμο, όχι απαραίτητα τον ίδιο, τότε ο μέσος όρος των στρατηγικών* $\bar{\mathbf{x}}^T = \frac{1}{T}\sum_{t=1}^{T}\mathbf{x}^t$ *θα συγκλίνει σε σημείο ισορροπίας.*

Ωστόσο αποδεικνύεται ότι στη γενική περίπτωση τα συνεξελικτικά παίγνια διαμόρφωσης άποψης δεν είναι κοινωνικά κοίλα γιατί το μοντέλο Friedkin Johnsen δεν ανήκει σε αυτή την υποκλάση.

1.4 Σύγκλιση σε Σημεία Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

Στα ασύμμετρα συνεξελικτικά παίγνια διαμόρφωσης άποψης κάθε παίκτης έχει μία ιδιωτική άποψη $s_i \in [0,1]$ που παραμένει σταθερή και μία δημόσια άποψη $x_i \in [0,1]$ που μπορεί να μεταβάλλεται. Επίσης έχει ένα παράγοντα αυτοεκτίμησης $\alpha_i \in [0,1]$ καθώς και βάρη $q_{ij}(\mathbf{x}_{-i})$ προς τους άλλους παίκτες που είναι συναρτήσεις των απόψεων των υπόλοιπων παικτών. Στο τέλος κάθε γύρου κάθε παίκτης i έχει κόστος

$$c_i(\mathbf{x}) = (1-\alpha_i)\sum_{j\neq i} q_{ij}(\mathbf{x}_{-i})(x_i - x_j)^2 + \alpha_i(x_i - s_i)^2,$$

όπου τα βάρη είναι κανονικοποιημένα, δηλαδή ισχύει $\sum_{j\neq i} q_{ij}(\mathbf{x}_{-i}) = 1$. Στο best response κάθε παίκτης ελαχιστοποιεί τη συνάρτηση κόστους του προηγούμενου γύρου. Επομένως έχουμε

$$x_i^{t+1} = \min_{x\in[0,1]} c_i(x, \mathbf{x}_{-i}^t) = (1-\alpha_i)\sum_{j\neq i} q_{ij}(\mathbf{x}_{-i}^t)x_j^t + \alpha_i s_i.$$

Το πρώτο ερώτημα που προκύπτει είναι κατά πόσο το best response συγκλίνει καθολικά στο σημείο ισορροπίας. Η απάντηση σε αυτό το ερώτημα είναι ότι δε συγκλίνει γιατί σε περιπτώσεις που υπάρχουν πολλά σημεία ισορροπίας στο παίγνιο μπορεί οι απόψεις κάποιων ατόμων να ταλαντώνονται διαρκώς. Η συμπεριφορά αυτή είναι εμφανής στο γράφημα 1.1α που έχει προκύψει από προσομοίωση ενός συγκεκριμένου παραδείγματος σε Python και δείχνει την μεταβολή των απόψεων δύο ατόμων καθώς η διαδικασία εξελίσσεται στο χρόνο. Επιπρόσθετα το best response μπορεί να μη συγκλίνει τοπικά σε κάποια σημεία ισορροπίας

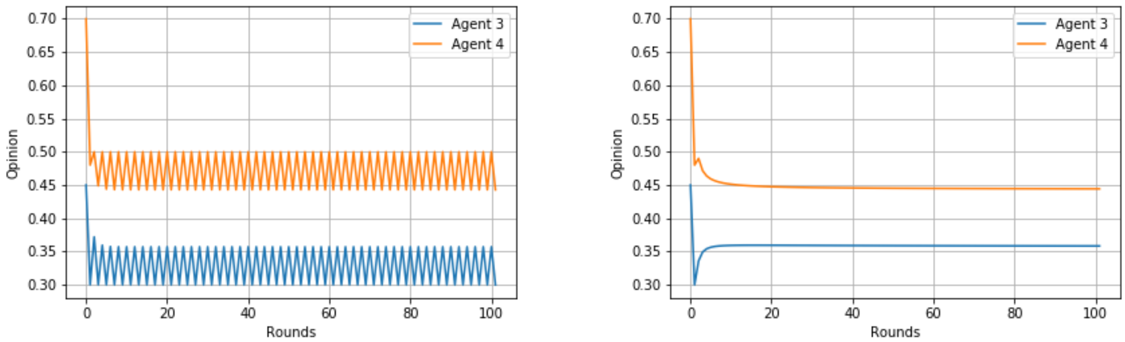

(a) best response (b) follow the leader

Figure 1.1: Σύγκριση των δύο δυναμικών για την ίδια περίπτωση

αν οι συναρτήσεις βάρους δεν είναι παραγωγίσιμες. Με άλλα λόγια όσο κοντά και αν είναι το διάνυσμα των απόψεων στο σημείο ισορροπίας αν όλα τα άτομα ακολουθούν το best response το σύστημα δε θα συγκλίνει σε εκείνο το σημείο ισορροπίας.

Ο αλγόριθμος follow the leader είναι η δεύτερη δυναμική που εξετάζουμε. Σε αυτήν κάθε άτομο ελαχιστοποιεί το συνολικό κόστος που είχε μέχρι τον προηγούμενο γύρο. Επομένως εκφράζει άποψη

$$x_i^{t+1} = \min_{x\in[0,1]} \sum_{\tau=1}^{t} c_i(x, \mathbf{x}_{-i}^{\tau})$$

ή συγκεκριμένα στα συνεξελικτικά παίγνια διαμόρφωσης άποψης

$$x_i^{t+1} = \frac{t}{t+1} x_i^t + \frac{1}{t+1} \left[(1-\alpha_i) \sum_{j\neq i} q_{ij}(\mathbf{x}_{-i}^t) x_j^t + \alpha_i s_i \right].$$

Παρατηρούμε ότι σε αυτή τη δυναμική το άτομο παίρνει ένα μέσο όρο της προηγούμενής του άποψης και των απόψεων των υπόλοιπων ατόμων με αποτέλεσμα η μεταβολή στην άποψη να είναι πιο ομαλή από ότι στο best response. Αυτό έχει ως αποτέλεσμα σε περιπτώσεις που οι απόψεις που προκύπτουν από το best response ταλαντώνονται να συγκλίνουν όταν τα άτομα ακολουθούν τον αλγόριθμο follow the leader, όπως φαίνεται και στο γράφημα 1.1β.

Μια καλή ιδιότητα του follow the leader για τα συγκεκριμένα παίγνια είναι ότι εξασφαλίζει no regret σε όποιο άτομο το ακολουθεί. Πιο συγκεκριμένα, αποδεικνύεται ότι ισχύει

$$\sum_{t=1}^{T} c_i(\mathbf{x}^t) - \min_{x\in[0,1]} \sum_{t=1}^{T} c_t(x, \mathbf{x}_{-i}^t) \leq H_T,$$

όπου H_T είναι ο T-οστός αρμονικός αριθμός. Επιπλέον κίνητρο για την περαιτέρω μελέτη αυτού του αλγορίθμου αποτελούν οι προσομοιώσεις που πραγματοποιήσαμε για διάφορους τύπους συναρτήσεων βαρών και παρατηρήσαμε ότι συγκλίνουν σε σημείο ισορροπίας, όπως στο γράφημα 1.2.

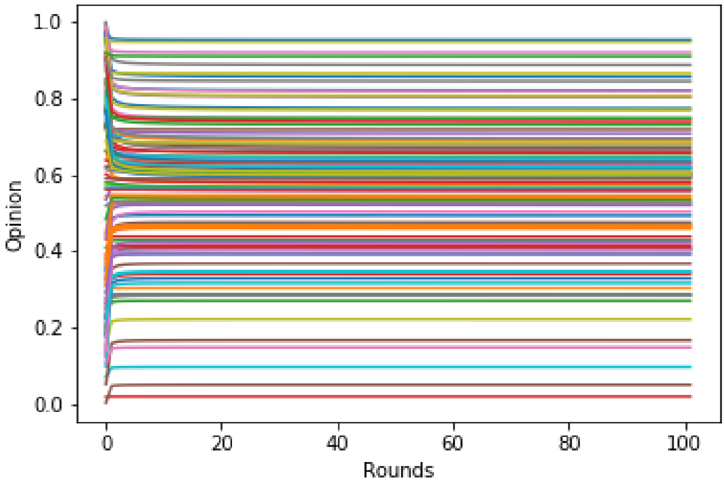

Figure 1.2: Η εξέλιξη των απόψεων σε στιγμιότυπο με 100 άτομα που ακολουθούν το follow the leader

Αν προσπαθήσουμε να χρησιμοποιήσουμε τη συνάρτηση δυναμικού που αναφέραμε για τα διαγώνια αυστηρά κοίλα παίγνια έχουμε

$$g(\mathbf{x}) = \begin{pmatrix} (1-\alpha_1)\sum_{j\neq 1} q_{1j}(\mathbf{x}_{-1})x_j + \alpha_1 s_1 - x_1 \\ \vdots \\ (1-\alpha_n)\sum_{j\neq n} q_{nj}(\mathbf{x}_{-n})x_j + \alpha_n s_n - x_n \end{pmatrix}.$$

Παρατηρούμε ότι ο κανόνας του follow the leader μπορεί να γραφτεί με τη μορφή $\mathbf{x}^t = \mathbf{x}^{t-1} + \frac{1}{t}g(\mathbf{x}^{t-1})$. Ισοδύναμα με το δυναμικό σύστημα συνεχούς χρόνου μπορούμε να παράξουμε το ίδιο θεώρημα για ένα διακριτό σύστημα.

Θεώρημα 4. *Αν η $\|g(\mathbf{x})\|^2$ είναι L-λεία και ο $(G(\mathbf{x}) + G(\mathbf{x})^T)$ είναι αρνητικά ορισμένος, τότε η δυναμική διαδικασία*

$$\mathbf{x}^t = \mathbf{x}^{t-1} + \frac{1}{t}g(\mathbf{x}^{t-1})$$

συγκλίνει στο σημείο ισορροπίας.

Το προηγούμενο θεώρημα δεν ισχύει μόνο για τα συνεξελικτικά παίγνια διαμόρφωσης άποψης που ικανοποιούν τις συνθήκες του αλλά είναι γενικά για τα κοίλα πάιγνια n ατόμων. Επομένως θα μπορούσε να έχει εφαρμογή και σε άλλα είδη παιγνίων που ανήκουν σε αυτή την κλάση.

Chapter 2

Introduction

Humans as social beings, interact with each other constantly and this dynamic process affects many aspects of their behavior, including their opinions regarding issues they are interested in. An opinion can be considered as a variable that either has discrete values, as for example the preferred candidate in an election, or values that lie in a continuous interval, with the prediction of the temperature at a specific time and place in the future being such an example. Although this dynamic procedure has been a well defined field of research for years, the advent of the Internet and the extensive use of social media have facilitated the monitoring of human interactions because social media users declare their friendships and general relationships publicly. In parallel, a group possibly interested in opinion formation would be that of advertisers, who attempt to predict and influence the opinions of the public on a specific product.

A series of opinion formation models, starting with DeGroot's consensus model [15], share some common elements. More specifically, they all describe a process that repeats in discrete steps, during which people, in order to formulate their opinion, take into consideration the weighted average of a subset of opinions. Depending on the model, the definition of the weights and the subset of opinions that influence a person differentiate. Another factor that varies between different models proposed is whether this averaging procedure happens randomly or in a deterministic manner. Therefore, these opinion formation dynamics lie in a greater class of models that imitate intricate natural processes through averaging systems. Other natural procedures displaying similar characteristics relate to animal behavior such as in birds and fish [34, 35, 10]. Additionally, artificial systems such as sensor networks might as well rely on an averaging process that involves communication between agents [9, 8].

The DeGroot model comprises of a fixed network, where every agent communicates with everyone else and studies the attainment of consensus, where eventually all the agents agree in one opinion or opinion disparity, in which case the network splits into groups of similarly thinking agents. This model was extended by Hegselmann and Krause in [23] by introducing the sense of neighborhood that consists of agents with similar opinions. Agents adopt the unweighted average of the opinions of their neighbors including their own opinions. A randomized version of this model is the Deffuant Weisbuch [14, 42], where agents meet in pairs and exchange opinions only if their opinions are close in comparison with a threshold value. A similar model is proposed in [13]. In [25, 26, 16] the models proposed extend the aforementioned models by including the evolution of friendships between people simultaneously with that of their opinion. In [26] the term coevolution was used to describe this phenomenon.

In 1990 Friedkin and Johnsen introduced a model [20] that studied opinion formation instead of the reaching of consensus in the sense that people usually do not agree, but rather form a personal opinion after being influenced by their environment. This process was modeled with the inclusion of a private opinion for each person along with their publicly expressed opinion. Bindel, Kleinberg and Oren presented a game theoretic equivalent of the FJ model in [5] and studied the social cost of disagreement through a measure called Price of Anarchy, that compares the cost of a game in an equilibrium to the optimal cost. This work in PoA bounds was extended for more general cases in [3, 12, 11]. In most models mentioned the opinions are signified by real numbers. However, there are also variants of the FJ model with binary opinions [43, 4, 12], usually 0 and 1.

A question that naturally occurs from the game theoretic view of the FJ model is how will the game evolve if the weights between the agents change as time progresses. In reality, apart from the actual opinions, even the relationships between people are influenced by the opinions they express. In this way, people usually become more fond of individuals with whom they share similar opinions and as a result they are greatly affected by those people's opinions in comparison to other's. The coevolutionary opinion formation game introduced in [3] extends the FJ model, taking into account these varying parameters. In particular, the weight of the influence of agent j to agent i becomes a function of the private opinion of i and the public opinions of all the other agents.

2.1 Opinion Formation Models

As people engage in conversation with their friends or relatives their opinion is shaped gradually. In opinion dynamics the process is modeled as a network of agents who are linked to each other and possess numerical opinions that change according to averaging rules. Multiple models have been proposed, each trying to simulate several natural characteristics of the opinion formation process.

The Friedkin Johnsen Model

In the FJ model [20] there exists an underlying network graph $G(V, E, w)$, where V is the set of the agents, E is the set of the links between pairs of agents and w are the weights of these links. For a pair of agents i and j the weight w_{ij} signifies how much agent j influences agent i during this procedure. It is assumed that all weights between distinct agents are non-negative, whereas for any agent i w_i denotes a positive self-confidence weight. The opinions expressed by the agent i at round t of the process is denoted by $x_i^t \in [0, 1]$. In parallel, each agent i maintains an internal opinion $s_i \in [0, 1]$, that does not change in time. All opinions of the agents are updated at every round according to

$$x_i^t = \frac{\sum_{j \neq i} w_{ij} x_j^{t-1} + w_i s_i}{\sum_{j \neq i} w_{ij} + w_i}.$$

In [21] it was illustrated that this dynamic converges linearly in time to a unique stable point.

Bindel, Kleinberg and Oren proposed in [5] a corresponding game, whose parallel best response is the same as the averaging update rule of the FJ model. In particular, at round t each agent expresses opinion x_i^t, which is her strategy, and incurs a cost

$$c_i(\mathbf{x}^t) = \sum_{j \neq i} w_{ij}(x_i^t - x_j^t)^2 + w_i(x_i^t - s_i)^2.$$

The agents acting selfishly are interested in minimizing the cost they suffer. Indeed, the opinion that minimizes the quadratic equation of the cost is the one suggested by the FJ dynamic. Consequently, the Nash equilibrium of this game coincides with the stable point of the relative dynamic.

The Hegselmann Krause Model

An interesting variation of the previous dynamic is the Hegselmann Krause model introduced in [23]. This model takes into account that people usually are influenced by others that share similar opinions with them. In fact, people tend to form groups of like-minded individuals that affect one another. In the HK model there are n agents, each with an initial opinion $x_i^0 \in [0, 1]$. The dynamic evolves again in discrete time steps. At each round every agent calculates her neighborhood, which includes all the other agents whose opinion is at most $\epsilon > 0$ away from her opinion, formally denoted by $N_i^t = \{j \neq i : |x_i^{t-1} - x_j^{t-1}| < \epsilon\}$. The parameter ϵ symbolizes how open or narrow minded the agents are. Similarly to the FJ model, the new opinion of agent i is derived from the averaging of the opinions of her neighbors. Hence, the opinion update rule for agent i at round t is

$$x_i^t = \frac{\sum_{j \in N_i^t} x_j^{t-1} + s_i}{|N_i^t| + 1}.$$

This dynamic has an infinite number of stable points that correspond to a separation of the agents into groups, such that any two agents of two different groups have opinions that are more than ϵ apart. Moreover, the averaging process always converges to such a stable point in a bounded number of steps ([2],[24], [28], [40], [30], [29]).

The Deffuant Weisbuch Model

The incentive behind the Deffuant Weisbuch model ([14], [42]) is that agents usually exchange opinions in pairs and an agent would discuss only with those who already possess opinions close to hers. The model consists of n agents and each agent i has an opinion $x_i \in [0, 1]$. At each round t two agents i and j are chosen at random. If their opinions are closer than a positive constant d, they update their opinions according to

$$x_i^t = x_i^{t-1} + \mu(x_j^{t-1} - x_i^{t-1})$$

and

$$x_j^t = x_j^{t-1} + \mu(x_i^{t-1} - x_j^{t-1}),$$

where μ is a convergence parameter in $[0, \frac{1}{2}]$. The dynamic process of the DW model always converges to a stable point. For high values of d the opinions converge to an average opinion, whereas for lower values the agents are partitioned in opinion clusters.

2.2 Questions and objectives

The DeGroot and the Friedkin Johnsen models are static in the sense that the underlying network which dictates the interactions between the agents is fixed. These static dynamics converge linearly to a unique equilibrium. Even though the Hegselmann Krause and Deffuant Weisbuch models are dynamic because each agent gets influenced by the other agents in a different degree in every round, they as well converge to an equilibrium relatively fast. However, it is not known whether the asymmetric coevolutionary opinion formation games, whose network also evolves in time, possess the same property.

This thesis focuses on studying whether reasonable dynamics, such as best response and no-regret dynamics, can converge to an equilibrium of an asymmetric coevolutionary opinion formation game. The main incentive for this approach originates in the fact that both best response and no-regret dynamics converge when all the averaging weights of the cost functions are fixed, which is a condition that reduces the game to a Friedkin Johnsen model. The existence of an equilibrium in asymmetrical coevolutionary opinion formation games is ensured by Rosen's theorem of equilibrium existence in concave n-person games [36], because each player's cost function is convex in her opinion. In chapter 4 we will present previous results about convergence properties of dynamics in concave n-person games. In particular, we will present the notion of diagonal strict concavity which ensures that the equilibrium of a concave n-person game is unique and study a continuous dynamic system that converges to the equilibrium of a concave n-person game under the assumption that the payoff functions are diagonally strictly concave. One question that we will attempt to answer in 5 is whether there exists an equivalent discrete dynamic that converges to the equilibrium under similar assumptions. Indeed, we will show that if the payoff functions are diagonally strictly concave and an additional smoothness condition holds, then a dynamic which resembles gradient descent converges to the unique equilibrium. Moreover, in 4 we will study a subclass of concave n-person games, called socially concave games, for which no-regret dynamics converge to an equilibrium. However, we will show that in general asymmetric coevolutionary opinion formation games are not in this subclass.

In chapter 3 we are interested in the theoretical background relevant to the techniques that are commonly used to prove the existence of an equilibrium and the convergence of a dynamic. Chapter 5 is devoted to studying the behavior of dynamics in asymmetric coevolutionary games and more specifically of best response and follow the leader, which is an algorithm commonly used in online convex optimization. We will show using specific examples that unlike in the Friedkin Johnsen model, best response does not converge globally for asymmetric coevolutionary opinion formation games that have multiple equilibria. Furthermore, we will demonstrate that best response does not even converge locally if the set of possible weight functions is unconstrained. Particularly, if the weight functions are not differentiable then best response might not converge to an equilibrium no matter how close to it the starting point of the game is. Regarding follow the leader, we will show that it is no-regret for these games as well as that it is equivalent to the dynamic that converges when the payoff functions of the game are diagonally strictly concave. Finally, we will present a set of simple conditions that would ensure that an asymmetric coevolutionary opinion formation game is diagonally strictly concave.

2.3 Overview of the chapters

This thesis comprises of three chapters. Chapter 3 includes all the basic notions, required for the better understanding of the main problem we consider in this thesis. Firstly, we have presented some basic definitions regarding correspondences, as they are used in the proof of existence of Nash equilibria in particular classes of games. In order to study the convergence of dynamics to equilibria, we present certain commonly used methods starting with basic linear algebra definitions and properties of vectors and square matrices and continue with the presentation of the basic concepts of convex optimization. We are particularly interested in the analysis of the gradient descent algorithm that calculates the optimal value of a convex function, the best response and follow the leader methods as natural strategy choosing behaviors of players in games, as well as the definition of no-regret algorithms.

In chapter 4 we describe a class of games called concave n-person games, which include the coevolutionary opinion formation games we are interested in. In this set of games there are n players, each with a payoff function which is concave in her strategy and continuous. First, we cite the result of [36] according to which these games always admit a Nash equilibrium, thus implying that the coevolutionary opinion formation games have this property too. The proof we present makes use of the Kakutani fixed point theorem. We continue by defining the notion of diagonal strict concavity, which involves the players' payoff functions and ensures that the equilibrium of a concave game is unique. Additionally, we provide a sufficient condition regarding the first derivatives of the payoff functions that guarantees that diagonal strict concavity holds. Following that, we illustrate that there exists a continuous time dynamic that always converges to a Nash equilibrium of such a game under the requirement that the latter condition holds. This is analyzed in steps, firstly showing that the strategies produced by the dynamic remain in the feasible strategy set, secondly proving that the dynamic converges to a stable point and finally showing that the stable point of the dynamic and the Nash equilibrium of the corresponding game coincide. We also present a subclass of concave games called socially concave games, for which it is illustrated that if all the players follow a no-regret algorithm, possibly each player a different one, then the average of the strategies of the players converges to a Nash equilibrium. Finally, we define the coevolutionary opinion formation games and explain why this particular group of games is concave.

In chapter 5 we have included all of our results, focusing on the analysis of the best response and follow the leader algorithms concerning whether they converge to an equilibrium of a coevolutionary opinion formation game. The first question that arose was whether a Nash equilibrium of such a game is unique. Providing simple examples we illustrate that this is not true. In the following sections we are concerned with the behavior of the two algorithms for different types of weight functions. In particular, we start by examining the global convergence of best response and show that it does not hold for either non differentiable or differentiable weight functions. A reasonable following question is whether best response converges locally to a Nash equilibrium, meaning that if the vector of opinions is close to the equilibrium then the dynamic will converge to it. We show that for the example with the non-differentiable function local convergence does not hold, but it does for the second example. As a response to the instability of best response we present the results of simulations showing that opinions expressed according to follow the leader do not fluctuate but rather converge to a stable point. Furthermore, we prove that follow the leader ensures no regret for agents in this game. In the following section we present our main result, which is not limited to the coevolutionary opinion formation games and can be applied to any concave games. Because we observed that the update rule of follow the leader in coevolutionary opinion formation games is equivalent to gradient descent with step size inverse of time, we proved that if the game is diagonally strictly concave and under a Lipschitness assumption this algorithm converges to the unique Nash equilibrium. Finally, we provide sufficient conditions for stict diagonal concavity of a coevolutionary opinion formation game.

Chapter 3

Theoretical Background

In this chapter we will provide the basic toolbox of definitions and theorems, that have been used for the study of the opinion formation models and are required for the comprehension of the subsequent chapters. More specifically, we are interested in the different methods and approaches of studying whether a dynamic converges to an equilibrium point.

3.1 Equilibria

The problems we are going to discuss throughout this thesis have a specific structure. In particular, like in the aforementioned models, there are n agents who adopt a certain value in every discrete round of a process according to some rule. If the values of the agents remain stable, then we say that the system has reached an **equilibrium**.

From a game theoretic point of view, there are n players who want to win in a game that is played in discrete rounds. At each round all the players choose simultaneously each one her own strategy for this round. The state of the game at round t is described by the strategies the players' chose. At the end of each round every player suffers a cost, which is dependent on the strategies of all the players. Every player is knowledgeable about her cost function. We suppose that all the players are selfish and they want to minimize the cost they incur.

In terms of notation, we denote the strategy set from which player i can choose her actions by S_i. The set of the strategy combinations is the Cartesian product of the individual strategy sets, i.e. $S = S_1 \times \ldots \times S_n$. The strategy vector of player i is denoted by $\mathbf{x}_i$ and that of the rest of the players by $\mathbf{x}_{-i}$. Finally, the cost function of player i is signified by $c_i(\mathbf{x})$.

An important notion that is regularly used in this setting is the Nash equilibrium [32]. This is a state of the game in which no player has an incentive to unilaterally deviate from her chosen strategy. Namely, a strategy combination $\mathbf{x}^* \in S$ is a **Nash equilibrium** if for every player i and for any $\mathbf{x}_i \in S_i$

$$c_i(\mathbf{x}^*) \leq c_i(\mathbf{x}_i, \mathbf{x}^*_{-i}).$$

A **best response strategy** of a player i is the one that minimizes her cost, given that the rest of the player will not change their strategies. More formally $\mathbf{x}_i \in S_i$ is a best response to $\mathbf{x}_{-i}$ if for any strategy $\mathbf{s}_i \in S_i$ we have

$$c_i(\mathbf{x}) \leq c_i(\mathbf{s}_i, \mathbf{x}_{-i}).$$

As we observe the notions of the Nash equilibrium and the best response strategy are linked because a joint strategy is a Nash equilibrium if it is the best response strategy to itself for all the players.

A more general notion of equilibrium is that of the **ϵ-Nash equilibrium** that attempts to loosen the conditional inequality. Certainly, there exists the corresponding best response. A strategy combination $\mathbf{x}^* \in S$ is a ϵ-Nash equilibrium if for a non-negative ϵ, for every player i and any $\mathbf{x}_i \in S_i$

$$c_i(\mathbf{x}^*) \leq c_i(\mathbf{x}_i, \mathbf{x}^*_{-i}) - \epsilon.$$

A strategy of a player i $\mathbf{x}_i \in S_i$ is an **ϵ-best response** to $\mathbf{x}_{-i}$ if for a non-negative ϵ and for any strategy $\mathbf{s}_i \in S_i$ we have

$$c_i(\mathbf{x}) \leq c_i(\mathbf{s}_i, \mathbf{x}_{-i}) - \epsilon.$$

Naturally, an ϵ-Nash equilibrium is an ϵ-best response to itself for all the players. Moreover, the Nash equilibrium is a 0-Nash equilibrium.

3.2 Correspondences

An equilibrium of the game could be defined as the fixed point of the update rule. In other words, when the system is at the equilibrium point then the players will adopt the exact same values as in the previous round. Therefore, if we want to examine the convergence of the dynamic to an equilibrium, we first have to prove that such a point exists. One method of proving the existence of a fixed point is by using correspondences.

In some cases there are multiple values that maximize or minimize an expression. For instance, the function $|x - a| + |x + a|$, where a is positive, is minimum for $x \in [-a, a]$. A correspondence maps a point of set X to a subset of set Y. In the previous example we mapped the real value of a to the interval $[-a, a]$. Since we will study optimization problems, we are interested in the structure of the set of optimal solutions. Formally, a **correspondence** ϕ from X to Y associates each point x in X with a subset of Y denoted by $\phi(x)$.

Since the calculation of equilibria in games can also be considered as an optimization problem, correspondences are crucial in determining the existence of an equilibrium. In particular, the Kakutani fixed point theorem is used to prove the existence of a Nash equilibrium. In the case that we will apply it we will also need the **Berge maximum theorem**, that is specific for maximization problems. It states that if $f : X \times Y \to \mathbb{R}$ is a continuous function on $X \times Y$ that we wish to maximize and S is a non-empty compact set that describes the feasible solutions then we can define the correspondence ϕ from X to Y

$$\phi(x) = \arg\max_{y \in S(x)} \{f(x, y)\}$$

The correspondence ϕ is upper semicontinuous, closed and compact-valued. If f is concave in y and S is convex, then ϕ is convex-valued. The **Kakutani fixed point theorem** ([27]) states that if a correspondence from $\mathbf{x} \in S$ to $\phi(\mathbf{x}) \subseteq S$ is upper semicontinuous and the set S is bounded, closed and convex, then there exists $\mathbf{x}^* \in S$ such that $\mathbf{x}^* \in \phi(\mathbf{x}^*)$. More detailed description of the notion of correspondence and the theorems mentioned can be found in [6].

3.3 Linear Algebra

One approach of examining the convergence of a game to an equilibrium would be to check if the distance between the vector of values that the players have chosen and some equilibrium point is decreasing as time progresses. The basic tool for this method is linear algebra, as it defines the notion of distance through norm and it provides helpful inequalities that bound the value of a norm.

First, we will define some basic notions regarding vector and matrix norms. We typically are concerned with vectors and matrices of real numbers. Vectors will be mostly denoted by $\mathbf{x}$ and $\mathbf{y}$ and their coordinates by $x_1, \ldots, x_n$ or $y_1, \ldots, y_n$ respectively. A **vector norm** is a mapping of a vector in $\mathbb{R}^n$ to a real number. The norm of vector $\mathbf{x} \in \mathbb{R}^n$ is denoted by $\|\mathbf{x}\|$ and has the following properties:

1. $\|\mathbf{x}\| \geq 0$
2. $\|c\mathbf{x}\| = |c| \cdot \|\mathbf{x}\|$, where c is a real number
3. $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = 0$
4. $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$, where $\mathbf{y}$ is also a vector in $\mathbb{R}^n$

The norm we will mostly use is the Euclidean norm that is defined as

$$\|\mathbf{x}\|_2 = (\mathbf{x}^T\mathbf{x})^{\frac{1}{2}} = \sqrt{x_1^2 + \ldots + x_n^2}.$$

Unless it is specified otherwise, it is implied that the norm used is the Euclidean. This norm satisfies the Cauchy-Schwartz inequality that is useful for bounding vector inner products.

$$|\mathbf{x}^T\mathbf{y}| \leq \|\mathbf{x}\|_2\|\mathbf{y}\|_2.$$

The **matrix norm** is a mapping of a square matrix with elements that are real numbers to a real number. The matrix norm we will use is defined in correspondence with a vector norm. If $\|\cdot\|$ is a vector norm, then the induced matrix norm of $n \times n$-dimensional matrix A is defined as:

$$\|A\| = \max_{\{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| = 1\}} \|A\mathbf{x}\|.$$

Similarly to the Cauchy-Schwartz inequality, it can be inferred directly from the previous definition that for matrix $A \in \mathbb{R}^{n \times n}$ and $\mathbf{x} \in \mathbb{R}^n$

$$\|A\mathbf{x}\| \leq \|A\|\|\mathbf{x}\|.$$

The maximum magnitude of the eigenvalues of a square matrix A is called the **spectral radius** and is denoted by $\rho(A)$. The spectral radius can be useful when a square matrix of real numbers is also symmetric, since then

$$\|A\|_2 = \rho(A).$$

In this case we can compute the norm of the matrix just by calculating its eigenvalues.

Additionally, the eigenvalues of a symmetric matrix are useful in determining several properties of the matrix. One such property is the definiteness a matrix. A $n \times n$ symmetric matrix A of real numbers is **positive definite** if:

1. for all $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{x} \neq 0$ we have $\mathbf{x}^T A \mathbf{x} > 0$, or

2. all the eigenvalues of A are positive

It is **positive semidefinite** if:

1. for all $\mathbf{x} \in \mathbb{R}^n$ we have $\mathbf{x}^T A \mathbf{x} \geq 0$, or

2. all the eigenvalues of A are non-negative .

Similarly, matrix A is **negative definite** if $-A$ is positive definite and **negative semidefinite** if $-A$ is positive semidefinite.

When a square matrix A is not symmetric we can study the definiteness of matrix $(A + A^T)$, that is symmetric, because $\mathbf{x}^T A \mathbf{x} = \frac{1}{2}\mathbf{x}^T (A + A^T)\mathbf{x}$. Sometimes we are not just interested in whether the eigenvalues are positive or negative, but we need more exact bounds. When all the eigenvalues of a symmetric matrix A are greater than or equal to l we write $A \succeq lI$. Similarly, when all the eigenvalues of A are less than or equal to L we write $A \preceq LI$. One way of bounding the eigenvalues of a square matrix with only non-negative elements is by calculating the row sums of the matrix. In particular, the largest eigenvalue of A is upperbounded by its largest row sum. Books [38] and [1] explain these concepts in further detail.

3.4 Convex Optimization

Another method of showing that a game converges to an equilibrium point involves the use of a potential function. More specifically, a **potential function** is a function that maps the strategies' vector to a real value and decreases when a player chooses a strategy that strictly decreases her cost. Therefore, a local minimum of the potential function signifies an equilibrium point of the game. Hence, showing that by following a specific update rule the players' strategies will converge to an equilibrium is equivalent to proving that the potential function will converge to a local minimum.

Convex optimization provides the tools to study different algorithms that optimize the value of functions that have specific "good" properties. If the potential satisfies such conditions then we can examine whether well known convex optimization algorithms converge to the minimum of the potential function. The term convex may refer to either a set or a function. In this section we start by defining both and continue with explaining notions, necessary for the analysis of the optimization of a convex function in a convex set, as presented in [7], [22] and [41].

A **set** $K \subseteq \mathbb{R}^n$ is **convex** if for every two points $\mathbf{x}$ and $\mathbf{y}$ in K all points described by $\mathbf{x}(\lambda) = \lambda\mathbf{x} + (1-\lambda)\mathbf{y}$, where $\lambda \in [0,1]$, are in set K. Intuitively, a set is convex when every line segment connecting any two points of the set lies entirely in the set.

There are multiple definitions of **convex functions** according to the level of smoothness of the function. The fundamental is the one following. It defines a function as convex if the line connecting any two points of its graph is above or on the graph. Namely, a function $f : K \to \mathbb{R}$, where K is a convex set, is convex if for any $\mathbf{x}, \mathbf{y} \in K$ and $\lambda \in [0,1]$

$$f(\lambda\mathbf{x} + (1-\lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1-\lambda) f(\mathbf{y}).$$

If the inequality above is strict for any $\mathbf{x} \neq \mathbf{y}$, then f is called **strictly convex**.

If the function is differentiable then its **gradient** is denoted by

$$\nabla f(\mathbf{x}) = \left(\frac{\partial f}{\partial x_1}(\mathbf{x}), \ldots, \frac{\partial f}{\partial x_n}(\mathbf{x}) \right)^T.$$

Geometrically, we can say that a function is convex if the tangent space of f at point x lies below the graph of f. The following definition states exactly this, but more formally. A differentiable function $f : K \to \mathbb{R}$, where K is a convex set, is convex if for any $\mathbf{x}, \mathbf{y} \in K$

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}).$$

The equivalent of the second derivative for multiple variable functions is the **Hessian** of f at $\mathbf{x}$, denoted by $\nabla^2 f(\mathbf{x})$.

$$\nabla^2 f(\mathbf{x}) = \begin{pmatrix} \frac{\partial^2 f}{\partial x_1^2}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_1 \partial x_2}(\mathbf{x}) & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n}(\mathbf{x}) \\ \frac{\partial^2 f}{\partial x_2 \partial x_1}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_2^2}(\mathbf{x}) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(\mathbf{x}) \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1}(\mathbf{x}) & \frac{\partial^2 f}{\partial x_n \partial x_2}(\mathbf{x}) & \cdots & \frac{\partial^2 f}{\partial x_n^2}(\mathbf{x}) \end{pmatrix}.$$

In this case, a twice differentiable function $f : K \to \mathbb{R}$, where K is convex and open, is convex if the Hessian $\nabla^2 f(\mathbf{x})$ is positive semidefinite.

Apart from convex functions we will study properties of a relative class of functions, **concave functions**. Suppose set K is convex. Function $f : K \to \mathbb{R}$ is concave if and only if function $-f$ is convex.

As we mentioned earlier, there are functions that are called strongly convex. These can be restricted further by determining how strongly convex they are. A differentiable function $f : K \to \mathbb{R}$, where K is convex, is **l-strongly convex** if for any $\mathbf{x}, \mathbf{y} \in K$

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) + \frac{l}{2}\|\mathbf{x} - \mathbf{y}\|^2.$$

There also exists an equivalent definition for functions that have Hessian matrix. More specifically, a twice differentiable function $f : K \to \mathbb{R}$, where K is convex and open, is l-strongly convex if

$$\nabla^2 f(\mathbf{x}) \succeq lI.$$

Another desirable property of functions is Lipschitz continuity, according to which their value does not change significantly for arguments that are close. Typically, a function is **G-Lipschitz** if for all $\mathbf{x}$, $\mathbf{y} \in K$ we have

$$\|f(\mathbf{y}) - f(\mathbf{x})\| \leq G\|\mathbf{y} - \mathbf{x}\|,$$

where G is a positive constant. If the function is differentiable it is equivalent to say that the gradient of the function is bounded, as the gradient shows how much the function changes at a neighborhood around a specific point. In such case, an equivalent condition is that

$$\|\nabla f(\mathbf{x})\| \leq G.$$

Correspondingly to strong convexity, there exists the notion of smoothness, which restricts the steepness of the function. In particular, a function is **L-smooth** if for any $\mathbf{x}$, $\mathbf{y} \in K$ we have

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) + \frac{L}{2}\|\mathbf{x} - \mathbf{y}\|^2,$$

where L is a positive constant. This is equivalent to saying that the gradient of f is L-Lipschitz and formally it holds that for any $\mathbf{x}$, $\mathbf{y} \in K$

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|.$$

By the conditions of Lipschitz continuity, if the Hessian of function f exists then we can also write that for all $\mathbf{x} \in K$

$$\nabla^2 f(\mathbf{x}) \preceq LI.$$

One of the main reasons we study convex functions is because of their valuable properties regarding minimization. For a convex function f a local minimum is also a global minimum and if f is defined in $\mathbb{R}^n$ and is differentiable then only at the global minimum it holds that $\nabla f(\mathbf{x}) = 0$. Intuitively, $-\nabla f(\mathbf{x})$ shows the direction towards which the function decreases, but if the gradient is zero then there is no such direction. In case function f is defined in a convex set K then only at the minimizer $\mathbf{x}$ it holds that for any $\mathbf{y} \in K$

$$\nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) \geq 0.$$

Again this shows that no matter where $\mathbf{y}$ is in K, if we move towards that point at any length the value of the function will increase.

Consequently, if we have a convex function f that we wish to minimize in $\mathbb{R}^d$, there is an iterative method to do so called **gradient descent**. We start at an initial point $\mathbf{x}_0$ and move with towards the opposite direction of the gradient of f at some specific length. More formally, if η_t is the step size, then

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \nabla f(\mathbf{x}_t).$$

Suppose that f is L-smooth and l-strongly convex. We will analyse the case where $\eta_t = \frac{1}{t}$, but the convergence properties are similar for other choices of η_t as well. Our goal is to show that the difference $f(\mathbf{x}_t) - f(\mathbf{x}^*)$, where $\mathbf{x}^*$ is the minimizer of the function, decreases in every step and converges to zero.

Due to the smoothness and the update rule we obtain

$$\begin{aligned} f(\mathbf{x}_t) &\leq f(\mathbf{x}_{t-1}) + \nabla f(\mathbf{x}_{t-1})^T(\mathbf{x}_t - \mathbf{x}_{t-1}) + \frac{L}{2}\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2 \\ &= f(\mathbf{x}_{t-1}) + \nabla f(\mathbf{x}_{t-1})^T\left(-\frac{1}{t}\nabla f(\mathbf{x}_{t-1})\right) + \frac{L}{2t^2}\|\nabla f(\mathbf{x}_{t-1})\|^2 \\ &= f(\mathbf{x}_{t-1}) - \frac{1}{t}\|\nabla f(\mathbf{x}_{t-1})\|^2 + \frac{L}{2t^2}\|\nabla f(\mathbf{x}_{t-1})\|^2. \end{aligned}$$

For $t \geq L$ it holds that $-\frac{1}{t} + \frac{L}{2t^2} \leq -\frac{1}{2t}$. Therefore, we obtain

$$f(\mathbf{x}_t) \leq f(\mathbf{x}_{t-1}) - \frac{1}{2t}\|\nabla f(\mathbf{x}_{t-1})\|^2.$$

Because of the strong concavity of f we have for any $\mathbf{x}$ and $\mathbf{y}$

$$\begin{aligned} f(\mathbf{y}) &\geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) + \frac{l}{2}\|\mathbf{y} - \mathbf{x}\|^2 \\ &\geq f(\mathbf{x}) - \frac{1}{2l}\|\nabla f(\mathbf{x})\|^2. \end{aligned}$$

This is derived from the minimization of the right term of the inequality over $\mathbf{y}$. Thus, for $\mathbf{x}_{t-1}$ and $\mathbf{x}^*$ we obtain

$$-\|\nabla f(\mathbf{x}_{t-1})\|^2 \leq -2l(f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*)).$$

Combining the above inequalities we have

$$\begin{aligned} f(\mathbf{x}_t) - f(\mathbf{x}^*) &\leq f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*) - \frac{1}{2t}\|\nabla f(\mathbf{x}_{t-1})\|^2 \\ &\leq f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*) - \frac{l}{2t}(f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*)) \\ &= (1 - \frac{l}{2t})(f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*)). \end{aligned}$$

We denote the difference of the value of f at t from the minimum value of f by $h(t) = f(\mathbf{x}_t) - f(\mathbf{x}^*)$ and by τ the minimum t which is at least L. Thus, from the inequality $(1 - x) \leq e^{-x}$ it holds that

$$\begin{aligned} h(t) &\leq \left(1 - \frac{l}{2t}\right)h(t-1) \\ &\leq e^{-\frac{l}{2t}}h(t-1) \\ &= e^{-\frac{l}{2}\sum_{i=\tau}^{t}\left(\frac{1}{i}\right)}h(\tau) \\ &\approx e^{-\frac{l}{2}\ln t}h(\tau) \\ &= \frac{1}{(\sqrt{t})^l}h(\tau). \end{aligned}$$

3.5 Online Convex Optimization

Sometimes intermediate steps might be necessary in order to prove convergence. The opinion formation games also fit in the setting of online convex optimization problems whose goal is to ensure that the regret of a player following a specific algorithm will be sublinear. Therefore, online convex optimization can be used for studying if the game converges to an equilibrium when all players follow a no-regret algorithm.

Online convex optimization is a framework that provides useful tools for solving a particular group of optimization problems with common properties. All these problems can be described as a game between a player and an adversary that evolves in discrete time. At round t the player chooses to play $\mathbf{x}_t$ from a convex set K. Then, the adversary chooses a function $f_t : K \to \mathbb{R}$ from a family of convex functions F and the player incurs cost $f_t(\mathbf{x}_t)$.

After T rounds the player has a regret corresponding to the actions chosen :

$$\text{regret}_T = \sup_{f_1,\dots,f_T} \left(\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in K} \sum_{t=1}^{T} f_t(\mathbf{x}) \right).$$

The aim of these techniques is to achieve regret that is a sublinear function of T. This means that in retrospect the average cost the player suffers is approximately equal to that incurred by choosing the best fixed action, no matter what the choices of the adversary were. More formally, an algorithm is **no-regret** if it achieves

$$\text{regret}_T = o(T).$$

A straightforward approach to this problem is **Follow The Leader**, also known as fictitious play. According to this method at round T the player chooses an $\mathbf{x}_T$ that would minimize the aggregated cost up to round $T-1$. Formally,

$$\mathbf{x}_T = \arg\min_{\mathbf{x}\in K} \sum_{t=1}^{T-1} f_t(\mathbf{x}).$$

Chapter 4

Concave n-person Games

In this chapter we will consider a specific class of games, the concave n-person games. Following [36] we will study the existence of pure Nash equilibria in these games as well as a condition called strict diagonal concavity that affects the equilibrium uniqueness. It shall be noted that the problem setting in [36] is relatively more general than the one presented here. Hence, we will adapt the theorems focusing on a more confined case where each player has her own separate strategy set, similar to that illustrated in [33]. Subsequently, we will study more specific subclasses of concave n-person games, namely socially concave and asymmetric coevolutionary opinion formation games.

In a concave n-person game there are n players. Specifically, player i, where i is an integer in $[1, n]$, chooses a strategy $\mathbf{x}_i$ from a compact and convex subset of $\mathbb{R}^{m_i}$ denoted by S_i, where m_i is a positive integer. If we define vector $\mathbf{x}$ to contain the strategies of all the players, then $\mathbf{x}$ is also in a convex and compact set S that results from the Cartesian product of $S_1, \ldots, S_n$ and is m-dimensional, where $m = \sum_{i=1}^{n} m_i$. In addition, every player i has a payoff function $u_i(\mathbf{x})$ which depends on the strategies of all the players involved in the game. In a concave n-person game all payoff functions $u_i(\mathbf{x})$ are continuous in $\mathbf{x}$ and concave in the strategy vector $\mathbf{x}_i$ of the same player for fixed $\mathbf{x}_{-i}$, which is the vector of strategies of all the other players.

A pure strategy Nash equilibrium of a concave n-person game is a point $\mathbf{x}^*$ in S, such that for every $i \in \{1, \ldots, n\}$

$$u_i(\mathbf{x}^*) = \max_{\mathbf{y}_i \in S_i} \{u_i(\mathbf{y}_i, \mathbf{x}^*_{-i})\}. \tag{4.1}$$

This condition indicates that no player has an incentive to change strategy given that the strategies of all the other players are fixed.

4.1 Existence of Equilibrium in Concave n-person Games

An interesting property of concave n-person games is that they always admit an equilibrium. The proof presented in [36] constructs a mapping from strategies to sets of strategies and uses the Kakutani fixed point theorem to show that the fixed point of the mapping is also an equilibrium of the relative game.

Theorem 4.1. *Every concave n-person game has an equilibrium point.*

Proof. As we mentioned previously a point $\mathbf{x}^*$ is an equilibrium if for all $i \in \{1, \dots, n\}$

$$\mathbf{x}_i^* \in \arg\max_{\mathbf{y}_i \in S_i} \{u_i(\mathbf{y}_i, \mathbf{x}_{-i}^*)\}.$$

Equivalently, we can write

$$\mathbf{x}^* \in \arg\max_{\mathbf{y} \in S} \left\{ \sum_{i=1}^{n} u_i(\mathbf{y}_i, \mathbf{x}_{-i}^*) \right\}$$

because if there existed an i such that the strategy of player i at the equilibrium was not in the set $\arg\max_{\mathbf{y}_i \in S_i} \{u_i(\mathbf{y}_i, \mathbf{x}_{-i}^*)\}$, then we would be able to augment this payoff function and therefore increase the whole sum by choosing a different strategy for i. However, then the equilibrium would not be in $\arg\max_{\mathbf{y} \in S} \{\sum_{i=1}^{n} u_i(\mathbf{y}_i, \mathbf{x}_{-i}^*)\}$, which is contradictory. The set of optimal solutions

$$\phi(\mathbf{x}) = \arg\max_{\mathbf{y} \in S} \left\{ \sum_{i=1}^{n} u_i(\mathbf{y}_i, \mathbf{x}_{-i}) \right\}$$

is a correspondence from S to a subset of S. The sum $\sum_{i=1}^{n} u_i(\mathbf{y}_i, \mathbf{x}_{-i})$ is continuous in $\mathbf{y}$ and $\mathbf{x}$. Since a player's payoff is concave in her own strategy, this sum is concave in $\mathbf{y}$. By the Berge maximum theorem ϕ is upper semicontinuous, closed, compact and convex and by the Kakutani fixed point theorem there exists a fixed point $\mathbf{x}^* \in \phi(\mathbf{x}^*)$, which clearly is a Nash equilibrium point of the game. □

4.2 Diagonal Strict Concavity and Uniqueness of Equilibrium

Concave n-person games may have multiple equilibrium points. As we observed previously, a point $\mathbf{x}^*$ in S that satisfies

$$\sigma(\mathbf{x}^*, \mathbf{r}) = \sum_{i=1}^{n} r_i u_i(\mathbf{x}^*) = \max_{\mathbf{y} \in S} \left\{ \sum_{i=1}^{n} r_i u_i(y_i, \mathbf{x}_{-i}^*) \right\},$$

where r_is are non-negative, is a Nash equilibrium of the game. A sufficient condition that ensures that the equilibrium of such a game is unique is a property of the payoff functions called diagonal strict concavity.

In the case where one player is responsible for maximizing her own payoff function, for instance if she has to choose a pair of $x_1, x_2 \in \mathbb{R}$ such that $u(x_1, x_2) = -(x_1 - 3)^2 - x_2^2$ is maximized, one way of proving that the maximum is unique is to show that the payoff function is concave. By the definition of concave functions we could either show that $u(\mathbf{y}) \leq u(\mathbf{x}) + \nabla u(\mathbf{x})^T(\mathbf{y}-\mathbf{x})$ for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^2$ or that $\nabla^2 u(\mathbf{x})$ is negative definite. However, in concave n-person games, where n is greater than 1, each player influences the payoff of everyone else and therefore, the notion of concavity needs to be generalized.

The definition of the equilibrium suggests that each player attempts to maximize her personal payoff by changing her strategy but has no say in the strategies of the others, that also influence her payoff. For this reason we will use a pseudogradient of $\sigma(\mathbf{x}, \mathbf{r})$

$$g(\mathbf{x}, \mathbf{r}) = \begin{pmatrix} r_1 \nabla_1 u_1(\mathbf{x}) \\ r_2 \nabla_2 u_2(\mathbf{x}) \\ \vdots \\ r_n \nabla_n u_n(\mathbf{x}) \end{pmatrix}.$$

Definition 4.1. *We call $\sigma(\mathbf{x}, \mathbf{r})$ diagonally strictly concave for $\mathbf{x} \in S$ and fixed $\mathbf{r}$ with non-negative elements, if for every pair of distinct points $\mathbf{x}^*, \bar{\mathbf{x}} \in S$ we have*

$$(\mathbf{x}^* - \bar{\mathbf{x}})^T g(\bar{\mathbf{x}}, \mathbf{r}) + (\bar{\mathbf{x}} - \mathbf{x}^*)^T g(\mathbf{x}^*, \mathbf{r}) > 0.$$

Using the Jacobian of vector $g(\mathbf{x}, \mathbf{r})$, denoted by $G(\mathbf{x}, \mathbf{r})$, it can be shown that function $\sigma(\mathbf{x}, \mathbf{r})$ is diagonally strictly concave. Notably, negative definiteness of matrix $(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)$ for $\mathbf{x} \in S$ and fixed $\mathbf{r} > 0$ provides a sufficient condition for $\sigma(\mathbf{x}, \mathbf{r})$ to be diagonally strictly concave.

Theorem 4.2. *If the symmetric matrix $(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)$ is negative definite for $\mathbf{x} \in S$ and fixed $\mathbf{r}$, whose elements are positive, then $\sigma(\mathbf{x}, \mathbf{r})$ is diagonally strictly concave.*

Proof. We can think of $G(\mathbf{x}, \mathbf{r})$ as an equivalent of a Hessian matrix whose definiteness determines the concavity of the initial function $\sigma(\mathbf{x}, \mathbf{r})$. Yet, we can not follow this reasoning in the proof because $g(\mathbf{x}, \mathbf{r})$ is not the actual gradient of $\sigma(\mathbf{x}, \mathbf{r})$. We begin with a property of G and we want to show a corresponding property of g. Because the Jacobian matrix G contains all the first-order partial derivatives of vector g, a way of getting the vector g from the matrix G is by integration.

Let $\mathbf{x}^*$, $\bar{\mathbf{x}}$ be two distinct points in set S for which we want to show strong diagonal concavity. Since S is a convex set, all the points on the line segment connecting them are in S and have $G(\mathbf{x}, \mathbf{r})$ negative definite. Formally these points can be written as

$$\mathbf{x}(\lambda) = \lambda \mathbf{x}^* + (1-\lambda)\bar{\mathbf{x}} \text{ , for } 0 \leq \lambda \leq 1.$$

The difference $g(\mathbf{x}^*, \mathbf{r}) - g(\bar{\mathbf{x}}, \mathbf{r})$ results from integrating on this line segment, i.e. $g(\mathbf{x}^*, \mathbf{r}) - g(\bar{\mathbf{x}}, \mathbf{r}) = \int_{\bar{\mathbf{x}}}^{\mathbf{x}^*} \frac{dg(\mathbf{x}(\lambda), \mathbf{r})}{d\lambda} d\lambda$. Since $G(\mathbf{x}, \mathbf{r})$ is the Jacobian of vector $g(\mathbf{x}, \mathbf{r})$, we can take the derivative of $g(\mathbf{x}(\lambda), \mathbf{r})$ with respect to λ and obtain

$$\frac{dg(\mathbf{x}(\lambda), \mathbf{r})}{d\lambda} = G(\mathbf{x}(\lambda), \mathbf{r}) \frac{d\mathbf{x}(\lambda)}{d\lambda} = G(\mathbf{x}(\lambda), \mathbf{r})(\mathbf{x}^* - \bar{\mathbf{x}}).$$

In order to get the condition wanted we need to multiply the difference of $g(\mathbf{x}^*, \mathbf{r}) - g(\bar{\mathbf{x}}, \mathbf{r})$ by $(\bar{\mathbf{x}} - \mathbf{x}^*)^T$.

$$
\begin{aligned}
(\bar{\mathbf{x}} - \mathbf{x}^*)^T(g(\mathbf{x}^*, \mathbf{r}) - g(\bar{\mathbf{x}}, \mathbf{r})) &= -\int_{\bar{\mathbf{x}}}^{\mathbf{x}^*} (\mathbf{x}^* - \bar{\mathbf{x}})^T G(\mathbf{x}(\lambda), \mathbf{r})(\mathbf{x}^* - \bar{\mathbf{x}}) d\lambda \\
&= -\frac{1}{2}\int_{\bar{\mathbf{x}}}^{\mathbf{x}^*} (\mathbf{x}^* - \bar{\mathbf{x}})^T [G(\mathbf{x}(\lambda), \mathbf{r}) + G(\mathbf{x}(\lambda), \mathbf{r})^T](\mathbf{x}^* - \bar{\mathbf{x}}) d\lambda > 0.
\end{aligned}
$$

For the last two steps we used $(\mathbf{x}^* - \bar{\mathbf{x}})^T G(\mathbf{x}(\lambda), \mathbf{r})(\mathbf{x}^* - \bar{\mathbf{x}}) = \frac{1}{2}(\mathbf{x}^* - \bar{\mathbf{x}})^T(G(\mathbf{x}(\lambda), \mathbf{r}) + G(\mathbf{x}(\lambda), \mathbf{r})^T)(\mathbf{x}^* - \bar{\mathbf{x}})$ and the fact that $G(\mathbf{x}(\lambda), \mathbf{r}) + G(\mathbf{x}(\lambda), \mathbf{r})^T$ is negative definite. As a result $\sigma(\mathbf{x}, \mathbf{r})$ is diagonally strictly concave. □

There are fairly simple examples where diagonal strict concavity does not hold. For instance, in a game with two players player 1 has strategy x_1, player 2 has strategy x_2 and their payoff functions are $u_1(x_1, x_2) = -x_1^2 + 6x_1x_2$ and $u_2(x_1, x_2) = -\frac{1}{2}x_2^2 + x_1x_2$ respectively. The function u_1 is concave in x_1 and u_2 is concave in x_2. For this game we have

$$
g(\mathbf{x}) = \begin{pmatrix} -2x_1 + 6x_2 \\ -x_2 + x_1 \end{pmatrix}
$$

and

$$
G(\mathbf{x}) = \begin{pmatrix} -2 & 6 \\ 1 & -1 \end{pmatrix},
$$

which is not negative definite.

Finally, we can analyse the main idea of this section, which is that diagonal strict concavity ensures the uniqueness of the equilibrium of a concave game.

Theorem 4.3. *If $\sigma(\mathbf{x}, \mathbf{r})$ is diagonally strictly concave for $\mathbf{x} \in S$ and fixed $\mathbf{r}$, whose elements are non-negative, then the equilibrium point of the game is unique.*

Proof. Let $\mathbf{x}^*$ be an equilibrium point of the game. In such case, point $\mathbf{x}^*$ is an optimal solution of the convex optimization problem for any $i = \{1, \ldots, n\}$:

$$
\begin{array}{cc}
\text{maximize} & u_i(\mathbf{y}_i, \mathbf{x}^*_{-i}) \\
\text{s.t.} & \mathbf{y}_i \in S_i.
\end{array}
$$

Since function $u_i(\mathbf{y}_i, \mathbf{x}^*_{-i})$ is concave in $\mathbf{y}_i$ and $\mathbf{x}^*_i$ is an optimal solution of the problem above, we can infer that at this point the inner product between the gradient and the direction towards an interior point of S_i is non-positive. The gradient at point $\mathbf{x}^*_i$ defines a hyperplane in S_i. If this inner product were to be positive that would mean that the two vectors are in the same side of the hyperplane and therefore, by following the direction of the projected gradient we would be able to further augment the objective function. However, this is not possible at a maximum point. Therefore, for any $\mathbf{y}_i \in S_i$ we have :

$$
\nabla_i u_i(\mathbf{x}^*)^T(\mathbf{y}_i - \mathbf{x}^*_i) \leq 0.
$$

Formally, this is a result that is derived from the Karush-Kuhn-Tucker conditions.

For a positive number r_i, we can multiply the inequality and obtain $r_i \nabla_i u_i(\mathbf{x}^*)^T(\mathbf{y}_i - \mathbf{x}_i^*) \leq 0$. The previous inequality is valid for any i. Hence, by summing for all $i = 1, \ldots, n$ we maintain the inequality sign.

$$\sum_{i=1}^{n} r_i \nabla_i u_i(\mathbf{x}^*)^T(\mathbf{y}_i - \mathbf{x}_i^*) \leq 0.$$

Let $\bar{\mathbf{x}}$ be also an equilibrium point of the game that differs from $\mathbf{x}^*$. The procedure above can be applied to this point as well. As a result,

$$\sum_{i=1}^{n} r_i \nabla_i u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_i - \mathbf{x}_i^*) + \sum_{i=1}^{n} r_i \nabla_i u_i(\bar{\mathbf{x}})^T(\mathbf{x}_i^* - \bar{\mathbf{x}}_i) \leq 0.$$

If we rewrite it in a more compact way we have

$$(\bar{\mathbf{x}} - \mathbf{x}^*)^T g(\mathbf{x}^*, \mathbf{r}) + (\mathbf{x}^* - \bar{\mathbf{x}})^T g(\bar{\mathbf{x}}, \mathbf{r}) \leq 0,$$

where $\mathbf{r}$ is the vector of elements r_i. By the assumptions of the theorem this is a contradiction. Therefore, diagonal strict concavity ensures the uniqueness of the equilibrium point of the game.

□

One example where diagonal strict concavity holds is the Friedkin Johnsen model. The theorem above provides an alternative proof for the uniqueness of its equilibrium. In particular, the pseudogradient of the payoff functions, if we consider as payoff the opposite of cost, is

$$g(\mathbf{x}) = \begin{pmatrix} 2(1-\alpha_1)\sum_{j \neq 1} w_{1j} x_j + 2\alpha_1 s_1 - 2x_1 \\ \vdots \\ 2(1-\alpha_n)\sum_{j \neq n} w_{nj} x_j + 2\alpha_n s_n - 2x_n \end{pmatrix}.$$

By differentiating the payoff functions once more we get the matrix G

$$G(\mathbf{x}) = \begin{pmatrix} -2 & 2(1-\alpha_1)w_{12} & \ldots & 2(1-\alpha_1)w_{1n} \\ 2(1-\alpha_2)w_{21} & -2 & \ldots & 2(1-\alpha_2)w_{2n} \\ \vdots & & & \\ 2(1-\alpha_n)w_{n1} & 2(1-\alpha_n)w_{n2} & \ldots & -2 \end{pmatrix}.$$

We can rewrite this matrix, using an auxiliary matrix W whose diagonal elements are zero and the others are $W_{ij} = (1-\alpha_i)w_{ij}$. Consequently, the matrix G is $G(\mathbf{x}) = 2(W - I)$, which is negative definite if all α_i are positive. More specifically, this derives from the fact that the rows of matrix W sum up to a value less than 1. Therefore, the Friedkin Johnsen model corresponds to a diagonally strictly concave game with unique equilibrium. The definiteness is relatively straightforward to study in this case. However, even if we use the first definition that includes only the vector g the condition is reduced to the same inequality. Specifically for this game the condition is for any $\mathbf{x}^* \neq \bar{\mathbf{x}}$

$$\sum_{i=1}^{n}(x_i^* - \bar{x}_i)\left[2(1-\alpha_i)\sum_{j\neq i} w_{ij}\bar{x}_j - 2\bar{x}_i - 2(1-\alpha_i)\sum_{j\neq i} w_{ij}x_j^* + 2x_i^*\right] > 0$$

or

$$\sum_{i=1}^{n}(x_i^* - \bar{x}_i)^2 > \sum_{i=1}^{n}(x_i^* - \bar{x}_i)\left(\sum_{j\neq i}(1-\alpha_i)w_{ij}(x_j^* - \bar{x}_j)\right),$$

which by definition is equivalent to $I \succ W$.

4.3 Stability of the Unique Equilibrium Point

In [36] the author suggests a dynamic model that describes a natural behavior of the players in a concave n-person game. The reasoning behind this model is that every player, given the strategies of the other players, attempts to maximize her payoff by choosing a feasible strategy from her strategy set. The rate according to which every player changes her strategy depends on the gradient of the function $u_i(\mathbf{x})$ she wishes to maximize over $\mathbf{x}_i$, since he gradient of the function shows the direction towards which the function increases. The feasible set of joint strategies can be defined as $S = \{\mathbf{x} | \forall j \in \{1, \ldots, k\} : h_j(\mathbf{x}) \geq 0\}$, where for any $j = 1, \ldots, k$ the function $h_j(\mathbf{x})$ is concave. If the strategy of the players is not in S, then it must move towards a direction that increases the value of the constraint functions that are not strictly satisfied. Consequently, a term is added to the strategy change rate for every constraint function not strictly satisfied, that is proportionate to the gradient of the relative constraint function. The differential equations of the dynamic system are for all $i = 1, \ldots, n$

$$\frac{d\mathbf{x}_i}{dt}(t) = \dot{\mathbf{x}}_i(t) = r_i \nabla_i u_i(\mathbf{x}(t)) + \sum_{j=1}^{k} \lambda_j(\mathbf{x}(t)) \nabla_i h_j(\mathbf{x}(t)), \tag{4.2}$$

where all $\lambda_j(\mathbf{x}(t))$ are non-negative and will be defined later and all r_i are positive constants. The first term of the right side of the equation expresses the maximization of the payoff function of player i and the second term ensures that strategy $\mathbf{x}$ is feasible.

The system above can be rewritten in a more compact way using the notation $H(\mathbf{x}) = [\nabla h_1(\mathbf{x}), \ldots, \nabla h_k(\mathbf{x})]$ and $\lambda(\mathbf{x}) = [\lambda_1(\mathbf{x}), \ldots, \lambda_k(\mathbf{x})]^T$

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \lambda(\mathbf{x}(t)), \mathbf{r}) = g(\mathbf{x}(t), \mathbf{r}) + H(\mathbf{x}(t))\lambda(\mathbf{x}(t)). \tag{4.3}$$

We will assume that both $g(\mathbf{x}, \mathbf{r})$ and $H(\mathbf{x})$ are continuous in $\mathbf{x}$.

The parameter vector $\lambda(\mathbf{x})$ lies in the set $\Lambda(\mathbf{x}) \subset \mathbb{R}^k$

$$\Lambda(\mathbf{x}) = \underset{l_j \geq 0, \text{ if } h_j(\mathbf{x}) \leq 0,\ l_j = 0, \text{ otherwise}}{\arg\min} \|f(\mathbf{x}, \mathbf{l}, \mathbf{r})\|.$$

The aim of this set is to ensure that the strategy vector is not moved more than needed towards a direction that would make it unfeasible. That is why we minimize $\dot{\mathbf{x}}(t)$ as much as possible. If one of the constraint functions is zero, this means that the strategy vector is on the bound of the set. Therefore, we still need to make sure it will not move outside of S. If $\mathbf{x}$ is a strictly feasible point, then there is no need to change the value of the constraint functions and the only concern of the players is to maximize their payoff. Hence, in this case $\Lambda(\mathbf{x}) = \{0\}$ and $f(\mathbf{x}, \lambda, \mathbf{r}) = g(\mathbf{x}, \mathbf{r})$.

If all the initial strategies are feasible, then the strategies described by the system of differential equations 4.3 always remain in S. Therefore, the constrain terms of the differential equations are successful in keeping the strategies in S.

Theorem 4.4. *If $g(\mathbf{x}, \mathbf{r})$ and $H(\mathbf{x})$ are continuous in $\mathbf{x}$ in compact set $\bar{S} \supset S$, then there exists a continuous solution $\mathbf{x}(t)$ of*

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \lambda(\mathbf{x}(t)), \mathbf{r}) = g(\mathbf{x}(t), \mathbf{r}) + H(\mathbf{x}(t))\lambda(\mathbf{x}(t)),$$

such that starting at a point $\mathbf{x}$ in S the trajectory of strategies $\mathbf{x}(t)$ will remain in S for all $t > 0$.

Proof. First, we need to show that this system of differential equations has a solution. Assuming that $\lambda(\mathbf{x}(t))$ is measurable in t and f is continuous in $\mathbf{x}$, by the Caratheodory's existence theorem there exists a continuous solution $\mathbf{x}(t)$ in $\bar{S}$ that satisfies the differential equations almost everywhere.

Second, we want to prove that starting from a strategy vector in S the system will never leave S. Let $\bar{\mathbf{x}}$ be the first point on the transit of the strategy vector that is external to S. For this point there exists a constraint that is not satisfied, i.e. there exists a $l \in \{1, \ldots, k\}$ such that $h_l(\bar{\mathbf{x}}) < 0$. Due to the continuity of $\mathbf{x}(t)$ there exists a "previous" point $\tilde{\mathbf{x}}$ that is on the border of S because $h_l(\bar{\mathbf{x}}) = 0$ and $\frac{d}{dt}h_l(\tilde{\mathbf{x}}) < 0$. If we assume that the system is at point $\tilde{\mathbf{x}}$ at time $t_1 > 0$, then by the chain rule we obtain

$$\frac{d}{dt}h_l(\mathbf{x}(t_1)) = \nabla h_l(\mathbf{x}(t_1))\frac{d}{dt}\mathbf{x}(t_1) < 0.$$

By substituting $\frac{d}{dt}\mathbf{x}(t_1)$ with $g(\mathbf{x}(t_1), \mathbf{r}) + H(\mathbf{x}(t_1)\lambda(\mathbf{x}(t_1)))$ we end up with

$$\frac{d}{dt}h_l(\mathbf{x}(t_1)) =$$

$$\nabla h_l(\mathbf{x}(t_1))g(\mathbf{x}(t), \mathbf{r}) + \sum_{j \neq l} \lambda_j(\mathbf{x}(t_1))\nabla h_l(\mathbf{x}(t_1))^T \nabla h_j(\mathbf{x}(t_1)) + \lambda_l(\mathbf{x}(t_1))\|\nabla h_l(\mathbf{x}(t_1))\|^2 < 0.$$

This implies that we could increase $\frac{d}{dt}h_l(\mathbf{x}(t_1))$ by choosing the appropriate λ_l and thereby there was no need to move to a not feasible strategy. We can see how the strategy change rate is influenced by the choice of λ_l following the steps below.

Formally, we have

$$\|\dot{\mathbf{x}}(t)\|^2 = \|g(\mathbf{x}(t))\|^2 + 2(H(\mathbf{x}(t))\lambda(\mathbf{x}(t)))^T g(\mathbf{x}(t)) + \|H(\mathbf{x}(t))\lambda(\mathbf{x}(t))\|^2.$$

The derivative of $\|\dot{\mathbf{x}}(t_1)\|^2$ over λ_l is

$$\frac{d\|\dot{\mathbf{x}}(t_1)\|^2}{d\lambda_l} = 2\nabla h_l(\mathbf{x}(t_1))^T [g(\mathbf{x}(t_1)) + H(\mathbf{x}(t_1))\lambda(\mathbf{x}(t_1))] = \nabla h_l(\mathbf{x}(t_1))\frac{d}{dt}\mathbf{x}(t_1) < 0.$$

Therefore, by increasing the value of λ_l we can decrease the $\|\dot{\mathbf{x}}(t)\|$. This is contradictory, because by the definition of $\Lambda(\mathbf{x}(t_1))$ we have already chosen the λ_l that minimizes $\|\dot{\mathbf{x}}(t)\|$. Consequently, a point $\bar{\mathbf{x}}$ such that $h_l(\bar{\mathbf{x}}) < 0$ can not exist in the trajectory of $\mathbf{x}(t)$.

□

Following our observations of the previous proof the appropriate λ should satisfy for any $j \in \{1, \ldots, k\}$, such that $h_j(\mathbf{x}(t)) \leq 0$

$$\frac{d\|\dot{\mathbf{x}}(t)\|^2}{d\lambda_j} = 2\nabla h_j(\mathbf{x}(t))^T[g(\mathbf{x}(t)) + H(\mathbf{x}(t))\lambda(\mathbf{x}(t))] = 0.$$

This means that if we denote by $\bar{H}$ the matrix that contains only the lines of matrix H that correspond to the constraints that are not strictly satisfied, then for the vector λ we have

$$\bar{H}(\mathbf{x}(t))^T \left[g(\mathbf{x}(t)) + \bar{H}(\mathbf{x}(t))\lambda(\mathbf{x}(t))\right] = 0$$

and therefore

$$\lambda(\mathbf{x}(t)) = -[\bar{H}(\mathbf{x}(t))^T\bar{H}(\mathbf{x}(t))]^{-1}\bar{H}(\mathbf{x}(t))^T g(\mathbf{x}(t))$$

which is non-negative.

According to basic optimization conditions, the value of g at a Nash equilibrium is zero, since at such a point every player has maximized her payoff in relation to the strategies of the other players. In fact, it is apparent that for this reason if the dynamic system reaches a Nash equilibrium then it will remain there.

We will focus on the case where the constraint functions are not necessary and thus the relative terms can be removed from the dynamic. Therefore, in order to show that regardless of the starting point the system will converge to a Nash equilibrium point of the corresponding game, we have to prove that the value of g converges to zero.

Theorem 4.5. *If $(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)$ is negative definite for any $\mathbf{x}$ in S then the solution of the system of differential equations $\dot{\mathbf{x}}(t) = g(\mathbf{x}(t), \mathbf{r})$ converges to a Nash equilibrium point for any initial point in S.*

Proof. We will show that the norm of $g(\mathbf{x}(t), \mathbf{r})$ always decreases in time if it is not zero.

$$\frac{1}{2}\frac{d}{dt}\|g(\mathbf{x}(t), \mathbf{r}))\|^2 = \frac{1}{2}\frac{d}{dt}(g(\mathbf{x}(t), \mathbf{r})^T g(\mathbf{x}(t), \mathbf{r}) = g(\mathbf{x}(t), \mathbf{r})^T \frac{d}{dt}(g(\mathbf{x}(t), \mathbf{r}).$$

By the definition of the Jacobian matrix it holds that $\frac{d}{dt}(g(\mathbf{x}(t), \mathbf{r}) = G(\mathbf{x}(t), \mathbf{r})\frac{d}{dt}\mathbf{x}(t)$. Furthermore, the differential equations state that $\frac{d}{dt}\mathbf{x}(t) = g(\mathbf{x}(t), \mathbf{r})$. As a result we obtain

$$\frac{1}{2}\frac{d}{dt}\|g(\mathbf{x}(t), \mathbf{r}))\|^2 = g(\mathbf{x}(t), \mathbf{r})^T G(\mathbf{x}(t), \mathbf{r}) g(\mathbf{x}(t), \mathbf{r}).$$

Because for any square $n \times n$ matrix A and any n-dimensional vector $\mathbf{x}$ we have $\mathbf{x}^T A\mathbf{x} = \mathbf{x}^T A^T \mathbf{x}$, it holds that

$$g(\mathbf{x}(t), \mathbf{r})^T G(\mathbf{x}(t), \mathbf{r}) g(\mathbf{x}(t), \mathbf{r}) = \frac{1}{2} g(\mathbf{x}(t), \mathbf{r})^T (G(\mathbf{x}(t), \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T) g(\mathbf{x}(t), \mathbf{r}),$$

which is negative because of the conditions of the theorem. Actually there exists a positive constant m, such that

$$g(\mathbf{x}(t),\mathbf{r})^T(G(\mathbf{x}(t),\mathbf{r})+G(\mathbf{x},\mathbf{r})^T)g(\mathbf{x}(t),\mathbf{r}) \leq -m\|g(\mathbf{x}(t),\mathbf{r})\|^2.$$

Combining all of the above we obtain

$$\frac{1}{2}\frac{d}{dt}\|g(\mathbf{x}(t),\mathbf{r}))\|^2 \leq -m\|g(\mathbf{x}(t),\mathbf{r})\|^2.$$

Let $h(t)$ be a real-valued function for non-negative t. If $\frac{d}{dt}h(t) = -mh(t)$, then $h(t) = h(0)e^{-mt}$ which means that $\lim_{t\to\infty} h(t) = 0$. Hence, $\lim_{t\to\infty} \|g(\mathbf{x}(t),\mathbf{r}))\|^2 = 0$, which means that in time the dynamic system will converge to a Nash equilibrium. This equilibrium is unique though, because of the negative definiteness of $G(\mathbf{x},\mathbf{r})^T + G(\mathbf{x},\mathbf{r})$. Therefore, the system converges to the unique Nash equilibrium of the corresponding game no matter the starting point. □

In a similar way in [36] it is proved that the constraints do not affect this convergence result.

4.4 Socially Concave Games

We previously studied diagonally strictly concave games, as a subclass of concave n-person games. In this section we will describe another subclass, called socially concave games. What is interesting about these games, is that according to [17] and [31] there exist natural discrete time dynamics that converge to the equilibrium point of the game. Particularly, if all players follow a no-regret algorithm, then the average vector of their strategies converges to a Nash equilibrium point and the average of the payoff of every player converges to her payoff at that point.

Socially concave games also consist of n players who are trying to maximize their payoff functions. Each player i plays a strategy $\mathbf{x}_i$ in $S_i \subseteq \mathbb{R}^{m_i}$ and gains payoff $u_i(\mathbf{x})$, which is continuous in $\mathbf{x} \in S_i$. Socially concave games specifically satisfy two conditions, the first one being that there exists a strict convex combination of the payoff functions that is concave in the strategy vector $\mathbf{x}$. To be precise, this reminds us of $\sigma(\mathbf{x},\mathbf{r}) = \sum_{i=1}^{n} r_i u_i(\mathbf{x})$ under the condition that all $r_i > 0$ and that they are normalized. This function was extensively used in strictly diagonally concave games. The second condition concerns each payoff function separately and states that all players' payoff functions are convex in the strategies of the other players. By this definition it can be inferred that socially concave games constitute a subclass of concave n-person games. Therefore, by theorem 4.1 they always admit a Nash equilibrium point.

A straightforward example of a socially concave game consists of two players with strategies x_1 and x_2. Their payoff functions are $u_1(x_1,x_2) = -x_1^2 + x_2^2$ and $u_2(x_1,x_2) = x_1^2 - x_2^2$. The function u_1 is concave in x_1 and convex in x_2, whereas the function u_2 is concave in x_2 and convex in x_1. Their convex combination $\frac{1}{2}u_1(x_1,x_2) + \frac{1}{2}u_2(x_1,x_2)$ is equal to zero and hence, because it is constant it is concave.

Theorem 4.6. *All socially concave games are concave n-person games.*

Proof. We need to show that each player's payoff function is concave in her own strategy. We will examine the behavior of all the payoff functions for the strategy of a fixed player $i \in \{1, \ldots, n\}$. All the payoff functions of the other players are convex in i's strategy. Therefore, $\sum_{j \neq i} r_j u_j(\mathbf{x})$ is convex in i's strategy. Directly by the first condition we obtain that $\sum_{j=1}^{n} r_j u_j(\mathbf{x})$ is concave in i's strategy. Therefore, the difference between the two terms $r_i u_i(\mathbf{x})$ is concave in strategy x_i. Given a socially concave game, all the players' payoff functions are concave in the same player's strategy. Hence, a socially concave game is also a concave n-person game. □

As we saw in chapter 3, there is a variety of online convex optimization algorithms, such as online gradient descent and follow-the-leader, that achieve no-regret when applied to the right setting. If every player follows a no-regret algorithm, not necessarily the same, then the average strategy vector will converge to a Nash equilibrium. Simultaneously, the average payoff of each player converges to the payoff at the equilibrium point the dynamic converges to.

Theorem 4.7. *If every player in a socially concave game follows a no-regret algorithm, then the average strategy vector in round T is an $\epsilon(T)$-Nash equilibrium, where $\epsilon(T) = \frac{1}{T}\sum_{i=1}^{n} \frac{r_i}{r_{min}} R_i(T)$ and $R_i(T)$ is the regret of player i at round T.*

Proof. Assume that at round t player i has a strategy denoted by x_i^t. Since she plays according to a no-regret algorithm, after round T we have

$$\frac{1}{T}\sum_{t=1}^{T} u_i(\mathbf{x}^t) \geq \frac{1}{T}\max_{\mathbf{x}_i \in S_i}\sum_{t=1}^{T} u_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) - \frac{R_i(T)}{T}.$$

The average strategy vector at this round is denoted by $\bar{\mathbf{x}}^T$. Following the definition of an ϵ-Nash equilibrium our aim is to show that for all i there exists an $\epsilon_i(T) \geq 0$, such that for any $x_i \in S_i$

$$u_i(\bar{\mathbf{x}}^T) \geq u_i(x_i, \bar{\mathbf{x}}_{-i}^T) - \epsilon_i(T).$$

We observe that the two inequalities are very similar. The first one upperbounds the difference between the average payoff if instead of the strategies indicated by the algorithm player i had chosen to play any fixed strategy for all rounds and the average payoff she actually achieved. Whereas, the second inequality attempts to bound the difference between the payoff of player i if she unilaterally deviated from the average strategy played up to this point and her payoff if everyone played according to the average strategy.

Since the average payoff is a convex combination we can use convexity arguments for the payoff functions to derive the second inequality. However, there is an extra degree of complexity because we know that $\sum_{i=1}^{n} r_i u_i(\mathbf{x})$ is concave and not $\sum_{i=1}^{n} u_i(\mathbf{x})$, which would make the proof relatively direct. In terms of intuition it does not make a significant difference. Starting from the initial inequality the convexity of $u_i(\mathbf{x})$ in the other players' strategies results in the right side of the final inequality lowerbounding that of the initial. Supposing we had that $\sum_{i=1}^{n} u_i(\mathbf{x})$ is concave in $\mathbf{x}$, which is not actually true, we would get the left side as an upper bound of the left side of the initial inequality.

That being so, we will multiply the initial inequality by r_i and take the sum for all players.

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} r_i u_i(\mathbf{x}^t) \geq \frac{1}{T}\max_{\mathbf{x}\in S}\sum_{t=1}^{T}\sum_{i=1}^{n} r_i u_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) - \sum_{i=1}^{n} r_i \frac{R_i(T)}{T}.$$

Due to the concavity of $\sum_{i=1}^{n} r_i u_i(\mathbf{x})$ in $\mathbf{x}$ we obtain

$$\sum_{i=1}^{n} r_i u_i(\bar{\mathbf{x}}^t) \geq \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{n} r_i u_i(\mathbf{x}^t).$$

Similarly, the convexity of $u_i(\mathbf{x})$ in $\mathbf{x}_i$ implies that

$$\frac{1}{T}\max_{\mathbf{x}\in S}\sum_{t=1}^{T}\sum_{i=1}^{n} r_i u_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) \geq \max_{\mathbf{x}\in S}\sum_{i=1}^{n} r_i u_i(\mathbf{x}_i, \bar{\mathbf{x}}_{-i}^t) \geq \sum_{i=1}^{n} r_i u_i(\mathbf{s}_i, \bar{\mathbf{x}}_{-i}^t),$$

where $\mathbf{s}_i$ denotes the best response of player i to $\bar{\mathbf{x}}_{-i}^t$. We now want to keep the payoff function of only one player j. Because $\sum_{i\neq j} r_i u_i(\bar{\mathbf{x}}^t) \leq \sum_{i\neq j} r_i u_i(\mathbf{s}_i, \bar{\mathbf{x}}_{-i}^t)$, we obtain for any $x_j \in S_j$

$$u_j(\bar{\mathbf{x}}^T) \geq u_j(\mathbf{x}_j, \bar{\mathbf{x}}_{-j}^T) - \frac{1}{r_j}\sum_{i=1}^{n} r_i \frac{R_i(T)}{T}.$$

Therefore, $\epsilon_i(T) = \frac{1}{r_i}\sum_{j=1}^{n} r_j \frac{R_j(T)}{T}$. The maximum value of $\epsilon_i(T)$ is $\frac{1}{r_{\min}}\sum_{j=1}^{n} r_j \frac{R_j(T)}{T}$. Thus, we conclude that for all the players

$$u_i(\bar{\mathbf{x}}^T) \geq u_i(x_i, \bar{\mathbf{x}}_{-i}^T) - \frac{1}{r_{\min}}\sum_{j=1}^{n} r_j \frac{R_j(T)}{T}.$$

□

Theorem 4.8. *If every player in a socially concave game follows a no-regret algorithm, then the average payoff of every player is close to her payoff for the average strategy vector, i.e.*

$$\left|\frac{1}{T}\sum_{t=1}^{T} u_i(\mathbf{x}^t) - u_i(\bar{\mathbf{x}}^T)\right| \leq \frac{1}{r_i}\sum_{j=1}^{n} r_j \frac{R_j(T)}{T},$$

where $R_j(T)$ is the regret of player j at round T.

Proof. The idea behind this proof is exactly the same as that of the previous theorem. The starting inequality is again

$$\frac{1}{T}\sum_{t=1}^{T} u_i(\mathbf{x}^t) \geq \frac{1}{T}\max_{\mathbf{x}_i\in S_i}\sum_{t=1}^{T} u_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) - \frac{R_i(T)}{T}.$$

Because of the convexity of $u_i(\mathbf{x}_i, \mathbf{x}_{-i})$ in $\mathbf{x}_{-i}$ we obtain

$$\frac{1}{T} \max_{x_i \in S_i} \sum_{t=1}^{T} u_i(\mathbf{x}_i, \mathbf{x}_{-i}^t) \geq \max_{\mathbf{x}_i \in S_i} u_i(\mathbf{x}_i, \bar{\mathbf{x}}_{-i}^T) \geq u_i(\bar{\mathbf{x}}^T).$$

We can also directly use the result of the previous proof

$$\frac{1}{T} \sum_{t=1}^{T} u_j(\mathbf{x}^T) \geq u_j(\bar{\mathbf{x}}^T) - \frac{1}{r_j} \sum_{i=1}^{n} r_i \frac{R_i(T)}{T}.$$

Combining the previous inequalities the result of the theorem ensues. □

By the two theorems above it immediately follows that if for all $i \in \{1, \ldots, n\}$ $R_i(T) = o(T)$ then the average strategy vector converges to a Nash equilibrium and that each players' average payoff converges to her payoff at this equilibrium.

An interesting question is what is the relation between strictly diagonally concave and socially concave games. In fact, these two classes of games overlap but neither of the two is a subset of the other.

On one side, the game corresponding to the FJ model is strictly diagonally concave but it is not socially concave. The cost functions of the players are

$$c_i(\mathbf{x}) = (1 - \alpha_i) \sum_{j \neq i} w_{ij}(x_i - x_j)^2 + \alpha_i(x_i - s_i)^2.$$

These functions are convex in not just x_i but also in any x_j, for $j \neq i$. Therefore, this game is not socially concave. On the other side, the definition of social concavity does not ensure the uniqueness of the equilibrium. However, if one of the conditions of a socially concave game is strict then the game is strictly diagonally concave.

Theorem 4.9. *A socially concave game is a strictly diagonally concave game if the convex combination of the payoff functions is strictly concave in the strategy vector or all the players' payoff functions are strictly convex in the strategies of the other players.*

Proof. In theorem 4.6 we proved that every player's payoff function is concave in the player's strategy. Since this results from the concavity in $\mathbf{x}_i$ of $\frac{1}{r_i}\left(\sum_{j=1}^{n} r_j u_j(\mathbf{x}) - \sum_{j \neq i} r_j u_j(\mathbf{x})\right)$, if any of the two conditions of social concavity is strict then $u_i(\mathbf{x})$ will be strictly concave in $\mathbf{x}_i$ because the sum of a concave and a strictly concave function is strictly concave. From now on we will not be concerned about which of the initial conditions was strict but we will use the strict concavity we proved.

Let $\mathbf{x}^*$ and $\bar{\mathbf{x}}$ be two distinct points. We want to construct $\sum_{i=1}^{n} r_i \nabla_i u_i(\bar{\mathbf{x}})^T (\mathbf{x}_i^* - \bar{\mathbf{x}}_i) + \sum_{i=1}^{n} r_i \nabla_i u_i(\mathbf{x}^*)^T (\bar{\mathbf{x}}_i - \mathbf{x}_i^*)$. Due to the concavity of $\sum_{i=1}^{n} r_i u_i(\mathbf{x})$ in $\mathbf{x}$ we have

$$\sum_{i=1}^{n} r_i u_i(\mathbf{x}^*) \leq \sum_{i=1}^{n} r_i u_i(\bar{\mathbf{x}}) + \sum_{i=1}^{n} r_i \nabla u_i(\mathbf{x}^*)^T (\mathbf{x}^* - \bar{\mathbf{x}}).$$

The same is true if we interchange the two points. Therefore, we obtain

$$\sum_{i=1}^{n} r_i \nabla u_i(\mathbf{x}^*)^T(\mathbf{x}^* - \bar{\mathbf{x}}) + \sum_{i=1}^{n} r_i \nabla u_i(\bar{\mathbf{x}})^T(\bar{\mathbf{x}} - \mathbf{x}^*) \geq 0.$$

Function $u_i(\mathbf{x})$ is convex in $\mathbf{x}_{-i}$ and strictly concave in $\mathbf{x}_i$. Subsequently,

$$u_i(\mathbf{x}_i^*, \bar{\mathbf{x}}_{-i}) \geq u_i(\mathbf{x}^*) + \nabla_{-i} u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_{-i} - \mathbf{x}_{-i}^*)$$

and

$$u_i(\mathbf{x}_i^*, \bar{\mathbf{x}}_{-i}) < u_i(\bar{\mathbf{x}}) + \nabla_i u_i(\bar{\mathbf{x}})^T(\mathbf{x}_i^* - \bar{\mathbf{x}}_i).$$

From these two we obtain

$$0 > u_i(\mathbf{x}^*) - u_i(\bar{\mathbf{x}}) + \nabla_{-i} u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_{-i} - \mathbf{x}_{-i}^*) - \nabla_i u_i(\bar{\mathbf{x}})^T(\mathbf{x}_i^* - \bar{\mathbf{x}}_i).$$

Of course the same is valid if we interchange the two points. This results to

$$\nabla_i u_i(\bar{\mathbf{x}})^T(\mathbf{x}_i^* - \bar{\mathbf{x}}_i) - \nabla_{-i} u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_{-i} - \mathbf{x}_{-i}^*) + \nabla_i u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_i - \mathbf{x}_i^*) - \nabla_{-i} u_i(\bar{\mathbf{x}})^T(\mathbf{x}_{-i}^* - \bar{\mathbf{x}}_{-i}) > 0.$$

This way we can eliminate all the terms that include the gradients of payoff functions over the strategies of the other players.

Consequently, we obtain that $\sum_{i=1}^{n} r_i \nabla_i u_i(\bar{\mathbf{x}})^T(\mathbf{x}_i^* - \bar{\mathbf{x}}_i) + \sum_{i=1}^{n} r_i \nabla_i u_i(\mathbf{x}^*)^T(\bar{\mathbf{x}}_i - \mathbf{x}_i^*) > 0$. Thus, the game is strictly diagonally concave. □

4.5 Asymmetric Coevolutionary Opinion Formation Games

A class of opinion formation games that fall into the category of concave n-person games is that of the asymmetric coevolutionary opinion formation games. The motivation for their study comes from the fact that the game corresponding to the FJ model is relatively limited, because the weights between the players are constant. A more realistic model of opinion formation would have weights that evolve as the opinions of the players change. In [3] the authors attempt to extend the FJ model towards this direction by defining the games described below.

Similarly to the FJ model, in the asymmetric coevolutionary games every player i possesses an intrinsic opinion $s_i \in [0, 1]$ and a self-confidence factor $\alpha_i \in [0, 1]$. Externally, she expresses to the other players an opinion $x_i \in [0, 1]$, that may differ from s_i. The affinity of the players is directly influenced by the opinions they express, that meaning that the weight player i attributes to player j is greater if the second player's external opinion is closer to player i's internal opinion. More specifically, for any $i \neq j$ weight $q_{ij}(\mathbf{x}_{-i})$ is a non-negative continuous function in $\mathbf{x}_i$ and it decreases as the difference $|x_j - s_i|$ increases. Additionally, the weights are normalized, i.e. $\sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}) = 1$. As a result, weight $q_{ij}(\mathbf{x}_{-i})$ is an increasing function of $|x_k - s_i|$ for $k \neq j$. In correspondence to the cost function of the FJ model, the players suffer cost

$$c_i(\mathbf{x}) = (1 - \alpha_i) \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i})(x_i - x_j)^2 + \alpha_i(x_i - s_i)^2.$$

Theorem 4.10. *Asymmetric coevolutionary opinion formation games always admit a pure Nash equilibrium.*

Proof. By theorem 4.1 all concave n-person games admit a pure Nash equilibrium. Thus, we only need to show that asymmetric coevolutionary games are a subclass of concave n-person games. The cost function $c_i(\mathbf{x})$ is continuous in $\mathbf{x}$, because all weights $q_{ij}(\mathbf{x}_{-i})$ are continuous in $\mathbf{x}$, as well as the rest of the terms, that are quadratic functions. At this point we should note that a convex minimization problem is equivalent to a concave maximization problem. Hence, we want to show that the cost functions are convex in x_i. In particular, $q_{ij}(\mathbf{x}_{-i})$ is not dependent on x_i and therefore $q_{ij}(\mathbf{x}_{-i})(x_i - x_j)^2$ is convex in x_i, as well as $\alpha_i(x_i - s_i)^2$. Consequently, all players' cost functions are convex in their own opinion. □

However, the asymmetric coevolutionary opinion formation games are not in general socially concave. This holds because the cost functions of the FJ model, which is the specific case of the coevolutionary for fixed weights, are convex in all the agent's opinions.

Chapter 5

Equilibrium Convergence in Coevolutionary Opinion Formation Games

In this chapter we are interested in the convergence properties specifically of the Asymmetric Coevolutionary Games. Our definition is less strict than that of [3] as the minimum requirements for the weight functions are that they are continuous in their arguments and normalized. In more detail, we have n-agents, each with a constant intrinsic opinion $s_i \in [0, 1]$, a constant self-confidence factor $\alpha_i \in [0, 1]$ and an external opinion $x_i \in [0, 1]$ that may vary as the game evolves. The cost player i incurs is calculated by the function

$$c_i(\mathbf{x}) = (1 - \alpha_i) \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i})(x_i - x_j)^2 + \alpha_i (x_i - s_i)^2,$$

where $q_{ij}(\mathbf{x}_{-i})$ is the weight between agent i and agent j that is dependent on the external opinions of all the agents but i and continuous in $\mathbf{x}_{-i}$. Furthermore, it holds that $\sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}) = 1$.

Our study focuses on two algorithms that are derived from intuitively natural behaviors of agents. The first one, called Best Response, determines that at every round the agent decides to express the opinion that minimizes the cost she incurs if the external opinions of all the other agents remain fixed. In mathematical terms, the opinion update rule at round $t + 1$ is

$$x_i^{t+1} = \min_{x \in [0,1]} c_i(x, \mathbf{x}_{-i}^t).$$

Since the cost function is twice differentiable in x_i, it is easy to see that the minimizer of the cost function is

$$x_i^{t+1} = (1 - \alpha_i) \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}^t) x_j^t + \alpha_i s_i.$$

The second algorithm is Follow-The-Leader, which makes use of the full history of the costs the agent has suffered. In particular, at round t the agent following this algorithm attempts to minimize the total sum of costs she suffers, meaning that she chooses the opinion that would minimize the cumulative cost up to this round if she had expressed this opinion at every round. Formally, the update rule is

$$x_i^{t+1} = \min_{x \in [0,1]} \sum_{\tau=1}^{t} c_i(x, \mathbf{x}_{-i}^{\tau}).$$

In this case, the agent is not as influenced by the last round as it is in Best Response. Thus, the change of opinions of one agent from one round to the next one can not be as dramatic. This is apparent in the following update rule, which is actually equivalent to the minimization of the sum of costs and can be expressed as a convex combination of the Best Response opinions.

$$x_i^{t+1} = \frac{t}{t+1} x_i^t + \frac{1}{t+1} \left[(1 - \alpha_i) \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}^t) x_j^t + \alpha_i s_i \right].$$

With the aid of examples simulated in Python we observed that Best Response does not generally converge in time, as the agent that follows it is susceptible to the changes of opinions of all the other agents and this can lead to oscillations. Even in terms of local convergence, the Best Response algorithm fails for some classes of weight functions such as those that do not have continuous first-order derivatives. On the opposite, Follow-The-Leader appears to converge to a Nash equilibrium in the examples we studied. We proved that under certain conditions that ensure that the Nash equilibrium is unique this holds. In addition, we showed that this algorithm is no-regret.

5.1 Global Convergence of Best Response

For general weight functions Best Response does not always converge to a specific point. In fact, the opinion vector might oscillate between multiple points. This is evident in two simple examples we constructed. In general, the minimum requirements for the weight functions is that they are continuous in the opinion vector. In the first example the weight functions are not differentiable, whereas in the second they are.

Both examples consist of a network of four agents. The two of them, agents 1 and 2, have fixed opinions, which are 0.3 and 0.5 respectively. In order to keep their opinions unchanged, both of them have self-confidence factors equal to 1. The other two agents display relatively symmetrical behavior. They both have self-confidence factors 0.2 and therefore, they are greatly influenced by the opinions of others. Particularly, agent 3 has 0.3 intrinsic opinion and is influenced only by agent 1 when the opinion of agent 4 is over 0.5 and equally by agents 1 and 4 if the opinion of agent 4 is at most 0.49. Likewise, agent 4 has an intrinsic opinion equal to 0.5 and is only influenced by agent 2 when the opinion of agent 3 is at most 0.3 and equally by agents 2 and 3 if the opinion of agent 3 is at least 0.31.

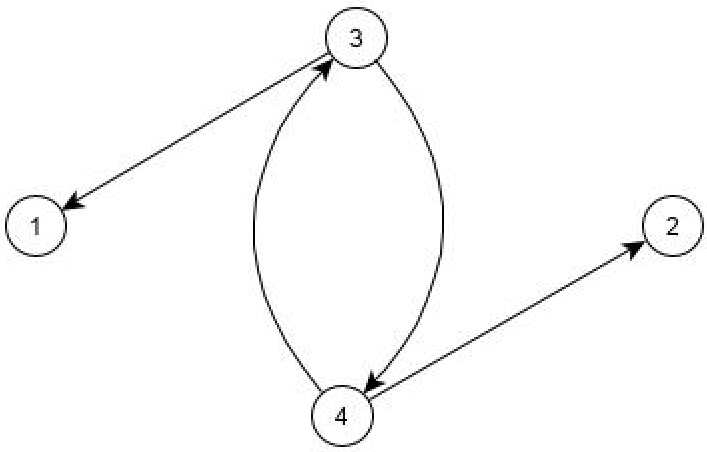

Figure 5.1: The graph of the examples network

Example 1

In the first example the weight functions do not have continuous first derivatives. Considering that agents 1 and 2 are not influenced by the opinions of the rest of the agents, there is no reason to define their weights. However, the weight functions of agents 3 and 4 are

$$q_{31}(x_4) = \begin{cases} 0.5, & \text{if } x_4 \leq 0.49 \\ 50x_4 - 24, & \text{if } 0.49 < x_4 < 0.5 \\ 1, & \text{if } x_4 \geq 0.5 \end{cases}$$

$$q_{32}(x_4) = 0$$

$$q_{34}(x_4) = \begin{cases} 0.5, & \text{if } x_4 \leq 0.49 \\ -50x_4 + 25, & \text{if } 0.49 < x_4 < 0.5 \\ 0, & \text{if } x_4 \geq 0.5 \end{cases}$$

$$q_{41}(x_3) = 0$$

$$q_{42}(x_3) = \begin{cases} 1, & \text{if } x_3 \leq 0.3 \\ -50x_3 + 16, & \text{if } 0.3 < x_3 < 0.31 \\ 0.5, & \text{if } x_3 \geq 0.31 \end{cases}$$

$$q_{43}(x_3) = \begin{cases} 0, & \text{if } x_3 \leq 0.3 \\ 50x_3 - 15, & \text{if } 0.3 < x_3 < 0.31 \\ 0.5, & \text{if } x_3 \geq 0.31 \end{cases}$$

.

By running simulations first, we noticed that this game has two equilibrium points $\mathbf{x}^* = [0.3, 0.5, 0.3, 0.5]$ and $\bar{\mathbf{x}} = [0.3, 0.5, 5/14, 31/70]$. Furthermore, we observed that the behavior of the system depends on the starting point. Notably, the distinct cases are

1. if none of the initial opinions of agents 3 and 4 are in the interval $(0.3, 0.5)$ then the game converges to the the equilibrium point [0.3, 0.5, 0.3, 0.5]

2. if only one of the initial opinions of agents 3 and 4 is in the interval $(0.3, 0.5)$ then the game eventually oscillates between the points [0.3, 0.5, 0.3, 31/70] and [0.3, 0.5, 5/14, 0.5]

3. if both initial opinions of agents 3 and 4 are in the interval $(0.3, 0.5)$ then the game converges to the equilibrium point [0.3, 0.5, 5/14, 31/70]

Specific results, that fall into these three cases and were produced by our simulations of the system, are depicted in 5.2.

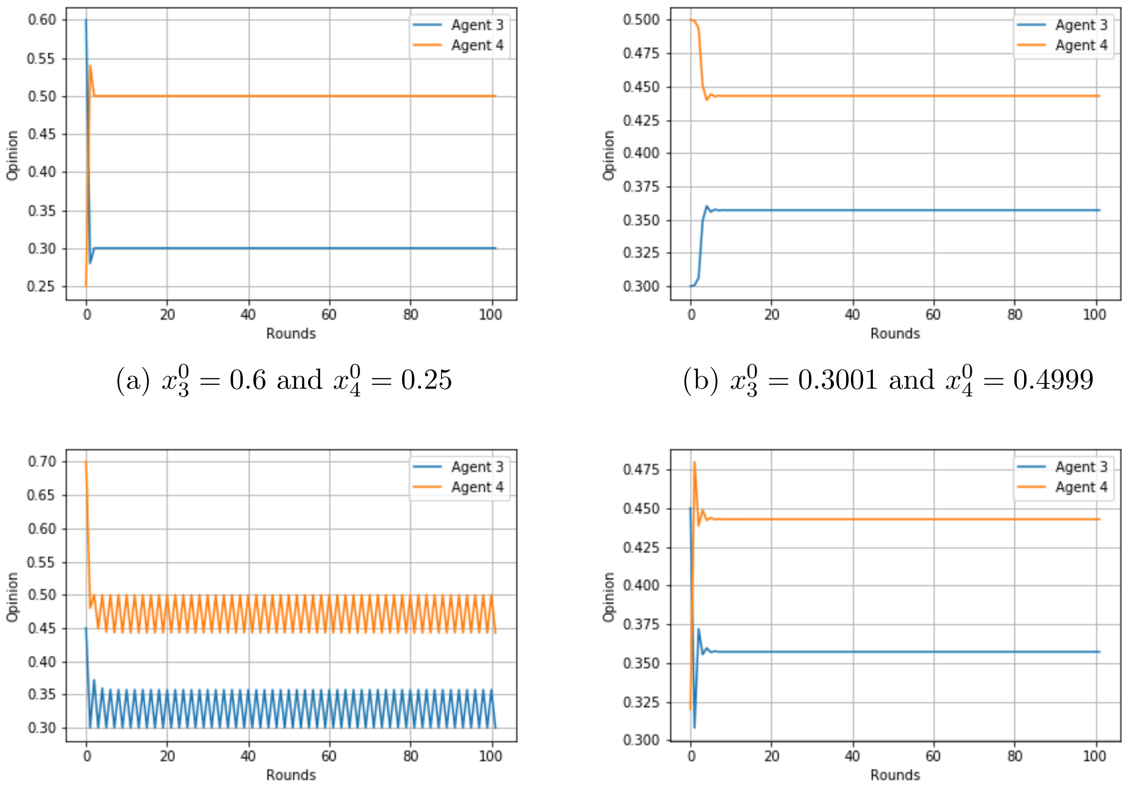

(a) $x_3^0 = 0.6$ and $x_4^0 = 0.25$

(b) $x_3^0 = 0.3001$ and $x_4^0 = 0.4999$

(c) $x_3^0 = 0.45$ and $x_4^0 = 0.7$

(d) $x_3^0 = 0.45$ and $x_4^0 = 0.32$

Figure 5.2: The evolution of the opinions of agents 3 and 4 for Best Response in example 1

The previous remarks can be proven mathematically. The general idea behind this proof is that at every round the agents 3 and 4 swap opinions in a way. Therefore, we can address the cases concerning the initial opinion of 3 only. The update rule for the opinions are :

$$x_3^t = 0.8\left[q_{31}(x_4^{t-1})x_1^{t-1} + q_{34}(x_4^{t-1})x_4^{t-1}\right] + \alpha_3 s_3,$$
$$x_4^t = 0.8\left[q_{40}(x_0^{t-1})x_0^{t-1} + q_{43}(x_3^{t-1})x_3^{t-1}\right] + \alpha_4 s_4.$$

Subsequently, we study the distinct cases of the initial opinion of agent 3.

1. If $x_3^t \leq 0.3$, then $x_4^{t+1} = 0.5$.

2. If $x_3^t \geq 0.5$, then after two rounds we have $x_3^{t+2} = 0.3$.
   The reason for this is that when $x_3^t \geq 0.5$, we obtain
   $$x_4^{t+1} = 0.8(0.5 \cdot 0.5 + 0.5 \cdot x_3^t) + 0.2 \cdot 0.5 \geq 0.5.$$
   By case 1 we have $x_3^{t+2} = 0.3$

3. If $x_3^t \in [0.31, 0.5)$, then $|x_4^{t+1} - \bar{x}_4| = 0.4|x_3^t - \bar{x}_3|$.
   The update rule for x_3^{t+1} is
   $$x_3^{t+1} = 0.8(0.5 \cdot 0.5 + 0.5x_2^t) + 0.2 \cdot 0.5.$$
   The same update rule applies to the equilibrium $\bar{\mathbf{x}}$. Thus, we have $\bar{x}_3 = 0.8(0.5 \cdot 0.5 + 0.5\bar{x}_2) + 0.2 \cdot 0.5$. As a result, it holds that $|x_3^{t+1} - \bar{x}_3| = 0.4|x_2^t - \bar{x}_2|$.

4. If x_3^t is in $(0.3, 0.31)$, then $|x_4^{t+1} - 0.5| > |x_3^t - 0.3|$.
   For simplicity we can consider that $x_3^t = 0.3 + \epsilon$, where $\epsilon \in (0, 0.01)$. By the best response update rule we have $x_4^{t+1} = 0.5 + 40\epsilon^2 - 8\epsilon$. Therefore the distance $0.5 - x_4^{t+1} = 8\epsilon - 40\epsilon^2 > \epsilon$ and x_4^{t+1} moves further away from 0.5.

Similarly, it can be showed that the initial opinion of agent 4 determines the opinion of agent 3 at the next round.

1. If $x_4^t \geq 0.5$, then $x_3^{t+1} = 0.3$.

2. If $x_4^t \leq 0.3$, then after two rounds we have $x_4^{t+2} = 0.5$.

3. If $x_4^t \in (0.3, 0.49]$, then $|x_3^{t+1} - \bar{x}_3| = 0.4|x_4^t - \bar{x}_4|$.

4. If x_4^t is in $(0.49, 0.5)$, then $|x_3^{t+1} - 0.3| > |x_4^t - 0.5|$.

Example 2

In the second example the weight functions are twice differentiable, but even with this requirement the Best Response might not converge. In particular, in our example the weight functions of agents 3 and 4 are

$$q_{31}(x_4) = \begin{cases} 0.5 & x_4 \leq 0.49 \\ 10^4(x_4 - 0.49)^2 + 0.5 & 0.49 < x_4 \leq 0.495 \\ 1 - 10^4(x_4 - 0.5)^2 & 0.495 < x_4 < 0.5 \\ 1 & x_4 \geq 0.5 \end{cases}$$

$$q_{32}(x_4) = 0$$

$$q_{34}(x_4) = \begin{cases} 0.5 & x_4 \leq 0.49 \\ 0.5 - 10^4(x_4 - 0.49)^2 & 0.49 < x_4 \leq 0.495 \\ 10^4(x_4 - 0.5)^2 & 0.495 < x_4 < 0.5 \\ 0 & x_4 \geq 0.5 \end{cases}$$

$$q_{41}(x_3) = 0$$

$$q_{42}(x_3) = \begin{cases} 1 & x_3 \leq 0.3 \\ 1 - 10^4(x_3 - 0.3)^2 & 0.3 < x_3 \leq 0.305 \\ 10^4(x_3 - 0.31)^2 + 0.5 & 0.305 < x_3 < 0.31 \\ 0.5 & x_3 \geq 0.31 \end{cases}$$

$$q_{43}(x_3) = \begin{cases} 0 & x_3 \leq 0.3 \\ 10^4(x_3 - 0.3)^2 & 0.3 < x_3 \leq 0.305 \\ 0.5 - 10^4(x_3 - 0.31)^2 & 0.305 < x_3 < 0.31 \\ 0.5 & x_3 \geq 0.31 \end{cases}$$

Again the game has the same two equilibrium points, $\mathbf{x}^* = [0.3, 0.5, 0.3, 0.5]$ and $\bar{\mathbf{x}} = [0.3, 0.5, 5/14, 31/70]$, as in example 1. Experimentally, we observe that the behavior of the system is also similar to that of example 1, meaning that the convergence of Best Response depends on the initial opinion vector. More specifically, if the opinions of agents 3 and 4 fall into an interval close to $\bar{\mathbf{x}}$ then the system converges to that point, if none of them are in the interval it converges to $\mathbf{x}^*$ and otherwise agents 3 and 4 oscillate between their opinions at the two Nash equilibria alternately. This is illustrated in the graphs of figure 5.3 that are similar to those of example 1. The difference between the two examples in the interval that determines the convergence is apparent in figures 5.2b and 5.3b where for the same initial opinions the two systems converge to different points.

A question that occurs from the previous examples is whether strong diagonal concavity, that ensures the uniqueness of the equilibrium, suffices for the convergence of Best Response in a Nash equilibrium.

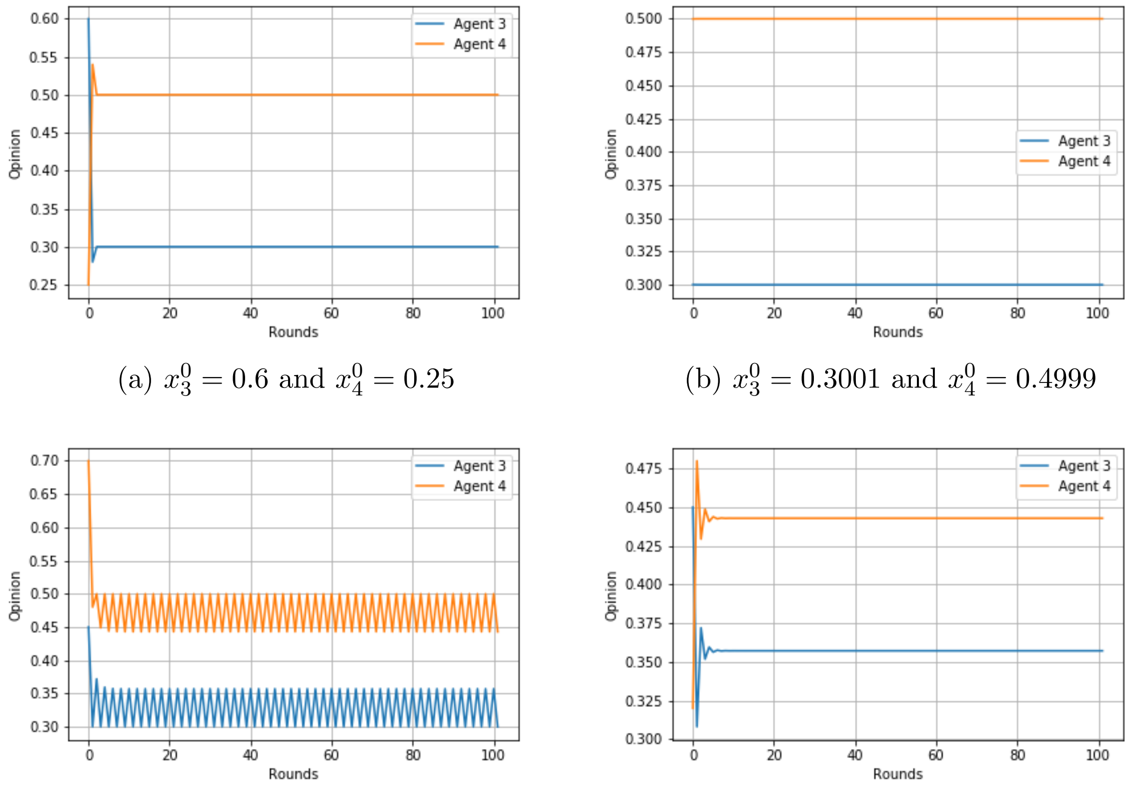

(a) $x_3^0 = 0.6$ and $x_4^0 = 0.25$ (b) $x_3^0 = 0.3001$ and $x_4^0 = 0.4999$

(c) $x_3^0 = 0.45$ and $x_4^0 = 0.7$ (d) $x_3^0 = 0.45$ and $x_4^0 = 0.32$

Figure 5.3: The evolution of the opinions of agents 3 and 4 for Best Response in example 2

5.2 Local Convergence of Best Response

A question that arose from the graph in 5.2b is whether Best Response always converges locally, regardless of the cost functions. Even though in the general case of cost functions Best Response does not converge for any initial opinion vector, we observed in the examples above that for each Nash equilibrium there is a non-empty subset of the joint strategy set that contains the initial opinion vectors that converge to it. Hence, the next logical step was to study the form of those subsets. In particular, we were interested in seeing whether there is a ball around each Nash equilibrium, such that if Best Response starts from any interior point of the ball it will converge to this equilibrium. This property is called local convergence. Remarkably, it does not hold for the example with the non differentiable weight functions. Nevertheless, the second example with the differentiable weight functions is locally convergent

Example 1

The definition of the example remains the same as in the previous section. One of the equilibrium points is $\mathbf{x}^* = [0.3, 0.5, 0.3, 0.5]$. We experimentally observed that if the initial point of the algorithm is of the type $[0.3, 0.5, 0.3+\epsilon_1, 0.5-\epsilon_2]$, where ϵ_1 and ϵ_2 are relatively small and positive, the system does not converge to $\mathbf{x}^*$. Formally, the proof for this is part of the proof that we presented in order to show that global convergence did not hold. Yet, for clarity reasons we will underline the specific part that refers to this issue.

Once more, the basic idea of the proof is that agents 3 and 4 swap opinions. Particularly, if $x_4^t = 0.5-\epsilon$, where ϵ is in $(0, 0.01)$, then $x_3^{t+1} - 0.3 = -40\epsilon^2 + 8\epsilon > \epsilon$. Correspondingly, if $x_3^t = 0.3+\epsilon$, where $\epsilon \in (0, 0.01)$, then we have $x_4^{t+1} = 0.5 + 40\epsilon^2 - 8\epsilon < 0.5 - \epsilon$. This means that the two agents' opinions move away from the Nash equilibrium alternately.

Example 2

In contrast, the results of the example 2 showed that for both Nash equilibria there is an area around them from where the Best Response algorithm converges to them. The point that is more interesting to study is again $\mathbf{x}^* = [0.3, 0.5, 0.3, 0.5]$, since the behavior of the algorithm around the point $\bar{\mathbf{x}} = [0.3, 0.5, 5/14, 31/70]$ is similar for both examples.

Interestingly, if $x_4^t = 0.5 - \epsilon$, where $\epsilon \in (0, \frac{20-\sqrt{395}}{200})$, then $x_3^{t+1} - 0.3 = 0.8[(1 - 10^4\epsilon^2)0.3 + 10^4\epsilon^2(0.5-\epsilon)] + 0.2 \times 0.3 - 0.3 = 1600\epsilon^2 - 8000\epsilon^3 < \epsilon$. Namely, if the opinion of agent 4 is a bit less than 0.5 but relatively close to it, then at the next round the opinion of agent 3 will be even closer to 0.3. A similar result holds if the opinion of agent 3 is ϵ over 0.3. In this case at the next round the opinion of agent 4 will move close to 0.5.

5.3 Examples of Convergence of Follow the Leader

In this section we cite the results of the simulations we ran as an incentive for studying the equilibrium convergence under Follow the Leader. As mentioned previously, this algorithm ensures more smooth changes in the opinion values than Best Response. Concerning the examples 1 and 2, the opinion vector always converges to a Nash equilibrium point regardless of the initial opinions if all the agents implement this algorithm. More specifically, figures 5.4 and 5.5 show for both the not differentiable and the differentiable weight functions that not only the oscillations are prevented but also the opinions change more smoothly from one round to the next one.

Even for more complicated weight functions and more agents, our simulations show that Follow The Leader converges to some point in S eventually. One such example is for weight functions that are exponential and increase as the distance between the other agent's opinion is closer to the agents' intrinsic opinion. More specifically, if we denote by $d_j^i = |x_j - s_i|$ the distance between the intrinsic opinion of agent i and the expressed opinion of agent j, then the weight between agent i and agent i is

$$q_{ij}(\mathbf{x}_{-i}) = \frac{e^{-d_j^j}}{\sum_{k \neq i} e^{-d_k^i}}.$$

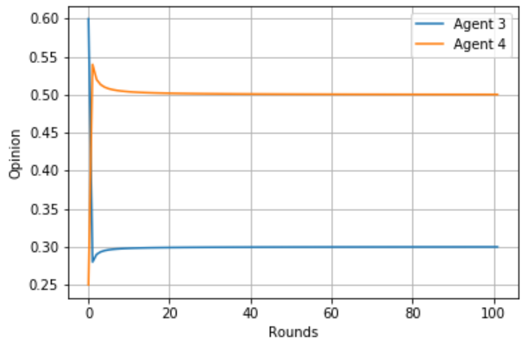

(a) $x_3^0 = 0.6$ and $x_4^0 = 0.25$

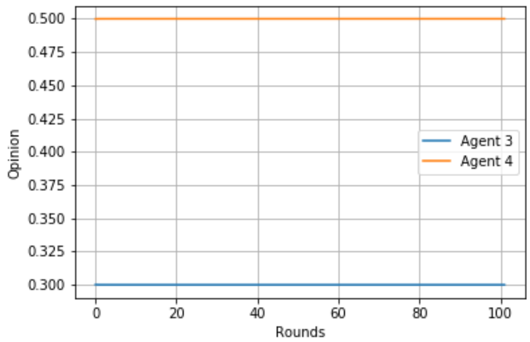

(b) $x_3^0 = 0.3001$ and $x_4^0 = 0.4999$

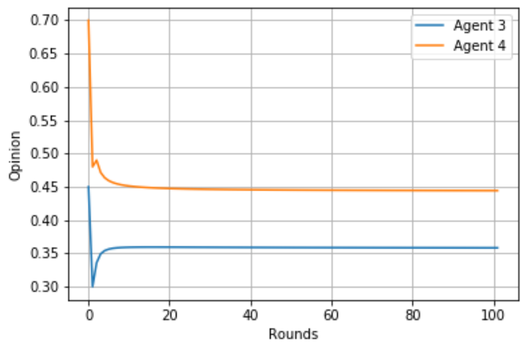

(c) $x_3^0 = 0.45$ and $x_4^0 = 0.7$

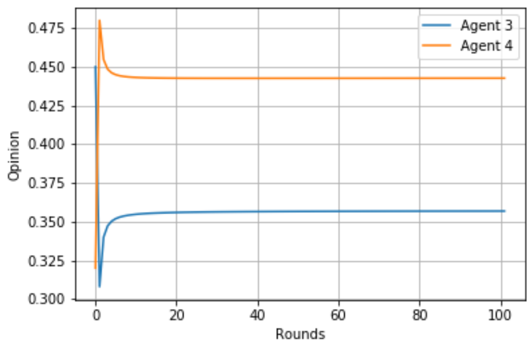

(d) $x_3^0 = 0.45$ and $x_4^0 = 0.32$

Figure 5.4: The evolution of the opinions of agents 3 and 4 for Follow The Leader in example 1

Figure 5.6 depicts the graphs of the evolution of the opinions of 20 and 100 agents respectively in time. In both cases the initial configuration, such as the intrinsic opinions and the self-confidence factors, were selected randomly. We observe that if all the agents choose their opinions according to Follow-The-Leader, then the system converges to a specific point.

Similarly, we consider the case where all the weight functions are

$$q_{ij}(\mathbf{x}_{-i}) = \frac{1}{d_j^i + \lambda} \times \left(\sum_k \frac{1}{d_k^i + \lambda}\right)^{-1},$$

where λ is a small positive constant, for example 10^{-10}. Figure 5.7 shows that for such a dynamic of 20 and 100 agents respectively, where the constant parameters are generated randomly and all the agents implement Follow The Leader, the opinion vector converges to a point eventually.

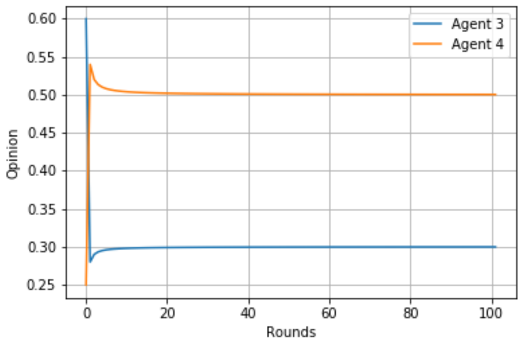

(a) $x_3^0 = 0.6$ and $x_4^0 = 0.25$

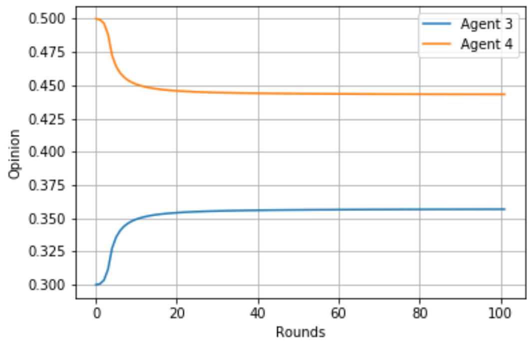

(b) $x_3^0 = 0.3001$ and $x_4^0 = 0.4999$

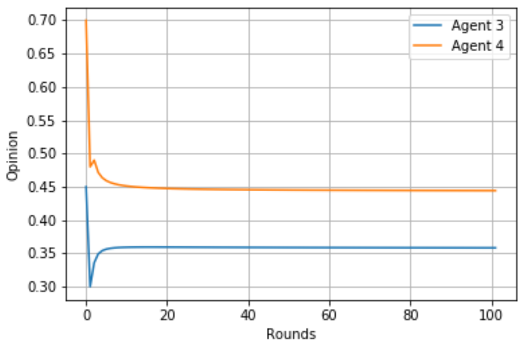

(c) $x_3^0 = 0.45$ and $x_4^0 = 0.7$

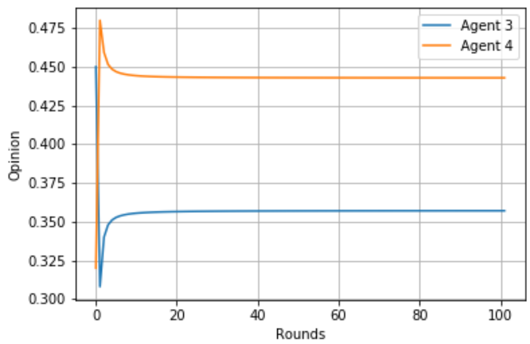

(d) $x_3^0 = 0.45$ and $x_4^0 = 0.32$

Figure 5.5: The evolution of the opinions of agents 3 and 4 for Follow The Leader in example 2

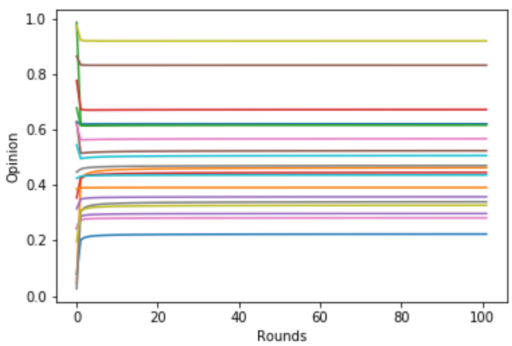

(a) 20 agents

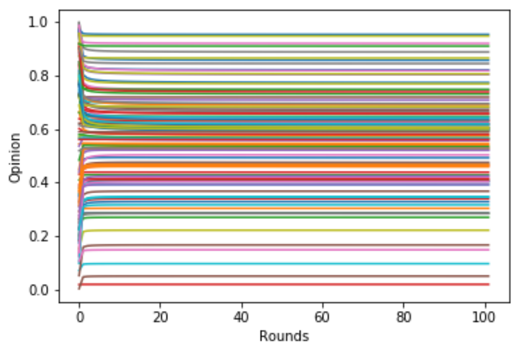

(b) 100 agents

Figure 5.6: The evolution of the opinions for exponential weight functions

5.4 Follow The Leader is no regret

In this section we will study the game from the point of view of a single agent, without knowing anything about the behavior of the rest of the agents, whose set is denoted by

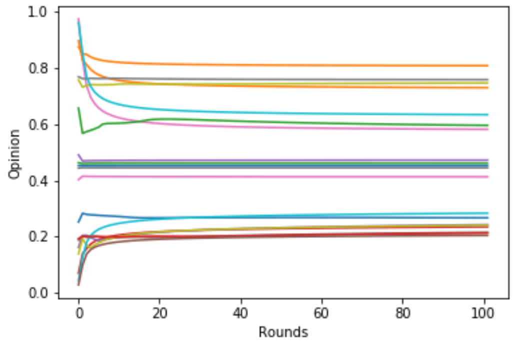

(a) 20 agents

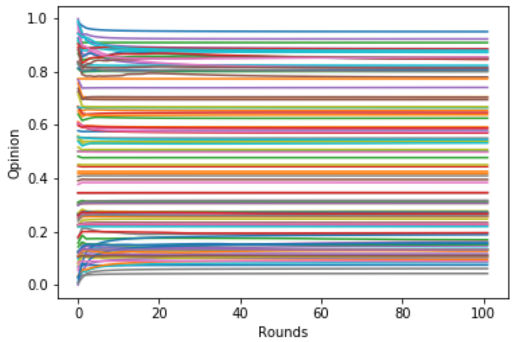

(b) 100 agents

Figure 5.7: The evolution of the opinions for weight functions

N. The agent we will focus on has an intrinsic opinion s, self-confidence factor α and a weight q_j for agent j. In this case, at round t the agent expresses opinion x^t, learns the opinions of the other agents denoted by $\mathbf{b}^t$ and incurs cost

$$c_t(x^t) = (1-\alpha)\sum_{j\in N} q_j(\mathbf{b}^t)(x^t - b_j^t)^2 + \alpha(x^t - s)^2.$$

Due to the strong convexity of the cost function in x_t we can show that if the agent plays according to the Follow-The-Leader algorithm, then he will have no regret. Applying this algorithm, in this specific case we have the following opinion change rule.

$$x^{t+1} = \frac{t}{t+1}x^t + \frac{1}{t+1}\left[(1-\alpha)\sum_{j\in N} q_j(\mathbf{b}^t)b_j^t + \alpha s\right].$$

Following the reasoning of [39] we will prove first a theorem that bounds the distance of the values that minimize two strongly convex one-variable functions who are close to each other. Closeness is indicated by the Lipschitzness of their difference. In other words, if we have two convex functions that for any $x \in [0,1]$ their values are relatively close to each other then the values of x that minimize them can not be far from one another. We can use this theorem show that the opinion of our agent will change less as the time progresses.

Theorem 5.1. *Let functions $f : [0,1] \to \mathbb{R}$ and $g : [0,1] \to \mathbb{R}$ be strongly convex, such that for all x the second derivatives $f''(x)$ and $g''(x) \geq \frac{1}{\eta}$ and their difference $h(x) = g(x) - f(x)$ is an L-Lipschitz function, i.e. $|h(x) - h(x')| \leq L|x - x'|$. Then if $x_f = \arg\min_{x\in[0,1]} f(x)$ and $x_g = \arg\min_{x\in[0,1]} g(x)$, it must hold that $|x_f - x_g| \leq \eta L$.*

Proof. The basic idea of this proof is that we can break the difference $h(x_f) - h(x_g)$ in the two separate contributions of the functions f and g.

Since f is $\frac{1}{\eta}$-strongly convex, we can lower bound the difference between the values of the function in x_g and x_f. The value $f(x_f)$ is the minimum of f, making the difference

$f(x_g) - f(x_f)$ positive if $x_g \neq x_f$. Specifically, we obtain

$$f(x_g) - f(x_f) \geq f'(x_f)(x_g - x_f) + \frac{1}{2\eta}(x_g - x_f)^2.$$

At point x_f and for all $x \in [0,1]$ it holds that $f'(x_f)(x - x_f) \geq 0$ because x_f minimizes the function f. Therefore, we have

$$f(x_g) - f(x_f) \geq \frac{1}{2\eta}(x_g - x_f)^2.$$

Similarly, we obtain $g(x_f) - g(x_g) \geq \frac{1}{2\eta}(x_f - x_g)^2$.

In order to bound the distance $|x_f - x_g|$ we will use the L-Lipschitzness of h.

$$\begin{aligned} L|x_f - x_g| \geq |h(x_f) - h(x_g)| &= |g(x_f) - f(x_f) - g(x_g) + f(x_g)| \\ &= g(x_f) - g(x_g) + f(x_g) - f(x_f) \\ &\geq \frac{1}{\eta}(x_f - x_g)^2. \end{aligned}$$

As a result, we have $|x_f - x_g| \leq \eta L$.

□

Proceeding, we will using the previous theorem we will prove that if the agent consults Follow-The-Leader during the game, then she will have incurred cost as little as that of the best fixed opinion on average.

Theorem 5.2. *The regret of the FTL algorithm is at most:*

$$\sum_{t=1}^{T} c_t(x_t) - \min_{x \in [0,1]} \sum_{t=1}^{T} c_t(x) \leq H_T \tag{5.1}$$

Proof. In order to show that the Follow-The-Leader is a no-regret algorithm for this game we will split the proof in three parts. First, we will bound the cumulative cost the agent incurs by playing according to FTL by the cost she would suffer if she was able to learn the opinions of the rest of the agents before expressing her own and therefore be one move ahead plus the sum of distances between the consecutive opinions she expresses. We could call the last one stability term, because it shows how often and how much the agent changes opinion. By the definition of the Follow The Leader the opinion the agent expresses at round t is

$$x^T = \arg\min_{x \in [0,1]} \sum_{t=1}^{T-1} c_t(x).$$

If she knew the rest of the opinions beforehand, which is not actually possible, she could be able to have at round t the opinion

$$\bar{x}^T = x^{T+1} = \arg\min_{x \in [0,1]} \sum_{t=1}^{T} c_t(x).$$

The relation between the cumulative costs incurred by the two methods is

$$\sum_{t=1}^{T} c_t(x^t) = \sum_{t=1}^{T} c_t(\bar{x}^t) + \sum_{t=1}^{T} \left(c_t(x^t) - c_t(\bar{x}^t)\right).$$

Specifically, it holds that

$$\begin{aligned} c_t(x^t) - c_t(\bar{x}^t) &= (1-\alpha)\sum_{j\in N} q_j(\mathbf{b}^t)(x^t - b_j^t)^2 + \alpha(x^t - s)^2 \\ &\quad - (1-\alpha)\sum_{j\in N} q_j(\mathbf{b}^t)(\bar{x}^t - b_j^t)^2 - \alpha(\bar{x}^t - s)^2 \\ &= (1-\alpha)\sum_{j\in N} q_j(\mathbf{b}^t)(x^t - \bar{x}^t)(x^t + \bar{x}^t - 2b_j^t) + \alpha(x^t - \bar{x}^t)(x^t + \bar{x}^t - 2s). \end{aligned}$$

All opinions are in $[0,1]$. Hence, we have that $(x^t + \bar{x}^t - 2b_j^t) \leq 2$ and $(x^t + \bar{x}^t - 2s) \leq 2$. Due to the fact that the weights are normalized, we obtain for any x and x'

$$\begin{aligned} c_t(x) - c_t(x') &\leq (1-\alpha)\sum_{j\in N} 2q_j(\mathbf{b}^t)|x - x'| + 2\alpha|x - x'| \\ &= 2|x - x'|. \end{aligned}$$

This result is useful in two ways as it shows that $c_t(x^t) - c_t(\bar{x}^t) \leq 2|x - x'|$ and that the cost functions are 2-Lipschitz.

Second, using theorem 5.1 we can bound this distance between consecutive opinions. In particular, x^T minimizes the function $\sum_{t=1}^{T} c_t(x)$ and $\bar{x}^T$ minimizes the function $\sum_{t=1}^{T+1} c_t(x)$. We now want to specify the values of the parameters of the theorem. Regarding the convexity we calculate the second derivative of the two functions.

$$c_t'(x) = (1-\alpha)\sum_{j\in N} 2q_j(\mathbf{b}^t)(x - b_j^t) + 2\alpha(x^t - s).$$

$$c_t''(x) = 2(1-\alpha) + 2\alpha = 2.$$

Thus, we obtain $\frac{d^2}{dx^2}\sum_{t=1}^{T} c_t(x) = 2T$ and $\frac{d^2}{dx^2}\sum_{t=1}^{T+!} c_t(x) = 2(T+1) \geq 2T$. As a result we can get $\eta = \frac{1}{2T}$. We have already showed that the difference $\sum_{t=1}^{T+1} c_t(x) - \sum_{t=1}^{T} c_t(x) = c_T(x)$ is 2-Lipschitz. Hence, we have

$$|x^T - \bar{x}^T| \leq \frac{1}{T}.$$

The third step is to find how $\sum_{t=1}^{T} c_t(\bar{x}^t)$ and $\min_{x\in[0,1]}\sum_{t=1}^{T} c_t(x) = \sum_{t=1}^{T} c_t(\bar{x}^T)$ are related. We will use the method of induction to show that

$$\sum_{t=1}^{T} c_t(\bar{x}^t) \leq \min_{x\in[0,1]}\sum_{t=1}^{T} c_t(x).$$

At $t = 1$ we clearly have $c_1(\bar{x}^1) \leq \min_{x\in[0,1]} c_1(x)$.

Let us suppose that at T it holds that $\sum_{t=1}^{T} c_t(\bar{x}^t) \leq \min_{x\in[0,1]} \sum_{t=1}^{T} c_t(x)$. At $T+1$ we can break the sum of costs into the cost of the last round and the sum of the rest of the costs which is bounded by induction. If instead of the best fixed opinion for the cumulative costs the agent expresses opinion $\mathbf{x}^{T+1}$, then the sum of costs from 1 until T increases. Formally, we have

$$\begin{aligned}\sum_{t=1}^{T+1} c_t(\bar{x}^t) &\leq c_T(\bar{x}^T) + \sum_{t=1}^{T} c_t(\bar{x}^t) \\ &= c_T(\bar{x}^T) + \min_{x\in[0,1]} \sum_{t=1}^{T} c_t(x) \\ &= \sum_{t=1}^{T+1} c_t(\bar{x}^T).\end{aligned}$$

Combining the results of the three parts we obtain

$$\sum_{t=1}^{T} c_t(x_t) - \min_{x\in[0,1]} \sum_{t=1}^{T} c_t(x) \leq \sum_{t=1}^{T} |x_t - x_{t+1}| = H_T,$$

where H_T is the T-th harmonic number and it is known that $H_T = O(logT)$. Subsequently, we conclude that the Follow-The-Leader algorithm is no-regret.

□

5.5 Convergence to the Unique Equilibrium point

If the equilibrium of the game is unique because strong diagonal concavity holds, according to 4.5 we have a continuous time dynamic system that converges to it. A question that occurs is whether this property can be extended to a discrete time dynamic, that would converge to the Nash equilibrium of the game. Apart from the negative definiteness of the matrix $(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)$, for the discrete equivalent we will need an extra condition that signifies how much the strategies can move at each step.

In order to translate these conditions in the setting of coevolutionary games we need to determine the vector

$$g(\mathbf{x}) = \begin{pmatrix} (1-\alpha_1)\sum_{j\neq 1} q_{1j}(\mathbf{x}_{-1})x_j + \alpha_1 s_1 - x_1 \\ \vdots \\ (1-\alpha_n)\sum_{j\neq n} q_{nj}(\mathbf{x}_{-n})x_j + \alpha_n s_n - x_n \end{pmatrix}.$$

Our goal creating this dynamic is to make the norm of $g(\mathbf{x}, \mathbf{r})$ equal to zero as time progresses, because the unique Nash equilibrium of the game is simultaneously the unique minimizer of $\|g(\mathbf{x}, \mathbf{r})\|$. Let

$$f(\mathbf{x}) = \frac{1}{2}\|g(\mathbf{x}, \mathbf{r})\|^2$$

be the function we attempt to minimize. Then by the definition of the Jacobian matrix we obtain

$$\nabla f(\mathbf{x}) = G(\mathbf{x}, \mathbf{r})^T g(\mathbf{x}, \mathbf{r}).$$

We assume that the function f is L-smooth, i.e. it holds that

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|^2.$$

The follow the leader update rule in this game is

$$x_i^{t+1} = \frac{t}{t+1}x_i^t + \frac{1}{t+1}\left[(1-\alpha_i)\sum_{j\neq i} q_{ij}(\mathbf{x}_{-i}^t)x_j^t + \alpha_i s_i\right],$$

which can also be written as

$$\mathbf{x}^{t+1} = \mathbf{x}^t + \frac{1}{t+1}g(\mathbf{x}^t)$$

which is the gradient descent algorithm with step size $\frac{1}{t}$. From the first update rule we can see that at any round the opinions remain in $[0, 1]$. In this case, we show that if all the players choose their strategies according to this update rule their strategies will converge to the unique Nash equilibrium of the game.

Theorem 5.3. *If $f(\mathbf{x})$ is L-smooth and $G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T$ is negative definite, then the opinion vector $\mathbf{x}^t$ produced by the update rule*

$$\mathbf{x}^t = \mathbf{x}^{t-1} + \frac{1}{t}g(\mathbf{x}^{t-1})$$

converges to the unique Nash equilibrium.

Proof. Because f is L-smooth we obtain

$$\begin{aligned}
f(\mathbf{x}^t) &\leq f(\mathbf{x}^{t-1}) + \nabla f(\mathbf{x}^{t-1})^T(\mathbf{x}^t - \mathbf{x}^{t-1}) + \frac{L}{2}\|\mathbf{x}^t - \mathbf{x}^{t-1}\|^2 \\
&= f(\mathbf{x}^{t-1}) + \nabla f(\mathbf{x}^{t-1})^T\left(\frac{1}{t}g(\mathbf{x}^{t-1}, \mathbf{r})\right) + \frac{L}{2}\left\|\frac{1}{t}g(\mathbf{x}^{t-1}, \mathbf{r})\right\|^2 \\
&= f(\mathbf{x}^{t-1}) + \frac{1}{t}g(\mathbf{x}^{t-1}, \mathbf{r})^T G(\mathbf{x}^{t-1}, \mathbf{r})g(\mathbf{x}^{t-1}, \mathbf{r}) + \frac{L}{2t^2}\left\|g(\mathbf{x}^{t-1}, \mathbf{r})\right\|^2.
\end{aligned}$$

Due to the fact that $G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})$ is negative definite for any $\mathbf{x}$ in S, there exists a positive constant m such that

$$g(\mathbf{x}, \mathbf{r})^T(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)g(\mathbf{x}, \mathbf{r}) \leq -m\left\|g(\mathbf{x}, \mathbf{r})\right\|^2.$$

Given that $g(\mathbf{x}, \mathbf{r})^T G(\mathbf{x}, \mathbf{r})g(\mathbf{x}, \mathbf{r}) = \frac{1}{2}g(\mathbf{x}, \mathbf{r})^T(G(\mathbf{x}, \mathbf{r}) + G(\mathbf{x}, \mathbf{r})^T)g(\mathbf{x}, \mathbf{r})$, it holds that

$$\begin{aligned}
f(\mathbf{x}^t) &\leq f(\mathbf{x}^{t-1}) - \frac{m}{2t}\left\|g(\mathbf{x}^{t-1}, \mathbf{r})\right\|^2 + \frac{L}{2t^2}\left\|g(\mathbf{x}^{t-1}, \mathbf{r})\right\|^2 \\
&= f(\mathbf{x}^{t-1}) + 2(-\frac{m}{2t} + \frac{L}{2t^2})f(\mathbf{x}^{t-1}).
\end{aligned}$$

Let α be a constant less than m, for example $\frac{m}{2}$. By the time t becomes at least $\frac{L}{m-a}$ we have that $-\frac{m}{t} + \frac{L}{t^2} \leq -\frac{\alpha}{t}$. We denote that time by τ. As a result, we obtain

$$\begin{aligned} f(\mathbf{x}^t) &\leq f(\mathbf{x}^{t-1}) - \frac{\alpha}{t} f(\mathbf{x}^{t-1}) \\ &\leq e^{-\frac{\alpha}{t}} f(\mathbf{x}^{t-1}) \\ &\leq e^{-a \sum_{i=\tau+1}^{t} \left(\frac{1}{i}\right)} f(\mathbf{x}^{\tau}) \\ &\approx e^{-\alpha \ln t + \alpha \sum_{i=1}^{\tau} \left(\frac{1}{i}\right)} f(\mathbf{x}^{\tau}) \\ &= c e^{-\alpha \ln(t)} f(\mathbf{x}^{\tau}) \\ &= \frac{c}{t^{\alpha}} f(\mathbf{x}^{\tau}). \end{aligned}$$

□

The proof above is general for any game that satisfies the conditions, as long as there is a way to ensure that the strategies that result from the dynamic will remain in the feasible strategy set.

5.6 Conditions for Diagonal Strict Concavity

In order to apply the result of the previous section we want to determine a set of conditions of the weight functions that are sufficient for proving that diagonal strict concavity holds in the game we are interested in. In fact, the definition of the diagonal strict concavity states that for every pair of distinct opinion vectors $\mathbf{x}^*$, $\bar{\mathbf{x}}$ we have

$$(\mathbf{x}^* - \bar{\mathbf{x}})^T (g(\bar{\mathbf{x}}) - g(\mathbf{x}^*)) > 0.$$

If we substitute g with its components, then we equivalently obtain

$$\sum_{i=1}^{n} (x_i^* - \bar{x}_i)^2 + \sum_{i=1}^{n} (x_i^* - \bar{x}_i)(1 - \alpha_i) \left(\sum_{j \neq i} (\bar{\mathbf{x}}_{-i}) \bar{x}_j - \sum_{j \neq i} q_{ij}(\mathbf{x}_{-i}^*) x_j^* \right) > 0.$$

We now define a weight matrix $Q(\mathbf{x})$ of the game, such that

$$Q(\mathbf{x})_{ij} = \begin{cases} 0, & \text{if } i = j \\ (1 - \alpha_i) q_{ij}(\mathbf{x}_{-i}), & \text{if } i \neq j \end{cases}$$

Therefore, the condition of diagonal strict concavity is written as

$$\|\mathbf{x}^* - \bar{\mathbf{x}}\|^2 > (\mathbf{x}^* - \bar{\mathbf{x}})^T Q(\mathbf{x}^*) \mathbf{x}^* - (\mathbf{x}^* - \bar{\mathbf{x}})^T Q(\bar{\mathbf{x}}) \bar{\mathbf{x}}$$

or equivalently

$$\|\mathbf{x}^* - \bar{\mathbf{x}}\|^2 > (\mathbf{x}^* - \bar{\mathbf{x}})^T Q(\mathbf{x}^*)(\mathbf{x}^* - \bar{\mathbf{x}}) + (\mathbf{x}^* - \bar{\mathbf{x}})^T (Q(\mathbf{x}^*) - Q(\bar{\mathbf{x}})) \bar{\mathbf{x}}. \tag{5.2}$$

This condition is very similar to the condition that ensures the convergence of the FJ model if we consider that the second term of the right side is very small, which means that the weight functions are close between the two opinion vectors. This formally is stated in the conditions

1. $\|Q(\mathbf{x}^*) - Q(\bar{\mathbf{x}})\| \leq \gamma \|\mathbf{x}^* - \bar{\mathbf{x}}\|$

2. $\|Q(\mathbf{x})\| \leq \rho$, for any $\mathbf{x} \in [0,1]^n$

where γ and ρ are positive constants.

By condition 1 we obtain $(\mathbf{x}^* - \bar{\mathbf{x}})^T Q(\mathbf{x}^*)(\mathbf{x}^* - \bar{\mathbf{x}}) \leq \rho \|\mathbf{x}^* - \bar{\mathbf{x}}\|^2$. Moreover, by the Cauchy-Schwartz inequality and the condition 2 we have

$$\begin{aligned}
|(\mathbf{x}^* - \bar{\mathbf{x}})^T Q(\mathbf{x}^*) - Q(\bar{\mathbf{x}}))\bar{\mathbf{x}}| \leq & \|\mathbf{x}^* - \bar{\mathbf{x}}\| \|(Q(\mathbf{x}^*) - Q(\bar{\mathbf{x}}))\bar{\mathbf{x}})\| \\
\leq & \|\mathbf{x}^* - \bar{\mathbf{x}}\| \|Q(\mathbf{x}^*) - Q(\bar{\mathbf{x}})\| \|\bar{\mathbf{x}}\| \\
\leq & \gamma \|\mathbf{x}^* - \bar{\mathbf{x}}\|^2 \sqrt{n}
\end{aligned}$$

As a result, if $\rho + \gamma\sqrt{n} < 1$, then the inequality 5.2 holds and the equilibrium of the game is unique.

5.7 Open Problems

In this section we mention a few immediate questions that remain unanswered. The first one, that would complete our initial efforts, would be to find a class of weight functions or even a specific one that satisfies the conditions of theorem 5.3. Granted this is realized, then we would have a subset of coevolutionary opinion formation games that have a dynamic that converges to a Nash equilibrium. However, it is also interesting to study the case where more than one Nash equilibria exist. One possible direction would be to check whether local convergence for best response holds for differential functions as well as if there is a connection between the areas where the matrix $G(\mathbf{x}) + G(\mathbf{x})^T$ is negative definite and local convergence.

Bibliography

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

[2] Arnab Bhattacharyya, Mark Braverman, Bernard Chazelle, and Huy L. Nguyen. On the convergence of the hegselmann-krause system. In *Innovations in Theoretical Computer Science*, ITCS 2013, pages 61–66, 2013.

[3] Kshipra Bhawalkar, Sreenivas Gollapudi, and Kamesh Munagala. Coevolutionary opinion formation games. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC 2013, pages 41–50, 2013.

[4] Vittorio Bilo, Angelo Fanelli, and Luca Moscardelli. Opinion formation games with dynamic social influences. *Web and Internet Economics*, pages 444–458, 2016.

[5] 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, pages 57–66, 2011.

[6] KC Border. Introduction to correspondences, November 2013.

[7] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[8] G. Bullo, J. Cortes, and S Martinez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.

[9] Bernard Chazelle. The dynamics of influence systems. In *proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS 2012, 2012.

[10] Bernard Chazelle. The convergence of bird flocking. *J. ACM*, 61(4), 2014.

[11] Po-An Chen, Yi-le Chen, and Chi-Jen Lu. Bounds on the price of anarchy for a more general class of directed graphs in opinion formation games. *Operational Research Letters*, 44(6):808–811, 2016.

[12] Flavio Chierichetti, Jon M. Kleinber, and Sigal Oren. On discrete preferences and coordination. In *ACM Conference on Electronic Commerce*, EC 2013, pages 233–250, 2013.

[13] P. Dandekar, A. Goel, and D. Lee. Biased assimilation, homophily and the dynamics of polarization. In *Workshop on Internet and Network Economics*, WINE 2012, 2012.

[14] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixed beliefs among interacting agents. *Advances in Complex Systems*, 3:87–98, 2001.

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

[16] R. Durrett, J. P. Gleeson, A. L. Lloyd, P. J. Mucha, F. Shi, D. Sivakoff, J. E. S. Socolar, and C. Varghese. Graph fission in an evolving voter model. *Proceedings of National Academy of Sciences*, 109(10):3682–3687, 2012.

[17] Eyal Even-Dar, Yishay Mansour, and Uri Nadav. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of computing*, STOC 2009, pages 523–532, 2009.

[18] Diodato Ferraioli, Paul W. Goldberg, and Carmine Ventre. Decentralized dynamics for finite opinion games. *Theor. Comput. Sci.*, 648(C), 2016.

[19] Dimirtis Fotakis, Vardis Kandiros, Vasilis Kontonis, and Stratis Skoulakis. Opinion dynamics with limited information. In *Web and Internet Economics - 14th International Conference*, WINE 2018, pages 282–296, 2018.

[20] Noah E. Friedkin and Eugene C. Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–205, 1990.

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

[22] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.

[23] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artifical Societies and Social Simulation*, 5(3), 2002.

[24] Julien M. Hendrickx and Vincent D. Blondel. Convergence of linear and non-linear versions of vicsek's model, 2006.

[25] A. D¿ Henry, Pralatm P, and C. Q. Zhang. Emergence of segregation in evolving social networks. *Proceedings of National Academy of Sciences*, 108:8605–8610, 2011.

[26] P. Holme and M. E. J. Newman. Nonequilibrium phase transitions in the coevolution of networks and opinions. *Physical review. E, Statistical, nonlinear, soft matter physics*, 74:056108, 2006.

[27] Shizuo Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3):457–459, 1941.

[28] Jan Lorenz. A stabilization theorem for dynamics of continuous opinions. *Physica A: Statistical Mechanics and its Applications*, 355, 2007.

[29] S. Martinez, F. Bullo, J. Cortes, and E. Frazzoli. On synchronous robotic networks part i: Models, tasks and complexity. 52(12):2199–2213, 2007.

[30] Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50:169–182, 2005.

[31] Uri Nadav. *Protocols for Selfish Agents*. PhD thesis, Tel Aviv University, 2009.

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

[33] Asu Ozdaglar. MIT game theory with engineering applications 6.254: Lecture 7, February 2010.

[34] Julia Parrish and William Hammer. Animal groups in three dimensions: How species aggregate, 1997.

[35] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH 1987, pages 25–34, 1987.

[36] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.

[37] Efstratios Panteleimon Skoulakis. *Natural and Efficient Dynamics through Convex Optimization*. PhD thesis, National Technical University of Athens, 2019.

[38] Gilbert Strang. *Linear Algebra and Its Applications*. Thomson Brooks/Cole, Belmont, California, 2006.

[39] Vasilis Syrgkanis. MIT algorithmic game theory and data science 6.853: Lecture 3, February 2017.

[40] B. Touri and A. Nedic. Discrete-time opinion dynamics. In *Conference Record of the Forthy fifth Asilomar Conference on Signals, Systems and Computers*, pages 1172–1176, 2011.

[41] Nisheeth K. Vishnoi. A mini-course on convex optimization, 2015.

[42] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal. Interacting agents and continuous opinions dyncamics. *Lecture Notes in Economics and Mathematical Systems*, 521:225–242, 2002.

[43] Mehmet Ercan Yildiz, Asuman E. Ozdaglar, daron Acemoglu, Amin Saberi, and Scaglione Anna. Binary opinion dynamics with stubborn agents. *ACM Trans. Economics and Comput.*