

Εθνικό Μετσόβιο Πολυτεχνείο

Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Αλγόριθμοι και Δυναμικές για Υπολογισμό Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

ΚΩΝΣΤΑΝΤΙΝΟΣ Α. ΤΣΟΠΕΛΑΣ

Επιβλέπων: Δημήτριος Φωτάκης

Αναπληρωτής Καθηγητής Ε.Μ.Π.



Εθνικό Μετσόβιο Πολυτεχνείο

Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Αλγόριθμοι και Δυναμικές για Υπολογισμό Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

ΚΩΝΣΤΑΝΤΙΝΟΣ Α. ΤΣΟΠΕΛΑΣ

Επιβλέπων: Δημήτριος Φωτάκης

Αναπληρωτής Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 28η Ιουνίου 2022.

Δημήτριος Φωτάκης Αριστείδης Παγουρτζής Παναγιώτης Μερτικόπουλος Αναπληρωτής Καθηγητής Ε.Μ.Π. Καθηγητής Ε.Μ.Π. Principal Researcher C.N.R.S.

Κωνσταντίνος Α	Α. Τσόπελας				
Διπλωματούχος Ε.Μ.Π.	Ηλεκτρολόγος	Μηχανικός	και	Μηχανικός	Υπολογιστών

Copyright © Κωνσταντίνος Α. Τσόπελας, 2022. Με επιφύλαξη παντός δικαιώματος. All rights reserved.

Απαγορεύεται η αντιγραφή, αποθήκευση και διανομή της παρούσας εργασίας, εξ ολοκλήρου ή τμήματος αυτής, για εμπορικό σκοπό. Επιτρέπεται η ανατύπωση, αποθήκευση και διανομή για σκοπό μη κερδοσκοπικό, εκπαιδευτικής ή ερευνητικής φύσης, υπό την προϋπόθεση να αναφέρεται η πηγή προέλευσης και να διατηρείται το παρόν μήνυμα. Ερωτήματα που αφορούν τη χρήση της εργασίας για κερδοσκοπικό σκοπό πρέπει να απευθύνονται προς τον συγγραφέα.

Οι απόψεις και τα συμπεράσματα που περιέχονται σε αυτό το έγγραφο εκφράζουν τον συγγραφέα και δεν πρέπει να ερμηνευθεί ότι αντιπροσωπεύουν τις επίσημες θέσεις του Εθνικού Μετσόβιου Πολυτεχνείου.

Περίληψη

Τα μοντέλα διαμόρφωσης άποψης προσπαθούν να κωδικοποιήσουν το πώς οι άνθρωποι διαμορφώνουν τις απόψεις τους αλληλεπιδρόντας ο ένας με τον άλλον. Οι άνθρωποι αναπαρίστανται, εν γένει, ως κόμβοι σε ένα (κοινωνικό) δίκτυο, οι ακμές του οποίου αναπαριστούν τις σχέσεις μεταξύ τους. Κατόπιν, κάθε άτομο υποτίθεται ότι διαμορφώνει την άποψή του βασιζόμενο σε κάποιου είδους μέσο όρο των απόψεων των φίλων του.

Σε αυτή την εργασία, μελετάμε ένα παιγνιοθεωρητικό μοντέλο διαμόρφωσης άποψης, που ονομάζεται συνεξελικτικά παίγνια διαμόρφωσης άποψης, το οποίο σημαίνει ότι αυτό το μοντέλο προσπαθεί να περιλάβει το ότι οι σχέσεις εξελίσσονται μαζί με τις απόψεις. Αυτά τα παίγνια ανήκουν στην κλάση των κοίλων παιγνίων, συνεπώς έχουν πάντα ισορροπίες Nash. Ο στόχος μας είναι να τις υπολογίσουμε αποδοτικά. Η προσέγγισή μας είναι να χρησιμοποιήσουμε ένα πρόσφατο αποτέλεσμα των Μερτικόπουλου και Zhou, στην τομή της θεωρίας παιγνίων και της κυρτής βελτιστοποίησης, το οποίο, αδρά, εγγυάται γρήγορη σύγκλιση ενός mirror descent αλγορίθμου (Dual Averaging) σε ισορροπία, σε κοίλα παίγνια όπου ισχύει η συνθήκη diagonal strict concavity του Rosen. Συνεπώς, κατά βάση η δική μας εργασία ήταν να αναζητήσουμε ενδιαφέρουσες υποκατηγορίες συνεξελικτικών παιγνίων διαμόρφωσης άποψης στις οποίες να ισχύει αυτή η συνθήκη. Η εργασία μας χωρίζεται σε δύο τμήματα.

Αρχικά, κάποια θεωρητικά αποτελέσματα, όπου δείχνουμε ότι περιορισμένα στιγμιότυπα των παιγνίων αυτών είναι diagonally strictly concave (συγκεκριμένα, όταν η αυτοπεποίθηση / αδράνεια των παικτών είναι επαρκώς μεγάλη).

Κατόπιν, λόγω της αυξημένης δυσκολίας στο να εκμαιεύσουμε περαιτέρω θεωρητικά αποτελέσματα, μελετάμε το μοντέλο μέσω προσομοιώσεων. Τρέχοντας τον αλγόριθμο Dual Averaging σε αρκετά στιγμιότυπα συνεξελικτικών παιγνίων διαμόρφωσης άποψης, βλέπουμε ότι έχει πολύ καλή επίδοση, φτάνοντας σε ισορροπία μετά από σχετικά λίγους γύρους. Επίσης, βρίσκουμε κάποια αριθμητικά αντιπαραδείγματα του diagonal strict concavity, αλλά πολύ λίγα. Το οποίο θα μπορούσε να δικαιολογεί το γιατί συναντάμε θεωρητικές δυσκολίες, παρόλο που ο αλγόριθμος δείχνει να αποδίδει πολύ καλά σε τυπικές περιπτώσεις.

Λέξεις κλειδιά

Διαμόρφωση Άποψης, Συνεξελικτικά Μοντέλα, Υπολογισμός Ισορροπιών, Κυρτή Βελτιστοποίηση, Μέθοδοι Κατοπτρικής Κλίσης

Abstract

Opinion formation models try to capture how people form their opinions through interacting with each other. In general, people are represented as nodes in a (social) network, edges represent the relationships between them. Then, each person is assumed to form their opinion based on some kind of average of the opinions of their friends.

In this work, we study a game theoretic model of opinion formation, called asymmetric coevolutionary opinion formation games, which means that this model tries to capture the fact that relationships evolve together with opinions. These games belong in the class of concave games, therefore always admit Nash equilibria; our goal is to compute them efficiently. Our approach is to use a recent result by Mertikopoulos and Zhou, on the intersection of game theory and convex optimization, which, crudely, guarantees fast convergence of a mirror descent algorithm (Dual Averaging) to equilibrium, in concave games where Rosen's diagonal strict concavity condition holds. Thus, the bulk of our efforts goes to seeking interesting subclasses of asymmetric coevolutionary opinion formation games where this condition holds. Our work can be separated into two main parts.

We begin with some theoretical results, where we show that restricted instances of asymmetric coevolutionary opinion formation games are diagonally strictly concave (specifically, when the self-confidence / stubborness of the players is sufficiently large).

Then, due to the increased difficulty in establishing theoretical results, we turn to studying the model via simulations. By running Dual Averaging on several instances of asymmetric coevlutionary opinion formation games, we see that the algorithm performs very well, reaching equilibrium in relatively few rounds. Moreover, we find some numerical counterexamples to diagonal strict concavity, but very few; which could explain why one encounters theoretical difficulties, while the algorithm runs very well in typical instances.

Key words

Opinion Formation, Coevolutionary Models, Equilibrium Computation, Convex Optimization, Mirror Descent Methods

Ευχαριστίες

Η ολοκλήρωση της διπλωματικής μου εργασίας σηματοδοτεί το τέλος της πορείας μου στη σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών Υπολογιστών. Σε αυτήν την πορεία και την ολοκλήρωσή της δεν ήταν λίγοι αυτοί που με βοήθησαν.

Κατόπιν τούτου, ευχαριστώ θερμά, καταρχάς, τον επιβλέποντα καθηγητή αυτής της διατριβής, κ. Δημήτρη Φωτάκη, για την καθοδήγηση και την εμπιστοσύνη που μου έδειξε, αλλά και για το ενδιαφέρον που μου εμφύσησε εξ αρχής για την θεωρητική πληροφορική μέσα από τα μαθήματα των Διακριτών Μαθηματικών και των Αλγορίθμων.

Επίσης, ευχαριστώ την οικογένειά μου, χωρίς την συμπαράσταση της οποίας δεν θα μπορούσα να είχα φτάσει εδώ που είμαι σήμερα, καθώς και τους φίλους μου, που με βοήθησαν εξίσου.

Κωνσταντίνος Α. Τσόπελας, Αθήνα, 28η Ιουνίου 2022

Περιεχόμενα

П	ερίλην	η	5
Al	strac		7
Ετ	χαριο	τίες	9
Па	εριεχό	ιενα	11
1.	Εκτ	ταμένη Ελληνική Περίληψη	13
	1.1	Εισαγωγή	13
		1.1.1 Κεντρικά Ερωτήματα	15
		1.1.2 Συνεισφορά	16
	1.2	Θεωρητικό Υπόβαθρο	17
		1.2.1 Θεωρία Παιγνίων και Κοίλα Παίγνια	17
		1.2.2 Κυρτή Βελτιστοποίηση	18
		1.2.3 Κυρτή Βελτιστοποίηση και Παίγνια	19
	1.3	Κοίλα Παίγνια και ο Αλγόριθμος Dual Averaging	20
		1.3.1 Σύγκλιση σε Ισορροπία	21
		1.3.2 Ταχύτητα Σύγκλισης	21
		1.3.3 Εφαρμογή στα Συνεξελικτικά Παίγνια	22
	1.4	Ο Αλγόριθμος Dual Averaging στα Συνεξελικτικά Παίγνια Διαμόρφωσης	
		Άποψης	24
		1.4.1 Θεωρητικά Αποτελέσματα στην Αυστηρή Διαγώνια Κοιλότητα	24
		1.4.2 Πειράματα και Προσομοιώσεις	28
Ka	είμενο	στα αγγλικά	37
2.	Intr	<mark>duction</mark>	37
	2.1	Contributions	38
	2.2	Related Work	41
	2.3	Overview of the chapters	41
3.	Tech	nical Background	43
	3.1	Games	43
	3.2	Nash Equilibrium	44
	3.3	Concave Games	45
		3.3.1 Diagonal Strict Concavity	46
	3.4	Convex Optimization	47

		3.4.1	Online Convex Optimization				 50
		3.4.2	Convex Optimization and Games				 53
	3.5	Socially	y Concave Games				
4.	Opir	nion For	mation				. 57
	4.1		ot model				
	4.2		n-Johnsen model				
	4.3		Theoretic viewpoint				
		4.3.1	Best Response Dynamics				
		4.3.2	Opinion Dynamics				
		4.3.3	Potential				
		4.3.4	Implications				
	4.4	Coevol	ution				
		4.4.1	Bounded Confidence - HK and DW				
		4.4.2	Asymmetric KNN Games				
		4.4.3	Asymmetric Coevolutionary Opinion Formation Games .				
5	Con	ooyo Co	mes and Dual Averaging				. 65
J.	5.1		onal Stability				
	5.2						
	3.2	5.2.1	veraging				
	5.3	•	Choice Map				
			al: Fenchel Coupling				
	5.4 5.5		gence				
	3.3		gence Speed				
			Equilibrium gap function				
	F (5.5.2	Running length				
	5.6	Applica	ation on Coevolutionary Games	•		•	 76
6.			ing on Coevolutionary Opinion Formation Games				
	6.1		tical Results on Diagonal Strict Concavity				
		6.1.1	Specialization for Normalized Independent Weights				
			Remarks				 88
	6.2		tions				
		6.2.1	Choice of Functions				
		6.2.2	Convergence in Small Instances				
		6.2.3	Convergence to Equilibrium for Large Numbers of Players				
		6.2.4	Diagonal Strict Concavity counterexample				
	6.3	Social (Concavity				 96
Βι	βλιογι	ραφία .		•	•	•	 . 99
По	ιράρτ	ημα					103
Α.	Omi	tted Pro	ofs and Calculations				. 103
R	Sim	ılation F	Plots				. 105
υ.	B.1		gence to Equilibrium for Large Numbers of Players				

Κεφάλαιο 1

Εκτεταμένη Ελληνική Περίληψη

Στο κεφάλαιο αυτό, θα εκθέσουμε περιληπτικά το μοντέλο και τα κύρια αποτελέσματα της εργασίας μας, παραλείποντας αρκετές λεπτομέρειες, όπως για παράδειγμα αποδείξεις.

1.1 Εισαγωγή

Ο τρόπος που οι άνθρωποι διαμορφώνουν τις απόψεις τους μελετάται εδώ και αρκετές δεκαετίες, από τη σκοπιά διαφόρων επιστημονικών πεδίων (π.χ. ψυχολογία, πολιτικές επιστήμες, κοινωνιολογία). Αυτό το επιστημονικό ενδιαφέρον έχει κορυφωθεί τα τελευταία χρόνια, ιδιαίτερα λόγω της προόδου του Διαδικτύου και των κοινωνικών δικτύων, τα οποία από τη μία μεριά κάνουν διαθέσιμα προς ανάλυση πάρα πολλά δεδομένα σχετικά με ανθρώπινες αλληλεπιδράσεις, και από την άλλη προσφέρουν διάφορες ευκαιρίες για κερδοφορία μέσω προβολής (π.χ. διαφημίσεις), όπου η γνώση της συμπεριφοράς των απόψεων του κοινού μπορεί να μεταφραστεί σε μεγαλύτερα κέρδη. Το γενικό πλαίσιο που μας ενδιαφέρει εμάς, λοιπόν, είναι οι τρόποι μαθηματικής μοντελοποίησης της διαδικασίας διαμόρφωσης άποψης, και το τι υπολογιστικές ιδιότητες έχουν αυτά τα μοντέλα.

Μία από τις πιο παλιές εργασίες στην κατεύθυνση αυτή είναι το μοντέλο του Morris Degroot [DeGr74]. Σε αυτό, υπάρχει ένας πεπερασμένος αριθμός N ανθρώπων (τους οποίους ονομάζουμε και πράκτορες) που αποτελούν το υποτιθέμενο κοινωνικό δίκτυο, και κάθε ένας τους διατηρεί έναν πραγματικό αριθμό (συνήθως στο [0,1]). Αυτός αναπαριστά την άποψή του, την οποία στη συνέχεια αναθεωρεί, σε διακριτούς γύρους, λαμβάνοντας υπόψη τις απόψεις όλων των άλλων πρακτόρων, μέσω ενός βεβαρυμένου μέσου όρου. Συγκεκριμένα:

$$x_{i,t+1} = \frac{\sum_{j=1}^{N} w_{ij} x_{j,t}}{\sum_{j=1}^{N} w_{ij}} = \sum_{j=1}^{N} q_{ij} x_{j,t}$$

όπου το $x_{i,t} \in [0,1]$ αναπαριστά την άποψη του πράκτορα i στο γύρο t. Τα βάρη $w_{ij} \geq 0$ αναπαριστούν την βαρύτητα που έχει για τον i η άποψη του j, την κρίση του για το ποιος από τους υπολοίπους είναι ο πιο έμπιστος / σοφός / πεπειραμένος κλπ. Τέλος, με $q_{ij} = \frac{w_{ij}}{\sum_{j=1}^N w_{ij}}$ συμβολίζουμε τα κανονικοποιημένα βάρη.

Το μοντέλο DeGroot, όπως λέγεται στη βιβλιογραφία, παρά την απλότητά του, παρουσιάζει αρκετές ενδιαφέρουσες ιδιότητες (π.χ. [Golu10, Amir21]), αλλά αυτό που μας ενδιαφέρει κυρίως είναι ότι θέτει πολλά από τα θεμέλια και για μετέπειτα μοντέλα. Το ότι οι απόψεις αναπαρίστανται από αριθμούς, το ότι κάθε πράκτορας αντιστοιχεί βάρη σε όλους τους άλλους, και το ότι η διαδικασία διαμόρφωσης άποψης συνίσταται από διακριτούς γύρους στους οποίους κάθε πράκτορας μετακινεί την άποψή του σε κάποιο μέσο όρο των απόψεων των υπολοίπων από τον προηγούμενο γύρο, όλα αυτά λίγο πολύ διατηρούνται σε πολλά από τα μοντέλα που ακολούθησαν. Ένα εξέχον παράδειγμα είναι το μοντέλο των Friedkin και Johnsen [Frie90], μία επέκταση του μοντέλου DeGroot, στην οποία κάθε πράκτορας i έχει επιπλέον μία σταθερή, κρυφή, εσωτερική άποψη, την οποία λαμβάνει υπόψη του κάθε φορά που υπολογίζει έναν μέσο όρο, και την οποία συμβολίζουμε με $s_i \in [0,1]$. Επιπλέον, διαθέτει και μία ακόμα σταθερά την οποία συνήθως ονομάζουμε παράγοντα "αυτοπεποίθησης", τη συμβολίζουμε με $\rho_i \in [0,1]$ και αναπαριστά το προς τα πού κλίνει ο συμβιβασμός που είναι αναγκασμένος να κάνει ο παίκτης μεταξύ της εσωτερικής του άποψης και των απόψεων των "φίλων" του (0 σημαίνει ότι τον ενδιαφέρει μόνο η κοινωνική συμφωνία, 1 ότι τον ενδιαφέρει μόνο η εσωτερική συμφωνία). Κατόπιν τούτου, ο κανόνας ανανέωσης των απόψεων στο μοντέλο FJ (Friedkin - Johnsen) είναι:

$$x_{i,t+1} = (1 - \rho_i) \sum_{j=1}^{N} q_{ij} x_{j,t} + \rho_i s_i$$

Ακόμη, αναφέρουμε επιγραμματικά ότι υπάρχει μία σειρά μοντέλων (π.χ. [Hegs02, Deff00]), στα οποία τα βάρη των αλληλεπιδράσεων δεν είναι πλέον σταθερά, αλλά μεταβάλλονται ανάλογα με την απόσταση μεταξύ των απόψεων των ατόμων. Αυτή η ιδέα υπάρχει και στο δικό μας μοντέλο, όπως θα δούμε.

Μία πιο πρόσφατη κατεύθυνση έρευνας συνδέει την διαδικασία διαμόρφωσης άποψης με τη θεωρία παιγνίων. Οι πράκτορες γίνονται παίκτες, καθένας εκ των οποίων θέλει να ελαχιστοποιήσει κάποιο κόστος, το οποίο εξαρτάται από τις απόψεις όλων των παικτών και εκφράζει τη δυσαρέσκεια του παίκτη με το εκάστοτε προφίλ απόψεων. Μία ενδιαφέρουσα και σχετική εργασία είναι αυτή των Bindel, Kleinberg και Oren [Bind11], οι οποίοι ορίζουν και μελετούν ένα παιγνιοθεωρητικό μοντέλο που αντιστοιχεί στο μοντέλο FJ, με την έννοια ότι η δυναμική βέλτιστης απόκρισης (best response, δηλαδή, κάθε παίκτης διαλέγει τη βέλτιστη άποψη με βάση τον προηγούμενο γύρο) σε αυτό το παίγνιο ταυτίζεται με την αρχική δυναμική του FJ μοντέλου.

Συγκεκριμένα, το κόστος κάθε παίκτη στο παιγνιοθεωρητικό αυτό μοντέλο είναι η παρακάτω, τετραγωνική στις απόψεις, συνάρτηση. Παρατηρείστε ότι το κόστος αυτό θα μπορούσε να ερμηνευθεί ως ένας μέσος δύο όρων, ενός που "χρεώνει" τον παίκτη όταν η εκφερόμενη άποψή του απομακρύνεται από τις απόψεις των υπολοίπων, και ειδικά των πιο "κοντινών" του, και ενός που τον χρεώνει όταν η εκφερόμενη άποψή του απομακρύνεται από την εσωτερική, πραγματική του, κατά μία ερμηνεία, άποψη.

$$c_i(x) = (1 - \rho_i) \sum_j q_{ij} (x_i - x_j)^2 + \rho_i (x_i - s_i)^2$$

την οποία μπορούμε να παραγωγίσουμε εύκολα και να βρούμε το σημείο ελαχίστου της, δηλαδή τη βέλτιστη απόκριση, που προκύπτει ίδια με την εξίσωση του FJ μοντέλου.

Αλλά δεν είναι και μόνο αυτό. Ο ίδιος κανόνας μπορεί κανείς, χωρίς πολύ κόπο, να δει ότι συμπίπτει με το να τρέχουν όλοι οι παίκτες κάθοδο κλίσης (gradient descent, θα τη δούμε λίγο αργότερα) με σταθερό βήμα! Επιπλέον, στο [Bind11] αποδεικνύεται ότι (για συμμετρικά δίκτυα, συμμετρικά βάρη δηλαδή), το παίγνιο είναι παίγνιο δυναμικού, το οποίο όπως θα δούμε σημαίνει ότι υπάρχει μία βαθμωτή συνάρτηση της οποίας τα ελάχιστα αντιστοιχούν ταυτίζονται με τα σημεία ισορροπίας του παιγνίου. Οπότε, αυτός είναι άλλος ένας τρόπος να δούμε την δυναμική του FJ, σαν να τρέχουμε κάθοδο κλίσης για την ελαχιστοποίηση μίας βαθμωτής συνάρτησης.

Το απάυγασμα εδώ είναι ότι υπάρχει μία στενή σχέση μεταξύ θεωρίας παιγνίων και βελτιστοποίησης, το οποίο βέβαια εν γένει είναι γνωστό, αλλά βλέπουμε εδώ και την σύγκλιση

των δύο αυτών κατευθύνσεων συγκεκριμένα στο πλαίσιο των μοντέλων διαμόρφωσης άποψης. Αυτό ακριβώς το τρίπτυχο είναι που θα προσπαθήσουμε να εκμεταλλευτούμε και εμείς, απλώς σε ένα πιο στριφνό παιγνιοθεωρητικό μοντέλο και χρησιμοποιώντας λίγο πιο προχωρημένες και πρόσφατες τεχνικές.

Κλείνοντας αυτό το μέρος της εισαγωγής, λοιπόν, ας ορίσουμε το μοντέλο μας. Πρόκειται για μία αρκετά φυσιολογική επέκταση του μοντέλου του [Bind11], η οποία εισήχθη από τους Bhawalkar, Gollapudi και Munagala στο [Bhaw13]. Η επέκταση είναι ότι τα βάρη του δικτύου δεν είναι πλέον σταθερά, αλλά μεταβάλλονται αλληλοεπηρεαζόμενα με τις απόψεις των παικτών. Συγκεκριμένα, οι συναρτήσεις κόστους των παικτών εδώ είναι:

$$c_i(x) = (1 - \rho_i) \sum_{j \neq i} q_{ij}(x_{-i}) \cdot (x_i - x_j)^2 + \rho_i(x_i - s_i)^2$$

όπου οι συναρτήσεις βαρών υποτίθεται ότι προκύπτουν ως $q_{ij}(x_{-i})=F_i(d^i_j,d^i_{-i-j})$, με $d^i_j=|x_j-s_i|$ και F_i μια συνεχή συνάρτηση φθίνουσα στο πρώτο όρισμα και αύξουσα στα υπόλοιπα. Με απλά λόγια, το βάρος που δίνει ο i στον j εξαρτάται από το πόσο κοντά βρίσκεται η άποψη του j από την εσωτερική άποψη του i, σε σχέση όμως με το πόσο κοντά βρίσκονται αντίστοιχα και των υπολοίπων παικτών.

1.1.1 Κεντρικά Ερωτήματα

Μία κεντρική ερώτηση, και αυτή που μας ενδιαφέρει, για οποιοδήποτε μοντέλο διαμόρφωσης άποψης είναι το κατά πόσο οι πράκτορες συγκλίνουν, τελικά, σε κάποιο σημείο ισορροπίας, υπό ποίες συνθήκες, αλλά και πόσο γρήγορα συμβαίνει αυτό.

Για τα μοντέλα που έχουμε αναφέρει, αυτά τα ερωτήματα έχουν εν πολλοίς απαντηθεί. Ήδη από το [DeGr74] είναι γνωστό ότι στο μοντέλο DeGroot οι πράκτορες τελικά συγκλίνουν σε ομοφωνία, δηλαδή όλοι στην ίδια άποψη (το ποια άποψη εξαρτάται από τα βάρη και τις αρχικές απόψεις). Στο μοντέλο FJ, οι πράκτορες και πάλι συγκλίνουν, αυτή τη φορά σε ένα πολύ συγκεκριμένο προφίλ απόψεων (δηλ. ανεξαρτήτως των αρχικών τους απόψεων). Και στις δύο περιπτώσεις, μετέπειτα δουλειές (π.χ. [Ghad13]) έχουν δείξει ότι η σύγκλιση είναι γρήγορη, συγκεκριμένα λογαριθμική στο $\frac{1}{\varepsilon}$ (όπου ε η απόσταση-στόχος από το τελικό όριο).

Μεταβαίνοντας στο πλαίσιο της θεωρίας παιγνίων, που είναι και αυτό που μας ενδιαφέρει, αντί για απλώς σταθερά σημεία, μπορούμε να αξιοποιήσουμε την κεντρική έννοια της ισορροπίας Nash. Για παράδειγμα, οι Bindel, Kleinberg και Oren δείχνουν στο [Bind11] ότι το παιγνιοθεωρητικό τους μοντέλο έχει μία και μοναδική ισορροπία Nash, η οποία πρέπει να ταυτίζεται με το σημείο ισορροπίας του FJ αφού συγκλίνει σε αυτήν η δυναμική βέλτιστης απόκρισης. Άρα, γνωρίζουμε ότι οι παίκτες συγκλίνουν στην ισορροπία, και μάλιστα γρήγορα.

Στο δικό μας μοντέλο, οι ιδιότητες αυτές δεν είναι δυνατό να αποδειχθούν με τον ίδιο τρόπο. Παρόλα αυτά, τα παίγνια που μας ενδιαφέρουν συνεχίζουν να είναι κοίλα παίγνια. Αυτό, όπως θα δούμε, σημαίνει ότι πάντα έχουν τουλάχιστον μία ισορροπία Nash, ο τρόπος υπολογισμού της, όμως, και η υπολογιστική της πολυπλοκότητα είναι ανοικτά προβλήματα. Υπάρχει, ωστόσο, μία πρόσφατη εργασία των Μερτικόπουλου και Zhou [Mert19], στην οποία αποδεικνύεται ότι για μία μεγάλη υποκατηγορία κοίλων παιγνίων (τα αυστηρώς διαγωνίως κοίλα, συγκεκριμένα) η ισορροπία μπορεί να υπολογιστεί αποδοτικά μέσω ενός αλγορίθμου κυρτής βελτιστοποίησης. Αυτό ακριβώς το αποτέλεσμα προσπαθούμε να εκμεταλλευτούμε στην παρούσα εργασία, υπό την έννοια ότι, αν καταφέρουμε να δέιξουμε για κάποια παίγνια ότι είναι αυστηρώς διαγωνίως κοίλα, τότε αμέσως έχουμε την εγγύηση ότι στα παίγνια αυτά μπορεί να υπολογιστεί αποδοτικά η ισορροπία.

1.1.2 Συνεισφορά

Καταρχάς, αποδεικνύουμε ότι στα συνεξελικτικά παίγνια διαμόρφωσης άποψης η αυστηρή διαγώνια κοιλότητα (concavity) είναι ισοδύναμη με μία συνθήκη ομαλότητας στις απεικονίσεις βέλτιστης απόκρισης των παικτών. Οι βέλτιστες αποκρίσεις είναι, και πάλι, ένας μέσος όρος των απόψεων των υπολοίπων παικτών, και η συνθήκη αυτή, λέει, αδρά, ότι όταν οι παίκτες μετακινούνται, οι μέσοι όροι δεν πρέπει να μετακινούνται αυθαίρετα πολύ.

Το παραπάνω έχει ως συνέπεια και μία συνθήκη Lipschitz η οποία είναι μόνο ικανή, αλλά πιο διαχειρίσιμη. Επιπλέον, η συνθήκη αυτή εξαρτάται από τους παράγοντες αυτοπεποίθησης των παικτών. Μία άμεση συνέπεια της συνθήκης αυτής είναι ότι εάν οι παίκτες είναι επαρκώς επίμονοι (αλλά όχι μόνο για $\rho_i=1$, που είναι εύκολο να δειχθεί), τότε το παίγνιο θα είναι αυστηρά διαγώνια κοίλο.

Στη συνέχεια, εξετάζουμε ελαφρώς πιο περιορισμένες (αλλά σίγουρα μη τετριμμένες) υποκλάσεις των παιγνίων αυτών και ποσοτικοποιούμε περαιτέρω τα παραπάνω αποτελέσματα, φράσσοντας κατάλληλα τις παραγώγους των μέσων όρων (ως συναρτήσεων των απόψεων). Κατόπιν τούτου, καταλήγουμε σε έναν τρόπο απόδειξης της αυστηρής διαγώνιας κοιλότητας για οποιοδήποτε παίγνιο στο οποίο η αυτοπεποίθηση των παικτών είναι πάνω από ένα κατώφλι. Η τιμή του κατωφλίου καθορίζεται από τα φράγματα που θα καταφέρει κανείς να βρει για τις συναρτήσεις βαρών και τις παραγώγους τους.

Οπότε, τελικά, η αυστηρή διαγώνια κοιλότητα ισχύει σε όλες αυτές τις περιπτώσεις, το οποίο με τη σειρά του σημαίνει ότι ο αλγόριθμος από το [Mert19] πράγματι συγκλίνει γρήγορα στη μοναδική ισορροπία αυτών των συνεξελικτικών παιγνίων διαμόρφωσης άποψης, τα οποία, απλώς, είναι κάπως περιορισμένα, υπό την έννοια ότι η αυτοπεποίθηση των παικτών πρέπει να είναι αρκετά κοντά στο 1.

Το δεύτερο μέρος των αποτελεσμάτων μας αφορά τις παρατηρήσεις που κάναμε όταν προσπαθήσαμε να μελετήσουμε το μοντέλο μέσω προσομοιώσεων (λόγω της δυσκολίας να δείξουμε ισχυρότερα αποτελέσματα). Μία άμεση παρατήρηση ήταν ότι, σε όλα τα παιχνίδια που εξετάσαμε, ο αλγόριθμος συγκλίνει σε ισορροπία, και μάλιστα αρκετά γρήγορα. Ενδεικτικά, προσομοιώσαμε παιχνίδια με έως 10000 παίκτες, και η σύγκλιση επαρκώς κοντά σε ισορροπία είχε πάντα επιτευχθεί μετά από 200 γύρους το πολύ.

Ένα ενδιαφέρον έυρημα ήταν ότι, ως προς την ℓ_∞ απόσταση, η ταχύτητα σύγκλισης (ο αριθμός των γύρων, συγκεκριμένα) του αλγορίθμου δεν έμοιαζε να εξαρτάται από το πλήθος των παικτών. Δηλαδή, ο κάθε παίκτης έφτανε κοντά στη βέλτιστή του απόκριση πάνω κάτω σε σταθερό πλήθος γύρων, ανεξάρτητο του πλήθους των παικτών. Από την άλλη, η ταχύτητα εξαρτάται αρκετά από την επιλογή της συνάρτησης βαρών, με την τελική απόσταση από την ισρροπία να κυμαίνεται από $\approx 10^{-9}$ έως $\approx 10^{-4}$, το οποίο, συγκριτικά, δείχνει ότι για κάποιες συναρτήσεις βαρών απαιτείται περισσότερος χρόνος για τη σύγκλιση. Λίγο πιο συγκεκριμένα, από τις συναρτήσεις που ελέγξαμε, αυτές που χρειάστηκαν το περισσότερο ήταν αυτές που μεταβάλλονται πιο πολύ με την απόσταση μεταξύ των απόψεων. Αξιοσημείωτο δε, ότι για αυτές τις συναρτήσεις ήταν και πιο εύκολο να βρούμε παραδείγματα όπου η απλή δυναμική βέλτιστης απόκρισης δεν πετύχαινε σύγκλιση.

Τέλος, ελέγχοντας αριθμητικά την αυστηρή διαγώνια κοιλότητα, βρήκαμε ότι υπάρχουν αποκλίσεις, αν και μικρές. Συγκεκριμένα, όταν τα βάρη της αλληλεπίδρασης μεταξύ των παικτών πέφτουν (περίπου) ως 1/d ως προς την απόσταση d μεταξύ δύο παικτών, τότε υπάρχουν σημεία στον χώρο των απόψεων όπου η αυστηρή διαγώνια κοιλότητα δεν ισχύει. Αυτό πιθανότατα είναι συνέπεια της απότομης μεταβολής του 1/d κοντά στο 0.

Κατόπιν τούτου, δεν είναι όλα τα συνεξελικτικά παίγνια διαμόρφωσης άποψης αυστηρά διαγώνια κοίλα, και αυτό εξηγεί, σε κάποιο βαθμό, και τις δυσκολίες που αντιμετωπίσαμε καθώς προσπαθούσαμε να δείξουμε αναλυτικά αποτελέσματα. Ταυτόχρονα, οι αποκλίσεις

που βρήκαμε είναι πάρα πολύ μικρές, το οποίο εξηγεί γιατί ο αλγόριθμος δούλεψε τόσο καλά (τουλάχιστον στα παίγνια που μελετήσαμε και προσομοιώσαμε). Επειδή, πρακτικά, η συνθήκη σχεδόν ισχύει.

1.2 Θεωρητικό Υπόβαθρο

Σε αυτή την παράγραφο, θα παρουσιάσουμε στον αναγνώστη μία πολύ σύντομη εισαγωγή σε στοιχεία της θεωρίας παιγνίων και της κυρτής βελτιστοποίησης, στο βαθμό που αυτό είναι απαραίτητο ώστε να είναι όσο το δυνατόν οικείες οι έννοιες που χρησιμοποιούμε.

1.2.1 Θεωρία Παιγνίων και Κοίλα Παίγνια

Ας ξεκινήσουμε με τις έννοιες του παιγνίου, του κοίλου παιγνίου και της ισορροπίας Nash. Τα παίγνια που μας ενδιαφέρουν αποτελούνται από έναν αριθμό παικτών N, καθένας εκ των οποίων επιλέγει μία στρατηγική από κάποιο σύνολο \mathcal{X}_i , το οποίο θεωρούμε πάντα ότι είναι **κυρτό** και **συμπαγές** υποσύνολο κάποιου Ευλείδιου χώρου \mathbb{R}^d (μπορούμε να το σκεφτόμαστε ως κάποιο ελλειψοειδές). Οι απολαβές του παίκτη κωδικοποιούνται μέσω μίας συνάρτησης κέρδους $u_i: \mathcal{X} \to \mathbb{R}$, όπου $\mathcal{X} = \prod_i \mathcal{X}_i$. Δηλαδή, για κάθε προφίλ x των στρατηγικών που έχουν διαλέξει όλοι οι παίκτες, ο παίκτης i κερδίζει $u_i(x)$.

Κοίλα παίγνια αποκαλούμε αυτά στα οποία η συνάρτηση $u_i(x_i;x_{-i})$ είναι κοίλη συνάρτηση της μεταβλητής x_i , για οποιαδήποτε συγκεκριμένη τιμή των x_{-i}^{-1} . Τα κοίλα παίγνια αποτελούν, κατά βάση, μία προσπάθεια να επεκταθεί η έννοια της κοίλης συνάρτησης (η οποία, όπως θα δούμε, είναι κεντρική για την ανάπτυξη μεθόδων βελτιστοποίησης) στο παιγνιοθεωρητικό πλαίσιο.

Κεντρική, φυσικά, είναι η έννοια της **ισορροπίας Nash**, η οποία ορίζεται ως οποιοδήποτε προφίλ στρατηγικών στο οποίο κανένας παίκτης δεν έχει συμφέρον να αλλάξει μονομερώς την στρατηγική του. Συμβολικά, το προφίλ $x^* \in \mathcal{X}$ είναι ισορροπία Nash αν για κάθε παίκτη i ισχύει:

$$u_i(x_i^*, x_{-i}^*) \ge u_i(x_i, x_{-i}^*)$$
 για κάθε $x_i \in \mathcal{X}_i$

Ένα πρώτο βασικό αποτέλεσμα, λοιπόν, που οφείλεται στη γνωστή εργασία του Rosen [Rose65], είναι ότι σε κάθε κοίλο παίγνιο, υπάρχει τουλάχιστον μία ισορροπία Nash.

Ο Rosen όμως δεν περιορίζεται σε αυτό, αλλά βλέπει τα κοίλα παίγνια και από την πλευρά της κυρτής βελτιστοποίησης, κατά κάποιον τρόπο. Ορίζει την έννοια της αυστηρής διαγώνιας κοιλότητας, η οποία γενικεύει την έννοια της αυστηρώς κοίλης συνάρτησης, επιβάλλοντας στο παίγνιο ένα αντίστοιχο της γνωστής συνθήκης κοιλότητας πρώτης τάξης $\langle \nabla f(x') - \nabla f(x), x' - x \rangle < 0$. Αυτή λέει, με απλά λόγια, ότι οι κατευθυνόμενες παράγωγοι της f κατά μήκος οποιασδήποτε γραμμής είναι γνησίως φθίνουσες. Αυτό που λέει, λοιπόν, η αυστηρή διαγώνια κοιλότητα είναι να ισχύει το ίδιο για την επονομαζόμενη ψευδοκλίση (pseudogradient) του παιγνίου, που δεν είναι τίποτα άλλο παρά το διάνυσμα των "διαγωνίων" κλίσεων των συναρτήσεων κέρδους, δηλαδή:

$$v(x) = (\nabla_{x_1} u_1(x), \dots, \nabla_{x_N} u_N(x))$$

Δηλαδή, το παίγνιο (ή οι συναρτήσεις κέρδους του) λέγεται αυστηρώς διαγωνίως κοίλο όταν ισχύει:

$$\sum_i \langle \nabla_{x_i} u_i(x') - \nabla_{x_i} u_i(x), x_i' - x_i \rangle < 0 \ \text{για κάθε} \ x \neq x'$$

 $^{^{1}}$ Συμβολίζουμε με $(x_{i};x_{-i})$ το προφίλ στρατηγικών x, απλά για τις περιπτώσεις όπου θέλουμε να τονίσουμε την στρατηγική που επιλέγει ο i έναντι των στρατηγικών που επιλέγουν όλοι οι άλλοι.

ή, χρησιμοποιώντας την ψευδοκλίση:

$$\langle v(x') - v(x), x' - x \rangle < 0$$
 για κάθε $x \neq x'$

όπου με $\langle \cdot, \cdot \rangle$ συμβολίζουμε το εσωτερικό γινόμενο.

Σε αυτή την περίπτωση, στο [Rose65] αποδεικνύεται ότι το παίγνιο έχει μοναδική ισορροπία Nash. Επιπλέον, υπάρχει δυναμική που συγκλίνει στην μοναδική αυτή ισορροπία, χωρίς, ωστόσο, να έχουμε κάποια εγγύηση για την ταχύτητα σύγκλισης.

1.2.2 Κυρτή Βελτιστοποίηση

Ένας συνήθης τρόπος προσέγγισης για το πρόβλημα του υπολογισμού ισορροπιών Nash είναι μέσω τεχνικών από την κυρτή βελτιστοποίηση (και άμεση - online - κυρτή βελτιστοποίηση). Αυτό ισχύει κατά μείζονα λόγο για τα παίγνια δυναμικού, όπου ο υπολογισμός ισορροπιών είναι εν τέλει ισοδύναμος με την έυρεση βελτίστων σημείων για κάποια συνάρτηση δυναμικού. Ωστόσο, χρήσιμες συνδέσεις και συμπεράσματα μπορούν να προκύψουν και για πιο γενικές κατηγορίες παιγνίων, όπως θα δούμε αργότερα.

Καταρχάς, ας ξεκινήσουμε με την έννοια της κυρτότητας. Πολύ σύντομα, δύο βασικές έννοιες κυρτότητας υπάρχουν. Ένα σύνολο λέγεται κυρτό όταν το ευθύγραμμο τμήμα που ενώνει δύο οποιαδήποτε σημεία του ανήκει εξ ολοκλήρου στο σύνολο (χαρακτηριστικό παράδειγμα κυρτου συνόλου είναι ένα ελλειψοειδές). Αλλά και μία συνάρτηση λέγεται κυρτή όταν η προσέγγιση (ανάπτυγμα Taylor) πρώτης τάξης γύρω από οποιοδήποτε σημείο της βρίσκεται εξ ολοκλήρου κάτω από την συνάρτηση.

Σε γενικές γραμμές, η η χρησιμότητα των κυρτών συναρτήσεων έγκειται στο ότι μπορούμε να υπολογίσουμε τα ελάχιστά τους αποδοτικά, καθώς έχουν την ιδιότητα ότι κάθε τοπικό ελάχιστο είναι και ολικό (έχετε στο μυαλό σας την $\|x\|_2$ ή την τετραγωνική $\|x\|_2^2$). Οπότε, ακολουθώντας κατά βάση την αντίθετη κατεύθυνση από την κλίση (το gradient) της συνάρτησης, που όπως γνωρίζουμε είναι η κατεύθυνση μεγίστης μείωσης της συνάρτησης, καταλήγουμε τελικά στην "καταβόθρα" που είναι το ολικό της ελάχιστο. Η πιο απλή εκδοχή της επονομαζόμενης καθόδου κλίσης, σε διακριτά βήματα, είναι η εξής:

$$x_{t+1} = x_t - \gamma_t \nabla f(x_t)$$

όπου f είναι η (κυρτή) συνάρτηση που θέλουμε να ελαχιστοποιήσουμε, και $\gamma_t>0$ είναι το μήκος του βήματος κατά το οποίο μετακινείται ο αλγόριθμος κατά την κατεύθυνση της κλίσης.

Φυσικά, χρειάζονται κάποιες προσαρμογές για περιπτώσεις όπου το επιτρεπτό σύνολο των x δεν είναι όλος ο χώρος, αλλά η βασική ιδέα είναι αυτή, και, παρά την απλότητά του, ο αλγόριθμος αυτός μπορεί να αποδειχθεί ότι για μεγάλες υποκλάσεις των κυρτών συναρτήσεων, και αν το μήκος του βήματος επιλεγεί κατάλληλα, μπορεί να φτάσει αυθαίρετα κοντά σε ελάχιστο της συνάρτησης, σε χρόνο πολυωνυμικό.

Να σημειώσουμε εδώ, επειδή ίσως να υπάρχει μία μικρή σύγχυση με τις εναλλαγές μεταξύ ελαχιστοποίησης και μεγιστοποίησης, ότι αντίστοιχα με την έννοια της **κυρτής** συνάρτησης ορίζεται και η έννοια της **κοίλης** συνάρτησης, δηλαδή πρόκειται ουσιαστικά για μία συνάρτηση που η αντίθετή της είναι κυρτή. Η αντιστοιχία, λοιπόν, από εκεί και πέρα είναι ελαχιστοποίηση - κυρτότητα - συνάρτηση κόστους και μεγιστοποίηση - κοιλότητα - συνάρτηση κέρδους².

² Εν γένει, στην βιβλιογραφία περί βελτιστοποίησης προτιμώνται τα κόστη, ενώ στη βιβλιογραφία περί θεωρίας παιγνίων τα κέρδη, αλλά με εξαίρεση αυτά που είδαμε και αφορούν τη διαμόρφωση άποψης, όπου πάλι συνήθως είναι κόστη. Αυτός είναι και ο λόγος που στην παρούσα εργασία εναλλασσόμαστε συχνά μεταξύ των δύο.

Τέλος, να κάνουμε μία μικρή αναφορά και στην άμεση κυρτή βελτιστοποίηση, η οποία αφορά πλέον περιπτώσεις όπου δεν έχουμε μία μόνο συνάρτηση, αλλά υπάρχει μία διαδικασία που αποτελείται από διακριτούς γύρους, και σε κάθε γύρο μας δίνεται και μία διαφορετική συνάρτηση κόστους, με τελικό στόχο να πετύχουμε όσο το δυνατόν μικρότερο συνολικό κόστος σε όλους τους γύρους. Είναι προφανές ότι αυτό σε πλήρη γενικότητα δεν μπορεί να έχει υπερβολικά καλές εγγυήσεις (π.χ. σύγκριση με τη βέλτιστη εκ των υστέρων λύση), παρόλα αυτά είναι αξιοσημείωτο ότι, για κυρτές συναρτήσεις, μπορεί να επιτευχθεί μία αναπάντεχη εγγύηση, μέσω αλγορίθμων αντιστοίχων με την κάθοδο κλίσης. Συγκεκριμένα, έγκειται στη σύγκριση του κόστους του αλγορίθμου μας με το κόστος της βέλτιστης εκ των υστέρων στατικής απόκρισης, δηλαδή του σημείου που, αν ξέραμε εξ αρχής τι συναρτήσεις θα μας δωθούν, αλλά ήμασταν υποχρεωμένοι να επιλέξουμε ένα και μόνο σημείο, αυτό θα ήθαν το βέλτιστο. Αυτή η μετρική ονομάζεται μεταμέλεια (regret) και ορίζεται ως ακριβώς η διαφορά μεταξύ των δύο συνολικών ζημιών:

$$R(T) = \max_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x_t) - f_t(x)$$

όπου T είναι ο εκάστοτε χρονικός ορίζοντας της διαδικασίας.

Κατόπιν τούτου, ένας αλγόριθμος θεωρούμε ότι συμπεριφέρεται καλά και τον αποκαλούμε χωρίς-μεταμέλεια (no-regret), όταν η μέση μεταμέλεια καθώς το $T\to\infty$ τείνει στο 0, ή ισοδύναμα, R(T)=o(T).

Ένας αλγόριθμος που έχει καλές εγγυήσεις υπό παρόμοιες προϋποθέσεις με την απλή κάθοδο κλίσης είναι η άμεση κάθοδος κλίσης, που δεν είναι παρά:

$$x_{t+1} = x_t - \gamma_t \nabla f_t(x_t)$$

Για τον αλγόριθμο αυτό, αποδεικνύεται ότι $R(T) = O(\sqrt{T})$ υπό ορισμένες προϋποθέσεις, και κατ' επέκταση είναι χωρίς-μεταμέλεια αλγόριθμος. Ο χωρίς-μεταμέλεια αλγόριθμος που θα μας ενδιαφέρει εμάς, ωστόσο, είναι άλλος, και θα τον δούμε στην επόμενη ενότητα.

Τόσο για την κυρτή βελτιστοποίηση, όσο και για την άμεση κυρτή βελτιστοποίηση, για περισσότερες λεπτομέρεις παραπέμπουμε τον αναγνώστη σε κάποιο από τα πολλά ποιοτικότατα συγγράμματα που υπάρχουν στην βιβλιογραφία (π.χ. [Vish21, Shal12]).

1.2.3 Κυρτή Βελτιστοποίηση και Παίγνια

Κλείνουμε με μία μικρή αναφορά στην κλάση των παιγνίων δυναμικού, με σκοπό να αναδείξουμε την ισχυρή σύνδεση που υπάρχει, σε κάποιες περιπτώσεις, μεταξύ της βελτιστοποίησης και της θεωρίας παιγνίων (συγκεκριμένα, τον υπολογισμό ισορροπιών).

Definition 1.1 (Παίγνιο Δυναμικού). Ένα παίγνιο $\mathcal{G}(N, (\mathcal{X}_i)_{i \in [N]}, (u_i)_{i \in [N]})$ καλείται παίγνιο δυναμικού (potential game) εάν υπάρχει κάποια συνάρτηση $\Phi \colon \mathcal{X} \to \mathbb{R}$ τέτοια ώστε:

$$u_i(y_i, x_{-i}) - u_i(z_i, x_{-i}) = \Phi(y_i, x_{-i}) - \Phi(z_i, x_{-i})$$

Όταν υπάρχουν παράγωγοι, ο παραπάνω ορισμός είναι ισοδύναμος με την κάπως πιο άμεση σχέση:

$$\nabla \Phi(x) = v(x)$$

δηλαδή η ψευδοκλίση του παιγνίου ταυτίζεται με την κλίση της Φ.

Κατόπιν τούτου, μπορούμε να συμπεράνουμε ότι οι ισορροπίες Nash του παιγνίου αντιστοιχούν στα τοπικά μέγιστα της συνάρτησης δυναμικού (για συναρτήσεις κόστους, τοπικα

ελάχιστα). Οπότε, το πρόβλημα υπολογισμού ισορροπιών μετατρέπεται σε πρόβλημα μεγιστοποίησης μιας (βαθμωτής) συνάρτησης, για το οποίο φυσικά, αν υπάρχει οι κατάλληλες προϋποθέσεις, μπορούν να χρησιμοποιηθούν οι διαφόρων ειδών μέθοδοι κλίσης που γνωρίζουμε από την κυρτή βελτιστοποίηση.

Ένα ερώτημα, λοιπόν, μετά είναι το τι θα συμβεί αν πάμε και εφαρμόσουμε τέτοιους αλγορίθμους και σε παίγνια που δεν έχουν δυναμικό, απλώς χρησιμοποιώντας την ψευδοκλίση στη θέση της κλίσης. Μία πιθανή απάντηση σε αυτό είναι τα αποτελέσματα του [Mert19], τα οποία θα δούμε στη συνέχεια.

1.3 Κοίλα Παίγνια και ο Αλγόριθμος Dual Averaging

Στην ενότητα αυτή, θα συνοψίσουμε τα αποτελέσματα του [Mert19]. Φυσικά, τα αποτελέσματα είναι πάρα πολλά, οπότε επικεντρωνόμαστε σε αυτά τα οποία μας ενδιαφέρουν και τα χρησιμοποιούμε στην δική μας εργασία. Με λίγα λόγια, αυτό που κατά βάση μας ενδιαφέρει είναι ότι σε όλα τα αυστηρώς διαγωνίως κοίλα παίγνια (μάλιστα, ένα υπερσύνολο αυτών), υπάρχει ένας αλγόριθμος που ανήκει στην γενική ομπρέλα των mirror descent αλγορίθμων και συγκλίνει στην (μοναδική) ισορροπία Nash σε πολυωνυμικό χρόνο.

Ξεκινάμε με την περιγραφή του αλγορίθμου αυτού. Πρόκειται για έναν σχετικά απλό αλγόριθμο, ευρέως γνωστό στη βιβλιογραφία περί άμεσης βελτιστοποίησης (πρώτη εμφάνιση στο [Nest09]), ο οποίος προσαρμόστηκε από το [Mert19] στο παιγνιοθεωρητικό πλαίσιο. Υπό ορισμένες χαλαρές προϋποθέσεις, είναι χωρίς-μεταμέλεια (π.χ. [Xiao09, Shal12]) και, ενώ δεν ασχοληθήκαμε ιδιαίτερα με αυτό στην εργασία μας, είναι μία σημαντικη ιδιότητα, η οποία εγγυάται καλές συνολικές απολαβές για τον εκάστοτε παίκτη και, κατά συνέπεια, είναι μία τακτική που είναι πολύ πιθανόν να την εφαρμόσουν πραγματικοί έξυπνοι παίκτες.

Ο αλγόριθμος λειτουργεί ως εξής. Κάθε παίκτης κρατάει δύο σημεία ανά πάσα στιγμή: την τρέχουσα στρατηγική του $x_{i,t} \in \mathcal{X}_i$, που είναι η στρατηγική την οποία θα παίξει σε κάθε γύρο t του παιγνίου, και ένα άλλο σημείο $y_{i,t}$ το οποίο μετακινείται ελεύθερα στον περιβάλλοντα χώρο. Σε κάθε γύρο, υπολογίζει το $y_{i,t+1}$ κάνοντας ένα ελεύθερο βήμα στην κατεύθυνση της κλίσης $\nabla_{x_i}u_i(x_t)$ (χωρίς να λαμβάνει υπόψη αν καταλήγει μέσα στο σύνολο των στρατηγικών του ή όχι), και το $x_{i,t}$ προβάλλοντας (υπό μία γενική έννοια) το $y_{i,t}$ στο σύνολο των στρατηγικών.

Με άλλα λόγια, θα μπορούσαμε να πούμε ότι ο αλγόριθμος μοιάζει με την προβεβλημένη κάθοδο κλίσης (projected gradient descent), με τη διαφορά ότι το $y_{i,t}$ δεν το διαγράφει τελείως μετά από κάθε βήμα κρατώντας μόνο την προβολή, αλλά το κρατάει, και κάνει το επόμενο βήμα με ορμητήριο αυτό το σημείο.

Τυπικά, ξεκινώντας με κάποιο αυθαίρετο y_1 , το γενικό αλγοριθμικό σχήμα περιγράφεται από τις εξισώσεις:

$$x_{i,t} = Q_i(y_{i,t})$$

$$y_{i,t+1} = y_{i,t} + \gamma_t \nabla_{x_i} u_i(x_t)$$
(DA)

όπου:

• Τα Q_i είναι γενικευμένες προβολές, που ονομάζονται *choice maps*. Για εμάς, θα είναι οι απλές ευκλείδιες προβολές³, δηλαδή:

$$Q_i(y_i) = \operatorname*{argmin}_{x_i \in \mathcal{X}_i} \|x_i - y_i\|_2$$

³ Οι choice maps που περιγράφονται στο [Mert19] είναι πολύ πιο γενικές, και βασίζονται σε έννοιες της κυρτής ανάλυσης όπως η κυρτή συζυγής μίας συνάρτησης. Οι Ευκλείδιες προβολές είναι ειδική τους περίπτωση, αλλά για τους σκοπούς μας αρκεί.

- Τα γ_t είναι μία φθίνουσα ακολουθία μηκών των βημάτων, τυπικά της μορφής $1/t^\beta$ για κάποιο $\beta \in (0,1].$
- Το $\nabla_{x_i} u_i(x_t)$ είναι η επονομαζόμενη "ιδιοκλίση" του παίκτη i στο προφίλ στρατηγικών x_t (μερικές φορές ίσως το συμβολίζουμε και ως $v_i(x_t)$). Επιπλέον, συμβολίζουμε $V_* = \max_{x \in \mathcal{X}} \|v(x)\|_*$. Παρατηρήστε ότι αυτό το μέγιστο υπάρχει πάντα, αφού έχουμε υποθέσει ότι το \mathcal{X} είναι συμπαγές και η ψευδοκλίση v(x) είναι συνεχής συνάρτηση.

1.3.1 Σύγκλιση σε Ισορροπία

Το βασικό αποτέλεσμα εδώ είναι το ακόλουθο:

Theorem 1.2. Σε ένα αυστηρώς διαγωνίως κοίλο παίγνιο, αν όλοι οι παίκτες εφαρμόζουν τον αλγόριθμο (DA), τότε το προφίλ στρατηγικών x_t θα συγκλίνει στη (μοναδική) ισορροπία Nash του παιγνίου (υπό χαλαρές προϋποθέσεις για τα μήκη βήματος γ_t).

Το αποτέλεσμα αυτό είναι, φυσικά, πολύ ενδιαφέρον και αποτελεί ένα πρώτο βήμα, ωστόσο αυτό που πραγματικά μας ενδιαφέρει είναι σε πεπερασμένο χρόνο τι γίνεται. Μπορούμε να πετύχουμε εγγυήσεις αντίστοιχες με αυτές που υπάρχουν για τις διάφορες μεθόδους καθόδου κλίσης στην απλή κυρτή βελτιστοποίηση (π.χ. ότι φτάνουμε ε -κοντά στο βέλτιστο / στην ισορροπία σε χρόνο $O(1/\varepsilon)$);

1.3.2 Ταχύτητα Σύγκλισης

Υπάρχουν δύο μετρικές που χρησιμοποιούνται στο [Mert19] για να δειχθεί η αποδοτικότητα του αλγορίθμου Dual Averaging. Η μία είναι η μέση απόσταση από την ισορροπία και η άλλη το running length.

Η μέση απόσταση από την ισορροπία στον γύρο t ορίζεται ως:

$$\bar{r}_t = \frac{\sum_{\tau=1}^t \gamma_\tau ||x_\tau - x^*||_2}{\sum_{\tau=1}^t \gamma_\tau}$$

Αποτελεί, λοιπόν, τον μέσο όρο των αποστάσεων κάθε σημείου από το οποίο περνάει ο αλγόριθμος από την ισορροπία, στην οποία τελικά συγκλίνει.

Από την άλλη, το running length στο γύρο t ορίζεται ως:

$$\ell_t = \sum_{\tau=1}^{t-1} \|x_{\tau+1} - x_{\tau}\|_2$$

Κατά συνέπεια, πρόκειται κυριολεκτικά για το συνολικό "μήκος" που έχει διαγράψει η τροχιά του αλγορίθμου στον χώρο $\mathcal X$ έως και τον γύρο t.

Τέλος, θα χρειαστούν δύο ακόμα έννοιες για να μπορέσουμε να διατυπώσουμε τις εγγυήσεις του [Mert19]. Εν συντομία:

• $\Omega = \max_{x \in \mathcal{X}} \frac{1}{2} \|x\|_2^2 - \min_{x \in \mathcal{X}} \frac{1}{2} \|x\|_2^2$. Η σταθερά αυτή προκύπτει από τον τρόπο με τον οποίο ορίζονται τα γενικά choice maps (και με τον οποίο δεν θα ασχοληθούμε εδώ). Είναι ένας από τους δύο παράγοντες που περιέχουν την επίδραση της διάστασης στα παρακάτω φράγματα (ο άλλος είναι το L που θα δούμε αμέσως μετά).

Απαιτείται μία ελαφρώς πιο ισχυρή μορφή της αυστηρής διαγώνιας κοιλότητας, την οποία ονομάζουμε ισχυρή διαγώνια κοιλότητα (strong diagonal concavity)⁴, και δεν είναι παρά η προσαρμογή της αντίστοιχης συνθήκης για τις ισχυρά κοίλες βαθμωτές συναρτήσεις. Έτσι, το παίγνιο λέμε ότι είναι L ισχυρώς διαγωνίως κοίλο αν ισχύει:

$$\langle v(x') - v(x), x' - x \rangle \leq L \|x' - x\|_2^2$$
 για κάθε $x \neq x'$

Δεδομένων, λοιπόν, αυτών των ορισμών, τα φράγματα που μας εγγυάται το [Mert19] έχουν ως εξής:

Theorem 1.3. Τρέχοντας τον αλγόριθμο (DA) σε ένα ισχυρά διαγωνίως κοίλο παίγνιο, με σταθερό μήκος βήματος $\gamma = \frac{1}{V_*} \sqrt{\frac{2\Omega}{T}}$, για T βήματα, θα ισχύει ότι:

$$\bar{r}_T^2 \le \frac{V_*}{L} \sqrt{\frac{2\Omega}{T}}$$

Theorem 1.4. Τρέχοντας τον αλγόριθμο (DA) σε ένα ισχυρά διαγωνίως κοίλο παίγνιο, με την προϋπόθεση $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$, θα ισχύει ότι:

$$\ell_{t_{\varepsilon}} \le \frac{V_*}{L} \frac{\Omega + \frac{V_*^2}{2} \sum_{t=1}^{\infty} \gamma_t^2}{\varepsilon^2}$$

όπου $t_{\varepsilon} = \inf\{t \geq 0 : \|x_t - x^*\|_2 \leq \varepsilon\}$ είναι ο πρώτος γύρος στον οποίο το προφίλ στρατηγικών βρίσκεται ε -κοντά στην ισορροπία.

Τα δύο αυτά αποτελέσματα μπορούμε να τα θεωρήσουμε και δυϊκά, κατά μία έννοια, καθώς το πρώτο φράσσει απόσταση από την ισορροπία συναρτήσει του "πόσο έτρεξε" ο αλγόριθμος, και το δεύτερο φράσσει κάτι σαν το "πόσο έτρεξε" ο αλγόριθμος συναρτήσει της απόστασης-στόχου από την ισορροπία. Επιπροσθέτως, η τάξη μεγέθους του πρώτου φράγματος είναι $O(1/\sqrt{T})$, ενώ του δεύτερου $O(1/\varepsilon^2)$. Κατ' αυτήν την έννοια εννοούμε την δυϊκότητα, καθώς όπως γνωρίζουμε από αντίστοιχες περιπτώσεις βελτιστοποίησης, το $1/\sqrt{T}$ και το $1/\varepsilon^2$ είναι δύο διαφορετικοί ισοδύναμοι τρόποι να εκφραστεί η ταχύτητα σύγκλισης ενός τέτοιο αλγορίθμου.

Ας σημειώσουμε, επίσης, ότι το πρώτο αποτέλεσμα έχει το μειονέκτημα ότι πρέπει να έχει αποφασιστεί εξ αρχής το πλήθος των γύρων για το οποίο θα τρέξει ο αλγόριθμος, και το μήκος βήματος εξαρτάται από αυτό. Το μειονέκτημα, βέβαια, του δευτέρου είναι ο όρος $\sum_{t=1}^{\infty} \gamma_t^2, \text{ ο οποίος μπορεί δυνητικά να είναι μεγάλος (αλλά και η εξάρτηση από τη διάσταση, το οποίο θα γίνει λίγο πιο εμφανές όταν θα εφαρμόσουμε τα αποτελέσματα αυτά στα δικά μας παίγνια).$

1.3.3 Εφαρμογή στα Συνεξελικτικά Παίγνια

Κλείνουμε την ενότητα αυτή εφαρμόζοντας τον γενικό αλγόριθμο των Μερτικόπουλου και Zhou στα συνεξελικτικά παίγνια διαμόρφωσης άποψης, και γράφοντας ρητά τις αντίστοιχες εξισώσεις. Ο σκοπός είναι να δούμε πώς μοιάζει ο αλγόριθμος αυτός στα δικά μας

⁴ Στην πραγματικότητα, στο [Mert19] χρησιμοποιείται μία πιο γενική συνθήκη αντί της ισχυρής, αλλά και της αυστηρής διαγώνιας κοιλότητας, που λέγεται variational stability και αντιστοιχεί στο unimodality των βαθμωτών συναρτήσεων, που είναι πιο γενικό από την αυστηρή τους κοιλότητα. Γενικά, είναι αρκετά ενδιαφέρουσα συνθήκη, αλλά εδώ δεν χρειάζεται να μπούμε σε τέτοιες λεπτομέρειες.

παίγνια, τι απλοποιήσεις μπορούν δυνητικά να γίνουν κλπ. Κατόπιν, στην επόμενη και τελευταία ενότητα, εκθέτουμε τα αποτελέσματά μας, που αφορούν το κατά πόσο μπορούμε να πούμε ότι ο αλγόριθμος υπολογίζει αποδοτικά την ισορροπία σε αυτά τα παίγνια.

Ξεκινάμε υπενθυμίζοντας τις συναρτήσεις κόστους των παικτών σε ένα συνεξελικτικό παίγνιο διαμόρφωσης άποψης, τις οποίες, επιπλέον, μετατρέπουμε σε συναρτήσεις κέρδους αντιστρέφοντας το πρόσημό τους:

$$u_i(x) = -(1 - \rho_i) \sum_{j \neq i} q_{ij}(x_{-i})(x_i - x_j)^2 - \rho_i(x_i - s_i)^2$$

Η κλίση, τώρα, της συνάρτησης αυτής δεν θα είναι παρά η παράγωγός της:

$$\frac{\partial u_i(x)}{\partial x_i} = -2(x_i - (1 - \rho_i) \sum_{j \neq i} q_{ij}(x_{-i}) x_j - \rho_i s_i) = -2(x_i - BR_i(x_{-i}))$$

όπου ο συντελεστής του x_i είναι μονάδα λόγω του ότι τα βάρη του κάθε παίκτη i έχουμε υποθέσει εξ αρχής ότι είναι κανονικοποιημένα, δηλαδή αθροίζουν στο 1. Επισης, συμβολίζουμε με $\mathrm{BR}_i(x_{-i})$ την ποσότητα που αποτελεί την βέλτιστη απόκριση του παίκτη i στις στρατηγικές x_{-i} , δηλαδή τον βεβαρυμένο μέσο όρο που λέγαμε στην εισαγωγή. Τυπικά:

$$BR_{i}(x_{-i}) = (1 - \rho_{i}) \sum_{j \neq i} q_{ij}(x_{-i})x_{j} + \rho_{i}s_{i}$$

Επιπλέον, όπως είπαμε, η choice map που χρησιμοποιούμε είναι η απλή Ευκλείδια προβολή, δηλαδή $Q_i(y_i) = \operatorname{argmin}_{x_i \in \mathcal{X}_i} \|y_i - x_i\|_2$. Από τη στιγμή που $\mathcal{X}_i = [0,1]$, η προβολή θα είναι απλώς:

$$Q_i(y_i) = egin{cases} 0, & ext{an } y_i < 0 \ 1, & ext{an } y_i > 1 \ y_i, & ext{διαφορετικά} \end{cases}$$

Αντικαθιστώντας, λοιπόν, τα παραπάνω στην (DA), θα πάρουμε:

$$\begin{split} x_{i,t} &= \min\{\max\{y_{i,t}, 0\}, 1\} \\ y_{i,t+1} &= y_{i,t} - 2\gamma_t(x_{i,t} - \mathrm{BR}_i(x_{-i})) \\ &= y_{i,t} + 2\gamma_t(\mathrm{BR}_i(x_{-i}) - x_{i,t}) \end{split}$$

Τέλος, για να απλοποιήσουμε περαιτέρω τον αλγόριθμο, κάνουμε δύο (αθώες, από όσο μπορούμε να καταλάβουμε) υποθέσεις:

- 1. $y_{i,1} \in [0,1]$.
- 2. $\gamma_t < 1/2$. Αφού τα γ_t είναι φθίνοντα, αρκεί $\gamma_1 < 1/2$.

Με αυτές τις υποθέσεις, και από τη στιγμή που $\mathrm{BR}_i(x_{-i}) \in [0,1]$, ως κυρτός συνδυασμός των x_{-i}, s_i , δεν είναι δύσκολο να διαπιστώσει κανείς ότι κάθε $y_{i,t}$ θα ανήκει σίγουρα στο [0,1]. Επομένως, θα ισχύει πάντα $x_{i,t}=y_{i,t}$ και ο αλγόριθμος ανάγεται στον πιο απλό κανόνα:

$$x_{i,t+1} = x_{i,t} + \gamma_t (BR_i(x_{-i,t}) - x_{i,t})$$
(1.1)

όπου αντικαταστήσαμε $\gamma_t \leftarrow 2\gamma_t$ για απλούστευση.

Οπότε, εν κατακλείδι, ο αλγόριθμος (DA) δεν κάνει κάτι το φοβερό στην περίπτωσή μας. Απλά, κάθε παίκτης, αντί να επιλέγει κάθε φορά τη βέλτιστή του απόκριση $BR_i(x_{-i,t})$,

επιλέγει ένα ενδιάμεσο σημείο μεταξύ της βέλτιστης απόκρισης και της στρατηγικής που επέλεξε στον προηγούμενο γύρο.

Το πιο σημαντικό που μας προσφέρει η εργασία των Μερτικόπουλου και Zhou, από εκεί και πέρα, είναι τα αναλυτικά εργαλεία (η αυστηρή διαγώνια κοιλότητα ως ικανή συνθήκη) και τις αντίστοιχες εγγυήσεις για την απόδοση αυτού του αλγορίθμου. Οπότε, εμείς θα αρκούσε, ιδανικά, να ασχοληθούμε μόνο με το να δείξουμε για όποια παίγνια μας ενδιαφέρουν ότι ισχύει η αυστηρή διαγώνια κοιλότητα.

1.4 Ο Αλγόριθμος Dual Averaging στα Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης

Σε αυτή την ενότητα, θα δούμε εν συντομία τα αποτελέσματα της εργασίας μας.

1.4.1 Θεωρητικά Αποτελέσματα στην Αυστηρή Διαγώνια Κοιλότητα

Μία κεντρική προσέγγιση στην δουλειά μας ήταν η προσπάθεια να βρούμε ενδιαφέρουσες υποκατηγορίες των συνεξελικτικών παιγνίων διαμόρφωσης άποψης, για τις οποίες να ισχύει η ισχυρή διαγώνια κοιλότητα. Τότε, όπως είδαμε στην προηγούμενη ενότητα, η δυναμική Dual Averaging θα συγκλίνει στη μοναδική ισορροπία σχετικά γρήγορα.

Μία τέτοια προσπάθεια η οποία έφερε κάποια αποτελέσματα ήταν να ξεκινήσουμε από τον ορισμό της ισχυρής διαγώνιας κοιλότητας, να αντικαταστήσουμε τις συγκεκριμένες πληροφορίες (τις παραγώγους) που προκύπτουν από τα κόστη των συνεξελικτικών παιγνίων και να προχωρήσουμε με απλές πράξεις έως κάποιο τελικό, πιθανώς ενδιαφέρον αποτέλεσμα. Συγκεκριμένα, κατορθώσαμε να φτάσουμε σε μία δυνητικά ενδιαφέρουσα ικανή συνθήκη για την αυστηρή διαγώνια κοιλότητα ενός συνεξελικτικού παιγνίου διαμόρφωσης άποψης.

Ας ξεκινήσουμε με τον ορισμό μία έννοιας που μας φάνηκε αρκετά χρήσιμη στα μετέπειτα αποτελέσματα. Πρόκειται για το τμήμα της βέλτιστης απόκρισης του παίκτη i που εμπλέκει μόνο τις απόψεις των υπολοίπων παικτών. Αυτό το καλούμε centroid, και το ορίζουμε ως εξής:

Definition 1.5 (Centroids). Για κάθε παίκτη i και κάθε προφίλ απόψεων $x \in \mathcal{X}$, ορίζουμε το centroid των γειτόνων του i ως:

$$\operatorname{ctr}_i(x) = \sum_{j \neq i} q_{ij}(x_{-i}) x_j$$

Επιπλέον, συμβολίζουμε με ctr: $\mathcal{X} \to \mathcal{X}$ την απεικόνιση που συγκεντρώνει τα centroids όλων των παικτών σε ένα διάνυσμα, δηλαδή:

$$\mathsf{ctr}(x) = (\mathsf{ctr}_1(x), \dots, \mathsf{ctr}_N(x))$$

Η έννοια του centroid πρωτοεμφανίστηκε στο [Fane21], όπου χρησιμοποιείται για την εύρεση προσεγγιστικών ισορροπιών στα συνεξελικτικά παίγνια διαμόρφωσης άποψης. Εμείς εδώ το προσαρμόζουμε για τους σκοπούς μας.

Παρατηρείστε ότι το centroid ενός παίκτη είναι απλά ένας κυρτός συνδυασμός (με άλλα λόγια, ένας βεβαρυμένος μέσος όρος) των απόψεων που εκφράζουν όλοι οι άλλοι παίκτες, βεβαρυμένων με τα ίδια βάρη που αναπαριστούν την ισχύ της αλληλεπίδρασης μεταξύ του

i και των γειτόνων του. Επιπροσθέτως, ας θυμηθούμε ότι η βέλτιστη απόκριση του παίκτη i ισούται με:

$$BR_i(x_{-i}) = (1 - \rho_i) \sum_{j \neq i} q_{ij}(x_{-i}) x_j + \rho_i s_i = (1 - \rho_i) \operatorname{ctr}_i(x) + \rho_i s_i$$

Οπότε, το centroid αναπαριστά ένα μέσο σημείο των απόψεων όλων των παικτών, εκτός του i, τέτοια ώστε ο παίκτης i ουσιαστικά ενδιαφέρεται μόνο για αυτό το σημείο, και όχι για το πλήρες προφίλ των απόψεων των υπολοίπων παικτών. Τουλάχιστον, ως προς την επιλογή της βέλτιστης απόκρισης δεδομένων των απόψεων των υπολοίπων.

Έχοντας τα παραπάνω στο μυαλό μας, μπορούμε τώρα να αποδείξουμε το ακόλουθο λήμμα, το οποίο προσαρμόζει τον ορισμό της αυστηρής διαγώνιας κοιλότητας στα συνεξελικτικά παίγνια διαμόρφωσης άποψης.

Lemma 1.6. Ένα συνεξελικτικό παίγνιο διαμόρφωσης άποψης είναι αυστηρώς διαγωνίως κυρτό αν και μόνο αν:

$$\sum_{i=1}^{N} (1 - \rho_i)(\mathsf{ctr}_i(x^1) - \mathsf{ctr}_i(x^2))(x_i^1 - x_i^2) < \|x^1 - x^2\|_2^2$$

για κάθε $x^1, x^2 \in \mathcal{X} = [0, 1]^N$.

Αν, επιπλέον, όλοι οι παίκτες έχουν τον ίδιο παράγοντα αυτοπεποίθησης $\rho_i = \rho \in [0, 1)$, τότε το παραπάνω είναι επίσης ισοδύναμο με:

$$\langle \operatorname{ctr}(x^1) - \operatorname{ctr}(x^2), x^1 - x^2 \rangle < \frac{1}{1 - \rho} ||x^1 - x^2||_2^2$$

Θυμηθείτε ότι, για μία βαθμωτή συνάρτηση $f: \mathcal{X} \to \mathbb{R}$, η συνθήκη:

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \leq K \|x - y\|^2$$
 για κάθε $x, y \in \mathcal{X}$

είναι ισοδύναμη με το ότι η f είναι μία K-λεία συνάρτηση (άλλα ονόματα: ισχυρά λεία, Lipschitz λεία). Οπότε, το λήμμα αυτό (ιδιαίτερα το δεύτερο μέρος) λέει ότι η ψευδοκλίση v(x) του παιγνίου μοιάζει με την κλίση μίας αυστηρά κοίλης (ή κυρτής, για κόστη) συνάρτησης (δηλαδή, ισχύει η αυστηρή διαγώνια κοιλότητα) αν και μόνο αν το διανυσματικό πεδίο $\operatorname{ctr}: \mathcal{X} \to \mathcal{X}$ μοιάζει με την κλίση μίας G-λείας συνάρτησης (για επαρκώς μικρή τιμή του G).

Υπό αυτή την ερμηνεία, εάν το διανυσματικό πεδίο ctr είναι Lipschitz, τότε θα περιμέναμε η συνθήκη να ισχύει, μια και για βαθμωτές συναρτήσεις το να είναι Lipschitz η κλίση συνεπάγεται το να είναι ισχυρά λεία η συνάρτηση. Και πράγματι, αυτό ισχύει και εδώ, όπως βλέπουμε στο ακόλουθο θεώρημα.

Theorem 1.7. Εάν η απεικόνιση ctr: $[0,1]^N \to [0,1]^N$ είναι G-Lipschitz ως προς την Eν-κλείδια νόρμα, τότε το συνεξελικτικό παίγνιο διαμόρφωσης άποψης με ίδιους παράγοντες αυτοπεποίθησης ρ για όλους τους παίκτες είναι αυστηρώς διαγωνίως κοίλο για όλες τις τιμές του ρ στο $(1-\frac{1}{G},1]$. Εάν $G\leq 1$, τότε το παίγνιο είναι πάντα αυστηρώς διαγωνίως κοίλο.

Μία άμεση συνέπεια του παραπάνω είναι ότι πάντα υπάρχει κάποια επαρκώς υψηλή τιμή αυτοπεποίθησης (αλλά μικρότερη από 1) για την οποία το παίγνιο γίνεται αυστηρώς διαγωνίως κοίλο. Μοναδική προϋπόθεση είναι τα βάρη $q_{ij}(x_{-i})$ να είναι συνεχώς διαφορίσιμες

συναρτήσεις. Πράγματι, σε αυτή την περίπτωση, τα ${\rm ctr}_i$ θα είναι και αυτά συνεχώς διαφορίσιμες συναρτήσεις. Γνωρίζουμε, όμως, ότι μία συνεχώς διαφορίσιμη συνάρτηση με συμπαγές πεδίο ορισμού (όπως εδώ το $[0,1]^N$) είναι πάντα Lipschitz, για κάποια επαρκώς μεγάλη σταθερά. Οπότε, για αρκετά μεγάλο G η υπόθεση του παραπάνω θεωρήματος ικανοποιείται!

Αυτό δεν είναι τελείως απροσδόκητο, από τη στιγμή που για $\rho=1$ το παίγνιο είναι αυστηρώς διαγωνίως κοίλο κατά τετριμμένο τρόπο, οπότε δεν φαντάζει παράλογο αυτή η συμπεριφορά να έχει μία, έστω τοπική, "συνέχεια", για τιμές του ρ κοντά στο 1 δηλαδή.

Για να κλείσουμε αυτή τη σειρά των αποτελεσμάτων, δίνουμε το παρακάτω θεώρημα, που είναι μία απλή εφαρμογή των αποτελεσμάτων του [Mert19] για τα δικά μας παίγνια, στην περίπτωση που ισχύουν οι προϋποθέσεις των συνθηκών που δώσαμε πιο πάνω 5 .

Theorem 1.8. Εστω ένα συνξελικτικό παίγνιο διαμόρφωσης άποψης, με κοινούς παράγοντες αυτοπεποίθησης $\rho_i = \rho$, και υποθέτοντας, επιπλέον, ότι η απεικόνιση ctr: $\mathcal{X} \to \mathcal{X}$ είναι *G-Lipschitz*.

Ας υποθέσουμε, τώρα, ότι οι παίκτες τρέχουν τον αλγόριθμο (DA) όπως τον δείζαμε στο τέλος της προηγούμενης ενότητας, δηλαδή:

$$x_{i,t+1} = x_{i,t} + \gamma_t((1-\rho)\operatorname{ctr}_i(x) + \rho s_i - x_{i,t})$$

Τότε, για όλα τα $\rho \in (1-\frac{1}{G},1]$, τα x_t συγκλίνουν στην μοναδική ισορροπία x^* του παιγνίου. Επιπλέον, ισχύουν οι ακόλουθες εγγυήσεις που αφορούν την ταχύτητα σύγκλισης:

1. Αν ο αλγόριθμος τρέξει για T γύρους με σταθερό βήμα $\gamma = \frac{1}{2\sqrt{T}}$, τότε για τη μέση απόσταση από την ισορροπία, $\bar{r}_t = \sum_{\tau=1}^t \gamma_\tau \|x_\tau - x^*\|_2 / \sum_{\tau=1}^t \gamma_\tau \ \theta$ α ισχύει:

$$\bar{r}_T^2 \leq \frac{N}{1 - G(1 - \rho)} \frac{1}{\sqrt{T}}$$

2. Αν $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$, το running length μέχρι το x_t να φτάσει ε-κοντά στην ισορροπία (όπως αυτό ορίστηκε στο (1.4)) φράσσεται ως:

$$\ell_{t_{\varepsilon}} \le \frac{N\sqrt{N}}{1 - G(1 - \rho)} \frac{\frac{1}{2} + 2\sum_{t=1}^{\infty} \gamma_t^2}{\varepsilon^2}$$

Ένα ερώτημα που παραμένει είναι το πώς θα μπορούσαν να εφαρμοστούν τα αποτελέσματα που έχουμε αναφέρει μέχρι στιγμής, είτε γενικά (το οποίο μοιάζει αρκετά υπεραισιόδοξο), είτε για πιο περιορισμένες υποκατηγορίες συνεξελικτικών παιγνίων διαμόρφωσης άποψης. Μια και η τελευταία συνθήκη που δώσαμε, λόγου χάρη, είναι συνθήκη Lipschitz, εύκολα αναρωτιέται κανείς αν θα μπορούσε να εφαρμοστεί βρίσκοντας φράγματα στις παραγώγους. Μία τέτοια ανάλυση δίνουμε στη συνέχεια, και κλείνουμε με εφαρμογή σε ένα συγκεκριμένο παράδειγμα συναρτήσεων βαρών.

Εξειδίκευση για Κανονικοποιημένα Ανεξάρτητα Βάρη

Εισάγουμε μία ειδική κατηγορία συναρτήσεων βαρών (με την οποία θα ασχοληθούμε αρκετά), τα οποία ονομάζουμε "κανονικοποιημένα ανεξάρτητα βάρη" και δεν είναι τίποτα

⁵ και ορισμένων προσαρμογών τους για την περίπτωση της επονομαζόμενης ισχυρής διαγώνιας κοιλότητας (strong diagonal concavity), με τις λεπτομέρειες των οποίων δεν χρειάζεται να ασχοληθούμε στην παρούσα περίληψη.

άλλο παρά βάρη που εξαρτώνται μόνο από την απόσταση των παικτών i και j, τα οποία στη συνέχεια κανονικοποιούνται στη μονάδα, όπως απαιτεί το μοντέλο μας:

$$q_{ij}(x) = F_i(d_j^i, d_{-i-j}^i) = \frac{f_i(d_j^i)}{\sum_{k \neq i} f_i(d_k^i)}$$
(1.2)

όπου $d_j^i = |x_j - s_i|$ και οι συναρτήσεις $f_i \colon [0,1] \to \mathbb{R}$ υποθέτουμε ότι είναι θετικές, συνεχείς (για να είναι κοίλο το παίγνιο) και φθίνουσες (για να πέφτουν τα βάρη με την απόσταση).

Με αυτά τα δεδομένα, καταφέραμε να αποδείξουμε το ακόλουθο αποτέλεσμα.

Theorem 1.9. Σε ένα συνεζελικτικό παίγνιο διαμόρφωσης άποψης με κανονικοποιημένα ανεζάρτητα βάρη, εάν $\min_{d \in [0,1]} f_i(d) = m > 0$ και $\max_{d \in [0,1]} f_i(d) + |f_i'(d)| = M$, τότε το etr είναι $G = \frac{M}{m} \sqrt{\frac{N}{N-1}} \approx \frac{M}{m}$ Lipschitz και, συνεπώς, το παίγνιο είναι ισχυρώς διαγωνίως κοίλο για όλες τις τιμές του $\rho > 1 - 1/G$.

Ως ένα τελικό πόρισμα, προσαρμόζουμε το θεώρημα (1.8) σε αυτή την περίπτωση:

Corollary 1.10. Σε ένα συνεζελικτικό παίγνιο διαμόρφωσης άποψης με κανονικοποιημένα ανεζάρτητα βάρη, εάν $\min_{d \in [0,1]} f_i(d) = m > 0$ και $\max_{d \in [0,1]} f_i(d) + |f_i'(d)| = M$, τότε για κάθε $\rho > 1 - \frac{m}{M} \sqrt{\frac{N-1}{N}}$, ο αλγόριθμος (1.1) συκλίνει στη μοναδική ισορροπία του παιγνίου, και έχουμε τα ίδια φράγματα που μας εγγυάται το θεώρημα (1.8), τα οποία μετά την αντικατάσταση $G = \frac{M}{m} \sqrt{\frac{N}{N-1}} \approx \frac{M}{m}$ γίνονται:

$$\begin{split} \bar{r}_T^2 &\leq \frac{N}{1 - (1 - \rho) \frac{M}{m}} \frac{1}{\sqrt{T}} \\ \ell_{t_{\varepsilon}} &\leq \frac{N\sqrt{N}}{1 - (1 - \rho) \frac{M}{m}} \frac{\frac{1}{2} + 2\sum_{t=1}^{\infty} \gamma_t^2}{\varepsilon^2} \end{split}$$

Να σημειωθεί ότι παρόλο που τυπικά ο όρος $\sqrt{\frac{N}{N-1}}$ χρειάζεται, τον έχουμε παραλείψει για να είναι λίγο πιο καθαρά τα αποτελέσματα και επειδή καθώς το N μεγαλώνει η συνεισφορά του γίνεται αμελητέα.

Παράδειγμα: εκθετικά βάρη Κλείνουμε με ένα παράδειγμα μίας συνάρτησης βαρών που είναι αρκετά ομαλή θα μπορούσε να πει κανείς, όπου $f_i(d)=e^{-d^2}$ (το τετράγωνο είναι για να ομαλοποιηθεί περαιτέρω η συνάρτηση). Σε αυτή την περίπτωση:

$$\min_{d \in [0,1]} f_i(d) = \min_{d \in [0,1]} e^{-d^2} = e^{-1} = m$$

$$\max_{d \in [0,1]} f_i(d) + |f_i'(d)| = \max_{d \in [0,1]} e^{-d^2} + 2de^{-d^2} = 1 + 2 = 3 = M$$

Οπότε, εφαρμόζοντας το παραπάνω θεώρημα, το ctr είναι G-Lipschitz, όπου:

$$G = \frac{M}{m} \sqrt{\frac{N}{N-1}} \approx \frac{M}{m} = 3e$$

το οποίο είναι < 10 για κάθε N, και < 9 για $N \ge 6$.

Σε κάθε περίπτωση, αυτό σημαίνει ότι για όλες τις τιμές του ρ μεγαλύτερες από 1-1/G, το οποίο είναι κάτι λίγο μικρότερο από 0.9, το παίγνιο είναι αυστηρώς διαγωνίως κοίλο και ο αλγόριθμος Dual Averaging συγκλίνει στην ισορροπία σχετικά γρήγορα.

Παρατηρήσεις

Καταρχάς, με τα αποτελέσματα της ενότητας αυτής θέλουμε να καταδείξουμε, όσο αυτό είναι δυνατόν, ότι το να δείξει κανείς την αυστηρή διαγώνια κοιλότητα για αυτά τα περίπλοκα, εξαιρετικά μη γραμμικά παίγνια δεν είναι τελείως απροσπέλαστο, ακόμα και αν αυτό αφορά κάπως περιορισμένες περιπτώσεις.

Δεύτερον, ουσιαστικά με τα παραπάνω έχουμε καταλήξει σε μία μέθοδο για την απόδειξη παρόμοιων αποτελεσμάτων για άλλες συναρτήσεις: βρίσκουμε κάτω φράγμα για την f_i και άνω φράγμα για την $f_i+|f_i'|$, και αυτομάτως λαμβάνουμε ένα εύρος τιμών της αυτοπεποίθησης ρ για τις οποίες το παίγνιο είναι αυστηρώς διαγωνίως κοίλο. Φυσικά, αυτή η διαδικασία δεν είναι εγγυημένο ότι θα δώσει αξιοσημείωτη τιμή για το ρ , αλλά είναι πιθανόν, όπως είδαμε στο παραπάνω παράδειγμα.

Τέλος, τα τελευταία μας αποτελέσματα θα μπορούσε να πει κανείς ότι αυτό που απαιτούν είναι, ουσιαστικά, άνω και κάτω φράγμα στην f_i , τα οποία να είναι όσο πιο κοντά και όσο πιο ψηλά γίνεται, και ένα άνω φράγμα στην $|f_i'|$, δηλαδή στον ρυθμό μεταβολής της f_i . Κατά συνέπεια, πιστεύουμε ότι τα αποτελέσματα αυτά μπορούν να ερμηνευθούν ώς μία ποσοτικοποίηση του ότι αν το δυναμικό / εξελισσόμενο κοινωνικό δίκτυο που έχουμε δεν μεταβάλλεται πολύ, δηλαδή τα βάρη είναι κοντά στο να είναι σταθερά, τότε και το μοντέλο πλησιάζει στο να συμπεριφέρεται σαν τα στατικά μοντέλα, δηλαδή ουσιαστικά σαν το FJ (για το οποίο είναι γνωστό ότι το αντίστοιχο παίγνιο των Bindel, Kleinberg και Oren [Bind11] είναι αυστηρώς διαγωνίως κοίλο, όπως αποδεικνύεται στο [Bair19], ενότητα 4.2, για παράδειγμα).

1.4.2 Πειράματα και Προσομοιώσεις

Πέρα από τα αποτελέσματα της προηγούμενης ενότητας, το να δείξουμε άλλα αξιοσημείωτα θεωρητικά αποτελέσματα αποδείχτηκε αρκετά δύσκολο.

Δεδομένης αυτής της δυσκολίας, λοιπόν, είναι φυσικό να γεννηθούν ορισμένα ερωτήματα. Ισχύει πραγματικά η αυστηρή διαγώνια κοιλότητα; Μήπως μπορούμε να βρούμε αντιπαραδείγματα; Αν τρέξουμε πραγματικά τον αλγόριθμο σε στιγμιότυπα συνεξελικτικών παιγνίων, τι θα δούμε; Θα συγκλίνει σε ισορροπία, και πόσο συχνά; Και, γενικά, θα θέλαμε να κατανοήσουμε λίγο καλύτερα ποια είναι η συμπεριφορά του αλγορίθμου σε τυπικά τέτοια παίγνια.

Για τον λόγο αυτό, αλλάξαμε την προσέγγισή μας και αρχίσαμε να μελετάμε το μοντέλο μέσω προσομοιώσεων, με σκοπό να διαπιστώσουμε, για έναν αριθμό από στιγμιότυπα, τις απαντήσεις στα παραπάνω ερωτήματα.

Το πλαίσιο στο οποίο δουλέψαμε ήταν και εδώ τα συνεξελικτικά παίγνια με κανονικοποιημένα ανεξάρτητα βάρη, δηλαδή, τα βάρη παράγονται από συναρτήσεις της μορφής (1.2).

Για να ολοκληρώσουμε το βασικό μας μοντέλο, παρακάτω δίνουμε μία απαρίθμηση των συγκεκριμένων συναρτήσεων βαρών που χρησιμοποιήσαμε. Ας σημειώσσουμε εδώ ότι σε όλους τους παίκτες δόθηκε, κάθε φορά, η ίδια συνάρτηση $f_i=f$, για να απλοποιήσουμε λίγο το μοντέλο μας.

- 1. Απλή εκθετική: $f(d_i^i) = e^{-d_j^i}$.
- 2. Εκθετική προσαρμοσμένη στο διάστημα [0,1]: $f(d_j^i) = e^{-20(d_j^i-5)}$.
- 3. Γραμμική μοναδιαίας κλίσης: $f(d_i^i) = 1 d_i^i$.

4. Απλή λογαριθμική: $f(d_i^i) = -\ln d_i^i$.

Φυσικά, η συνάρτηση αυτή παράγει βάρη που δεν είναι συνεχή, αφού απειρίζονται στο 0

Παρόλα αυτά, ήταν από τη μια μεριά εύκολο να την ελέγξουμε και αυτή και, από την άλλη, δεν έτυχε να συναντήσουμε κάποιο παθολογικό σημείο (με κάποια x_j, s_i ακριβώς ίσα).

5. Λογαριθμική προσαρμοσμένη στο διάστημα [0,1]:

$$f(d_j^i) = -\ln\left(\frac{999}{1000}d_j^i + \frac{1}{1000}\right)$$

.

6. Mía logistic: $f(d_j^i) = \frac{1}{1 + e^{20(d_j^i - 0.5)}}$.

7. Η αντίστροφη: $f(d_{j}^{i}) = \frac{1}{d_{j}^{i}}$.

Όπως και ο λογάριθμος, δεν είναι συνεχής, αλλά στις προσομοιώσεις την συμπεριλάβαμε.

- 8. Προσαρμοσμένη αντίστροφη: $f(d_j^i) = \frac{1}{d_j^i + \frac{1}{10000}}$.
- 9. Τετραγωνική ρίζα: $f(d^i_j) = -\sqrt{d^i_j}$.

Μικρά Στιγμιότυπα

Ξεκινάμε με μερικές απλές προσομοιώσεις σε μικρά στιγμιότυπα συνεξελικτικών παιγνίων. Τόσο για διευκόλυνση του αναγνώστη, όσο και για να εκμεταλλευτούμε το γεγονός ότι για μικρά στιγμιότυπα μπορούμε να σχεδιάσουμε την πλήρη χρονική εξέλιξη του αλγορίθμου σε σχετικά ευδιάκριτες εικόνες. Στην εικόνα 1.1, λοιπόν, έχουμε τρέξει τον αλγόριθμο για τρία στιγμιότυπα του παιγνίου, καθένα με μια διαφορετική συνάρτηση βαρών.


Αυτό που βλέπουμε σε αυτές τις περιπτώσεις, λοιπόν, είναι ότι επιτυγχάνεται σύγκλιση αρκετά γρήγορα, με τους παίκτες να "ηρεμούν" στην ισορροοπία Nash μέσα σε 100 περίπου βήματα. Αυτό μοιάζει να ισχύει για όλες τις αρχικοποιήσεις και τις εσωτερικές απόψεις που έχουμε δοκιμάσει. Επιπλέον, να σημειώσουμε ότι δεν έχουμε συναντήσει στιγμιότυπο με πάνω από ένα οριακό σημείο (δηλαδή, ισορροπία Nash).

Μία άλλη ενδιαφέρουσα παρατήρηση απεικονίζεται στην εικόνα 1.2, όπου έχουμε τρέξει το ίδιο παίγνιο δύο φορές, μία με τους παίκτες να τρέχουν απλή δυναμική Βέλτιστης Απόκρισης (best response) και μία με το Dual Averaging.


Εδώ, βλέπουμε ότι η Βέλτιστη Απόκριση παρουσιάζει ταλαντώσεις και δεν συγκλίνει σε ισορροπία⁶. Αντίθετα, το Dual Averaging συγκλίνει σε ισορροπία, και μάλιστα σχετικά γρήγορα.

Με αυτό το παράδειγμα θέλαμε να καταδείξουμε ότι ο αλγόριθμος Dual Averaging υπερτερεί, στα παίγνια αυτά, έναντι απλών αλγορίθμων όπως η Βέλτιστη Απόκριση.

⁶ Στην πραγματικότητα, είναι ήδη γνωστό ότι η Βέλτιστη Απόκριση δεν εγγυάται τη σύγκλιση σε ισορροπία για γενικά ασύμμετρα συνεξελικτικά παίγνια διαμόρφωσης άποψης (βλ. [Bair19])



Σχήμα 1.1: Ο αλγόριθμος Dual Averaging στο συνεξελικτικό παίγνιο με 5 παίκτες, παράγοντα αυτοπεποίθησης 0.5 για όλους τους παίκτες, εσωτερικές απόψεις s=(0.4,0.5,0.6,0.127,0.897) και τρεις διαφορετικές συναρτήσεις βαρών. Κάθε χρώμα αντιπροσωπεύει μία διαφορετική τυχαία αρχική ανάθεση απόψεων.




Σχήμα 1.2: Σύγκριση Βέλτιστης Απόκρισης με Dual Averaging: Το παίγνιο έχει N=50 παίκτες, παράγοντα αυτοπεποίθησης 0.5 για όλους και για συνάρτηση βαρών την αντίστροφη: $f(d_j^i)=\frac{1}{d_j^i+\frac{1}{10000}}$. Οι εσωτερικές απόψεις s_i των παικτών και οι αρχικές τους $x_{i,1}$ έχουν επιλεγεί τυχαία.


Σύγκλιση για Μεγάλα Πλήθη Παικτών

Ο στόχος μας εδώ ήταν να δούμε πώς η σύγκλιση συμπεριφέρεται για μεγαλύτερα πλήθη παικτών, και αν συνεχίζει η καλή εικόνα που είχαμε στις απλές περιπτώσεις ή όχι. Για το σκοπό αυτό, προσομοιώσαμε στιγμιότυπα με έως 10000 παίκτες και προσπαθήσαμε να αξιολογήσουμε αν και πόσο γρήγορα συγκλίνουν σε ισορροπία.

Φυσικά, αυτές οι τάξεις μεγέθους των παικτών δεν είναι ιδιαίτερα μεγάλες, και το πλήθος των στιγμιοτύπων που τρέξαμε επίσης δεν είναι εξαντλητικά μεγάλος. Ο λόγος είναι ο φυσικός περιορισμός του χρόνου εκτέλεσης, μια και κάθε γύρος του αλγορίθμου απαιτεί χρόνο (²), για να υπολογιστούν όλα τα βάρη. Από την άλλη μεριά, δεν μας ενδιέφερε τόσο να κάνουμε μία σοβαρή αξιολόγηση της πραγματικής, πρακτικής επίδοσης του αλγορίθμου όσο να δούμε κάποιες ενδείξεις σχετικά με την συμπεριφορά του.

Το πρώτο αποτέλσμα που θα δείξουμε εδώ είναι, για έναν αριθμό στιγμιοτύπων (συγκεκριμένα, ένα για κάθε πλήθος παικτών), τη διαφορά μεταξύ των συναρτήσεων κόστους των παικτών και των ελαχίστων τους (δηλαδή, στη βέλτιστη απόκριση κάθε παίκτη), την από-





- (a) Τελική ℓ_{∞} απόσταση από τα ελάχιστα κόστη (b) Εξέλιξη με το χρόνο του ιδίου μεγέθους, για μετά από 200 γύρους του αλγορίθμου.
 - ένα επιλεγμένο υποσύνολο των στιγμιοτύπων.

Σχήμα 1.3: Σύγκλιση του αλγορίθμου (DA) για εκθετικά βάρη: $f_i(d_i^i) = e^{-d_j^i}$.

σταση των οποίων μετράμε με την ℓ_{∞} νόρμα, μετά από 200 γύρους του αλγορίθμου. Τυπικά, η ποσότητα που μετράμε είναι:

$$\max_{1 \le i \le N} c_i(x_t) - c_i(\mathbf{BR}_i(x_{-i,t}), x_{-i,t})$$

στο t = 200.

Στην εικόνα 1.3, βλέπουμε ένα παράδειγμα ενός τέτοιου συνόλου προσομοιώσεων (για εκθετικά βάρη, συγκεκριμένα).


Από τις εικόνες αυτές, φάινεται ότι ο αλγόριθμος έχει και πάλι αρκετά καλή συμπεριφορά, αφού μέσα σε μόλις 200 βήματα οι παίκτες έχουν ήδη φτάσει αρκετά κοντά σε ισορροπία (μια τιμή σφάλματος 10^{-6} σημαίνει ότι κάθε παίκτης είναι το πολύ 10^{-3} μακριά από τη βέλτιστή του απόκριση).

Να σημειώσουμε ότι έχουμε κάνει το ίδιο με παραπάνω για όλες τις συναρτήσεις βαρών που αναφέραμε, και οι αντίστοιχες εικόνες (στα Αγγλικά) βρίσκονται στο παράρτημα Β.

Επιπλέον, ας σημειώσουμε άλλες δύο ενδιαφέρουσες παρατηρήσεις. Πρώτον, όπως μπορούμε να δούμε και στην 1.3, πέρα από κάποια στιγμιότυπα με λίγους παίκτες, κατά τα άλλα δεν φαίνεται να εξαρτάται το ℓ_{∞} σφάλμα από το πλήθος των παικτών. Φυσικά, αν και δεν το έχουμε δείξει εδώ, το ℓ_1 και το ℓ_2 αυξάνονται με το N στις περισσότερες περιπτώσεις, όπως θα περίμενε κανείς. Ωστόσο, για το ℓ_{∞} αυτό σημαίνει ότι το πόσο κοντά φτάνει ο κάθε ένας παίκτης στη βέλτιστη απόκρισή του δεν φαίνεται να εξαρτάται από το πλήθος των παικτών, αλλά μόνο από το πλήθος των γύρων που τρέχει ο αλγόριθμος.

Δεύτερον, ο κύριος παράγοντας που φαίνεται να επηρεάζει την τελική απόκλιση από ισορροπία είναι η συνάρτηση βαρών. Από αυτές που δοκιμάσαμε εμείς, βλέπουμε να υπάρχουν, κατά βάση, τρεις ομάδες:

1. Η αντίστροφη και η προσαρμοσμένη αντίστροφή είχαν την χειρότερη επίδοση, με την ℓ_{∞} απόσταση να "κλείνει" στο 10^{-4} περίπου. Κάτι όχι παράλογο, αφού πρόκειται για τις πιο ακραία μεταβαλλόμενες συναρτήσεις βαρών, αφού δίνουν πολύ μεγάλο βάρος σε απόψεις που είναι κοντά στην εσωτερική του παίκτη.



Σχήμα 1.4: Running length και πλήθος γύρων μέχρι οι παίκτες να φτάσουν ε κοντά στην ισορροπία. Το παίγνιο που απεικονίζουμε εδώ έχει 8000 παίκτες, οι οποίοι χρησιμοποιούν εκθετικά βάρη: $f_i(d_i^i) = e^{-d_j^i}$.

- 2. Λίγο καλύτερες ήταν η λογαριθμική, η προσαρμοσμένη λογαριθμική και η προσαρμοσμένη εκθετική (τάξη μεγέθους $10^{-6}-10^{-7}$), που είναι επίσης απότομες, αλλά όχι τόσο όσο οι προηγούμενες.
- 3. Τέλος, η τετραγωνική ρίζα, η γραμμική, η απλή εκθετική και η logistic, ως οι πιο ομαλές, είχαν την καλύτερη επίδοση, με το σφάλμα να είναι της τάξης του 10^{-9} .

Ταχύτητα Σύγκλισης Κλείνουμε το κομμάτι των προσομοιώσεων με κάποιες επιπλέον μετρήσεις, με το σκεπτικό να καταφέρουμε ίσως να αξιολογήσουμε λίγο καλύτερα την ταχύτητα σύγκλισης και να τη συγκρίνουμε με τις εγγυήσεις του [Mert19].

Συγκεκριμένα, για τα ίδια στιγμιότυπα και τα ίδια τρεξίματα του αλγορίθμου όπως και πριν, έχουμε υπολογίσει δύο μετρικές, το running length έως να φτάσουμε ε κοντά σε ισορροπία (όπως το έχουμε ορίσει στην προηγούμενη ενότητα), αλλά και το πλήθος των βημάτων που απαιτείται για τον ίδιο σκοπό.

Απεικονίζουμε ένα παράδειγμα αυτών των μετρήσεων, και πάλι για την εκθετική συνάρτηση βαρών, στην εικόνα 1.4.

Στη θέση της πραγματικής ισορροπίας έχουμε χρησιμοποιήσει το τελικό προφίλ απόψεων των παικτών, αφού ήδη είδαμε ότι είναι κοντά σε ισορροπία.

Μία αξιοσημείωτη παρατήρηση είναι ότι το running length φαίνεται να είναι φραγμένο. Όπως έχει σημειωθεί στο [Mert19], στην παράγραφο 6.2, το running length σε εκείνη την

περίπτωση δεν περιμένουμε να είναι φραγμένο λόγω της τυχαιότητας που υπάρχει στο πλήρες μοντέλο του [Mert19]. Από τη στιγμή, όμως, που εδώ δεν έχουμε τυχαιότητα, δεν είναι παράλογο που το συνολικό running length είναι φραγμένο. Το μόνο μειονέκτημα είναι ότι είναι κάπως δύσκολο να το συγκρίνουμε με το φράγμα που δίνει το [Mert19].

Τέλος, ας σημειώσουμε ότι και από αυτές τις μετρήσεις βγάζουμε τα ίδια συμπεράσματα όπως και πριν, ότι ο αλγόριθμος γενικά συγκλίνει γρήγορα, και υπάρχει η ίδια περίπου διάκριση μεταξύ των επιδόσεων για τις διαφορετικές συναρτήσεις βαρών.

Αντιπαράδειγμα στην Αυστηρή Διαγώνια Κοιλότητα

Αν η αρχική μας ελπίδα ίσως να ήταν τα συνεξελικτικά παίγνια διαμόρφωσης άποψης να είναι πάντα αυστηρώς διαγωνίως κοίλα, δυστυχώς κάτι τέτοιο φαίνεται να μην ισχύει. Σε αυτή την παράγραφο, δίνουμε ένα σχετικά μικρό αντιπαράδειγμα.

Εξετάζουμε το παίγνιο με τις εξής παραμέτρους:

- 1. N = 5 παίκτες.
- 2. $\rho_i = \frac{1}{2}$ για όλους τους παίκτες.
- 3. s = (0.9, 0.84, 0.83, 0.2, 0.18)
- 4. κανονικοποιημένα ανεξάρτητα βάρη με εκθετικές συναρτήσεις βαρών, δηλαδή:

$$f_i(d) = \frac{1}{d + 0.0001}$$

Επιπλέον, εξετάζουμε τα ακόλουθα δύο προφίλ απόψεων των παικτών:

$$x^1 = (0.84, 0.89, 0.8, 0.49, 0.3)$$

 $x^2 = (0.83, 0.87, 0.8, 0.49, 0.3)$

Τέλος, για να ελέγξουμε γραφικά την αυστηρή διαγώνια κοιλότητα, αντί να καταπιαστούμε με αριθμητικές πράξεις, μένει να παρατηρήσουμε για ακόμα μία φορά ότι ο ορισμός της είναι άμεσα ισοδύναμος με το εξής: οι προβολές της ψευδοκλίσης πάνω σε οποιαδήποτε γραμμή, οποιοδήποτε ευθύγραμμο τμήμα του χώρου των απόψεων είναι φθίνουσες (αύξουσες αν θεωρήσουμε τα κόστη αντί για τα κέρδη).

Κατά συνέπεια, μπορούμε να ελέγζουμε γραφικά την αυστηρή διαγώνια κοιλότητα απεικονίζοντας την ακόλουθη μονοδιάστατη συνάρτηση:


$$g(t) = \langle v(x^1 + t(x^2 - x^1)), x^2 - x^1 \rangle$$


 $\mathrm{gia}\ t\in[0,1].$

Αυτήν, λοιπόν, ακριβώς την συνάρτηση έχουμε σχεδιάσει στην εικόνα 1.5a, για το παίγνιο και τα δύο σημεία του χώρου των απόψεων που δώσαμε παραπάνω. Εδώ, έχουμε χρησιμοποιήσει τις συναρτήσεις κόστους των παικτών, οπότε θέλουμε η g(t) να είναι αύξουσα.

Είναι εμφανές, όμως, από την εικόνα αυτή, ότι μεταξύ των δύο συγκεκριμένων σημείων η g(t) δεν είναι αύξουσα, εξ ου το αντιπαράδειγμα στο οποίο αναφερθήκαμε. Για πληρότητα, έχουμε κάνει και μία ανάλυση με αριθμητικές πράξεις, που βρίσκεται στο παράρτημα \mathbf{A} , αν θέλει ο αναγνώστης να την δει.

Η επόμενη παρατήρηση που θα πρέπει να κάνουμε, ωστόσο, είναι ότι ενώ στην εικόνα 1.5a φαίνεται να υπάρχει μεγάλη απόκλιση, στην πραγματικότητα δεν είναι ακριβώς έτσι.





- (a) Εστίαση στο ευθύγραμμό τμήμα με άκρα τα x^1, x^2 .
- (b) Ολόκληρη η ευθεία, έως σχεδόν το σύνορο του υπερκύβου.

Σχήμα 1.5: Προβολές της ψευδοκλίσης πάνω στη γραμμή x^1 - x^2 .

Παρατηρήστε την εικόνα 1.5b, όπου έχουμε επεκτείνει το ευθύγραμμο τμήμα μεταξύ των x^1, x^2 μέχρι σχεδόν το σύνορο του υπερκύβου. Στην εικόνα αυτή, είναι πιο εμφανές ότι η g(t) είναι σχεδόν αύξουσα, η τοπική αυτή απόκλιση είναι πολύ μικρή.

Επίσης, θα πρέπει να σημειώσουμε ότι, αν παρατηρήσει κανείς το s και τα x^1, x^2 , βλέπει ότι κάποιες συνιστώσες τους είναι σχεδόν ίσες. Οπότε, στην πραγματικότητα, το αντιπαράδειγμά μας βασίζεται καίρια στο ότι ο παρονομαστής της αντίστροφης είναι πολύ κοντά στο 0, δηλαδή βασίζεται στην συμπεριφορά της αντίστροφης στο 0, δεν είναι ότι καταφέραμε να βρούμε αντιπαράδειγμα σε μια περιοχή όπου η συνάρτηση είναι πιο ομαλή.

Παρόλα αυτά, η συγκεκριμένη συνάρτηση είναι πλήρως συμβατή με το μοντέλο μας (συνεχής, φθίνουσα, παραγωγίσιμη), οπότε, αν μη τι άλλο, το συγκεκριμένο αντιπαράδειγμα, όσο οριακή περίπτωση κι αν είναι, μας δείχνει ότι είναι αδύνατον να αποδείξει κάποιος απλά ότι οποιοδήποτε συνεξελικτικό παίγνιο διαμόρφωσης άποψης είναι αυστηρώς διαγωνίως κοίλο, χωρίς καμία επιπλέον υπόθεση ως προς τις συναρτήσεις των βαρών.



Chapter 2

Introduction

The way people form their opinions has been of interest since many decades ago, from the perspective of several scientific fields. For example, psychologists may have been studying how opinions affect and are affected by the state of a person's psyche, while political scientists may have been interested specifically in how political opinions propagate, while sociologists were interested in the way opinions influence and are influenced by the interaction between people or groups of people (e.g. cultural norms).

This interest has grown significantly in the last few years, particularly due to the advent of the Internet and social networks. On one hand, the same phenomena of opinion formation have taken on much larger proportions, since there are now thousands or even millions of people forming some kind of connections, groups etc (interactions, in general) and, moreover, this information is publicly available. This is unprecedented, and it means that we can now observe how peoples' behaviour, interaction and opinion formation translate to such large scales, but also that there are much more data available to analyze and draw conclusions from. On the other hand, with the use of social networks for advertising, political campaigns etc., there is a great deal of profit to be made by knowing how opinions are formed. Advertisers, for example, are particularly interested in how consumer opinions are shaped, so that they can, possibly, be swayed to allow for easier or more widespread adoption of products or services.

One of the earliest works on opinion formation is the one by Morris DeGroot ([DeGr74]), which studies a simple model, namely one where there are some people (a.k.a. agents) comprising a social network, and each one has a number. This represents the opinion of the agent, which is then revised at discrete rounds, by taking into account the opinions of all other agents, formally by taking a weighted average of all opinions:

$$x_{i,t+1} = \frac{\sum_{j=1}^{N} w_{ij} x_{j,t}}{\sum_{j=1}^{N} w_{ij}}$$

where $x_{i,t} \in [0,1]$ represents the opinion of agent i at round t. The weights $w_{ij} \ge 0$ represent the judgement of the agent, so to speak, as to who of the others is the most knowledgable / experienced / trustworthy, etc.

Simple as it may seem, the DeGroot model has spawned a substantial amount of research. A recent interesting example is on the so-called "social learning" line of work, where the opinions are seen as the agent's inexact perception of some ground truth about the world. Then, the opinion formation process essentially represents the attempts of the people to combine their incomplete knowledge with the incomplete knowledge of other people, in order to arrive at a new perception, which is as close as possible to the truth.

On this subject, [Golu10] first showed that, in a setting where the agents' initial opinions are noisy samples of a ground truth number, the simple DeGroot dynamics lead to a consensus which is very close to the actual number, when the network is large.

Building on this, a more recent result by Amir et al. [Amir21] shows that, additionally, weaknesses of the DeGroot model (sensitivity to "bots" and to specific types of miscommunications) can be overcome by simply discretizing the opinion space with sufficient granularity. This means that, with a slight compromise with respect to precision, the dynamics is made much more robust.

This is only an example of works related to the DeGroot model. And, after all, even though it is a relatively simple model compared to others, it already lays much of the foundation (the norms, if you will) as to how opinion formation models are constructed. Opinions are, in general, represented by numbers, each agent assigns a weight to all other agents (and herself, in some cases) and the way that the agents' opinions evolve depends solely on an interaction process between them. This interaction process amounts to repeated revision rounds, where in each round, each agent updates their opinion to a function of the opinions of all agents (including their own) from the previous round, in most cases some kind of weighted average. All these aspects are more or less preserved by the models that followed afterwards.

A prominent example is the model of Friedkin and Johnsen, which extends the DeGroot model by also assigning hidden, intrinsic opinions to the agents, who then take those into account whenever they perform an average. The models of Hegselmann and Krause ([Hegs02]) and Deffuant and Weisbuch ([Deff00]), on the other hand, first introduced the concept of dynamic (co-)evolution of interactions, with the strength of the interaction between two agents changing dynamically, based on the proximity of their opinions.

A more recent line of work has connected the opinion formation procedure to game theory. Under this light, the agents are now players, who incur a cost if, for example, their expressed opinion is too far away from their peers', or if it is too far away from their own intrinsic, hidden opinion (a "cognitive dissonance" of sorts); all this information is encoded in the players' cost functions, which, in general, depend on the expressed opinions of all players. The opinion formation procedure is, then, the repeated version of this game over many discrete rounds; and the purpose of the agents is to update their opinions at each round in such a way as to minimize their total cost across all rounds that the game is played.

An exemplary work is the one of Bindel, Kleinberg and Oren ([Bind11]), who formed a game theoretic formulation for the FJ (Friedkin - Johnsen) model, by assigning quadratic cost functions to the player, in such a way that their best response coincides with the update rule of the original FJ model. In other words, the opinion dynamic of the FJ model is exactly the simultaneous best-response dynamic in this game. [Bind11] focuses on the social cost of disagreement of these games, proving bounds for the Price of Anarchy (introduced as a concept in [Kout99]), which measures the worst-case social cost at a Nash equilibrium versus the optimal social cost achievable. They also study how the edges of the social network can be manipulated in order to reduce the social cost at equilibrium.

This has also spawned a number of works that extended the Proof of Anarchy bounds to more general cases, for continuous but also for discrete binary opinions (usually in $\{0,1\}$) [Bhaw13, Chen16, Bil \Box 18, Chie18, Yild13], which is also an interesting setting, since, despite its seeming simplicity, it models some interesting scenarios, such as the choice of a electoral candidate, or the choice between two social media platforms by a user.

2.1 Contributions

A central question for any opinion formation model is whether it eventually converges to a stable point. Also, under what conditions and how fast does this happen.

For example, it is already known even from [DeGr74] that the DeGroot model eventually settles on a consensus and, depending on the weights and the initial opinions, any opinion in [0,1] can be reached. Later works also showed that the convergence is fast [Ghad13].

In the FJ model, on the other hand, agents do not reach consensus, due to their intrinsic opinions. Rather, they converge to a specific, unique stable point, regardless of the initial opinions. Moreover, [Ghad13] again shows that the convergence is linear.

As for the HK and DW models, they exhibit what can be called *polarization*, where the agents converge (again, relatively fast) to clusters, where agents of the same cluster have the same opinion, which is somewhat far away from the opinions of the other clusters.

Now, the opinion formation model that we are interested in is game theoretic, like the one in [Bind11]. In the context of games, instead of stable points, we can directly talk about the central concept of Nash equilibriua. Many questions can be asked about the Nash equilibria of a game (the following list is non exhaustive). Do they always exist? Do players eventually converge to them? Under which dynamics? How quickly? How much worse is the total social welfare at an equilibrium, compared to the centrally computed optimum (price of anarchy)? Our main concern here is about the convergence to Nash equilibrium (i.e. the computation of Nash equilibria), and its efficiency.

In general, the computation of Nash equilibria is a vast research domain, with overwhelming amount of works. There are negative, hardness results (e.g. a known early example is that mixed equilibria in finite games are PPAD-complete [Dask09, Chen09]), but also several positive ones (e.g. the widely known convergence of no-regret dynamics to coarse correlated equilibria [Fost97, Freu99, Hart00]), even for classes of games where not only the computation of equilibria is plausible, but it can be achieved by the players themselves with very simple, natural dynamics (e.g. best-response in congestion games [Even03, Polj95], and no-regret we mentioned before).

The discipline that interests us the most, however, is the one intersecting with convex optimization. The research is again plentiful, especially in the area of no-regret algorithms and their behaviour with respect to equilibria. For example, in [Even09], it is shown that in a large class of games, all no-regret dynamics converge ergodically to Nash equilibrium at the same rate as the regret decay. In [Gian21b], it is shown that, in finite games, a known no-regret algorithmic scheme (FTRL) converges only to strict equilibria, and [Gian21a] also gives convergence rates. In potential games, which intersects heavily with convex optimization, as the problem of equilibrium computation translates to an optimization problem, there are works that make heavy use of optimization techniques such as mirror descent, extra-gradients, acceleration etc. [Vu21, Anto21, Hsie21].

In this work, we focus on the question of efficient equilibrium computation, for a specific game model of opinion formation, called asymmetric coevolutionary opinion formation games. First introduced by Bhawalkar, Gollapudi and Munagala ([Bhaw13]), "coevolutionary" refers to the fact that, as the HK and DW models, the bond between two players is dynamic, and it depends on the proximity of their opinions. It is also intrinsically asymmetric, because each player possesses a hidden opinion, which affects the weight they assign to other players. In other words, this whole model is the same as the FJ model, but with dynamic weights, instead of static.

In these games, equilibria always exist, due to Rosen's theorem [Rose65], because each player's cost is convex in her own opinion. However, the computational complexity of finding such equilibria is still an open problem, and this is the motivation of our work. More

¹ as a matter of fact, it is the same, but with the interaction weights changing depending on the players' opinions

specifically, we seek to find an algorithm which, if employed by the players of such a game, will guarantee convergence to equilibrium, and in a reasonable number of steps.

To that end, our main approach is to employ the recent results by Mertikopoulos and Zhou [Mert19]. Their work deals with a more general setting, that of concave games, and it says that, under a slightly stronger version of concavity, there is a suitable algorithm that does exactly what we want. The condition is called diagonal strict concavity, and the algorithm is the well known in online optimization Dual Averaging. Thus, the main focus of our work was to study if and how this result can be applied to the asymmetric coevolutionary opinion formation games.

Our results First of all, we show, via straightforward calculations, that, in these games, diagonal strict concavity is equivalent to a smoothness condition on the players' best response mappings², which are, again, an average, in this case of the expressed opinions of all other players and the player's intrinsic opinion; crudely, this condition says that, when players move, their averages shouldn't move arbitrarily much.

This (as one may imagine), then, also translates to a (stronger) Lipschitzness condition which is sufficient for diagonal strict concavity. Moreover, this condition is critically dependent on the players' so-called stubbornness / self-confidence factors, which are numbers $\rho_i \in [0,1]$ that express what compromise the player is going to make between being close to her intrinsic opinion and to the opinions of her "friends" (0 means she cares only about social conformity, 1 that she cares only about cognitive dissonance). An immediate consequence of this condition is that, basically, if players are stubborn enough (but not only for $\rho_i = 1$, which is easy to show), then the game will be diagonally strictly concave.

We are then able to go to slightly more specific, albeit nontrivial, subclasses of these games and quantify this even further, by appropriately bounding the derivatives of the averages. Thus, we procure a way to show diagonal strict concavity for any game in which the players' self-confidence is above a threshold; this threshold is determined by the bounds one may manage to work out for the weight functions and their derivatives.

So, in the end, diagonal strict concavity holds for those cases, which, in turn, means that the algorithm from [Mert19] does indeed converge to the unique equilibrium of those asymmetric coevolutionary opinion formation games, which, however, are a bit constrained, in the sense that the players' self-confidence needs to be somewhat close to 1.

The second part of our results concerns the observations we made when we tried to explore the games of our model via simulations (due to the difficulty in establishing stronger results). An immediate observation was that, on all games we saw, the algorithm converges to equilibrium, and somewhat fast at that. We simulated games with up to 10000 players, and convergence sufficiently close to equilibrium was always achieved after at most 200 rounds.

An interesting find was that, with respect to the ℓ_{∞} distance, the convergence speed (in number of *rounds*) of the algorithm did not seem to depend on the number of players. In other words, any given player reached close to her best response in more or less the same number of rounds. On the other hand, the speed does seem to depend on the choice of the weight functions, with the distance to equilibrium after the same number of steps ranging from $\approx 10^{-9}$ to $\approx 10^{-4}$, which, comparatively, shows that for some weight functions it takes much longer to converge. Specifically, from the functions we tested, those that took the longest were the ones that changed the most extremely with the distance between the players' opinions. Interestingly, for these functions it was also easier to procure examples where the simple best response dynamics failed to converge (and oscillated instead).

² i.e. best response correspondences, but the best response is actually always unique

Finally, by testing experimentally the diagonal strict concavity condition, we have found that there are deviations, even if small ones. Specifically, when the interaction weights drop (approximately) as 1/d with the distance d between two players, then there are points on the opinion space where diagonal strict concavity fails. This most probably comes as a consequence of 1/d being too aggressively changing. We should note that, as the reader may have realized, this weight function was one of those that resulted in the slowest convergence, that we mentioned before.

This, in turn, implies that not all coevolutionary games are diagonally strictly concave and, up to some degree, explains the difficulties we have faced while trying to approach this analytically. At the same time, the discrepancies we found are very small, which explains why (at least in the games we studied and simulated) the algorithm performed so well; because, practically, the condition does hold.

2.2 Related Work

There are a number of works which try to show convergence to either stable points or equilibria, and are tangential to our goals. For example, [Lore05] shows that there is a wide range of opinion dynamics in which weights are dynamically changing, for which convergence to polarized groups is guaranteed, as long as the self-confidence of the agents is high enough. The crucial deviation from our model is that this does not have intrinsic opinions, i.e. it more closely resembles the DeGroot model with dynamic weights, rather than the FJ model with dynamic weights. This difference seems to be crucial, as, indicatively, one can see (both in [Frie90] and [Ghad13]) that the dominant term for convergence is the one that involves the intrinsic opinions, the other simply vanishes. So, we do not expect the results of [Lore05] to give any insight into our games, in general.

From another perspective, [Fane21] studies specifically the same games as us; the central idea here is that, by slightly modifying the cost functions, he derives a new game, and it is then shown that either a consensus or the vector of intrinsic opinions must be a 2-approximate Nash equilibrium. Not only in the modified game, but also in the original. Here, the truth is that we are interested in exact equilibria, but this is a very interesting result.

Finally, our work is in many ways a follow-up of [Bair19], where diagonal strict concavity again makes an appearance, as it is shown that a dynamic which coincides with both follow the leader and a version of gradient descent does converge, if diagonal strict concavity holds. The main difference here is that we now know from [Mert19] that a similar dynamic always converges fast to the unique equilibrium of any diagonally strictly concave game, thus we focus all our efforts into studying how diagonal strict concavity can or cannot be shown for asymmetric coevolutionary opinion formation games, and then testing experimentally whether the algorithm seems to indeed converge or not.

2.3 Overview of the chapters

In chapter 3, we begin by providing some background material needed for the sequel. There are two main pillars: convex optimization and game theory (concave games, in particular). We give a brief overview of central concepts from convex optimization and online convex optimization. Then, we focus on the concept of games and Nash equilibrium, following with some known results (mostly of Rosen - [Rose65]) for the special, generally well-behaved class of concave games, in which all games we discuss lie.

Then, in chapter 4, we give a more detailed account of various opinion formation models that have been studied in the literature, while also underlying the connection between opinion formation and optimization that is found in a lot of cases.

In chapter 5, we give an overview of the results of [Mert19]. We introduce the algorithm of Dual Averaging, and prove that it converges to the unique equilibrium of any diagonally strictly concave game. We also show that it does so in polynomial number of steps.

Finally, in chapter 6, we exhibit the results of our work. We start with our main theoretical result, which shows diagonal strict concavity and, hence, convergence to equilibrium, albeit for a limited subclass of coevolutionary opinion formation games. We then continue with the results of our simulations, which provide some insight into the practical performance of the algorithm for some characteristic weight functions, but also show that there are valid, even if extreme, counterexamples to diagonal strict concavity, meaning that proving diagonal strict concavity in full generality is unlikely to be possible.

Chapter 3

Technical Background

3.1 Games

All the games we will study are finite-player, continuous-action games. Specifically, there are N players $(N \in \{1, 2, 3, ...\})$ and each player $i \in [N] = \{1, ..., N\}$ selects his strategy / action from a compact convex set \mathcal{X}_i , generally assumed to be a subset of some Euclidean space \mathbb{R}^{m_i} . The game's (combined) action space is then denoted as $\mathcal{X} = \prod_{i=1}^N \mathcal{X}_i$.

The payoff to each player is modeled by a payoff function dependent on all players' chosen actions, i.e. $u_i \colon \mathcal{X} \to \mathbb{R}$, with $u_i(x) \in \mathbb{R}$ representing the payoff received by player i when the players choose the profile $x = (x_1, \dots, x_N)$ of actions. Often, the action profile x is denoted as (x_i, x_{-i}) , to highlight the action x_i of player i against the actions x_{-i} of all other players.

One way to think of continuous-action games is as one possible generalization of the classical (continuous) optimization setting, where we have one function over some continuous domain and we seek to find its extreme points, either maxima or minima. Then, continuous games can be seen as a generalization of this to more than one (but finite) functions, but, instead of regularization, Pareto optimality or other generalizations, here we associate each function i with a player who controls the component i of $x \in \mathcal{X}$ and seeks to optimize this function.

All the above are the definitions of the basic concepts involved in a game. Below, we define some more refined classes of games, which satisfy specific conditions, and which are relevant for us.

Definition 3.1 (Continuous Games). We call a game continuous when, for all players $i \in [N]$, u_i is continuous in $x \in \mathcal{X}$.

Definition 3.2 (Concave Games). A continuous game is further called concave when, for all players $i \in [N]$, $u_i(x_i, x_{-i})$ is concave in x_i for every fixed value of $x_{-i} \in \mathcal{X}_{-i} = \prod_{j \neq i} \mathcal{X}_j$.

Definition 3.3 (Smooth Games). A game is called smooth when, for all players $i \in [N]$, $u_i(x_i, x_{-i})$ continuously differentiable in x_i (for any fixed value of x_{-i}), i.e. each players payoff function is continuously differentiable in his own action.

In this case, we will denote by $v_i(x) = \nabla_{x_i} u_i(x_i, x_{-i})$ the individual gradient of player i, and we will call the following ensemble:

$$v(x) = (v_1(x), v_2(x), \dots, v_N(x))$$

the pseudo-gradient of the game at the action profile $x \in \mathcal{X}$.

Most of the games we will deal with will be concave and smooth (we may call them simply concave for brevity). We will later look into concave games in more detail, as they are of great interest.

3.2 Nash Equilibrium

A central solution concept of game theory, first introduced by Nash ([Nash50], [Nash51]). A (pure) Nash Equilibrium is any action profile in which no player has incentive to unilaterally deviate from their individual action. Formally:

Definition 3.4 (Nash Equilibrium). *An action profile* $x^* \in \mathcal{X}$ *is called a* Nash equilibrium *if*:

$$u_i(x_i^*, x_{-i}^*) \ge u_i(x_i, x_{-i}^*), \text{ for all } i \in [N], x_i \in \mathcal{X}_i$$

Furthermore, since at equilibrium each player is maximizing their payoff, we get the following first-order necessary condition for equilibrium:

$$\langle v_i(x_i^*, x_{-i}^*), x_i - x_i^* \rangle \le 0, \text{ for all } i \in [N]$$
 (3.1)

This is a well known generalization of simple stationarity ($\nabla f(x^*) = 0$) for extreme points of functions, and it has a simple geometric interpretation.

Consider a differentiable function $f : C \to \mathbb{R}$. Then, as we know, the gradient $\nabla f(x) \in C$ is a vector that is perpendicular to the contour of f on x, and it points to the direction of steepest ascent. Moreover, the directional derivative of f along any line starting from x and parallel to, let's say, the vector h, can be expressed as the inner product $\langle \nabla f(x), h \rangle$.

Under this light, if, for any point x of the domain it holds that $\langle \nabla f(x^*), x - x^* \rangle > 0$, then this means that the directional derivative of f at x^* towards x is positive, which means that f is locally strictly increasing at x^* , on the line from x^* to x. But then, x^* cannot be a maximum point.

This is why this is always a necessary condition for optimality. Plus, it is also sufficient if the function is concave (for more details on this, see, for example, [Vish21], theorem 3.14, or [Boyd14], section 4.2.3).

Applying this fact to the payoff function $u_i(\cdot, x_{-i}^*)$, which we know is maximized at the Nash equilibrium strategy of player i, we conclude that (3.1) for every player $i \in [N]$, is a necessary condition for a Nash equilibrium. Moreover, in the case of concave (or pseudo concave) games, it is also sufficient.

Furthermore, it can be concisely encoded to the following proposition, encompassing all players at once:

Theorem 3.5. If x^* is a Nash equilibrium, then:

$$\langle v(x^*), x - x^* \rangle \le 0 \tag{3.2}$$

for all $x \in \mathcal{X}$.

If the game is concave, then the reverse is also true.

Proof. If x^* is a Nash equilibrium, then we can simply add (3.1) for all $i \in [N]$, and we get (3.2).

For the converse, if (3.2) holds for all $x \in \mathcal{X}$, then we can substitute $x = (x_i, x_{-i}^*)$, and then all components of the inner product vanish, except for the *i*-th component, so we get:

$$\langle v_i(x^*), x_i - x_i^* \rangle \le 0$$

which, since the game is concave, implies that player i is best responding. This holds for any i, so x^* must be a Nash equilibrium.

3.3 Concave Games

Concave games are an especially interesting class of games. As is the case for concave maximization, there are several results concerning Nash equilibria (properties, computation etc.) of concave games or subclasses thereof.

A major part of our work concerns concave games or subclasses. Here, we begin the journey by giving a brief overview of some interesting results from [Rose65], which is one of the first works that pioneered the study of concave games.

We start with the existence of Nash equilibria, which is always guaranteed for any concave game.

Theorem 3.6 (Rosen). In a concave game, a Nash equilibrium is guaranteed to exist.

We give a small proof for completeness, although we believe there are much better specimens out there (see [Rose65] itself, for example).

Proof. Define the best response correspondence of player $i \in [N]$ as the set of all best responses of player i to a profile x_{-i} of all other players' actions, i.e. the set of optimal actions player i can choose to play, actions that maximize her payoff, if she knows that all other players will play x_{-i} . Formally:

$$\begin{aligned} & \mathrm{BR}_i \colon \mathcal{X}_{-i} \to 2^{\mathcal{X}_i} \\ & \mathrm{BR}_i(x_{-i}) = \{x_i^* \in \mathcal{X}_i \colon u_i(x_i^*, x_{-i}) \ge u_i(z_i, x_{-i}) \text{ for all } z_i \in \mathcal{X}_i \} \end{aligned}$$

Then, we define the combined best response correspondence as the ensemble of the above:

CBR:
$$\mathcal{X} \to 2^{\mathcal{X}}$$

CBR $(x) = BR_1(x_{-1}) \times BR_2(x_{-2}) \times \cdots \times BR_N(x_{-N})$

It is not difficult to see that a fixed point of the above correspondence, i.e. a $x^* \in \mathcal{X}$ such that $x^* \in \operatorname{CBR}(x^*)$, has to consist, for each player i, of an action x_i^* that is best responding against x_{-i}^* . But this means that x^* is a Nash equilibrium.

As such, the proof mainly comes down to an application of Kakutani's fixed point theorem ([Kaku41]) on the point-to-set correspondence CBR. The form of the theorem we use here is as follows.

Suppose \mathcal{X} is a nonempty, compact and convex subset of some Euclidean space and $\phi \colon \mathcal{X} \to 2^{\mathcal{X}}$ is a point-to-set correspondence on \mathcal{X} . Then, ϕ is guaranteed to have a fixed point (i.e. $x \in \phi(x)$), if the following two conditions hold:

- $\phi(x)$ is nonempty and convex for all $x \in \mathcal{X}$.
- ϕ has a closed graph. This means that for all sequences $\{x_n\}_{n\in\mathbb{N}}$ and $\{y_n\}_{n\in\mathbb{N}}$ such that $x_n \to x$, $y_n \to y$ and $y_n \in \phi(x_n)$ for all n, it holds that $y \in \phi(x)$.

Here, the action space \mathcal{X} is already assumed to be non-empty, compact and convex.

 $\mathrm{BR}_i(x)$ is always nonempty and convex as the set of maximizers of the function u_i which is continuous and concave. The same holds for $\mathrm{CBR}(x)$, as it is the cartesian product of such sets.

Finally, CBR has a closed graph, since for any sequences $\{x_n\}_{n\in\mathbb{N}}$, $\{y_n\}_{n\in\mathbb{N}}$ such that $x_n \to x$, $y_n \to y$ and $y_n \in CBR(x_n)$ it holds that:

$$\begin{aligned} y_{n,i} &\in \mathrm{BR}_i(x_{n,-i}) \\ \Rightarrow & u_i(z_i,x_{n,-i}) \leq u_i(y_{n,i},x_{n,-i}) \text{ for all } z_i \in \mathcal{X}_i \\ \Rightarrow & \forall z_i \in \mathcal{X}_i, u_i(z_i,x_{-i}) = \lim_n u_i(z_i,x_{n,-i}) \leq \lim_n u_i(y_{n,i},x_{n,-i}) = u_i(y_i,x_{-i}) \\ \Rightarrow & y_i \in \mathrm{BR}_i(x_{-i}) \end{aligned}$$

In other words, $y \in CBR(x)$.

Thus, the conditions for Kakutani's fixed point theorem are satisfied.

3.3.1 Diagonal Strict Concavity

Moving on from simple concavity of the game, there is an additional condition which, when imposed on a concave game, can give us guarantees regarding the computation of an equilibrium. [Rose65] calls this condition **diagonal strict concavity** and shows that, for a diagonally strictly concave game, not only is the equilibrium unique, but there are also both continuous-time and discrete-time dynamics that can be followed by each player and converge to that equilibrium.

Formally, it is defined as follows.

Definition 3.7 (Diagonal Strict Concavity). A game is called diagonally strictly concave if

$$\langle v(x') - v(x), x' - x \rangle < 0$$
 (DSC)

П

holds for every $x \neq x'$.

In other words, a game is diagonally strictly concave if its pseudogradient "behaves" like the gradient of a strictly concave function, in the sense of the first order condition $\langle \nabla f(x') - \nabla f(x), x' - x \rangle < 0$, which says that the directional derivative of f along any line is decreasing. So, here, it is something similar, but for the pseudogradient.

Lemma 3.8. A diagonally strictly concave game is a concave game.

Proof. In the above definition of diagonal strict concavity, let $x'_{-i} = x_{-i}$. Then:

$$\langle v_i(x_i', x_{-i}) - v_i(x_i, x_{-i}), x_i' - x_i \rangle < 0$$

which implies (strict) concavity of the function $u_i(\cdot, x_{-i})$, which in turn means that the game is concave.

Theorem 3.9. If a game is diagonally strictly concave, then it has a unique Nash equilibrium.

Proof. Since the game is concave (from the above lemma), we already know that there is some equilibrium, call it $x^* \in \mathcal{X}$. Now, suppose that there is also another, $x^{**} \neq x^*$. Since the game is diagonally strictly concave, we have:

$$\langle v(x^{**}) - v(x^{*}), x^{**} - x^{*} \rangle < 0 \Rightarrow \langle v(x^{**}), x^{**} - x^{*} \rangle < \langle v(x^{*}), x^{**} - x^{*} \rangle \le 0$$

where the last inequality comes from (3.2).

Thus, $\langle v(x^{**}), x^* - x^{**} \rangle > 0$, which in turn violates the necessary condition (3.2) for x^{**} , so it cannot be an equilibrium.

As is the case for concave scalar functions, there is a sufficient second-order condition for the diagonal strict concavity of a game, which involves the jacobian matrix of the game's pseudogradient; similarly to concave functions, where the matrix involved is the Hessian, which coincides, as is known, with the jacobian of the function's gradient.

The game's $Hessian\ G\colon \mathcal{X}\to\mathbb{R}^{m\times m}$, where m is the dimension of the ambient Euclidean space of \mathcal{X} (equal to the sum of all dimensions of the players actions), is thus defined as the Jacobian of $v(\cdot)$, i.e. the j-th column of G(x) is $\frac{\partial}{\partial x_j}v(x)$, $j=1,\ldots,m$.

Theorem 3.10. If the symmetric matrix $G(x) + G(x)^T$ is negative definite, then the game is diagonally strictly concave.

For the proof, we defer the reader to [Rose65].

Finally, we state the following two theorems regarding dynamics that converge to equilibrium. The proofs are omitted, as they are not so straightforward and of no great interest to us. For details, see [Rose65].

Theorem 3.11. There is a continuous-time dynamical system, i.e. a system of differential equations, of the form $\dot{x}_i = f_i(x)$, for which:

- 1. There always exists a continuous solution x(t) which remains in \mathcal{X} for all t > 0.
- 2. If $G(x) + G(x)^T$ is negative definite, then, under these dynamics, the unique equilibrium x^* is globally asymptotically stable.

Theorem 3.12. There is a discrete-time dynamical system under which the unique equilibrium of a diagonally strictly concave game (for which, specifically, $G(x) + G(x)^T$ is negative definite) is globally asymptotically stable.

Specifically, the discrete-time system is the forward Euler discretization of the continuous-time system, i.e. $x_{t+1} = x_t + \gamma_t f(x_t)$, with the step size γ_t chosen appropriately.

The key takaway from this is that [Rose65] has shown that equilibria of diagonally strictly concave games can be computed with some dynamics.

A natural following question is what is the complexity of computing these equilibria, since [Rose65] does not address how long it takes these dynamics to approach the equilibrium. This is addressed in [Mert19], which we will review later on.

3.4 Convex Optimization

One common way that equilibrium computation can be tackled is with techniques from (online) convex optimization. This holds first and foremost for the class of potential games, where equilibrium computation turns out to be equivalent to the optimization of some potential function. However, fruitful conclusions can be derived even for more general classes of games, as we will later see.

In this section, we provide definitions for fundamental concepts of (online) convex analysis and we give a brief overview of some straightforward results from (online) convex optimization, to lay the groundwork for what is to come and to showcase the analysis of an optimization method in simple cases.

We begin, as is the norm, with the definition of convexity. Take note that the term "convex" is used to describe two things: convex sets and convex functions.

Definition 3.13 (Convexity). A set $C \subseteq \mathbb{R}^n$ is called convex if for any two points $x, y \in C$ the line segment connecting them is inside of C, i.e. $x + \lambda(y - x) \in C$ for all $\lambda \in [0, 1]$.

A function $f: C \to \mathbb{R}$ is called convex if C is convex and for any two points $x, y \in C$ it holds that $f(\lambda x + (1 - \lambda)y) \le \lambda f(x) + (1 - \lambda)f(y)$.

Intuitively, a function is convex if any line connecting two points on the graph of f lies above the graph. There are, however, two equivalent, and perhaps more straightforward, characterizations of differentiable convex functions.

The first is that they lie below their first-order (linear) approximation at any point. Formally, for all $x, y \in dom(f)$, it holds that:

$$f(y) \ge f(x) + \langle \nabla f(x), y - x \rangle$$

The second is that the directional derivative of a convex function along any line is increasing. Formally, for all $x, y \in dom(f)$, it holds that:

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \ge 0$$

Of course, this should remind you of diagonal strict concavity. Indeed, diagonal strict concavity is a generalization of the above condition to the pseudogradient of a game. As a matter of fact, for a potential game, the game being diagonally strictly concave and its potential being strictly (strict inequality in the above) concave are both the same as the above.

Finally, there is yet another characterization of convexity for twice differentiable functions. This says that a function $f: C \to \mathbb{R}$, where C is convex and open, is convex if its Hessian $\nabla^2 f(x)$ is positive semidefinite for all $x \in C$.

These are all fundamental properties, variations of which we will encounter in all of the following work.

Now, the next concept we should mention is further conditions which guarantee the convergence of convex optimization algorithms. Because, the truth is, convexity by itself is not enough to prove convergence and efficiency for almost any optimization algorithm.

We start with strong convexity, which basically says that the function has a quadratic lower bound. Formally:

Definition 3.14 (Strong Convexity). A differentiable function $f: C \to \mathbb{R}$ is called l-strongly convex if:

$$f(y) \ge f(x) + \langle \nabla f(x), y - x \rangle + \frac{l}{2} ||x - y||^2$$

for all $x, y \in C$.

In a similar manner to convexity, strong convexity of a twice differentiable function is also equivalent to $\nabla^2 f(x) \succeq lI$.

Next, we introduce Lipschitz continuity, which imposes a bound on the rate of change of a function.

Definition 3.15 (Lipschitz Continuity). *A function* $f: C \to \mathbb{R}$ *is called G-Lipschitz continuous if:*

$$|f(x) - f(y)| \le G||x - y||$$

for all $x, y \in C$.

which, for convex functions, is equivalent to $\|\nabla f(x)\|_* \leq G$.

Finally, when the gradient of a function is itself Lipschitz, the function is called smooth:

Definition 3.16 (Smoothness). A function $f: C \to \mathbb{R}$ is called L-smooth if its gradient is L-Lipschitz continuous, i.e.:

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

for all $x, y \in C$.

Moreover, for convex functions, this condition can be proved equivalent to:

$$f(y) \le f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} ||x - y||^2$$
(3.3)

for all $x, y \in C$.

As such, it is also a somewhat dual property to strong convexity, in that it provides a quadratic upper bound on the function.

Finally, we should mention the definition of concave functions:

Definition 3.17 (Concavity). A function $f: C \to \mathbb{R}$ is called concave if -f is convex.

Obviously, all concepts and results of convex analysis apply for concave functions too, with trivial adjustments (e.g. reversing the direction of an inequality), and all subsequent results of convex minimization also apply as-is for concave maximization. All one has to do is reverse the sign and, thus, turn a payoff function to a cost function.

Gradient Descent A simple and rather archetypal (unconstrained) convex optimization algorithm is gradient descent (GD). At each step, we follow the direction of the steepest descent, which is, as is well known, the negative gradient. Specifically, the algorithm is:

$$x_{t+1} = x_t - \gamma_t \nabla f(x_t)$$

where $\gamma_t > 0$ is the step size.

Here, we shall provide the convergence analysis for when the function is L-smooth and the step size is constant, $\gamma_n = \gamma$.

From the equivalent formulation of smoothness (3.3), we get:

$$f(x_{t+1}) \le f(x_t) + \langle \nabla f(x_t), x_{t+1} - x_t \rangle + \frac{L}{2} ||x_{t+1} - x_t||_2^2$$
$$= f(x_t) - \gamma ||\nabla f(x_t)||_2^2 + \frac{L\gamma^2}{2} ||\nabla f(x_t)||_2^2$$

As a function of γ , the RHS of the above is minimized for step size $\gamma = \frac{1}{L}$, for which we get:

$$f(x_{t+1}) \le f(x_t) - \frac{1}{2L} \|\nabla f(x_t)\|_2^2 \Rightarrow R_{t+1} \le R_t - \frac{1}{2L} \|\nabla f(x_t)\|_2^2$$

where $R_t = f(x_t) - f(x^*)$ denotes the distance of the current value to the optimal value. Then, since f is convex, we get:

$$R_t = f(x_t) - f(x^*) \le \langle \nabla f(x_t), x_t - x^* \rangle$$

Combining the last two inequalities, we have:

$$\begin{split} R_{t+1} & \leq \langle \nabla f(x_t), x_t - x^* \rangle - \frac{1}{2L} \| \nabla f(x_t) \|_2^2 \\ & = -\frac{1}{2L} (\| \nabla f(x_t) \|_2^2 - 2L \langle \nabla f(x_t), x_t - x^* \rangle) \text{ now, we can complete the square} \\ & = -\frac{1}{2L} (\| \nabla f(x_t) \|_2^2 - 2L \langle \nabla f(x_t), x_t - x^* \rangle + L^2 \| x_t - x^* \|_2^2 - L^2 \| x_t - x^* \|_2^2) \\ & = -\frac{1}{2L} (\| \nabla f(x_t) - L(x_t - x^*) \|_2^2 - L^2 \| x_t - x^* \|_2^2) \\ & = -\frac{L^2}{2L} (\| \frac{1}{L} \nabla f(x_t) - x_t + x^* \|_2^2 - \| x_t - x^* \|_2^2) \\ & = -\frac{L}{2} (\| -x_{t+1} + x^* \|_2^2 - \| x_t - x^* \|_2^2) = \frac{L}{2} (\| x_t - x^* \|_2^2 - \| x_{t+1} - x^* \|_2^2) \end{split}$$

Thus, if we telescope the above:

$$\sum_{t=1}^{T} R_t \le \frac{L}{2} (\|x_0 - x^*\|_2^2 - \|x_T - x^*\|_2^2) \le \frac{L}{2} \|x_0 - x^*\|_2^2$$

Finally, since R_t is decreasing, we conclude that:

$$R_T \le \frac{1}{T} \sum_{t=1}^{T} R_t \le \frac{L \|x_0 - x^*\|_2^2}{2T} \le \frac{LD^2}{2T}$$

where D can be interpreted as either an upper bound on $||x_0 - x^*||_2$, which must be a starting condition for the algorithm, or as an upper bound for the diameter of the feasible set. Of course, in the latter case, gradient descend must be adapted to work for constrained optimization. This is not difficult; for a demonstration of something along these lines, check out the following section.

Here, the final result is that gradient descent reaches a point x_T with $R_T \leq \varepsilon$ in at most $O(\frac{LD^2}{\varepsilon})$ steps.

3.4.1 Online Convex Optimization

In general, convex optimization by itself is enough only for the case of potential games, where there is an underlying potential function and, whenever a player moves to improve his payoff, the potential increases; through such arguments, convergence to Nash equilibria by, e.g., gradient descent employed by all players can be shown.

However, for non potential games, convex optimization is no longer enough. Since we have no underlying single function, here we must say that each player, whenever she plays, sees a different payoff function, specifically $u_i(x_i; x_{-i})$, and what she wishes is to maximize her cumulative gains along the whole play of the game, if the game will be played repeatedly.

This, as seen from the perspective of one player, is basically the setting of *online* convex optimization. As a result, online convex optimization methods and algorithms are generally regarded to be "natural" tactics, which an intelligent player of a game is highly likely to employ; convergence results of such algorithms to equilibria are highly desirable, and, as a matter of fact, it is one such algorithm that is employed in [Mert19] to compute equilibria in diagonally strictly concave games, and which we will discuss in a later chapter.

In this section, we will briefly introduce fundamental concepts and an algorithm from online convex optimization, both to establish useful concepts (basically, the regret) and to get a first feeling of the analysis of online convex optimization algorithms.

First of all, let us describe exactly the general setting, which is, as the name suggests, a generalization of convex optimization to the online case:

Definition 3.18 (Online Convex Optimization). *An online convex optimization problem is* an iterative process conducted between some agent and some (possibly adversarial) environment, in which, at each step $t \in \mathbb{N}$:

- 1. The agent chooses an action $x_t \in \mathcal{X}$, where \mathcal{X} is a compact convex set, subset of some Euclidean space.
- 2. The environment selects a (convex) cost function $f_t: \mathcal{X} \to \mathbb{R}$.
- 3. The agent incurs a loss of $f_t(x_t)$.
- 4. The agent receives any feedback that is to be received.

One can easily understand that the above process is somewhat unfair, and that it is too much to ask for guarantees comparing, for example, the effectiveness of the agent to the optimal offline solution. For this reason, the merit of success here is taken to be the difference between the cost that the agent manages to incur and the minimum cost that she could have incurred by playing *the same* action on all rounds. This is called the agent's (external) regret, and it is defined below.

Definition 3.19 (Regret). In an online convex optimization problem, the agent's regret is defined as:

$$R(T) = \max_{x \in \mathcal{X}} \sum_{t=1}^{T} f_t(x_t) - f_t(x)$$

where T is the horizon of play.

An algorithm is then called no-regret if it guarantees vanishing average regret for the agent, or, equivalently, R(T) = o(T).

Online Gradient Descent One example of an algorithm that achieves no-regret (under suitable conditions) is the online version of gradient descent, appropriately called *online gradient descent (OGD)*. The algorithm is the following:

$$x_{t+1} = x_t - \gamma_t \nabla f_t(x_t)$$

We will also need the projected version of OGD, as we have to work with a bounded action set \mathcal{X} . The only difference is that each new point given by the above recursion is projected onto the action set:

$$x_{t+1} = \Pi_{\mathcal{X}}(x_t - \gamma_t \nabla f_t(x_t))$$

where $\Pi_{\mathcal{X}}(x) = \operatorname{argmin}_{x' \in \mathcal{X}} \|x - x'\|_2$ is the projection mapping.

Theorem 3.20 (OGD No Regret). Let \mathcal{X} (subset of some Euclidean space) be a compact convex set with diameter D, i.e. $\max_{x,y\in\mathcal{X}} \|x-y\|_2 = D$. Also, let $(f_t)_{t\in\mathbb{N}}$ be a sequence of differentiable convex functions. Then, (projected) OGD enjoys the following regret bound:

$$R(T) \le \frac{D^2}{2\gamma_T} + \sum_{t=1}^T \frac{\gamma_t}{2} \|\nabla f_t(x_t)\|_2^2$$

If, moreover, all f_t are assumed to be L-Lipschitz and OGD is run with step size $\gamma_t = \frac{D}{L\sqrt{t}}$, then:

$$R(T) \leq \frac{3}{2} L D \sqrt{T}$$

Proof. Firstly, we use the convexity of f_t to bound the difference $f_t(x_t) - f_t(x)$ (for any $x \in \mathcal{X}$):

$$f_t(x) \ge f_t(x_t) + \langle \nabla f_t(x_t), x - x_t \rangle$$

$$\Rightarrow f_t(x_t) - f_t(x) \le \langle \nabla f_t(x_t), x_t - x \rangle$$

Then, we will bound the quantity $\langle \nabla f_t(x_t), x_t - x \rangle$ using the update rule.

$$||x_{t+1} - x||_2^2 = ||\Pi_{\mathcal{X}}(x_t - \gamma_t \nabla f_t(x_t)) - x||$$

Now, it is a well known fact and not hard to see geometrically that $\|\Pi_{\mathcal{X}}(y) - x\|_2 \le \|y - x\|_2$, for any $x \in \mathcal{X}$. So, we can continue:

$$||x_{t+1} - x||_{2}^{2} \leq ||x_{t} - \gamma_{t} \nabla f_{t}(x_{t}) - x||_{2}^{2}$$

$$= ||x_{t} - x||_{2}^{2} + \gamma_{t}^{2} ||\nabla f_{t}(x_{t})||_{2}^{2} - 2\gamma_{t} \langle \nabla f_{t}(x_{t}), x_{t} - x \rangle$$

$$\Rightarrow 2\langle \nabla f_{t}(x_{t}), x_{t} - x \rangle \leq \frac{1}{\gamma_{t}} (||x_{t} - x||_{2}^{2} - ||x_{t+1} - x||_{2}^{2} + \gamma_{t}^{2} ||\nabla f_{t}(x_{t})||_{2}^{2})$$

Now, we are ready to bound the regret (we use $R_x(T)$ to denote the regret with respect to the specific fixed action $x \in \mathcal{X}$):

$$\begin{split} 2R_x(T) &= 2\sum_{t=1}^T (f_t(x_t) - f_t(x)) \leq \sum_{t=1}^T 2\langle \nabla f_t(x_t), x_t - x \rangle \\ &\leq \sum_{t=1}^T \frac{1}{\gamma_t} (\|x_t - x\|_2^2 - \|x_{t+1} - x\|_2^2 + \gamma_t^2 \|\nabla f_t(x_t)\|_2^2) \\ &= \sum_{t=1}^T \frac{\|x_t - x\|_2^2 - \|x_{t+1} - x\|_2^2}{\gamma_t} + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \\ &= \frac{\|x_1 - x\|_2^2}{\gamma_1} + \sum_{t=2}^T \|x_t - x\|_2^2 \left(\frac{1}{\gamma_t} - \frac{1}{\gamma_{t-1}}\right) - \frac{1}{\gamma_T} \|x_{T+1} - x\|_2^2 + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \\ &\leq \frac{\|x_1 - x\|_2^2}{\gamma_1} + \sum_{t=2}^T \|x_t - x\|_2^2 \left(\frac{1}{\gamma_t} - \frac{1}{\gamma_{t-1}}\right) + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \\ &\leq \frac{D^2}{\gamma_1} + D^2 \sum_{t=2}^T \left(\frac{1}{\gamma_t} - \frac{1}{\gamma_{t-1}}\right) + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \\ &= \frac{D^2}{\gamma_1} + D^2 \left(-\frac{1}{\gamma_1} + \frac{1}{\gamma_T}\right) + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \\ &= \frac{D^2}{\gamma_T} + \sum_{t=1}^T \gamma_t \|\nabla f_t(x_t)\|_2^2 \end{split}$$

So, finally:

$$R(T) = \max_{x \in \mathcal{X}} R_x(T) \le \frac{D^2}{2\gamma_T} + \sum_{t=1}^{T} \frac{\gamma_t}{2} \|\nabla f_t(x_t)\|_2^2$$

The second part of the theorem is derived easily from the first, if we substitute $\gamma_t = \frac{D}{L\sqrt{t}}$ and we use:

- 1. that a convex differentiable function is L-Lipschitz if and only if $\|\nabla f(x)\| \leq L$.
- 2. that $\sum_{t=1}^{T} \frac{1}{\sqrt{t}} \le 2\sqrt{T}$.

3.4.2 Convex Optimization and Games

In this section, we utilize the example of *potential games* to showcase the strong connection that exists, in some cases, between optimization and games (specifically, equilibrium computation).

Definition 3.21 (Potential Game). A game $\mathcal{G}(N, (\mathcal{X}_i)_{i \in [N]}, (u_i)_{i \in [N]})$ is called an (exact) potential game if there exists some function $\Phi : \mathcal{X} \to \mathbb{R}$ such that:

$$u_i(y_i, x_{-i}) - u_i(z_i, x_{-i}) = \Phi(y_i, x_{-i}) - \Phi(z_i, x_{-i})$$

For smooth games, an equivalent, more straightforward formulation is for the function Φ to hold that:

$$\nabla \Phi(x) = v(x)$$

i.e. the gradient of Φ must be equal to the pseudogradient of the game.

Given this, it is easy to see that the Nash equilibria of the game correspond to the local maxima of the potential function. If, specifically, the game is also concave, then the equilibria and the local (which are also global) maxima coincide.

Thus, in a concave potential game, any concave optimization algorithm can be employed, using the gradient of the potential or, equivalently, the pseudogradient of the game.

An idea, then, is what happens if we use the pseudogradient for such an algorithm / dynamic, but for general, and not only potential games. One possible answer are the results of [Mert19], which we review later in this work.

3.5 Socially Concave Games

[Even09] discusses another interesting subclass of concave games, with respect to equilibrium computation. The definition is the following:

Definition 3.22 (Socially Concave Games). *A game is called socially concave if the following two conditions hold:*

- 1. there is some strict convex combination of the payoff functions which is concave. Formally, there are $\lambda_i > 0$ with $\sum_i \lambda_i = 1$, such that $g(x) = \sum_{i=1}^N \lambda_i u_i(x)$ is a concave function in x.
- 2. for every $i \in [N]$ and $a_i \in \mathcal{X}_i$ (fixed), $u_i(a_i, x_{-i})$ is convex in x_{-i} .

It is not difficult to show that a socially concave game is always a concave game.

The main result here is that, for socially concave games, not only are equilibria computable, but they can be approached in average whenever all players employ some no-regret algorithm (not necessarily the same).

This is a considerable improvement compared to the aforementioned classes, as no-regret algorithms are arguably simple, natural and generally viable to be employed by real, incentive-driven players. Not only this, but the result does not constrain the players to use some specific no-regret algorithm. On the contrary, they can use whatever algorithm / strategy they want, as long as it is no-regret.

To state exactly this result of [Even09], we first define an approximate notion of equilibrium:

Definition 3.23. An $x_i \in \mathcal{X}_i$ is an ε -best response to x_{-i} if player i can score no more than ε better against x_{-i} than what she earns by playing x_i , formally $u_i(x_i, x_{-i}) \geq u_i(y_i, x_{-i}) - \varepsilon$ for all $y_i \in \mathcal{X}_i$.

Then, an ε -equilibrium is simply any action profile where each player i is ε -best responding to x_{-i} .

Theorem 3.24. Suppose we have an N player concave game, which is also socially concave. If each player acts in a way that guarantees a regret bound $R_i(T)$, then at any time step t > 0:

1. The average strategy vector $\bar{x}_t = \frac{1}{t} \sum_{k=1}^t x_k$ is an ε_t -Nash equilibrium, where $\varepsilon_t = \frac{1}{\lambda_{min}} \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t}$ and $\lambda_{min} = \min_{i \in [N]} \lambda_i$.

2. The average payoff $\bar{u}_{i,t} = \frac{1}{t} \sum_{k=1}^{t} u_i(x_k)$ of each player is close to her payoff at \bar{x}_t , in the sense that:

$$|\bar{u}_{i,t} - u_i(\bar{x}_t)| \le \frac{1}{\lambda_i} \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t}$$

Proof. The proof basically comes down to properly rewriting the definitions for regret, convexity of u_i in x_{-i} and concavity of $g(x) = \sum_{i=1}^{N} \lambda_i u_i(x)$ in x, applied on the points x_t .

1. $R_i(t)$ is an upper bound on the regret of player i, thus, for every $x_i \in \mathcal{X}_i$:

$$R_{i}(t) \ge \sum_{k=1}^{t} \left[u_{i}(x_{i}, x_{-i,k}) - u_{i}(x_{i,k}, x_{-i,k}) \right] \Rightarrow \bar{u}_{i,t} \ge \frac{1}{t} \sum_{k=1}^{t} u_{i}(x_{i}, x_{-i,k}) - \frac{1}{t} R_{i}(t)$$
(3.4)

2. $u_i(x_i, x_{-i})$ is convex in x_{-i} , thus, for any $x_i \in \mathcal{X}_i$:

$$u_i(x_i, \frac{1}{t} \sum_{k=1}^t x_{-i,k}) \le \frac{1}{t} \sum_{k=1}^t u_i(x_i, x_{-i,k})$$
(3.5)

3. from the concavity of g(x):

$$g(\bar{x}_t) \ge \frac{1}{t} \sum_{k=1}^t g(x_k) \Rightarrow \sum_{i=1}^N \lambda_i u_i(\bar{x}_t) \ge \sum_{i=1}^N \lambda_i \bar{u}_{i,t}$$
 (3.6)

By combining (3.4) and (3.5), we get:

$$\bar{u}_{i,t} \ge u_i(x_i, \frac{1}{t} \sum_{k=1}^t x_{-i,k}) - \frac{1}{t} R_i(t) = u_i(x_i, \bar{x}_{-i,t}) - \frac{1}{t} R_i(t)$$

Finally, combining the above with (3.6), we get:

$$\sum_{i=1}^{N} \lambda_{i} u_{i}(\bar{x}_{t}) \ge \sum_{i=1}^{N} \lambda_{i} u_{i}(x_{i}, \bar{x}_{-i,t}) - \sum_{i=1}^{N} \frac{\lambda_{i} R_{i}(t)}{t}$$

Now, almost all results can be derived from this inequality. Indeed, since x_i can be freely chosen, we can choose $x_j = \bar{x}_{i,t}$ for all $j \neq i$ and keep x_i free. Then, we get:

$$\begin{split} \lambda_i u_i(\bar{x}_t) + \sum_{j=1, j \neq i}^N \lambda_j u_j(\bar{x}_t) &\geq \lambda_i u_i(x_i, \bar{x}_{-i,t}) + \sum_{j=1, j \neq i}^N \lambda_j u_i(\bar{x}_t) - \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t} \\ \Rightarrow & \lambda_i u_i(\bar{x}_t) \geq \lambda_i u_i(x_i, \bar{x}_{-i,t}) - \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t} \\ \Rightarrow & u_i(\bar{x}_t) \geq u_i(x_i, \bar{x}_{-i,t}) - \frac{1}{\lambda_i} \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t} \geq u_i(\bar{x}_t) \geq u_i(x_i, \bar{x}_{-i,t}) - \frac{1}{\lambda_{min}} \sum_{i=1}^N \frac{\lambda_i R_i(t)}{t} \end{split}$$

In other words, \bar{x}_t is an ε_t -Nash equilibrium.

For the proof of the second part, which is equally straightforward, we defer the reader to [Even09].

Chapter 4

Opinion Formation

The topic of opinion formation is relatively large, with much research having been conducted through the decades, in fields ranging from mathematical sociology to game theory. By no means do we intend here to summarize the whole topic, but we give some brief overviews of models and approaches that are relevant to our setting, and which our model stems from.

Network representation In general, the models we deal with represent the interacting, opinion forming entities as nodes in a (social) network, using edges between them to represent some kind of bond or interaction (e.g. Facebook friends). We generally denote this network as a graph G(V, E) (following standard graph notation) or G(V, E, w) for the case of weighted graphs. We generally deal with directed graphs, unless explicitly specified otherwise.

The agents' opinions are represented as single numbers on the real line, usually assumed to lie in [0, 1].

In addition to the above, each model specifies some opinion dynamics, i.e. update rules which dictate how each agent's opinion evolves over time, while interacting with the other agents. There are many variations of those and, additionally, it is not uncommon for a network or game model to specify and study several dynamics. This is not unusual; these should be thought of as two perpendicular, albeit related, aspects of every model. On one hand, we have the specifications for how agents are stationed in the world, what is their opinion, what interests them and holds the most weight for them. On the other hand, we have the specifications for how players think, decide and take action. We can think of it as a function to be minimized and the optimization method that is used to minimize it (which is indeed the case on many occasions, as we shall see).

4.1 DeGroot model

First specified in [DeGr74], this model posits that, at each round, each agent updates their opinion to some (fixed) convex combination of the others' opinions. Formally:

$$x_{t+1} = Ax_t = A^t x_1 (4.1)$$

with A a stochastic matrix, i.e. it has nonnegative entries and all its rows sum to 1.

In [DeGr74], it is shown that the agents' opinions converge to a consensus (i.e. all agents to the same opinion), under some mild assumptions on the Markov chain corresponding to the stochastic transition probability matrix A (specifically, irreducibility and aperiodicity).

Furthermore, it is shown in [Ghad13] that the time it takes for the agents to reach ε close to the equilibrium is of order $\Theta(\frac{1}{1-\rho_2}\log\frac{1}{\varepsilon})$, where ρ_2 is the second largest eigenvalue modulus of A (the largest is 1, since A is a stochastic matrix).

It is interesting to notice that in [Ghad13], which otherwise does not refer to optimization whatsoever, the way they resolve one case which does not converge under the simple best response dynamics (specifically, bipartite graphs) is to utilize a variation of best response in which, at each turn, each player chooses a point somewhere between their best response and their previous point. They call it "noisy best response dynamics", but it is not much different than what a GD-like optimization algorithm would do.

4.2 Friedkin-Johnsen model

First studied in [Frie90], in this model, there are still no games, but only N agents with opinions who, again, at each step of an iterative process, update their opinion to some weighted sum of all opinions (with fixed weights). The difference now is that an extra term is added, that can be interpreted as their intrinsic, unchanging opinions ([Frie90] interprets it as the effect of exogenous variables on the opinions, but, as long as it's static, there is no substantial difference). Formally:

$$x_{t+1} = Ax_t + Bs \tag{4.2}$$

where A is some matrix of interaction weights.

It is assumed that each agent updates their opinion to a weighted average of all agents' opinions from the previous round and their own intrinsic opinion.

Thus, the matrix of interaction weights A is substochastic and the matrix B is a diagonal matrix, whose entries b_{ii} are linked to the agents' "self-confidence".

Moreover, assuming there is at least one agent for whom $b_{ii} > 0$ (otherwise, we are back in the DeGroot model), A has at least one row with sum strictly less than 1. Additionally, A is assumed to be irreducible, which is equivalent to the graph of the agents being connected (otherwise, one can work on each connected component independently).

Under these assumptions, the spectral radius of A, $\rho(A) = \max_i |\lambda_i(A)|$, is guaranteed to be less than 1. But then, we know that A^k behaves like a geometric progression with rate less than 1, in the sense that:

$$\lim_{k \to \infty} A^k = 0$$

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

Moreover, from (4.2) we recursively get that:

$$x_{t+1} = A^t x_1 + (\sum_{k=0}^t A^k) Bs$$

As such, and as is shown also in [Frie90] (under some different but related conditions), the opinion vector converges to:

$$x_{\infty} = (I - A)^{-1} Bs$$

As for the convergence speed, [Ghad13] show that a similar result holds as in the DeGroot model. Specifically, the time until an ε -neighbourhood of the limit is reached is bounded by $\Theta(\frac{1}{1-\rho(A)}\log(\frac{1}{\varepsilon}))$, where $\rho(A)=\max_i \lambda_i(A)$ is the spectral radius of A.

4.3 Game Theoretic viewpoint

In [Bind11] (and [Ghad13], as a matter of fact), the above regime is reinterpreted in terms of game theory. Specifically, the averaging by a player of other players' opinions and their own intrinsic opinion is seen as the player's best response in a game with suitable quadratic costs.

Specifically, each player i is assigned a cost function of the form:

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

where $w_{i,j}$ is the weight of the edge (i, j) (in other words, the strength of their interaction) and N(i) is the set of all direct neighbours of i in G(V, E).

The first term of this sum expresses the cost of disagreement with the player's own intrinsic opinion, while the second term represents the cost of disagreement with their "friends".

Now, given the above, it is easy to find the best-response correspondence of a player, which, in this case, is a simple mapping, since the cost is a simple quadratic function and the best response is unique. All we need to do is differentiate:

$$\frac{\partial c_i}{\partial x_i}(x_i^*) = 2(x_i^* - s_i) + 2\sum_{j \in N(i)} w_{i,j}(x_i^* - x_j) = 0$$

$$\Rightarrow x_i^* = BR_i(x) = \frac{s_i + \sum_{j \in N(i)} w_{i,j} x_j}{1 + \sum_{j \in N(i)} w_{i,j}}$$

This quantity is a convex combination of s_i , x_j (for all $j \neq i$) and, as such, it always lies within [0,1]. So, it is indeed always the best response of player i (i.e. we do not have to consider boundary cases; the quadratic function is constructed in such way that its minimum point over \mathbb{R} is always in [0,1]).

In other words, a necessary and sufficient condition for x^* to be (the) Nash equilibrium in these games is $v(x^*) = 0$. Now, if we try to iterate on that a little, the pseudo-gradient of these games is:

$$v(x) = 2Lx + 2x - 2s = 2(L+I)x - 2s$$

where L is the Laplacian matrix of the weighted graph, i.e. $L_{ii} = \sum_{j \in N(i)} w_{ij}$ and $L_{ij} = -w_{ij}$.

This is simply the partial derivative we wrote above, but in matrix form.

Then, the condition $v(x^*) = 0$ becomes:

$$v(x^*) = 0 \Leftrightarrow 2Lx^* + 2x^* - 2s = 0 \Leftrightarrow (L+I)x^* = s \Leftrightarrow x^* = (L+I)^{-1}s$$

where L + I is invertible, since it is strictly diagonally dominant (easily checked).

4.3.1 Best Response Dynamics

Given the above expression for the best response of a player, the best response dynamics for this game is:

$$x_{i,t+1} = BR_i(x_t) = \frac{s_i + \sum_{j \in N(i)} w_{i,j} x_{j,t}}{1 + \sum_{j \in N(i)} w_{i,j}}$$

4.3.2 **Opinion Dynamics**

It is not difficult to see that the best response dynamics described above fall into the opinion dynamics model of [Frie90] that we described, with:

$$A_{ij} = \frac{w_{ij}}{1 + \sum_{j \in N(i)} w_{ij}}$$
$$B_{ii} = \frac{s_i}{1 + \sum_{j \in N(i)} w_{ij}}$$

As a result, all that we said previously still holds: the dynamics converge to $(I-A)^{-1}Bs$ and the convergence speed is bounded by the same quantity.

Moreover, since they coincide with the best response dynamics of the game, this unique limit must also be the unique Nash equilibrium of the game (i.e. $(L+I)^{-1}s = (I-A)^{-1}Bs$).

4.3.3 **Potential**

For the case of undirected graphs, the above game is additionally a potential game, with the potential function:

$$\phi(x) = \sum_{i=1}^{N} (x_i - s_i)^2 + x^T L x$$
$$\nabla \phi(x) = 2Lx + 2x - 2s = v(x)$$
$$\frac{\partial^2 \phi}{\partial x^2}(x) = 2(L+I)$$

Since the graph is undirected, the Laplacian is symmetric. Moreover, L can be easily seen to be positive semidefinite, as $x^T L x = \sum_{\{i,j\} \in E} w_{ij} (x_i - x_j)^2 \ge 0$. As a result, ϕ has a positive definite Hessian. Thus, it is strictly convex; and strongly

convex, for that matter, since $L \geq 0 \Rightarrow L + I \geq I$.

As a result, methods of convex optimization can be applied to obtain its minimum, or equivalently, the Nash equilibrium of the game, in polynomial time.

But, if we write explicitly, let's say, gradient descent for this problem, we get:

$$x_{t+1} = x_t - \gamma_t \nabla \phi(x_t) = x_t - \gamma_t 2((L+I)x_t - s)$$

$$\Rightarrow x_{i,t+1} = x_{i,t} - 2\gamma_t ((1 + \sum_j w_{ij})x_{i,t} - \sum_{j \in N(i)} w_{ij}x_j - s_i)$$

$$= x_{i,t} - 2\gamma_t (1 + \sum_j w_{ij})(x_{i,t} - \frac{s_i + \sum_{j \in N(i)} w_{i,j}x_j}{1 + \sum_{j \in N(i)} w_{i,j}})$$

$$= x_{i,t} - 2\gamma_t (1 + \sum_j w_{ij})(x_{i,t} - BR_i(x_t))$$

In other words, what we end up with again resembles a lot the "noisy" best response dynamics of [Ghad13], in the sense that a player does not go straight to their best response at each step, but to an opinion between their best response and their previous opinion.

Furthermore, since all other stuff are constant, we can renormalize the step and say that:

$$\gamma_t \leftarrow 2\gamma_t (1 + \sum_j w_{ij})$$

Then, $x_{i,t+1} = x_{i,t} - \gamma_t(x_{i,t} - BR_i(x_t))$ and one can easily see that the best response dynamics we already talked about coincide with gradient descent if we take a constant step of 1.

4.3.4 Implications

What all the above tries to convey is that, for this class of games, there is a strong connection between the game, the opinion dynamics and the optimization methods that can be performed on them. Moreover, each perspective can be used in different ways to arrive at certain conclusions. An example here is that simultaneous best response dynamics converge, which, in general, does not hold, but, in this case, it can be shown both via the opinion dynamics and the potential function with gradient descent, because in both cases, the algorithms end up being identical.

Another example, which also portrays some advantages of optimization methods, is the case we discussed in DeGroot opinion dynamics, where simple best response dynamics do not converge, but a slight modification to a more GD-like dynamic results in convergence.

This warms up the field and points us to the direction of using optimization methods for equilibrium computation in not only those games, but maybe in generalizations of them.

4.4 Coevolution

The final consideration, and the one encompassing our model, is to consider that the network is not static, but it co-evolves with the opinions. For every player, the players whom she values the most, and whose opinions she will weigh more in her cost of disagreement, are those whose opinions are closer to her, and to her intrinsic opinion in particular.

The first two models are classic in this line of work. The rest are taken (either verbatim or slightly adjusted) from [Bhaw13].

4.4.1 Bounded Confidence - HK and DW

Hegselmann Krause One of the first approaches for this line of work was the one in [Hegs02]. In this opinion formation model, agents do not possess any hidden information (such as the s_i before), but each agent only "listens" to those other agents whose opinions are close enough to hers.

Specifically, there is a threshold parameter $\varepsilon > 0$ (which represents how "open-minded" the agents are) and each agent takes into account only the agents in the set:

$$I(i,x) = \{1 \le j \le N : |x_i - x_j| \le \varepsilon\}$$

in the sense of the following update rule:

$$x_{i,t+1} = \frac{\sum_{j \in I(i,x)} x_j(t)}{|I(i,x(t))|}$$

Obviously, the agent herself is always in I(i, x).

Moreover, notice that the HK model dynamic admits an infinite number of stable points, namely any partition of the agents into groups, such that the opinions of agents from different groups are more than ε apart.

Finally, concerning convergence properties of the model, it has been shown that the HK model dynamic converges to a stable opinion profile as described above in a finite number of steps [Hend06, Lore05, More05]. Moreover, there are both upper [Tour11, Bhat12] and lower [Wedi14] bounds on the number of rounds needed for this to happen.

Deffuant Weisbuch Another model which is on the same wavelength is the Deffuant Weisbuch (DW) model [Weis02, Deff00, Weis03]. Here, at each round t, two agents i, j are chosen at random to interact. If their opinions are closer than some threshold d, then they both update their opinions to a middle point, using the same coefficients:

$$x_{i,t+1} = x_{i,t} + \mu(x_{j,t} - x_{i,t})$$

$$x_{j,t+1} = x_{j,t} + \mu(x_{i,t} - x_{j,t})$$

where μ is a convergence parameter in [0, 1/2].

This model seems to be more difficult, with results being more scarce, but still, it has been shown that this opinion dynamic also converges to a stable point ([Lore05]) and relatively fast at that ([Zhan15]).

4.4.2 Asymmetric KNN Games

Going again into game theoretic models, here, each player again possesses an intrinsic opinion $s_i \in [0, 1]$. Then, she is connected, for any given profile $x \in \mathcal{X}$, only to the K other players whose expressed opinions are closest to her intrinsic opinion s_i (with ties broken arbitrarily but consistently). Formally:

$$N(i) \subseteq [N] \setminus \{i\}, |N(i)| = K, \text{ s.t. } \forall j \in N(i), \forall j' \notin N(i) \cup \{i\}, |x_j - s_i| \le |x_{j'} - s_i|$$

$$c_i(x) = \sum_{j \in N(i)} (x_i - x_j)^2 + \rho K(x_i - s_i)^2$$

where ρ is some number which represents the self-confidence of the player.

These games are, in general, difficult to work with. They are not concave and, as a matter of fact, there are very simple examples of such games in which Nash equilibria do not even exist.

Such an example is given in [Bhaw13]. Let N=3 players, with $s_1=0, s_2=0.5, s_3=1$ and K=1 and $\rho=1$. Then, player $i,i\in\{1,2,3\}$ has the cost function:

$$c_i(x) = (x_i - x_{\sigma(i)})^2 + (x_i - s_i)^2$$

where $\sigma(i) = \arg\min_{j \neq i} |x_j - s_i|$, and the best response of player i is:

$$x_i^* = \frac{x_{\sigma(i)} + s_i}{2}$$

Given this, it is relatively easy to see that a Nash equilibrium x^* must have $x_1^* \le x_2^* \le x_3^*$, since:

$$\begin{split} x_1^* &= \frac{\min\{x_2, x_3\} + 0}{2} = \frac{\min\{x_2, x_3\}}{2} \leq \min\{x_2, x_3\} \\ x_3^* &= \frac{\max\{x_1, x_2\} + 1}{2} \geq \max\{x_1, x_2\} \text{ (verifiable with simple calculations)} \end{split}$$

So, then,
$$x_1^* = \frac{x_2^* + s_1}{2} = \frac{x_2^*}{2}$$
 and $x_3^* = \frac{x_2^* + s_3}{2} = \frac{x_2^* + 1}{2}$. Finally, there are three cases for x_2^* :

- 1. $x_2^* \in [0, 0.5)$: then, $x_1^* \in [0, 0.25)$ and $x_3^* \in [0.5, 0.75)$, so $\sigma(2) = 3$ and x_2^* should be $x_2^* = \frac{x_3^* + 0.5}{2} \in [0.5, x_3^*]$, which is a contradiction.
- 2. $x_2^* \in (0.5, 1]$: completely symmetric, $x_1^* \in (0.25, 0.5], x_3^* \in (0.75, 1]$, so $\sigma(2) = 1$ and x_2^* should be $x_2^* = \frac{x_1^* + 0.5}{2} \in [x_1^*, 0.5]$, which is a contradiction.
- 3. $x_2^*=0.5$, in which case $x_1^*=0.25$ and $x_3^*=0.75$, thus 1 and 3 are tied for $\sigma(2)$. But, no matter which one we choose, the best response of player 2 cannot be 0.5, as it should either be $\frac{0.25+0.5}{2}=0.375$ or $\frac{0.5+0.75}{2}=0.625$.

So, this game does not have any Nash equilibria, and there isn't much hope in terms of equilibrium computation, at least in general.

4.4.3 Asymmetric Coevolutionary Opinion Formation Games

Finally, this is the model that we are mainly concerned with. First of all, we explicitly constrain the individual action space of each player to $\mathcal{X}_i = [0, 1]$. Further, each player is assigned a cost function of the form:

$$c_{i}(x_{i}, x_{-i}) = (1 - \rho_{i}) \sum_{j \neq i} q_{ij}(x_{-i}; s_{i})(x_{i} - x_{j})^{2} + \rho_{i}(x_{i} - s_{i})^{2}$$

$$= x_{i}^{2} - 2\beta(x_{-i}; s_{i})x_{i} + \gamma(x_{-i}; s_{i}), \text{ with } \begin{cases} \beta(x_{-i}; s_{i}) = (1 - \rho_{i}) \sum_{j \neq i} x_{j}q_{ij}(x_{-i}; s_{i}) + \rho_{i}s_{i} \\ \gamma(x_{-i}; s_{i}) = (1 - \rho_{i}) \sum_{j \neq i} x_{j}^{2}q_{ij}(x_{-i}; s_{i}) + \rho_{i}s_{i}^{2} \end{cases}$$

where

- $q_{ij}(x_{-i}; s_i) = F_i(d_j^i, d_{-i-j}^i)$, with $d_i^j = |x_j s_i|$
- F_i is a continuous function, nonincreasing in the first argument and nondecreasing in all the rest.
- $q_{ij}(x_{-i}; s_i) \ge 0$, $\sum_{j \ne i} q_{ij}(x_{-i}; s_i) = 1$, hence the equivalent form where c_i is a quadratic with unit coefficient in x_i .

Here, q_{ij} tries to capture the role of the (normalized) edge weights that we had before, with the fundamental difference that these weights change dynamically depending on the players' opinions vector.

The second condition stems exactly from that. It basically says that if player j goes away from player i (and all other players stay in the same spot), then the weight of the edge (i, j) is expected to decrease (or at least not increase).

Nash Equilibrium Existence It is easy to see that the above game is concave, since each player's cost is a continuous function and convex in the player's action. As a result, the existence of at least one Nash equilibrium is guaranteed by [Rose65].

Consequently, as noted in [Bhaw13] (section 4.3), it is an interesting open question to decide the computational complexity of Nash equilibria in this case, and this is exactly the ultimate goal of the current line of work.

To that end, in the following chapter we are going to discuss the Dual Averaging dynamics, which is a known algorithm in convex optimization, and which has been shown in [Mert19] to converge fast to equilibria in a large subclass of concave games. And, basically, as we will see, all we do is try to make the results of [Mert19] applicable on asymmetric coevolutionary games.

Properties We summarize here some properties of the above game that may be of use later.

Individual gradient / derivative

$$\frac{\partial}{\partial x_i} c_i(x_i, x_{-i}) = 2(1 - \rho_i) \sum_{j \neq i} (x_i - x_j) q_{ij}(x_{-i}; s_i) + 2\rho_i(x_i - s_i)$$
$$= 2x_i - 2\beta(x_{-i}; s_i)$$

Best responses Since the payoff of player i is a simple quadratic with respect to x_i , we can calculate the best response to a strategy profile of the other players in a straightforward manner:

$$x_i^* = \operatorname*{argmin}_{y \in \mathcal{X}_i} c_i(y; x_{-i}) = \beta(x_{-i}; s_i) = (1 - \rho_i) \sum_{j \neq i} x_j q_{ij}(x_{-i}; s_i) + \rho_i s_i$$

and the minimum cost that is attained by choosing this strategy is:

$$c_i(x_i^*, x_{-i}) = -\beta(x_{-i}; s_i)^2 + \gamma(x_{-i}; s_i)$$

Chapter 5

Concave Games and Dual Averaging

In this chapter, we give an overview of the main results of [Mert19], that are also the most relevant to our work.

If we had to summarize the most relevant and interesting results, it would be that for (a superclass of) diagonally strictly concave games, as those were introduced in [Rose65], there is an algorithm (as a matter of fact, a somewhat general algorithmic scheme) which converges to the unique equilibrium in polynomial time.

5.1 Variational Stability

A notion used widely in [Mert19], one way to see it is as a generalization of unimodality.

Another way, if it seems more relevant to the reader, which is also how the author puts it, is that it plays the same role for learning in games with continuous action spaces as evolutionary stability (first introduced by [Smit73]) plays for evolution in games with a continuum of players.

Formally, variational stability is defined with respect to an action profile¹:

Definition 5.1. An action profile $x^* \in \mathcal{X}$ is called (globally²) variationally stable state if:

$$\langle v(x), x - x^* \rangle < 0$$
 for all $x \in \mathcal{X} \setminus \{x^*\}$ (5.1)

We will often call a variationally stable state simply stable.

Geometrically, stability means that the pseudogradient vector at any point "points toward" x^* , in the sense that v(x) has positive inner product with the vector starting from x and ending on x^* . Thus, intuitively, we generally expect maximization algorithms that follow the gradient to approach the equilibrium x^* .

Two noteworthy properties concerning variational stability are the following:

A variationally stable point is the game's unique Nash Equilibrium. We omit the formal proof, but it is straightforward: the definition, on one hand, implies that on any line from x_i to x_i*, u(·, x_{-i}*) is increasing, thus x_i* is a maximizer, and, on the other hand, (5.1) is an exact counterexample of the necessary condition (3.2) for x, so any x ≠ x* cannot be a Nash equilibrium.

¹ To encompass cases such as, for example, potential games whose potential is concave, but not strictly concave (there is a plateau), [Mert19] also gives a more general definition that applies to sets of profiles. We will not deal with stable sets in this work, as it did not seem to apply to our games.

² The original definition of [Mert19] is even more general, defining locally stable states, in the sense that the above inequality holds for x in some neighbourhood of x^* . We will not deal with that either.

Diagonal strict concavity implies variational stability, so that the (unique) Nash equilibrium of a diagonally strictly concave game is always variationally stable. Unsurprising, since, essentially, variational stability is to diagonal strict concavity what unimodality is to strict concavity.

Formally, it follows directly from the definition (5.1) plus condition (3.2):

$$\langle v(x) - v(x^*), x - x^* \rangle < 0 \Rightarrow \langle v(x), x - x^* \rangle < \langle v(x^*), x - x^* \rangle \le 0$$

5.2 Dual Averaging

Dual averaging is a fairly simple, GD like algorithm, widely known in online optimization literature (introduced in [Nest09]), adapted in [Mert19] to the game theoretic setting. It is a well-known no-regret algorithm under mild assumptions (see for example [Xiao09] and [Shal12]), and, even though it is not exactly in our agenda to talk about this, it is an important property, which guarantees a decent total payoff for the player and, thus, as we have mentioned, it contitutes a strategy very likely to be employed by intelligent players.

The algorithm can be described as follows. Each player keeps track of two points: their current action $x_{i,t} \in \mathcal{X}_i$, which is essentially the action they choose to play at each round of the game, and one point $y_{i,t}$ which moves freely in the ambient space³. Then, at each step, $y_{i,t+1}$ is found by taking a free step along the gradient $v_i(x_t)$ (with no regard as to whether we end up in the action set or not), while $x_{i,t}$ is found by projecting $y_{i,t}$ onto the action set.

In other words, it resembles projected GD, the difference being that $y_{i,t}$ are "persistent". We don't simply project an infeasible point onto the feasible set and keep the projection, but it is the infeasible point from which we will take the next gradient step, in the following round.

Formally, starting with some arbitrary y_1 , the general algorithmic framework can be described by the following update rule:

$$x_{i,t} = Q_i(y_{i,t}) y_{i,t+1} = y_{i,t} + \gamma_t v_i(x_t)$$
(DA)

where:

- Q_i is a generalized projection, called the *choice map* (we shall see the formal definition later).
- γ_t is a nonincreasing step sequence, typically of the form $1/t^{\beta}$ for some $\beta \in (0,1]$ (one typical choice is $\gamma_t = 1/\sqrt{t}$).
- $v_i(x_t)$ is the individual gradient, as we have already said. Furthermore, we denote $V_* = \max_{x \in \mathcal{X}} \|v(x)\|_*$. Note that this maximum always exists, since we have already imposed that \mathcal{X} be compact and v(x) be continuous.

5.2.1 Choice Map

Later in this work, we will mainly utilize the Euclidean projection, that is, the minimum distance projection of a point onto a (convex) set:

$$Q(y) = \operatorname*{argmin}_{x \in \mathcal{X}} \|y - x\|_2^2$$

³ More accurately, the dual space of the ambient vector space. This makes [Mert19] very general, but, for the purposes of this work, we decided that it was unnecessary algebraic bloating.

Nevertheless, the more generic choice map we will present shortly gives useful insight that is connected to concepts from convex analysis, that would perhaps not be as apparent if we focused only on a specific choice map.

To start thinking about the choice map, consider the points $y_{i,t}$. From (DA), it can be seen that $y_{i,t}$ aggregate gradient steps, in a similar way to a simple gradient descent. Ideally, we would like to be able to choose $y_{i,t}$ as the actual actions to be played; of course, y_t are not guaranteed to lie inside the action space, and it is the role of Q_i to project them back, as we already said. For this reason, we want Q_i to choose an action that is as closely aligned as possible with $y_{i,t}$. Thus, an obvious choice would be $Q_i(y_i) = \operatorname{argmax}_{x_i \in \mathcal{X}_i} \langle y_i, x_i \rangle$.

The problem with this choice is that it is too aggresive and it cannot possibly converge to interior points, because, generically, the maximizer will be on the boundary. As an example, let's try to apply the algorithm with this choice map to a simple 1-dimensional optimization problem (i.e. single "player") with:

$$u(x) = -x^2, x \in [-1, 1]$$

$$\Rightarrow u'(x) = -2x$$

Then, the algorithm will be:

$$\begin{aligned} x_t &= Q(y_t) = \operatorname*{argmax}_{x \in [-1,1]} y_t x = \begin{cases} 1 & y_t > 0 \\ -1 & y_t < 0 \\ \text{arbitrary} & y_t = 0 \end{cases} \\ y_{t+1} &= y_t - 2\gamma_t x_t = \begin{cases} y_t - 2\gamma_t & x_t = 1 \Leftrightarrow y_t > 0 \\ y_t + 2\gamma_t & x_t = -1 \Leftrightarrow y_t < 0 \end{cases} \end{aligned}$$

Let's say, now, that y_1 is chosen to be some positive value (negative would be essentially the same, but we do want to avoid 0; besides, it is only one out of infinite possible initializations, so the generic case is the one we study). Then, on the next step, $y_2 = y_1 - 2\gamma_1$. If $y_2 > 0$, on the next $y_3 = y_2 - 2\gamma_2 = y_1 - 2(\gamma_1 + \gamma_2)$. As long as $y_t > 0$, this will continue, so that the next y will be $y_{t+1} = y_1 - 2\sum_{\tau=1}^t \gamma_\tau$. Of course, this cannot continue forever, due to the requirement $\sum_{t=1}^\infty \gamma_t = \infty$. Eventually,

Of course, this cannot continue forever, due to the requirement $\sum_{t=1}^{\infty} \gamma_t = \infty$. Eventually, y_t will either reach exactly 0 or pass over to the other side. Then, the same process will continue, with $y_{t+1} = y_{t_0} - 2\sum_{\tau=t_0}^{t} \gamma_{\tau}$, where t_0 is the step when we passed through 0. So, we can see that y_t will continuously move towards 0, bypassing it if the sums of γ_t

So, we can see that y_t will continuously move towards 0, bypassing it if the sums of γ_t do not equal 0 exactly, but then again moving immediately in the other direction, toward 0 again. By the way, $\sum_{t=1}^{\infty} \gamma_t = \infty$ means that there is no way that y_t eventually remain on any one side of 0 forever (because that would mean that $\sum_{t=t_0}^{\infty}$ is finite).

Now, if y_1 and γ_t are such, that the pathological case of $y_t = 0$ is avoided for all t (or, more concretely, if one chose to use 1 or -1 as the arbitrary maximizer of the zero function, since, let's say, -1 and 1 are the results for the two other cases), then it is easy to see that x_t will continuously alternate between 1 and -1, never converging (or even going near, for that matter) to the maximum point 0.

This is the reason (or, at least, one of the reasons) why, instead of a hard max, one must take a soft max approach, in the sense of:

$$y_i \mapsto \operatorname*{argmax}_{x_i \in \mathcal{X}_i} \{ \langle y_i, x_i \rangle - h_i(x_i) \}$$

where h_i is some strongly convex regularizer.

Connection to FTRL Actually, there is also one other interesting way to look at the Dual Averaging algorithm that we know of, and that is as a linearized and weighted version of the Follow The Regularized Leader (FTRL) algorithm. This is one of the fundamental algorithms in online optimization. It works directly on the online maximization setting that we discussed in the previous chapter, trying to minimize regret by selecting for the agent the best possible action in hindsight, i.e.:

$$x_{t+1} = \underset{x \in \mathcal{X}}{\operatorname{argmax}} \sum_{\tau=1}^{t} f_{\tau}(x) + h(x)$$

plus h(x), which is the regularization term again.

Then, the idea is to approximate the functions f_t by their linear approximation around the point chosen in the respective turn, i.e.:

$$f_t(x) \approx f_t(x_t) + \langle \nabla f_t(x_t), x - x_t \rangle$$

If one then takes these approximations, puts weights γ_t on them, and substitutes the functions in FTRL with these approximations, the result is the update rule:

$$\begin{split} x_{t+1} &= \operatorname*{argmax}_{x \in \mathcal{X}} \sum_{\tau=1}^t \gamma_\tau (f_\tau(x_\tau) + \langle \nabla f_\tau(x_\tau), x - x_t \rangle) + h(x) \\ &= \operatorname*{argmax}_{x \in \mathcal{X}} \langle \sum_{\tau=1}^t \gamma_\tau \nabla f_\tau(x_\tau), x \rangle + h(x) \text{ (the other terms are independent of } x) \end{split}$$

Now, modulo the initialization, this is exactly (DA) for each player i, since:

$$\begin{aligned} y_{i,t+1} &= y_{i,t} + \gamma_t v_i(x_t) = y_{i,1} + \sum_{\tau=1}^t \gamma_\tau v_i(x_\tau) \\ x_{i,t+1} &= Q_i(y_{i,t+1}) = \underset{x_i \in \mathcal{X}_i}{\operatorname{argmax}} \langle y_{i,t+1}, x_i \rangle - h_i(x_i) \\ &= \underset{x_i \in \mathcal{X}_i}{\operatorname{argmax}} \langle y_{i,1} + \sum_{\tau=1}^t \gamma_\tau v_i(x_\tau), x_i \rangle - h_i(x_i) \end{aligned}$$

For more details on FTRL, see, for example, [Shal12] (section 2.3) and [Haza19] (chapter 5).

We now continue with the formal definitions of a regularizer function and the choice map.

Definition 5.2 (Regularizer). A function $h: C \to \mathbb{R}$, where C is compact and convex, is called a regularizer (or penalty function) if it is continuous and K-strongly convex (for some K > 0).

Notice that this will not only solve the issue of sticking to the boundary, but also the equally annoying issue of poorly defined argmax, since $\langle y_i, x_i \rangle - h_i(x_i)$ is strongly concave, and strongly concave functions have unique maximum points.

Now, the choice map Q is defined as follows:

Definition 5.3. Suppose we have a convex (action) set C and a regularizer $h: C \to \mathbb{R}$. Then, the choice map $Q: \mathcal{Y} \to C$ induced by h is defined as:

$$Q(y) = \operatorname*{argmax}_{x \in C} \{ \langle y, x \rangle - h(x) \}$$

One can already see the correlation between Q(y) and the convex conjugate h^* of h, even formally, since $h^*(y) = \max_{x \in C} \{\langle y, x \rangle - h(x) \}$.

Example As an example, the simple Euclidean projection can be derived as the choice map induced by the squared Euclidean norm, $h(x) = \frac{1}{2}||x||_2^2$, since:

$$\begin{aligned} & \underset{x \in C}{\operatorname{argmax}} \{ \langle y, x \rangle - \frac{1}{2} \| x \|_2^2 \} = \underset{x \in C}{\operatorname{argmax}} \{ 2 \langle y, x \rangle - \| x \|_2^2 \} = - \underset{x \in C}{\operatorname{argmin}} \{ \| x \|_2^2 - 2 \langle y, x \rangle \} \\ & = - \underset{x \in C}{\operatorname{argmin}} \{ \| y \|_2^2 + \| x \|_2^2 - 2 \langle y, x \rangle \} \text{ (adding a constant makes no difference)} \\ & = - \underset{x \in C}{\operatorname{argmin}} \{ \langle y - x, y - x \rangle \} = - \underset{x \in C}{\operatorname{argmin}} \{ \| y - x \|_2^2 \} \\ & = - \underset{x \in C}{\operatorname{argmin}} \{ \| y - x \|_2 \} \text{ since square is increasing} \end{aligned}$$

Finally, some interesting results can be derived, which are, in fact, difficult to see in some cases for specific functions. There is no reason to dig deeper into these results for the current work, since we do not play around with choice maps ourselves, but we list some of these facts here, as they are needed for the discussion on convergence.

- 1. x = Q(y) if and only if $y \in \partial h(x) \Leftrightarrow h(x') h(x) \ge \langle y, x' x \rangle, \forall x' \in C$; in particular, im $Q = \operatorname{dom} \partial h$.
- 2. h^* is differentiable on y and $\nabla h^*(y) = Q(y)$ for all y.
- 3. Q is (1/K)-Lipschitz continuous.

These three points are quite well known in convex analysis; for proofs, see, for example, [Zhou18].

Adaptation to Games Up until this point, we basically discussed choice maps in a convex optimization context, with no mention of multiple players.

The adaptation is straightforward: each player $i \in [N]$ will be endowed with an individual regularizer h_i which is K_i strongly convex and induces a choice map Q_i .

One thing to note, however, is that we will often need to refer to the players as an ensemble, with one choice map, i.e. $x \in \mathcal{X}$ and Q(x). This is done by considering $h(x) = \sum_i h_i(x_i)$ as the aggregate penalty function, which is K-strongly convex, where $K = \min_i K_i$. The induced choice map in this case is simply $Q(x) = (Q_1(x_1), \dots, Q_N(x_N))$.

5.3 Potential: Fenchel Coupling

The quantity that is used as a potential in order to show convergence is the Fenchel coupling, a relatively known primal-dual measure of distance which can be used to measure the divergence between a primal and a dual point (by dual we mean the points $Y_{i,n}$). As a demonstrative example, the Fenchel coupling for the entropic regularizer, which is the one used in the case of simplex strategy sets, is exactly the KL-divergence.

For one action profile $x \in \mathcal{X}$ and one point in the dual space y, the Fenchel coupling tries to measure their distance in terms of the defining characteristics of the choice map, i.e. the convex conjugate of h:

$$F(x,y) = h^*(y) + h(x) - \langle y, x \rangle$$

F(x,y) as defined above collects all the terms of Fenchel inequality, hence the name (first used in [Mert16]). This means that F(x,y) is always nonnegative.

There are two basic properties of the Fenchel coupling which will be of use to us. The first one captures the fact that we can use it as a potential function, in the sense that if it goes to 0, then the dynamics converge to equilibrium. Formally:

$$F(x,y) \ge \frac{1}{2}K||Q(y) - x||^2$$

This is a direct consequence of the strong convexity of h, since:

$$F(x,y) = h^*(y) + h(x) - \langle y, x \rangle = \langle y, Q(y) \rangle - h(Q(y)) + h(x) - \langle y, x \rangle$$
$$= h(x) - h(Q(y)) - \langle y, x - Q(y) \rangle \ge \frac{K}{2} \|Q(y) - x\|^2$$

and this inequality is equivalent to h being strongly convex (see [Zhou18] for a proof of this), because $y \in \partial h(Q(y))$ as we said before.

The second property is an inequality, which is the basic bound used to show that the Fenchel coupling under dual averaging dynamics goes to 0. Formally:

$$F(x, y') \le F(x, y) + \langle y' - y, Q(y) - x \rangle + \frac{1}{2K} ||y' - y||_*^2$$

This, on the other hand, is a direct consequence of the duality between strong convexity and smoothness (or lipschitz continuous gradient), showcased in [Zhou18], for example, which implies that, for a K-strongly convex h, h^* will be 1/K-smooth. Because of this, we get:

$$\begin{split} F(x,y') &= h^*(y') + h(x) - \langle y', x \rangle \\ &\leq h^*(y) + \langle y' - y, \nabla h^*(y) \rangle + \frac{1}{2K} \|y' - y\|_*^2 + h(x) - \langle y', x \rangle \\ &= h^*(y) + \langle y' - y, Q(y) \rangle + \frac{1}{2K} \|y' - y\|_*^2 + h(x) - \langle y', x \rangle \\ &= (h^*(y) + h(x) - \langle y, x \rangle) + \langle y, x \rangle - \langle y', x \rangle + \langle y' - y, Q(y) \rangle + \frac{1}{2K} \|y' - y\|_*^2 \\ &= F(x,y) + \langle y' - y, Q(y) - x \rangle + \frac{1}{2K} \|y' - y\|_*^2 \end{split}$$

As a final note, if this inequality by itself seems a bit stiff, try and look at its application on (DA) below. It may seem much more intuitive.

5.4 Convergence

With regard to convergence, the basic result which we are interested in is the following:

Theorem 5.4. If the (necessarily unique) Nash equilibrium of a game is variationally stable, then (DA) converges to it.

We will try to give an outline of the proof here, which we went to great lengths to make as simple and easy as possible. It is not an easy result (and, as a matter of fact, the results of [Mert19] that encompass local stability and noise are even harder).

Intuition from continuous time Firstly, let us examine the much simpler case of the continuous time dynamics, to see that the Fenchel coupling is indeed a suitable potential function and that (DA) indeed tends to converge to the stable equilibrium.

A mean field approximation of (DA) is be the following:

$$\dot{y} = v(x)$$
$$x = Q(y)$$

From this, we can take the time derivative of the Fenchel coupling:

$$\frac{d}{dt}F(x^*,y) = \frac{d}{dt}[h^*(y) + h(x^*) - \langle y, x^* \rangle]
= \langle \dot{y}, \nabla h^*(y) \rangle - \langle \dot{y}, x^* \rangle
= \langle v(x), Q(y) - x^* \rangle = \langle v(x(t)), x(t) - x^* \rangle < 0$$

where the last inequality stems from the fact that x^* is stable.

So, by Lyapunov's second method, we know the above system is (globally) asymptotically stable, which means that $F(x^*, y(t)) \to 0$ and, as a result, $x(t) \to x^*$.

The Proof The proof tries to recreate the above Lyapunov argument in the discrete time. To do so, the pivotal property it uses is the final Fenchel inequality we mentioned, applied to the current and the next step of the algorithm, i.e.:

$$F(x^*, y_{t+1}) \leq F(x^*, y_t) + \gamma_t \langle v(x_t), x_t - x^* \rangle + \frac{\gamma_t^2}{2K} \|v(x_t)\|_*^2$$

$$\leq F(x^*, y_t) + \gamma_t \langle v(x_t), x_t - x^* \rangle + \frac{\gamma_t^2 V_*^2}{2K}$$

Since x^* is variationally stable, the second term is, as in continuous time, always negative; if we only had that, then we would immediately get $F(x^*, y_{t+1}) < F(x^*, y_t)$ and convergence of $F(x^*, y_t)$ to 0 would follow instantly.

Alas, in discrete time, arises the extra final term, which is the hurdle that must be overcome to show convergence. So, to overcome this, the proof tries to control all three terms, so that, in the end, their sum always remains (eventually) in acceptable levels, in the sense that if $F(x^*, y_t) < \varepsilon$, then $F(x^*, y_{t+1}) < \varepsilon$. For this, there are two cases.

If x_t is sufficiently far from x^* , then the directional derivative from x_t towards x^* (which is always positive, by variational stability) is big enough to push us in the right direction, despite the extra positive term in the above. In this case, we get $F(x^*, y_{t+1}) < F(x^*, y_t)$, and we are good to go.

If, on the other hand, x_t is close to the equilibrium, then, unfortunately, we cannot guarantee $F(x^*,y_{t+1}) < F(x^*,y_t)$. Nevertheless, the variational stability is still enough to keep it from straying too far, because something else favorable happens; $F(x^*,y_t)$ is itself very small. Thus, we can still control $F(x^*,y_{t+1})$ so that it remains within a (arbitrarily small) neighbourhood of 0, and convergence will still hold.

More formally, we distinguish the two cases as follows:

1. $\frac{\varepsilon}{2} \leq F(x^*, y_t) < \varepsilon$. For this case, we rewrite the above inequality as:

$$F(x^*, y_{t+1}) \le F(x^*, y_t) + \gamma_t \left(\langle v(x_t), x_t - x^* \rangle + \frac{\gamma_t V_*^2}{2K} \right)$$

Now, since $F(x^*, y_t) \ge \frac{\varepsilon}{2}$, x_t is bounded away from x^{*4} , so the variational stability quantity, as a continuous function, has a maximum, which is now strictly negative. As a result:

$$\langle v(x_t), x_t - x^* \rangle \le -c(\varepsilon) < 0$$

$$\Rightarrow F(x^*, y_{t+1}) \le F(x^*, y_t) + \gamma_t \left(-c(\varepsilon) + \frac{\gamma_t V_*^2}{2K} \right)$$

If, for the moment, we assume that the step γ_t is small enough (specifically, $\gamma_t \leq c(\varepsilon) \frac{2K}{V_*^2}$), then we have indeed managed to make the parenthesized quantity negative, and we get:

$$F(x^*, y_{t+1}) \le F(x^*, y_t) < \varepsilon$$

2. $F(x^*, y_t) < \frac{\varepsilon}{2}$. For this case, we can immediately use the variational stability and rewrite the above inequality as:

$$F(x^*, y_{t+1}) \le F(x^*, y_t) + \frac{\gamma_t^2 V_*^2}{2K} < \frac{\varepsilon}{2} + \frac{\gamma_t^2 V_*^2}{2K}$$

Again, if, for the moment, we assume that the step γ_t is small enough (specifically, $\gamma_t^2 \leq \frac{\varepsilon}{2} \frac{2K}{V_*^2}$), then we indeed get:

$$F(x^*, y_{t+1}) < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon$$

Consequently, even if $F(x^*, y_t)$ are not monotonically decreasing, it is still guaranteed that, for any $\varepsilon > 0$, if some $F(x^*, y_t) < \varepsilon$, then all subsequent ones will also be $< \varepsilon$.

The proof is basically over here. There is only one loose end to be tied: guarantee that for every ε there will come a time when $F(x^*, y_t) < \varepsilon$ and γ_t is sufficiently small.

To that end, [Mert19] shows something slightly more general: that every neighbourhood of x^* is recurrent under (DA), i.e. x_t will reach the neighbourhood from any starting point; then it may walk away for a while, but it will again visit the neighbourhood eventually, and this will go on indefinitely. In other words, and more concisely, what he shows here is $\lim\inf_{t\to\infty}F(x^*,y_t)=0$.

Due to this, we are sure to eventually find ourselves in any neighbourhood of x^* with step size sufficiently small.

Moreover, this is not difficult to show. All we need to do is take the same inequality, this time telescoping it:

$$F(x^*, y_{t+1}) \le F(x^*, y_1) + \sum_{\tau=1}^t \gamma_\tau \langle v(x_\tau), x_\tau - x^* \rangle + \frac{V_*^2}{2K} \sum_{\tau=1}^t \gamma_\tau^2$$

Now, if some neighbourhood of x^* is not recurrent, then x_t must, eventually, stay permanently outside of it, so, the same way as in the first case above, the variational stability

⁴ for more accurate details, see the respective proof of [Mert19]

quantity must be bounded:

$$\langle v(x_{\tau}), x_{\tau} - x^* \rangle \le -c$$

$$\Rightarrow F(x^*, y_{t+1}) \le F(x^*, y_1) - c \sum_{\tau=1}^{t} \gamma_{\tau} + \frac{V_*^2}{2K} \sum_{\tau=1}^{t} \gamma_{\tau}^2$$

$$= F(x^*, y_1) - \sum_{\tau=1}^{t} \gamma_{\tau} \left(c - \frac{V_*^2}{2K} \frac{\sum_{\tau=1}^{t} \gamma_{\tau}^2}{\sum_{\tau=1}^{t} \gamma_{\tau}} \right)$$

for all considered step sizes, $\sum_{\tau=1}^t \gamma_{\tau}^2 / \sum_{\tau=1}^t \gamma_{\tau} \to 0$, so in the end we get:

$$F(x^*, y_{t+1}) \to -\infty$$

which is a contradiction, since $F(x, y) \ge 0$.

Thus, we have proved that x_t visit every neighbourhood of x^* an infinite number of times. Our proof is now complete.

5.5 Convergence Speed

[Mert19] uses two different, albeit similar, approaches / measures to bound the convergence speed of (DA) to equilibrium.

5.5.1 Equlibrium gap function

Consider a stable equilibrium $x^* \in \mathcal{X}$. Then, the equilibrium gap function is defined as:

$$\epsilon(x) = \langle v(x), x^* - x \rangle$$

By the definition of variational stability, $\epsilon(x) \geq 0$, with equality iff $x = x^*$.

Consequently, the equilibrium gap function can be viewed as a distance measure from x^* , even if not quite as good as the Euclidean distance itself.

Moreover, we call an $x^* \in \mathcal{X}$ strongly stable if, for some L > 0:

$$\langle v(x), x - x^* \rangle \le -L \|x - x^*\|^2 \Leftrightarrow \epsilon(x) \ge L \|x - x^*\|^2$$

i.e. the equilibrium gap function grows at least quadratically with the distance between x^* and x.

Now, denoting with $\bar{\epsilon}_t = \sum_{\tau=1}^t \gamma_{\tau} \epsilon(x_{\tau}) / \sum_{\tau=1}^t \gamma_{\tau}$ the average equilibrium gap, [Mert19] proves the following bounds.

Theorem 5.5. For the average equilibrium gap it holds that:

$$\bar{\epsilon}_t \le \frac{F_1 + V_*^2/(2K) \sum_{\tau=1}^t \gamma_\tau^2}{\sum_{\tau=1}^t \gamma_\tau}$$

where $F_1 = F(x^*, y_1)$.

Proof. The proof comes as a result of the fundamental inequality for Fenchel coupling, the same one we also used to show convergence:

$$0 \leq F(x^*, y_{t+1}) \leq F(x^*, y_t) + \gamma_t \langle v(x_t), x_t - x^* \rangle + \frac{\gamma_t^2 V_*^2}{2K}$$

$$\leq F(x^*, y_1) + \sum_{\tau=1}^t \gamma_\tau \langle v(x_\tau), x_\tau - x^* \rangle + \frac{V_*^2}{2K} \sum_{\tau=1}^t \gamma_\tau^2$$

$$\implies \sum_{\tau=1}^t \gamma_\tau \langle v(x_\tau), x^* - x_\tau \rangle \leq F(x^*, y_1) + \frac{V_*^2}{2K} \sum_{\tau=1}^t \gamma_\tau^2$$

$$\implies \frac{\sum_{\tau=1}^t \gamma_\tau \langle v(x_\tau), x^* - x_\tau \rangle}{\sum_{\tau=1}^t \gamma_\tau} = \bar{\epsilon}_t \leq \frac{F(x^*, y_1) + \frac{V_*^2}{2K} \sum_{\tau=1}^t \gamma_\tau^2}{\sum_{\tau=1}^t \gamma_\tau}$$

For concreteness, if we choose, for example, a step size policy of $\gamma_t = \frac{1}{\sqrt{t}}$ (which, as described in [Mert19], is optimal out of all $\frac{1}{t^{\beta}}$), then we get $\sum_{\tau=1}^{t} \gamma_{\tau}^{2} = \sum_{\tau=1}^{t} \frac{1}{\tau} \approx \log t$ (harmonic series), $\sum_{\tau=1}^{t} \gamma_{\tau} = \sum_{\tau=1}^{t} \frac{1}{\sqrt{\tau}} \approx \sqrt{t}$ (by integration), and, as a result, the average equilibrium gap $\bar{\epsilon}_t$ is $O(\frac{\log t}{\sqrt{t}})$. Which is pretty nice! As a small corollary of the above, if x^* is strongly stable, then:

$$\bar{\epsilon}_t = \frac{\sum_{\tau=1}^t \gamma_{\tau} \epsilon(x_{\tau})}{\sum_{\tau=1}^t \gamma_{\tau}} \ge L \frac{\sum_{\tau=1}^t \gamma_{\tau} \|x_{\tau} - x^*\|^2}{\sum_{\tau=1}^t \gamma_{\tau}} \ge L \left(\frac{\sum_{\tau=1}^t \gamma_{\tau} \|x_{\tau} - x^*\|}{\sum_{\tau=1}^t \gamma_{\tau}}\right)^2 = L \bar{r}_t^2$$

and thus

$$\bar{r}_t^2 \le \frac{1}{L} \frac{F(x^*, y_1) + \frac{V_*^2}{2K} \sum_{\tau=1}^t \gamma_\tau^2}{\sum_{\tau=1}^t \gamma_\tau}$$
 (5.2)

where $\bar{r}_t = \sum_{\tau=1}^t \gamma_\tau \|x_\tau - x^*\| / \sum_{\tau=1}^t \gamma_\tau$ is the average distance to equilibrium.

Corollary 5.6. If (DA) is run for T iterations with $y_1 = 0$ and constant step size $\gamma = V_*^{-1} \sqrt{2K\Omega/T}$, where $\Omega = \max h - \min h$, then:

$$\bar{\epsilon}_T \leq V_* \sqrt{2\Omega/KT}$$

If, further, x^* is L-strongly stable, then

$$\bar{r}_T \le \left(\frac{2V_*^2}{L^2} \cdot \frac{\Omega}{KT}\right)^{1/4}$$

where $\bar{r}_t = \frac{\sum_{\tau=1}^t \gamma_{\tau} \|x^* - x_{\tau}\|}{\sum_{\tau=1}^t \gamma_{\tau}}$ is again the average distance to equilibrium.

Proof. First of all, we note that:

$$\begin{split} F(x^*, y_1) &= F(x^*, 0) = h^*(0) + h(x^*) - \langle 0, x^* \rangle \\ &= \max_{x \in \mathcal{X}} \{ \langle 0, x \rangle - h(x) \} + h(x^*) = -\min_{x \in \mathcal{X}} h(x) + h(x^*) \\ &\leq -\min_{x \in \mathcal{X}} h(x) + \max_{x \in \mathcal{X}} h(x) = \Omega \end{split}$$

Then, we use directly the previous theorem, plugging in $\gamma_t = \gamma$:

$$\bar{\epsilon}_T \le \frac{F(x^*, y_1) + \frac{V_*^2}{2K}T\gamma^2}{T\gamma} \le \frac{\Omega}{T\gamma} + \frac{V_*^2\gamma}{2K}$$

Now, it is not hard to see, analytically, that the above quantity is minimized for $\gamma = V_*^{-1} \sqrt{2K\Omega/T}$, for which the average equilibrium gap bound becomes:

$$\bar{\epsilon}_T \le V_* \sqrt{\frac{2\Omega}{KT}}$$

Finally, similarly as before, $\bar{r}_T^2 \leq L^{-1}\bar{\epsilon}_T$ and, as such:

$$\bar{r}_T \leq \sqrt[4]{\frac{V_*^2}{L^2} \frac{2\Omega}{KT}}$$

We should note that the advantage of this last result over the previous theorem is that it is more straightforward, and the error guarantee is more tight, since here it is actually $O(\frac{1}{\sqrt{t}})$ (without the log factor).

The downside, of course, is that, here, the error goal needs to be known from the beginning, and the number of iterations is predetermined, based on that goal. This means that the process is not cumulative, in the sense that if one wants to improve an already calculated equilibrium, they must start the whole process anew.

5.5.2 Running length

The second measure is the running length of (DA), which is defined as:

$$\ell_t = \sum_{\tau=1}^{t-1} \|x_{\tau+1} - x_{\tau}\|$$

Now, if x_t converges to some x^* , then a shorter length signifies less oscillations of x_t around x^* . Thus, in a certain way, ℓ_t is a more refined convergence criterion than the equilibrium gap function.

The main result here is that the running length of (DA) until players reach an ε -neighbourhood of a strongly stable state is at most $O(1/\varepsilon^2)$. Formally:

Theorem 5.7. Suppose (DA) is run with step size such that $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$ and $\sum_{t=1}^{\infty} \gamma_t = \infty$. Also, let $t_{\varepsilon} = \inf\{t \geq 0 : \|x^* - x_t\| \leq \varepsilon\}$ be the stopping time until an ε -neighbourhood is reached, and $\ell_{\varepsilon} = \ell_{t_{\varepsilon}}$ denote the respective running length. If x^* is L-strongly stable, then:

$$\ell_{\varepsilon} \leq \frac{V_*}{KL} \frac{F_1 + (2K)^{-1} V_*^2 \sum_{t=1}^{\infty} \gamma_t^2}{\varepsilon^2}$$

Proof. We will utilize results from the previous section, and specifically the bound on the average distance to equilibrium. Since for all $t < t_{\varepsilon}$, $||x_t - x^*|| > \varepsilon$, the average distance to equilibrium up to that point cannot be less than ε , i.e. $\bar{r}_{t_{\varepsilon}-1} \ge \varepsilon$. Combining this with (5.2), we get:

$$\varepsilon^{2} \leq \bar{r}_{t_{\varepsilon}-1}^{2} \leq \frac{1}{L} \frac{F_{1} + \frac{V_{*}^{2}}{2K} \sum_{\tau=1}^{t_{\varepsilon}-1} \gamma_{\tau}^{2}}{\tau_{t_{\varepsilon}-1}} \Rightarrow \tau_{t_{\varepsilon}-1} \leq \frac{F_{1} + \frac{V_{*}^{2}}{2K} \sum_{\tau=1}^{\infty} \gamma_{\tau}^{2}}{L\varepsilon^{2}}$$

$$(5.3)$$

Finally, we use the fact that the choice map Q is 1/K-Lipschitz continuous:

$$\ell_{\varepsilon} = \sum_{\tau=1}^{t_{\varepsilon}-1} \|x_{\tau+1} - x_{\tau}\| \le \frac{1}{K} \sum_{\tau=1}^{t_{\varepsilon}-1} \|y_{\tau+1} - y_{\tau}\|_{*} = \frac{1}{K} \sum_{\tau=1}^{t_{\varepsilon}-1} \gamma_{\tau} \|v(x_{\tau})\|_{*} \le \frac{V_{*}}{K} \sum_{\tau=1}^{t_{\varepsilon}-1} \gamma_{\tau} = \frac{V_{*}}{K} \tau_{t_{\varepsilon}-1}$$

The original claim follows by combining the last two inequalities.

The above is already a more than sufficient argument as to the time efficiency of dual averaging. Nevertheless, we found it a nice testament to the above that even concrete bounds on the stopping time itself can be calculated for specific step size sequences. For this reason, we give an example of such a bound below:

Corollary 5.8. If (DA) is run with step size $\gamma_t = 1/t^{\frac{3}{5}}$ then $t_{\varepsilon} = O(1/\varepsilon^5)$.

Proof. $\sum_{t=1}^{\infty} \gamma_t$ of course diverges, and $\gamma_t^2 = 1/t^{\frac{6}{5}}$, so $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$, and it is equal to some constant c > 0. So, first of all, the requirements of the theorem are satisfied.

Moreover, given the explicit form of γ_t , we can now calculate an explicit formula for τ_t as a function of t. Specifically, skipping minor details, we can approximate the sum by integrals, which yields the inequality:

$$\tau_t = \sum_{\tau=1}^t \frac{1}{\tau^{\frac{3}{5}}} \ge \int_1^{t+1} \frac{dx}{x^{\frac{3}{5}}} = \frac{5}{2} \left[x^{\frac{2}{5}} \right]_1^{t+1} = \frac{5}{2} ((t+1)^{\frac{2}{5}} - 1)$$

which we can now combine with (5.3) to receive:

$$\frac{5}{2}(t_{\varepsilon}^{\frac{2}{5}} - 1) \le \tau_{t_{\varepsilon} - 1} \le \frac{F_1 + \frac{V_*^2}{2K}c}{L\varepsilon^2} \Rightarrow t_{\varepsilon} \le \left(1 + \frac{2}{5}\frac{F_1 + \frac{V_*^2}{2K}c}{L\varepsilon^2}\right)^{\frac{5}{2}} = O\left(\frac{1}{\varepsilon^5}\right)$$

which may not be as small, but it is comparably efficient to the original $\frac{1}{\varepsilon^2}$.

5.6 Application on Coevolutionary Games

We close the chapter by introducing the application of the (DA) algorithm on asymmetric coevolutionary opinion formation games and writing explicitly the respective formulas. Then, the main purpose of the following chapter will be to exhibit our results toward establishing that the Dual Averaging algorithm converges to equilibrium on these games.

We first recall (DA):

$$x_{i,t} = Q_i(y_{i,t})$$

$$y_{i,t+1} = y_{i,t} + \gamma_t v_i(x_t)$$

Since c_i are cost functions, the payoff functions will be $-c_i$. As a result:

$$v_i(x_t) = -\frac{\partial}{\partial x_i} c_i(x_t) = -2(x_{i,t} - (1 - \rho_i) \sum_{j \neq i} q_{ij}(x_{-i,t}; s_i) x_{j,t} - \rho_i s_i)$$
$$= -2(x_{i,t} - \beta(x_{-i,t}; s_i))$$

Furthermore, as we said previously, the choice map we will use will be the Euclidean projection, i.e. $Q_i(y_i) = \arg\min_{x_i \in \mathcal{X}_i} \|y_i - x_i\|_2^2$. Since $\mathcal{X}_i = [0, 1]$, the projection will simply be:

$$Q_i(y_i) = \begin{cases} 0, & \text{if } y_i < 0\\ 1, & \text{if } y_i > 1\\ y_i, & \text{otherwise} \end{cases}$$

Substituting the above on (DA), we get:

$$\begin{split} x_{i,t} &= \min\{\max\{y_{i,t}, 0\}, 1\} \\ y_{i,t+1} &= y_{i,t} - 2\gamma_t(x_{i,t} - \beta(x_{-i,t}; s_i)) \\ &= y_{i,t} + 2\gamma_t(\beta(x_{-i,t}; s_i) - x_{i,t}) \end{split}$$

Finally, to simplify things even more, we make two (innocent, as far as we can tell) assumtions:

- 1. $y_{i,1} \in [0,1]$.
- 2. $\gamma_t < 1/2$. This is easy, since the step sequence γ_t is decreasing anyway, so all we need to do is make sure that $\gamma_1 < 1/2$. After all, the results of [Mert19] are independent of the initialization of the run.

Under these assumptions, and since $\beta(x_{-i}; s_i) \in [0, 1]$, as a convex combination of $x_{-i}; s_i$ which are all in [0, 1], it is not difficult to see, recursively, that every $y_{i,t}$ will be a convex combination of the previous point and β which will both be in [0, 1].

As a result, $y_{i,t}$ will always be in [0,1], so $x_{i,t} = y_{i,t}$ and the (DA) dynamics can be reduced to the following:

$$x_{i,t+1} = x_{i,t} + \gamma_t(\beta(x_{-i,t}; s_i) - x_{i,t})$$
(5.4)

where we have substituted $\gamma_t \leftarrow 2\gamma_t$ for clarification purposes.

So, after all, the (DA) dynamics, in our case, does not do something particularly fancy, it's easily understandable: at each step, player i does not play her best response, but some suitable compromise between her best response and her previous opinion⁵.

⁵ Not to be confused with the soft maximization in Q. By now, we have done away with Q. Here, essentially, we have a (discrete-time) gradient descent algorithm.

Chapter 6

Dual Averaging on Coevolutionary Opinion Formation Games

As we saw, the results of [Mert19] seem very promising for the efficient computation of equilibria in a very broad class of (concave) games, in games that satisfy (DSC), for example, wherein, as a matter of fact, lies our particular interest. Moreover, we know for a fact that asymmetric coevolutionary games are concave and, seemingly, they are fairly well behaved; for example, the cost functions are a simple quadratic in x_i . That is to say, they seem highly eligible to satisfy these slightly stricter requirements of [Mert19].

This is why, with our main goal being equilibrium computation on asymmetric coevolutionary opinion formation games, our main approach was to try and apply the results of [Mert19] to these games, by showing that they (or interesting subclasses) satisfy the needed requirements, mainly diagonal strict concavity.

Moreover, we should point out that the Dual Averaging algorithm studied in [Mert19] has several other desirable properties. For one, it is decentralized, which means that the players can be thought of as being pitted against each other in a repeated iteration of the game in question, and (given sufficient feedback) each player employs the algorithm independently of others, to select their action at each round of the game. In other words, from the player's perspective, the setting conforms to that of online optimization. Additionally, Dual Averaging is a well known no-regret algorithm, hence it is arguably "compatible" with the incentive of the players, which is to minimize their total cost.

Thus, if we manage to satisfy the requirements, we immediately get a simple, decentralized, natural, incentive-compatible algorithm / dynamic that converges to equilibrium fast. Hence our motivation for exploring how this can be achieved, and for what subclasses of asymmetric coevolutionary games.

6.1 Theoretical Results on Diagonal Strict Concavity

A central approach in our work was to try and find interesting subclasses of asymmetric coevolutionary opinion formation games, for which diagonal strict concavity holds. Then, as we saw in the previous chapter, any dual averaging dynamic will converge to the unique equilibrium of those games relatively quickly.

One such endeavor that bore some fruit theoretically was to start from the definition of diagonal strict concavity, substitute the specifics of asymmetric coevolutionary games and work our way until some interesting conclusion is reached. Specifically, we reach a potentially interesting sufficient condition for the diagonal strict concavity of an asymmetric coevolutionary opinion formation game.

First of all, a concept that will be of use is the part of the best response of a player that involves only the x_j of the other players. We call that the centroid, defined formally as:

Definition 6.1 (Centroids). For any player i and any opinion profile $x \in \mathcal{X}$, we define the centroid of i's neighbours as:

$$\operatorname{ctr}_i(x) = \sum_{j \neq i} q_{ij}(x_{-i}) x_j$$

Furthermore, we denote by ctr: $\mathcal{X} \to \mathcal{X}$ the mapping which is the ensemble of these centroids, and can also be written in matrix notation as below.

$$\operatorname{ctr}(x) = (\operatorname{ctr}_1(x), \dots, \operatorname{ctr}_N(x)) = Q(x) \cdot x$$

where $Q(x) = \{q_{ij}(x)\}$ is the matrix of all interaction weights (assigning $q_{ii}(x) = 0$).

The concept of the centroid has been taken from [Fane21], which uses it to find approximate equilibria of coevolutionary opinion formation games, and adapted for our purposes below.

Notice that the centroid of a player is simply a convex combination (in other words, a weighted average) of the expressed opinions of all other players, weighted by the same weights that represent the strength of the interaction between i and her neighbours. Furthermore, we recall that player i's best response is equal to:

$$\beta_i(x_{-i}) = (1 - \rho_i) \sum_{j \neq i} x_j q_{ij}(x_{-i}) + \rho_i s_i = (1 - \rho_i) \operatorname{ctr}_i(x) + \rho_i s_i$$

So, the centroid represents an aggregation of the players' opinions, except i's, such that player i is essentially only interested in this point, rather than the full profile x_{-i} of the other players; with respect to selecting the best possible action, of course.

With this in mind, we can now prove the following lemma, which essentially adapts the definition of diagonal strict concavity to the specific case of asymmetric cevolutionary opinion formation games.

Lemma 6.2. An asymmetric coevolutionary opinion formation game is diagonally strictly concave if and only if:

$$\sum_{i=1}^{N} (1 - \rho_i)(\operatorname{ctr}_i(x^1) - \operatorname{ctr}_i(x^2))(x_i^1 - x_i^2) < \|x^1 - x^2\|_2^2$$

for all $x^1, x^2 \in \mathcal{X} = [0, 1]^N$.

If, further, all players share the same self confidence factor $\rho_i = \rho \in [0, 1)$, then this is also equivalent to the following:

$$\langle \operatorname{ctr}(x^1) - \operatorname{ctr}(x^2), x^1 - x^2 \rangle < \frac{1}{1 - \rho} ||x^1 - x^2||_2^2$$

Proof. Recall the definition of diagonal strict concavity, which, here, we adapt for the case of cost functions, so all signs are reversed:

$$\langle v(x^1) - v(x^2), x^1 - x^2 \rangle > 0$$
 (6.1)

In asymmetric coevolutionary opinion formation games, starting with the cost of player i:

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

we can calculate the individual gradient (here, derivative) of player i as:

$$\begin{split} \frac{\partial c_i}{\partial x_i}(x) &= 2(1-\rho_i) \sum_{j \neq i} (x_i - x_j) q_{ij}(x_{-i}) + 2\rho_i(x_i - s_i) \\ &= 2(1-\rho_i) x_i \sum_{j \neq i} q_{ij}(x_{-i}) - 2(1-\rho_i) \sum_{j \neq i} x_j q_{ij}(x_{-i}) + 2\rho_i x_i - 2\rho_i s_i \\ &= 2(1-\rho_i) x_i + 2\rho_i x_i - 2(1-\rho_i) \sum_{j \neq i} x_j q_{ij}(x_{-i}) - 2\rho_i s_i, \text{ since } \sum_{j \neq i} q_{ij}(x) = 1 \\ &= 2\left(x_i - (1-\rho_i) \sum_{j=1}^N q_{ij}(x_{-i}) x_j - \rho_i s_i\right) \\ &= 2(x_i - (1-\rho_i) \operatorname{ctr}_i(x) - \rho_i s_i) \end{split}$$

Note that $q_{ij}(x_{-i})$ is independent of x_i , therefore it is treated as a constant.

Gathering the above for all players, we can write the pseudogradient of the game at any $x \in [0,1]^N$ in matrix notation, as:

$$v(x) = 2(x - \operatorname{diag}\{1 - \rho_i\}\operatorname{ctr}(x) - \operatorname{diag}\{\rho_i\}s)$$

Thus, we can compute the LHS of the definition (6.1):

$$\begin{split} &\langle v(x^1) - v(x^2), x^1 - x^2 \rangle \\ = & 2 \langle x^1 - \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^1) - \operatorname{diag}\{\rho_i\} s - x^2 + \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^2) + \operatorname{diag}\{\rho_i\} s, x^1 - x^2 \rangle \\ = & 2 \langle x^1 - \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^1) - x^2 + \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^2), x^1 - x^2 \rangle \\ = & 2 \|x^1 - x^2\|_2^2 - 2 \langle \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^1) - \operatorname{diag}\{1 - \rho_i\} \operatorname{ctr}(x^2), x^1 - x^2 \rangle \\ = & 2 \|x^1 - x^2\|_2^2 - 2 \sum_{i=1}^N (1 - \rho_i) (\operatorname{ctr}_i(x^1) - \operatorname{ctr}_i(x^2)) (x_i^1 - x_i^2) \\ = & 2 \|x^1 - x^2\|_2^2 - 2 (1 - \rho) \langle \operatorname{ctr}(x^1) - \operatorname{ctr}(x^2), x^1 - x^2 \rangle, \text{ if self confidence is uniform} \end{split}$$

From the two final equalities, combined with the definition, the two parts of the lemma follow immediately:

$$\langle v(x^{1}) - v(x^{2}), x^{1} - x^{2} \rangle$$

$$= 2\|x^{1} - x^{2}\|_{2}^{2} - 2\sum_{i=1}^{N} (1 - \rho_{i})(\operatorname{ctr}_{i}(x^{1}) - \operatorname{ctr}_{i}(x^{2}))(x_{i}^{1} - x_{i}^{2}) > 0$$

$$\Leftrightarrow \sum_{i=1}^{N} (1 - \rho_{i})(\operatorname{ctr}_{i}(x^{1}) - \operatorname{ctr}_{i}(x^{2}))(x_{i}^{1} - x_{i}^{2}) < \|x^{1} - x^{2}\|_{2}^{2}$$

and for uniform self-confidence:

$$\begin{split} &\langle v(x^1) - v(x^2), x^1 - x^2 \rangle \\ = &2\|x^1 - x^2\|_2^2 - 2(1-\rho)\langle \text{ctr}(x^1) - \text{ctr}(x^2), x^1 - x^2 \rangle > 0 \\ \Leftrightarrow &\langle \text{ctr}(x^1) - \text{ctr}(x^2), x^1 - x^2 \rangle < \frac{1}{1-\rho}\|x^1 - x^2\|_2^2 \end{split}$$

Recall that, for a function $f: \mathcal{X} \to \mathbb{R}$, the condition:

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \le K ||x - y||^2$$
 for all $x, y \in \mathcal{X}$

is equivalent to f being K-smooth (e.g. [Zhou18]). So, the above lemma (the second part, in particular) tells us that the pseudogradient v(x) of the game resembles the gradient of a strictly concave (convex) function (i.e. diagonal strict concavity holds) if and only if the vector field ctr: $\mathcal{X} \to \mathcal{X}$ resembles the gradient of a G-smooth function (for a sufficiently small value of G).

As a small note, should the reader be interested in complete accuracy, the equivalence of those two conditions is not absolutely true (the "if" part, specifically for $G = \frac{1}{1-\rho}$). There is the small complication of < in the definition of (DSC), which, in the definition of smoothness, is a \le . So, if we want to be accurate, it holds that:

- 1. If diagonal strict concavity holds, then ctr does satisfy the same inequality as the gradient of a $\frac{1}{1-\rho}$ -smooth function.
- 2. If ctr is like the gradient of a G-smooth function, for any $G < \frac{1}{1-\rho}$, then diagonal strict concavity holds for self-confidence level ρ .
- 3. What does not necessarily hold is that if ctr is like the gradient of a $\frac{1}{1-\rho}$ -smooth function, then diagonal strict concavity holds.

This makes a small appearance in the following results, where the value $1 - \frac{1}{G}$ is excluded, but it is, still, a very minor detail.

One thing to note is that it is one of those annoying complications / asymmetries that appear when one talks about "strict" convexity / concavity properties. These, however, usually do not appear when we move on to "strong" properties, and this is also the case here, as we will see in the results for strong stability, in which this issue vanishes.

Now, it is easy to prove the following sufficient condition, which, under the same light, says nothing more than that if the vector field ctr is *G*-Lipschitz, then it does indeed resemble the gradient of a *G*-smooth function; two concepts that, as we know, are indeed very close in the scalar case, where, for a convex function, smoothness and having a Lipschitz gradient are equivalent.

Theorem 6.3. If the mapping ctr: $[0,1]^N \to [0,1]^N$ is G-Lipschitz continuous with respect to the Euclidean norm, then the asymmetric coevolutionary opinion formation game with uniform self confidence levels ρ is diagonally strictly concave for all values of ρ in $(1-\frac{1}{G},1]$. If $G \leq 1$, then, evidently, the game is always diagonally strictly concave.

Proof. From Cauchy-Schwartz inequality, it holds that:

$$\langle \mathsf{ctr}(x^1) - \mathsf{ctr}(x^2), x^1 - x^2 \rangle \leq \|\, \mathsf{ctr}(x^1) - \mathsf{ctr}(x^2)\|_2 \|x^1 - x^2\|_2$$

and, since ctr is G-Lipschitz:

$$||ctr(x^1) - ctr(x^2)||_2 \le G||x^1 - x^2||_2$$
 (6.2)

Finally, combining the above:

$$\langle \mathsf{ctr}(x^1) - \mathsf{ctr}(x^2), x^1 - x^2 \rangle \leq G \|x^1 - x^2\|_2^2 < \frac{1}{1 - \rho} \|x^1 - x^2\|_2^2$$

for all values of ρ such that $G < \frac{1}{1-\rho} \Leftrightarrow \rho > 1 - \frac{1}{G}$; which, given the previous lemma, implies that the game is diagonally strictly concave.

What we have established so far is not as difficult as it may seem at first. In more plain words, we have begun with a vector field (v(x)) which, for our specific games, is equal to another vector field (ctr) plus an affine term. Then, we say that v is diagonally strictly concave if ctr is Lipschitz. But, given that DSC means $\langle v(x^1)-v(x^2),x^1-x^2\rangle<0$ and Lipschitzness implies almost immediately $\langle \operatorname{ctr}(x^1)-\operatorname{ctr}(x^2),x^1-x^2\rangle< G\|x^1-x^2\|_2^2$, this conclusion is not hard to reach.

One immediate consequence is the following. It is well known that continuously differentiable functions over a compact domain are always Lipschitz, for a large enough constant. If q_{ij} are continuously differentiable, it is easy to see that the same holds for ctr_i . Consequently, for continuously differentiable weights, there always exists some large enough (but less than 1) level of self confidence for which the game is diagonally strictly concave.

This is not entirely unexpected, since for $\rho=1$ the game is trivially diagonally strictly concave, plus, we have already acquired some intuition from other opinion formation models which says that players' self-confidence / stubbornness "helps" convergence, more or less.

We close this analysis by adapting the above results to account for strong variational stability (as described in the previous chapter). Now, strong variational stability is a only a slight reinforcement of the concept of variational stability, but, as we saw, it is necessary for some of the results with regard to the convergence speed of Dual Averaging, and this is why we study it here.

Lemma 6.4. The unique equilibrium of a diagonally strictly concave asymmetric coevolutionary opinion formation game is also L-strongly stable if:

$$\sum_{i=1}^{N} (1 - \rho_i)(\operatorname{ctr}_i(x^1) - \operatorname{ctr}_i(x^2))(x_i^1 - x_i^2) \le (1 - \frac{L}{2})\|x^1 - x^2\|_2^2$$

for all $x^1, x^2 \in \mathcal{X} = [0, 1]^N$.

If, further, all players have the same self confidence factor $\rho_i = \rho \in [0, 1)$, then this is equivalent to the following:

$$\langle \operatorname{ctr}(x^1) - \operatorname{ctr}(x^2), x^1 - x^2 \rangle \le \frac{1 - \frac{L}{2}}{1 - \rho} ||x^1 - x^2||_2^2$$

Proof. First of all, notice that the correlation between diagonal strict concavity and variational stability can be extended to strong stability and "strong" diagonal convexity, by which we mean:

$$\langle v(x^1) - v(x^2), x^1 - x^2 \rangle \ge L ||x^1 - x^2||_2^2$$

Indeed, if the above holds, then we have:

$$\langle v(x) - v(x^*), x - x^* \rangle \ge L \|x - x^*\|_2^2$$

$$\Leftrightarrow \langle v(x), x - x^* \rangle \ge \langle v(x^*), x - x^* \rangle + L \|x - x^*\|_2^2 \ge L \|x - x^*\|_2^2$$

$$\Rightarrow \langle -v(x), x^* - x \rangle \ge L \|x - x^*\|_2^2$$

where we used the necessary condition (3.2) (with reversed sign, for the case of cost functions). Thus, x^* is L-strongly stable.

Given this, it is easy to treat the above inequality the same way we did with (6.1) in the

previous lemma:

$$\begin{split} \langle v(x^1) - v(x^2), x^1 - x^2 \rangle &\geq L \|x^1 - x^2\|_2^2 \\ \Leftrightarrow &2 \|x^1 - x^2\|_2^2 - 2 \sum_{i=1}^N (1 - \rho_i) (\operatorname{ctr}_i(x^1) - \operatorname{ctr}_i(x^2)) (x_i^1 - x_i^2) \geq L \|x^1 - x^2\|_2^2 \\ \Leftrightarrow &\sum_{i=1}^N (1 - \rho_i) (\operatorname{ctr}_i(x^1) - \operatorname{ctr}_i(x^2)) (x_i^1 - x_i^2) \leq (1 - \frac{L}{2}) \|x^1 - x^2\|_2^2 \end{split}$$

And, of course, if $\rho_i = \rho$, this becomes:

$$\begin{split} &(1-\rho)\langle \text{ctr}(x^1) - \text{ctr}(x^2), x^1 - x^2 \rangle \leq (1-\frac{L}{2})\|x^1 - x^2\|_2^2 \\ \Leftrightarrow &\langle \text{ctr}(x^1) - \text{ctr}(x^2), x^1 - x^2 \rangle \leq \frac{1-\frac{L}{2}}{1-\rho}\|x^1 - x^2\|_2^2 \end{split}$$

Theorem 6.5. If the mapping $\operatorname{ctr}\colon [0,1]^N \to [0,1]^N$ is G-Lipschitz continuous with respect to the Euclidean norm, then the asymmetric coevolutionary opinion formation game with uniform self confidence levels ρ is diagonally strictly concave and its unique equilibrium is L-strongly stable for all values of $\rho \in [0,1), L \in [0,2)$ such that $\rho \geq 1-(1-\frac{L}{2})/G$. If $G \leq 1-\frac{L}{2}$, then, evidently, this holds for all $\rho \in [0,1)$.

Proof. The proof follows the same line of thought. From Cauchy-Schwartz and Lipschitzness we get:

$$\langle \mathsf{ctr}(x^1) - \mathsf{ctr}(x^2), x^1 - x^2 \rangle \le G \|x^1 - x^2\|_2^2 \le \frac{1 - \frac{L}{2}}{1 - \rho} \|x^1 - x^2\|_2^2$$

where the final inequality holds for any values of ρ , L such that $G \leq \frac{1-\frac{L}{2}}{1-\rho} \Leftrightarrow \rho \geq 1 - \frac{1-\frac{L}{2}}{G}$.

So, according to the previous lemma, if $\rho \geq 1 - \frac{1 - \frac{L}{2}}{G}$, the game is diagonally strictly concave and its unique equilibrium is L-strongly stable.

These two results could be seen as a compromise between the values of ρ and L, in that the larger the value of L, the closer ρ needs to be to 1 in order to guarantee L-strong stability.

A surprising implication here is that, for this class of games, one can get strong stability almost for free; in the sense that, if we know some coevolutionary game is diagonally strictly concave, then all we need to do is increase the self-confidence ever so slightly, and it will now surely be strongly diagonally concave, for some suitable L.

Finally, to tie in with equilibrium computation, we combine the above results with the results of [Mert19] (specifically, theorems 6.2 and 6.4) into the following theorem.

Theorem 6.6. Consider an asymmetric coevolutionary opinion formation game, with uniform self-confidence factors $\rho_i = \rho$, and assume that the mapping ctr: $\mathcal{X} \to \mathcal{X}$ is G-Lipschitz.

Suppose, now, that the players run the (DA) algorithm with Euclidean projections and $\gamma_1 < 1/2$, i.e. their opinions evolve according to the update rule:

$$x_{i,t+1} = x_{i,t} + \gamma_t((1-\rho)\operatorname{ctr}_i(x_t) + \rho s_i - x_{i,t})$$

Then, for all $\rho \in (1 - \frac{1}{G}, 1]$, x_t converges to the unique equilibrium x^* of the game. Moreover, we have the following results with regard to convergence speed. Note that x^* is L-strongly stable, where $L = 2(1 - G(1 - \rho))$.

1. The average equilibrium gap $\bar{\epsilon}_t = \frac{\sum_{\tau=1}^t \gamma_{\tau} \langle v(x_{\tau}), x_{\tau} - x^* \rangle}{\sum_{\tau=1}^t \gamma_{\tau}} decays as$:

$$\bar{\epsilon}_t \le \frac{N(\frac{1}{2} + 2\sum_{\tau=1}^t \gamma_\tau^2)}{\sum_{\tau=1}^t \gamma_\tau}$$

2. The average distance to the equilibrium, $\bar{r}_t = \sum_{\tau=1}^t \gamma_\tau ||x_\tau - x^*||_2 / \sum_{\tau=1}^t \gamma_\tau$ decays as:

$$\bar{r}_t^2 \le \frac{1}{1 - (1 - \rho)G} \frac{N(\frac{1}{4} + \sum_{\tau=1}^t \gamma_\tau^2)}{\sum_{\tau=1}^t \gamma_\tau}$$

3. Define the running length $\ell_t = \sum_{\tau=1}^{t-1} \|x_{\tau+1} - x_{\tau}\|_2$ and the stopping time $t_{\varepsilon} = \inf\{t \geq 0 : \|x_t - x^*\|_2 \leq \varepsilon\}$, as in theorem 5.7.

Then, the running length until x_t reaches an ε -neighbourhood of x^* is bounded as:

$$\ell_{t_{\varepsilon}} \le \frac{N\sqrt{N}}{1 - (1 - \rho)G} \frac{\frac{1}{2} + 2\sum_{\tau=1}^{\infty} \gamma_{\tau}^{2}}{\varepsilon^{2}}$$

$$(6.3)$$

In a sense, the first and third convergence speed bounds could be seen as dual statements. If, for example, one chooses a step size $\gamma_t \propto 1/\sqrt{t}$, then the equilibrium gap decays as $O(\frac{\log t}{\sqrt{t}})$, which is almost the same as $O(\frac{1}{\sqrt{t}})$. On the other hand, the running length until an ε -neighbourhood of the equilibrium is reached is on the order of $\frac{1}{\varepsilon}$. So, the "distance" to equilibrium decays as $1/\sqrt{t}$ and the steps until we reach distance ε are $1/\varepsilon^2$; as one would expect.

Now, the next thing one may ask how can the conditions we have given so far be satisfied, either generally (which seems a bit far fetched) or for more specific subclasses of the general asymmetric coevolutionary opinion formation game. Since the condition is, after all, that ctr be Lipschitz, what one could do is, roughly, give concrete bounds on the derivatives of the centroid mapping. We give such an analysis in the next section.

6.1.1 Specialization for Normalized Independent Weights

Here, we introduce a special class of weight generating functions F_i which will be our main focus for a significant part of the rest of this work, and we apply the results of the previous section on this class of weights.

We call the class in question "normalized independent weights", and it is the following:

$$q_{ij}(x) = F_i(d_j^i, d_{-i-j}^i) = \frac{f_i(d_j^i)}{\sum_{k \neq i} f_i(d_k^i)}$$
(6.4)

where, as we have said previously, $d_j^i = |x_j - s_i|$, and the functions $f_i : [0, 1] \to \mathbb{R}$ are assumed to be:

- 1. Positive, mostly for simplicity.
- 2. Continuous, so that the weights are continuous functions, and, thus, we have a valid concave game and the guarantee that Nash equilibria always exist, by [Rose65].
- 3. Nonincreasing, to fulfil the requirement that $F_i(\cdot)$ is nonincreasing in the first argument and nondecreasing in the rest.

Furthermore, for this section, we will assume that the function $g_i(x_j) = f_i(d_i^i) = f_i(|x_j|)$ s_i) is continuous and differentiable on [0, 1]. In other words, f_i is required to be differentiable and $f_i'(0) = 0$; so that, when it is mirrored and translated by $f_i(|x_i - s_i|)$, there is no corner at $x_j = s_i$.

Given this assumption, we will fulfil the Lipschitzness condition (6.2) by bounding the derivatives of the mapping ctr.

Theorem 6.7. In an asymmetric coevolutionary opinion formation game with normalized independent weights, if $f_i(d) \ge m > 0$ and $f_i(d) + |f'_i(d)| \le M$, for all $d \in [0, 1]$, then the centroid mapping ctr is $G = \frac{M}{m} \sqrt{\frac{N}{N-1}}$ Lipschitz, and, consequently, the game is diagonally strictly concave for all uniform $\rho > 1 - 1/G$.

Proof. We begin by rewriting the centroids:

$$\operatorname{ctr}_{i}(x) = \sum_{k \neq i} x_{k} q_{ik}(x) = \sum_{k \neq i} x_{k} \frac{g_{i}(x_{k})}{\sum_{k \neq i} g_{i}(x_{k})} = \frac{\sum_{k \neq i} x_{k} g_{i}(x_{k})}{\sum_{k \neq i} g_{i}(x_{k})}$$

where $g_i : [0,1] \to \mathbb{R}$, $g_i(x_j) = f_i(|x_j - s_i|)$, as we mentioned before.

Now, we can calculate their partial derivatives (of course, x_i does not appear in any $q_{ij}(x)$, so this partial derivative is 0).

$$\begin{split} \frac{\partial}{\partial x_j} \operatorname{ctr}_i(x) &= \frac{\partial}{\partial x_j} \left(\frac{\sum_{k \neq i} x_k g_i(x_k)}{\sum_{k \neq i} g_i(x_k)} \right) \\ &= -\frac{g_i'(x_j)}{(\sum_{k \neq i} g_i(x_k))^2} \sum_{k \neq i} x_k g_i(x_k) + \frac{g_i(x_j) + x_j g_i'(x_j)}{\sum_{k \neq i} g_i(x_k)} \\ &= -\frac{g_i'(x_j)}{\sum_{k \neq i} g_i(x_k)} \frac{\sum_{k \neq i} x_k g_i(x_k)}{\sum_{k \neq i} g_i(x_k)} + \frac{g_i(x_j) + x_j g_i'(x_j)}{\sum_{k \neq i} g_i(x_k)} \\ &= -\frac{g_i'(x_j)}{\sum_{k \neq i} g_i(x_k)} \operatorname{ctr}_i(x) + \frac{g_i(x_j) + x_j g_i'(x_j)}{\sum_{k \neq i} g_i(x_k)} \\ &= \frac{1}{\sum_{k \neq i} g_i(x_k)} (g_i(x_j) + x_j g_i'(x_j) - g_i'(x_j) \operatorname{ctr}_i(x)) \\ &= \frac{1}{\sum_{k \neq i} g_i(x_k)} (g_i(x_j) + g_i'(x_j)(x_j - \operatorname{ctr}_i(x))) \end{split}$$

Since x_j and $\operatorname{ctr}_i(x)$ are both in [0,1], $x_j - \operatorname{ctr}_i(x) \in [-1,1]$ and we get:

$$g_i(x_j) + g_i'(x_j)(x_j - \mathsf{ctr}_i(x)) \in [g_i(x_j) - |g_i'(x_j)|, g_i(x_j) + |g_i'(x_j)|]$$

$$\Rightarrow |g_i(x_j) + g_i'(x_j)(x_j - \mathsf{ctr}_i(x))| \le g_i(x_j) + |g_i'(x_j)|$$

because $g_i(x_i) > 0$ and, thus, $g_i(x_i) + |g_i'(x_i)|$ must be greater than $|g_i(x_i) - |g_i'(x_i)||$.

Now, g_i is nothing more than a reflection and translation of f_i , so they share the same derivative bounds (in absolute value). Thus, $g_i(x_j) + |g_i'(x_j)| \leq M$.

Moreover, $g_i(x_j) \ge m \Rightarrow \sum_{k \ne i} g_i(x_k) \ge (N-1)m$. Combining all the above, we can bound the partial derivatives of the centroid by:

$$\left| \frac{\partial}{\partial x_j} \operatorname{ctr}_i(x) \right| \le \frac{M}{(N-1)m}$$

The remaining steps are straightforward. First, use this to bound the gradient of ctr_i:

$$\|\nabla \operatorname{ctr}_i(x)\|_2 = \sqrt{\sum_{j \neq i} \left(\frac{\partial}{\partial x_j} \operatorname{ctr}_i(x)\right)^2} \le \sqrt{(N-1) \left(\frac{M}{(N-1)m}\right)^2} = \frac{M/m}{\sqrt{N-1}}$$

Since bounded gradient implies Lipschitzness with respect to the same norm (see, for example, [Vish21], section 6.2.2), ctr_i is $\frac{M/m}{\sqrt{N-1}}$ -Lipschitz.

Finally, we show that if all ctr_i are G-Lipschitz, then ctr is $G\sqrt{N}$ -Lipschitz:

$$\begin{split} \|\operatorname{ctr}(x^1) - \operatorname{ctr}(x^2)\|_2 &= \sqrt{(\operatorname{ctr}_1(x^1) - \operatorname{ctr}_1(x^2))^2 + \dots + (\operatorname{ctr}_N(x^1) - \operatorname{ctr}_N(x^2))^2} \\ &\leq \sqrt{(G\|x^1 - x^2\|_2)^2 + \dots + (G\|x^1 - x^2\|_2)^2} \\ &= G\sqrt{N}\|x^1 - x^2\|_2 \end{split}$$

Consequently, ctr will be
$$\frac{M/m}{\sqrt{N-1}}\sqrt{N} = \frac{M}{m}\sqrt{\frac{N}{N-1}}$$
 Lipschitz.

As a final corollary, let us apply theorem 6.6 to this case:

Corollary 6.8. In an asymmetric coevolutionary opinion formation game with normalized independent weights, if $f_i(d) \geq m > 0$ and $f_i(d) + |f_i'(d)| \leq M$, for all $d \in [0,1]$, then, for all $\rho > 1 - \frac{m}{M} \sqrt{\frac{N-1}{N}}$, the algorithm (5.4) converges to the unique equilibrium of the game, and we have the same bounds on $\bar{\epsilon}_t$, \bar{r}_t and $\ell_{t_{\varepsilon}}$ as in theorem 6.6, which, after substituting $G = \frac{M}{m} \sqrt{\frac{N}{N-1}}$, become:

$$\bar{\epsilon}_t \le \frac{N(\frac{1}{2} + 2\sum_{\tau=1}^t \gamma_\tau^2)}{\sum_{\tau=1}^t \gamma_\tau}$$

$$\bar{r}_t^2 \le \frac{N}{1 - (1 - \rho)\frac{M}{m}\sqrt{\frac{N}{N-1}}} \frac{\frac{1}{4} + \sum_{\tau=1}^t \gamma_\tau^2}{\sum_{\tau=1}^t \gamma_\tau}$$

$$\ell_{t_\varepsilon} \le \frac{N\sqrt{N}}{1 - (1 - \rho)\frac{M}{m}\sqrt{\frac{N}{N-1}}} \frac{\frac{1}{2} + 2\sum_{\tau=1}^\infty \gamma_\tau}{\varepsilon^2}$$

Example: exponential weights As an example, let us consider the case where $f_i(d) = e^{-d^2}$, thus $g_i(x_j) = e^{-(x_j - s_i)^2}$ and $g_i'(x_j) = -2(x_j - s_i)e^{-(x_j - s_i)^2}$. Then:

$$g(x_j) = e^{-(x_j - s_i)^2} \ge e^{-1} = m$$

$$g(x_j) + |g'(x_j)| = e^{-(x_j - s_i)^2} (1 + 2|x_j - s_i|) \le 3 = M$$

for all $x_i \in [0, 1]$.

Thus, applying the above theorem, we get that ctr is G-Lipschitz, where:

$$G = \frac{M}{m} \sqrt{\frac{N}{N-1}} = 3e\sqrt{\frac{N}{N-1}}$$

which is < 10 for all N, and < 9 for $N \ge 6$.

In any case, this means that, for all values of ρ greater than $1 - \frac{1}{G} < 0.9$, the game is guaranteed to be diagonally strictly concave, and the adapted Dual Averaging algorithm (5.4) converges to equilibrium with the speed guarantees of the above corollary.

6.1.2 Remarks

First of all, this section wants to demonstrate that showing diagonal strict concavity for these complex games is not completely unapproachable, even if for somewhat constrained settings.

Secondly, the two theorems demonstrate a method for showing similar results for other functions: find a lower bound of g and an upper bound of $g(x_j) + |g'(x_j)|$, and we have automatically a set of values for ρ for which the game is diagonally strictly concave. Of course, this procedure is not guaranteed to yield a decent value for G, but it may do so.

Thirdly, since we are talking about a lower bound on g and (something like) an upper bound on |g'(x)|, we believe that this can intuitively be interpreted as a quantification of the fact that if our coevolving network changes relatively little, i.e. the weights are close to being constant, then it basically behaves like the static models, e.g. the FJ model (in which case the game is known to be diagonally strictly concave, as shown in [Bair19], section 4.2, for example).

6.2 Simulations

Apart from the results of the previous section, obtaining other significant theoretical results proved to be rather difficult. In general terms, bounding the derivative of the weights gave us the bounds we have talked about so far, but, other than that, the highly nonlinear nature of the weights made it difficult to analyze, for example, the Hessian matrix of the game, to prove diagonal strict concavity for broader classes of games, variational stability or, more directly, that the Dual Averaging algorithm converges to equilibrium.

It was for this reason that we then shifted our approach to exploring the model by means of simulation, in order to ascertain, for a number of instances, whether the facts that we are aiming for seem to hold or not and, in general, to get a more direct intuition of the way the (DA) algorithm evolves in these games.

6.2.1 Choice of Functions

We have implemented and run the (DA) algorithm on several instances of asymmetric coevolutionary opinion formation games with normalized independent weights. The setup is as we previously discussed, i.e. weights are generated by functions of the form:

$$q_{ij}(x) = F_i(d_j^i, d_{-i-j}^i) = \frac{f_i(d_j^i)}{\sum_{k \neq i} f_i(d_k^i)}$$

where $f_i : [0,1] \to \mathbb{R}$ are positive, continuous and nonincreasing.

Here, we give a roundup of the weight functions f_i we used. We have tried, as far as possible, to choose functions that are simple but also representative of many ways a continuous, differentiable function over [0, 1] could look like.

We should also note that all players were given the same weight function $f_i = f$ for any given run, to keep it simple.

- 1. A simple exponential: $f(d_j^i) = e^{-d_j^i}$.
- 2. Exponential, but adjusted to the interval [0,1]: $f(d_j^i) = e^{-20(d_j^i 5)}$ (because the simple one actually closely resembles a linear function in [0,1]).

- 3. A unit slope linear function: $f(d_i^i) = 1 d_i^i$.
- 4. A simple logarithm: $f(d_i^i) = -\ln d_i^i$.

Of course, this function produces weights that are not continuous, since they go to infinity at distance 0.

Nevertheless, on one hand, it was easy to check and, on the other, while running the algorithm we did not seem to hit any pathological points $(x_j, s_i \text{ exactly equal})$. Thus, we include it.

5. A valid logarithmic function, adjusted to the interval [0, 1]:

$$f(d_j^i) = -\ln\left(\frac{999}{1000}d_j^i + \frac{1}{1000}\right)$$

which, of course, now does not diverge at 0.

- 6. A logistic function: $f(d_j^i) = \frac{1}{1+e^{20(d_j^i-0.5)}}$.
- 7. The reciprocal: $f(d_j^i) = \frac{1}{d_j^i}$.

 Just like the logarithm, it is not continuous, but we include it.
- 8. An adjusted reciprocal: $f(d_j^i) = \frac{1}{d_j^i + \frac{1}{10000}}$
- 9. The square root function: $f(d_j^i) = -\sqrt{d_j^i}$.

Now, we can provide the specific results of our simulations.

6.2.2 Convergence in Small Instances

Our first endeavor was to verify experimentally that the (DA) algorithm does indeed converge to equilibrium; or, if luck has it, to find counterexamples. Moreover, by simulating specific examples, we hoped to gain further insight on the actual practical performance of the algorithm on real instances.

To that end, we start with some simple small simulations, both to ease the reader into the process and to take advantage of the fact that, for small numbers of players, the time progress of the algorithm can be seen directly.

In figure 6.1, we have run the algorithm for three instances of the game, each with a different weight function.

In these cases, we can see that convergence is achieved somewhat fast, since the players settle to equilibrium in a matter of about 100 steps. This seems to hold for all initializations and intrinsic opinions we have tried so far. Convergence is achieved to an equilibrium relatively fast, and, moreover, we should note that we have not encountered any instances with multiple limit points (i.e. equilibria).

These things could further suggest that diagonal strict concavity may hold, at least in some cases.

Another interesting observation is depicted in figure 6.2, where we have run the same game twice, once with all players running simple Best Response Dynamics (i.e. $x_{i,t+1} = \beta_i(x_t)$) and once with Dual Averaging.




Figure 6.1: Dual Averaging on the coevolutionary opinion formation game with 5 players, self-confidence 0.5 for all players, intrinsic opinions s=(0.4,0.5,0.6,0.127,0.897), and three different weight functions. Each color represents a different, randomly chosen initialization.






Figure 6.2: Best Response VS Dual Averaging: The game has N=50 players, self-confidence 0.5 for all players and as weight generating the adjusted reciprocal $f(d_j^i) = \frac{1}{d_j^i + \frac{1}{10000}}$. Both the intrinsic opinions vector and the initialization have been chosen randomly.

Here, we see that the Best Response dynamics constantly fluctuate, never converging to equilibrium¹. On the contrary, the Dual Averaging dynamics converge to equilibrium, and relatively fast at that.

We have picked this experiment as a demonstration that Dual Averaging is indeed more powerful than simple algorithms such as Best Response. Moreover, the choice of the weight generating function (the reciprocal) is not completely arbitrary. As we shall later see, we have found games, with reciprocal weight functions, which fail to satisfy the diagonal strict concavity condition everywhere; a strong suggestion that these games are not so well behaved. In spite of that, we see that the Dual Averaging algorithm still manages to converge to equilibrium.

¹ As a matter of fact, Best Response is already known to not guarantee convergence for general asymmetric coevolutionary opinion formation games. See, for example, [Bair19]





- (a) Final ℓ_{∞} error from minimum costs after running the algorithm for 200 steps.
- **(b)** Distance from minimum costs evolution with time, for a select few of the games.

Figure 6.3: Convergence of (DA) algorithm for exponential weights $f_i(d_i^i) = e^{-d_j^i}$.

6.2.3 Convergence to Equilibrium for Large Numbers of Players

Our goal here was primarily to see how the convergence behaviour scales with the number of players. To that end, we simulated games with up to almost 10000 players, and we wanted to assess whether they converge to equilibrium and, if so, how fast they reach it, estimate their convergence speed.

First of all, let us comment on the computational complexity of the simulations. Note that it takes $O(N^2)$ time to compute all the interaction weights q_{ij} between the players at a given opinion profile. As a result, a single step of the algorithm takes quadratic time in N.

This means that we are somewhat restricted, computationally, for numbers of players that reach the thousands. The first consequence is, of course, that we went up to 10000 players and not more. However, we also needed to restrict our exploration, varying only the weight function; while choosing, for each simulation, only one initialization at random, and, as usual, s_i are chosen at random and the self confidence ρ_i is set to 0.5.

Now, firstly, we plot the difference / error between the actual costs and the minimum costs of all players (i.e. at their best responses), as measured by the l_{∞} norm, after 200 steps of the algorithm. Formally, we plot the quantity:

$$\max_{1 \le i \le N} c_i(x_t) - c_i(\beta_i(x_{-i,t}), x_{-i,t})$$

at t = 200 (the absence of absolute value is because the quantity is always non negative).

In figure 6.3, we see an example of such a group of runs. Moreover, to showcase a bit the full evolution of the error as the algorithm progresses, we have also plotted the time sequence of those errors, for a select few of the games.

It is already somewhat visible that, in all these games, the algorithm gets in very close to equilibrium. The final errors are somewhat small, considering that $c_i(\cdot; x_{-i})$ is a simple, unit quadratic, and thus we know that it increases around its minimum as x^2 ; which means that, for example, an error in objective value of 10^{-6} of a player translates directly to a distance to equilibrium of 10^{-3} and vice versa. Additionally, it is evident from the time series of the errors that they continuously decrease towards 0.

Plus, after all, 200 steps are not very much. Moreover, we should note that for a few games that we selected at random throughout this process and run for more steps, the error continues to decrease for as long as the algorithm keeps going. Remember, for example, the run we showed in figure 6.2, where the algorithm clearly converged, but needed more than 200 steps. So, overall, given our settings, we believe that the performance of the algorithm is highly satisfactory.

We have also done the same for all the other weight function in our pool. All the respective diagrams can be seen in B.

Another interesting observation is that there does not seem to exist a significant correlation between the number of players and the ℓ_{∞} objective-value error to equilibrium. Of course, although we have not shown it here, the ℓ_1 error, which sums all the players' errors, does, as expected, increase with N in most cases. However, this means that how close any single player gets to their best response seems to be dependent only on the number of rounds that have been run, and not on the number of players.

Another thing that seems to affect the final proximity to equilibrium is, of course, the choice of weight function. Out of the functions that we have chosen here, we have been able to make out basically three tiers, based on how large is the error at the end of the algorithm:

- 1. The reciprocal and adjusted reciprocal functions performed the worst out of all, with an ℓ_{∞} error on the order of 10^{-4} .
- 2. Slightly better performance was demonstrated by the logarithm, the adjusted logarithm, and the adjusted exponential function, which are on the order of $10^{-6} 10^{-7}$.
- 3. Finally, the square root, the linear (unit slope), the simple exponential and the logistic function performed the best, with the distance to best response being on the order of 10^{-9} .

Most of these are not unexpected, or at least can be given plausible explanations.

The reciprocal and the adjusted reciprocal were expected to be "difficult" cases from the beginning. They are very aggressive, they explode at 0 and the interaction weights resulting from them probably change too abruptly.

The same arguments hold for the two logarithms, albeit slightly less here, since the logarithm is less aggressive; consider, for example, that the adjusted logarithm at 0 gives a value of ≈ 6.9 , while the reciprocal 10000.

The adjusted exponential $(f_i(d_j^i) = e^{-20(d_j^i - 0.5)})$ may surprise at first, but, after all, even though it does not become infinite, it, too, is too aggressive. Consider that $f_i(0) \approx 22000$ and $f_i(1) \approx 4 \cdot 10^{-5}$. So, it also changes too fiercely.

Finally, the rest are all much smaller and change less aggressively, so their good behaviour is not unexpected.

Measuring Convergence Speed Now, since, in the above runs, the algorithm seems to converge, the next thing was to measure the convergence speed with a little more precision.

We focused on two measures. First, the running length metric, which is also used in [Mert19] and we can compare the respective results. Secondly, the number of steps needed until the algorithm reaches an ε -neighbourhood of the equilibrium, which is more straightforward. We remind the definitions:

1. Running length at step t: $\ell_t = \sum_{k=1}^{t-1} \|x_{k+1} - x_k\|_2$. Of special interest is the running length until an ε -neighbourhood is reached: $\ell_\varepsilon \equiv \ell_{t_\varepsilon}$ (see below).




Figure 6.4: Running length and number of steps until an ε -neighbourhood of the equilibrium is reached. The game used has 8000 players, who employ exponential weights $f_i(d_j^i) = e^{-d_j^i}$.

2. Number of steps until a ε -neighbourhood is reached: $t_{\varepsilon} = \inf\{t \geq 0 : ||x_t - x^*||_2 \leq \varepsilon\}$.

We use the same runs of the algorithm as previously, but now we keep track of these two measures. Specifically, for each game, we have plotted ℓ_{ε} and t_{ε} as a function of $1/\varepsilon$. An example of such a plot is given in figure 6.4.

Of course, in the place of the exact equilibrium, which is unknown, we have used the opinion profile on which the algorithm ended, since, as we already saw, it is very close to being a Nash equilibrium.

One observation is that the total running length of the algorithm seems to be bounded. This is even more apparent if we plot ℓ_t directly as a sequence over the step variable t, shown in figure 6.5 for the same game.

[Mert19] notes, in section 6.2, that the running length of the (DA) algorithm cannot be expected to be bounded, in general, but this is due to the stochastic nature of the setting of [Mert19], in which the noise does not vanish even near an equilibrium.

But, here, our setting is significantly more constrained, there is no noise and, as it seems by the experiments, the running length does, in fact, turn out to be bounded.

This, in turn, however, makes it somewhat difficult to compare the running length covered by our algorithm to the theoretical guarantees given in [Mert19], since these bounds (e.g. (6.3)) are of the form $O(1/\varepsilon^2)$ and thus, if tight, the running length ℓ_{ε} should be increasing towards ∞ .

More plots such as these, for other weight functions and numbers of players, can be found




Figure 6.5: Running length at all steps of the algorithm run on the game with 8000 players and exponential weights.

in B.

From there, we can see again that weight functions such as the reciprocal are at a clear disadvantage with respect to convergence speed, while it can also be seen that ℓ_{ε} as well as t_{ε} generally increase with the number of players; something not unexpected, since here the norm used is ℓ_2 and not ℓ_{∞} as before.

6.2.4 Diagonal Strict Concavity counterexample

As we have said, our first hope would be for the asymmetric coevolutionary opinion formation games to always be diagonally strictly concave (3.7).

Sadly, that does not seem to be the case, at least in general. In this section, we shall provide a (relatively) small counterexample.

We consider the game with the following parameters:

- N = 5 players
- $\rho_i = \frac{1}{2}$ for all i
- s = (0.9, 0.84, 0.83, 0.2, 0.18)
- normalized independent weights with reciprocal generating functions, i.e.:

$$f_i(d) = \frac{1}{d + 0.0001}$$


Furthermore, we consider the following two points of the opinion space:


$$x^{1} = (0.84, 0.89, 0.8, 0.49, 0.3)$$

 $x^{2} = (0.83, 0.87, 0.8, 0.49, 0.3)$

Now, notice that diagonal strict concavity is equivalent to the following $\mathbb{R} \to \mathbb{R}$ function being monotone increasing over $[0,1] \ni t$:

$$g(t) = \langle v(x^1 + t(x^2 - x^1)), x^2 - x^1 \rangle$$





- (a) Zoomed onto the line segment between x^1 and x^2 .
- **(b)** Whole line, up to almost the boundary of the hypercube.

Figure 6.6: Pseudogradient projections along the line connecting x^1 - x^2

for all $x^1, x^2 \in \mathcal{X}$.

Indeed, sufficiency is trivial, since for $t_1 = 0$ and $t_2 = 1$ we retrieve the original definition of diagonal strict concavity. The fact that diagonal strict concavity implies the monotonicity of g is also not difficult to see, since monotonicity is equivalent to:

$$\begin{split} &(g(t_1)-g(t_2))(t_1-t_2)>0 \text{ for all } t_1\neq t_2, \text{ of course} \\ &\Leftrightarrow (\langle v(x^1+t_1(x^2-x^1)), x^2-x^1\rangle - \langle v(x^1+t_2(x^2-x^1)), x^2-x^1\rangle)(t_1-t_2)>0 \\ &\Leftrightarrow \langle v(x^1+t_1(x^2-x^1))-v(x^1+t_2(x^2-x^1)), x^2-x^1\rangle(t_1-t_2)>0 \\ &\Leftrightarrow \langle v(x^1+t_1(x^2-x^1))-v(x^1+t_2(x^2-x^1)), (t_1-t_2)(x^2-x^1)\rangle>0 \end{split}$$

If we denote $x_{t_1} = x^1 + t_1(x^2 - x^1)$ and $x_{t_2} = x^1 + t_2(x^2 - x^1)$, then $x_{t_1} - x_{t_2} = x^1 + t_1(x^2 - x^1) - x^1 - t_2(x^2 - x^1) = (t_1 - t_2)(x^2 - x^1)$ and the above becomes:

$$\langle v(x_{t_1}) - v(x_{t_2}), x_{t_1} - x_{t_2} \rangle > 0$$

which is true, because of diagonal strict concavity. Therefore, the original is true as well, which means q is strictly increasing.

Now, to return back to the example, if we plot g(t) for the two x^1 and x^2 we specified above, then we get figure 6.6a.

It is evident that g(t) is not increasing between these two points. Thus, our counterexample. For completeness, we point the reader to appendix A, should the reader be interested in seeing the exact arithmetic calculations of $\langle v(x^1) - v(x^2), x^1 - x^2 \rangle$ between those two points.

Something else we should note is that, although in figure 6.6a it appears that the deviation is significant, in fact it is not. See, for example, figure 6.6b, where we have drawn the same graph, but for two points further apart on the line between x^1 and x^2 , i.e. for two points $x_1 + t(x_2 - x_1)$, one with t < 0 and one with t > 1.

With this extension, it can be seen that the discrepancy of x^1 and x^2 is only a minor perturbation of g(t), which is otherwise very clearly increasing.

Moreover, we should note that, by taking a closer look at s and s, s, the reader can notice that some s-components and some s components are almost equal. This means that the denominator of $1/d_j^i$ would be zero or very close to zero, and the only reason this does not happen is due to the adjustment we implemented on the reciprocal, so that it does not become infinite at 0. Thus, the example relies somewhat on the fact that the weight function almost "explodes" at 0, and, by this perspective, could be seen as an edge case.

Nevertheless, the weight function remains fully legitimate (continuous, nonincreasing, differentiable), so, at the very least, this counterexample, however crude, shows that it is impossible to simply prove that the game is diagonally strictly concave, with no further assumptions at all with regard to the weight functions.

6.3 Social Concavity

Since we know that, for all socially concave games, no-regret learning converges to Nash equilibrium, one of our first approaches was to try and connect socially concave games to our setting; mostly by trying to find interesting subclasses or specific instances of asymmetric coevolutionary games that are socially concave.

We did not draw decisive conclusions on this direction, however it seems highly challenging to find such a connection, so we thought that, as a final note, it would be interesting to record some of the difficulties.

First of all, notice that even in the (4.2) model, where the interaction weights are taken to be constant (i.e., in the FJ model), and which is a special case of asymmetric coevolutionary games, the cost functions are convex, and specifically quadratic, in all arguments. So, they are very far from being concave in x_{-i} .

Nevertheless, one could search beyond constant weights, in the hope of finding other classes of weight functions that make the game socially concave. To that end, let us examine the two conditions needed for social concavity.

Firstly, since in opinion formation games we have costs, but in socially concave games we have payoffs, we turn costs into payoffs by simply reversing the sign. Thus, in coevolutionary opinion formation games, each player i is assigned the payoff function:

$$u_i(x) = -c_i(x) = -(1 - \rho_i) \sum_{j \neq i} (x_i - x_j)^2 q_{ij}(x) - \rho_i (x_i - s_i)^2$$

Thus, for the game to be socially concave, we need one, to find some strict convex combination of the payoffs which is concave (equivalently, a combination of costs which is convex), and two, to assert that the payoff of each player i is convex in x_{-i} (equivalently, player i's cost is concave in x_{-i}).

For the first, we want the sum:

$$g(x) = \sum_{i=1}^{N} \lambda_i c_i(x) = \sum_{i=1}^{N} \lambda_i \rho_i (x_i - s_i)^2 + \sum_{i,j=1}^{N} \lambda_i (1 - \rho_i) (x_i - x_j)^2 q_{ij}(x)$$

to be convex in x (where we have taken $q_{ii}(x) = 0$).

It is obvious that, if $q_{ij}(x)$ are convex in x, then this requirement is satisfied. Indeed, this is true, however, in such a case, costs are convex in all opinions (as was the case for FJ), therefore the second requirement for social concavity cannot possibly hold. But, even if we ignore the second requirement, there are two important conflicts with the model we have given for asymmetric coevolutionary opinion formation games:

1. $\sum_{j\neq i}q_{ij}(x)=1$ for all $x\in\mathcal{X}$. This normalization of weights makes it, of course, impossible for all weights to be convex, unless linear, since $q_{ij}(x)=1-\sum_{k\neq i,j}q_{ik}(x)$ must be concave if all other weights are convex.

2. $q_{ij}(x) = F_i(d^i_j, d^i_{-i-j})$, as we have said, so $q_{i,j}$ actually needs to depend only on the absolute differences between the player's intrinsic opinion and other players' public opinions. Moreover, F_i should be nonincreasing in the first argument and nondecreasing in the rest. As a result, $q_{ij}(x_j, x_{-i-j})$, as a function of x_j , will most probably have a strict maximum at s_i and, therefore, not be convex.

Of course, forfeiting these conditions is not at all out of the question, but, in this work, this is the model we studied.

As for the second condition, for which we need each $c_i(x_i, x_{-i})$ to be concave in x_{-i} , we already mentioned a prominent counterexample, which is the constant weights of the FJ model.

These are the main reasons why we decided that this direction will probably not be very fruitful.

Βιβλιογραφία

- [Amir21] Gideon Amir, Itai Arieli, Galit Ashkenazi-Golan and Ron Peretz, "Robust Naive Learning in Social Networks", 2021.
- [Anto21] Kimon Antonakopoulos and Panayotis Mertikopoulos, "Adaptive first-order methods revisited: Convex optimization without Lipschitz requirements", in NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems, 2021.
- [Bair19] Konstantina Bairaktari, "Υπολογισμός και Σύγκλιση σε Σημεία Ισορροπίας σε Συνεξελικτικά Παίγνια Διαμόρφωσης Άποψης", Master's thesis, School of Electrical and Computer Engineering, National Technical University of Athens, 2019.
- [Bhat12] Arnab Bhattacharyya, Mark Braverman, Bernard Chazelle and Huy Nguyen, "On the Convergence of the Hegselmann-Krause System", *ITCS 2013 Proceedings of the 2013 ACM Conference on Innovations in Theoretical Computer Science*, 11 2012.
- [Bhaw13] Kshipra Bhawalkar, Sreenivas Gollapudi and Kamesh Munagala, "Coevolutionary Opinion Formation Games", in *Proc. of the 45th ACM Symposium on Theory of Computing*, ACM, 2013.
- [Bil□18] Vittorio Bilò, Angelo Fanelli and Luca Moscardelli, "Opinion formation games with dynamic social influences", *Theoretical Computer Science*, vol. 746, pp. 73–87, 2018.
- [Bind11] David Bindel, Jon Kleinberg and Sigal Oren, "How Bad is Forming Your Own Opinion?", in 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, pp. 57–66, 2011.
- [Boyd14] Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*, Cambridge University Press, 2014.
- [Chen09] Xi Chen, Xiaotie Deng and Shang-Hua Teng, "Settling the Complexity of Computing Two-Player Nash Equilibria", *J. ACM*, vol. 56, no. 3, may 2009.
- [Chen16] 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", *Operations Research Letters*, vol. 44, no. 6, pp. 808–811, 2016.
- [Chie18] Flavio Chierichetti, Jon Kleinberg and Sigal Oren, "On discrete preferences and coordination", *Journal of Computer and System Sciences*, vol. 93, pp. 11–29, 2018.

- [Dask09] Constantinos Daskalakis, Paul W. Goldberg and Christos H. Papadimitriou, "The Complexity of Computing a Nash Equilibrium", *SIAM Journal on Computing*, vol. 39, no. 1, pp. 195–259, 2009.
- [Deff00] Guillaume Deffuant, David Neau, Frédéric Amblard and Gérard Weisbuch, "Mixing Beliefs Among Interacting Agents", *Advances in Complex Systems*, vol. 3, pp. 87–98, 01 2000.
- [DeGr74] Morris H. DeGroot, "Reaching a Consensus", *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [Even03] Eyal Even-Dar, Alex Kesselman and Yishay Mansour, "Convergence Time to Nash Equilibria", in Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow and Gerhard J. Woeginger, editors, *Automata, Languages and Programming*, pp. 502–513, Berlin, Heidelberg, 2003, Springer Berlin Heidelberg.
- [Even09] 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 '09, p. 523–532, New York, NY, USA, 2009, Association for Computing Machinery.
- [Fane21] Angelo Fanelli, "On Coevolutionary Opinion Formation". Private Communication, 2021.
- [Fost97] Dean P. Foster and Rakesh V. Vohra, "Calibrated Learning and Correlated Equilibrium", *Games and Economic Behavior*, vol. 21, no. 1, pp. 40–55, 1997.
- [Freu99] Yoav Freund and Robert E. Schapire, "Adaptive Game Playing Using Multiplicative Weights", *Games and Economic Behavior*, vol. 29, no. 1, pp. 79–103, 1999.
- [Frie90] Noah E. Friedkin and Eugene C. Johnsen, "Social Influence and Opinions", *Journal of Mathematical Sociology*, 1990.
- [Ghad13] Javad Ghaderi and R. Srikant, "Opinion dynamics in social networks: A local interaction game with stubborn agents", in *2013 American Control Conference*, pp. 1982–1987, 2013.
- [Gian21a] Angeliki Giannou, Emmanouil Vasileios Vlatakis-Gkaragkounis and Panayotis Mertikopoulos, "The convergence rate of regularized learning in games: From bandits and uncertainty to optimism and beyond", in *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.
- [Gian21b] Angeliki Giannou, Emmanouil Vasileios Vlatakis-Gkaragkounis and Panayotis Mertikopoulos, "Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information", in *Proceedings of Thirty Fourth Conference on Learning Theory*, vol. 134 of *Proceedings of Machine Learning Research*, pp. 2147–2148, PMLR, 15–19 Aug 2021.
- [Golu10] Benjamin Golub and Matthew O. Jackson, "Naïve Learning in Social Networks and the Wisdom of Crowds", *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–49, February 2010.

- [Hart00] Sergiu Hart and Andreu Mas-Colell, "A Simple Adaptive Procedure Leading to Correlated Equilibrium", *Econometrica*, vol. 68, no. 5, pp. 1127–1150, 2000.
- [Haza19] Elad Hazan, "Introduction to Online Convex Optimization", *CoRR*, vol. abs/1909.05207, 2019.
- [Hegs02] Rainer Hegselmann and Ulrich Krause, "Opinion Dynamics and Bounded Confidence Models, Analysis and Simulation", *Journal of Artificial Societies and Social Simulation*, vol. 5, 07 2002.
- [Hend06] Julien M. Hendrickx and Vincent D. Blondel, "Convergence of linear and non-linear versions of vicsek's model", 2006.
- [Hsie21] Yu-Guan Hsieh, Kimon Antonakopoulos and Panayotis Mertikopoulos, "Adaptive Learning in Continuous Games: Optimal Regret Bounds and Convergence to Nash Equilibrium", in COLT '21: Proceedings of the 34th Annual Conference on Learning Theory, 2021.
- [Kaku41] Shizuo Kakutani, "A generalization of Brouwer's fixed point theorem", *Duke Mathematical Journal*, vol. 8, no. 3, pp. 457 459, 1941.
- [Kout99] Elias Koutsoupias and Christos H. Papadimitriou, "Worst-case Equilibria", in *STACS*, 1999.
- [Lore05] Jan Lorenz, "A stabilization theorem for dynamics of continuous opinions", *Physica A: Statistical Mechanics and its Applications*, vol. 355, no. 1, pp. 217–223, 2005. Market Dynamics and Quantitative Economics.
- [Mert16] Panayotis Mertikopoulos and William H. Sandholm, "Learning in Games via Reinforcement and Regularization", *Mathematics of Operations Research*, vol. 41, no. 4, pp. 1297–1324, 2016.
- [Mert19] Panayotis Mertikopoulos and Zhengyuan Zhou, "Learning in games with continuous action sets and unknown payoff functions", *Mathematical Programming*, vol. 173, pp. 465–507, 2019.
- [More05] Luc Moreau, "Stability of multiagent systems with time-dependent communication links", *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.
- [Nash50] John F. Nash, "Equilibrium points in n-person games", *Proceedings of the National Academy of Sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [Nash51] John F. Nash, "Non-Cooperative Games", *Annals of Mathematics*, vol. 54, p. 286, 1951.
- [Nest09] Yurii Nesterov, "Primal-dual subgradient methods for convex problems", *Mathematical Programming*, vol. 120, pp. 221–259, 2009.
- [Polj95] Svatopluk Poljak, "Integer Linear Programs and Local Search for Max-Cut", *SIAM Journal on Computing*, vol. 24, no. 4, pp. 822–839, 1995.
- [Rose65] J. B. Rosen, "Existence and Uniqueness of Equilibrium Points for Concave N-Person Games", *Econometrica*, vol. 33, 1965.

- [Shal12] Shai Shalev-Shwartz, "Online Learning and Online Convex Optimization", *Foundations and Trends in Machine Learning*, vol. 4, pp. 107–194, 2012.
- [Skou19] Panteleimon Efstratios Skoulakis, *Natural and Efficient Dynamics through Convex Optimization*, Ph.D. thesis, Electrical and Computer Engineering, 2019.
- [Smit73] J. Maynard Smith and G. R. Price, "The Logic of Animal Conflict", *Nature*, vol. 246, pp. 15–18, 1973.
- [Tour11] Behrouz Touri and Angelia Nedić, "Discrete-time opinion dynamics", 2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR), pp. 1172–1176, 2011.
- [Vish21] Nisheeth K. Vishnoi, *Algorithms for Convex Optimization*, Cambridge University Press, 2021.
- [Vu21] Dong Quan Vu, Kimon Antonakopoulos and Panayotis Mertikopoulos, "Fast routing under uncertainty: Adaptive learning in congestion games with exponential weights", in *NeurIPS '21: Proceedings of the 35th International Conference on Neural Information Processing Systems*, 2021.
- [Wedi14] Edvin Wedin and Peter Hegarty, "A quadratic lower bound for the convergence rate in the one-dimensional Hegselmann-Krause bounded confidence dynamics", 2014.
- [Weis02] Gérard Weisbuch, Guillaume Deffuant, Frédéric Amblard and Jean-Pierre Nadal, "Meet, discuss, and segregate!", *Complexity*, vol. 7, no. 3, pp. 55–63, 2002.
- [Weis03] Gérard Weisbuch, Guillaume Deffuant, Frédéric Amblard and Jean-Pierre Nadal, "Interacting Agents and Continuous Opinions Dynamics", *arXiv: Disordered Systems and Neural Networks*, pp. 225–242, 2003.
- [Xiao09] Lin Xiao, "Dual Averaging Method for Regularized Stochastic Learning and Online Optimization", in Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams and A. Culotta, editors, *Advances in Neural Information Processing Systems*, vol. 22, Curran Associates, Inc., 2009.
- [Yild13] Ercan Yildiz, Asuman Ozdaglar, Daron Acemoglu, Amin Saberi and Anna Scaglione, "Binary Opinion Dynamics with Stubborn Agents", *ACM Trans. Econ. Comput.*, vol. 1, no. 4, dec 2013.
- [Zhan15] Jiangbo Zhang and Ge Chen, "Convergence rate of the asymmetric Deffuant-Weisbuch dynamics", *Journal of Systems Science and Complexity*, vol. 28, pp. 773–787, 2015.
- [Zhou18] Xingyu Zhou, "On the Fenchel Duality between Strong Convexity and Lipschitz Continuous Gradient", 2018.

Παράρτημα Α

Omitted Proofs and Calculations

Arithmetic counterexample (from section 6.2.4) We begin by reminding the parameters of the game:

- N = 5 players
- $\rho_i = \frac{1}{2}$ for all i
- s = (0.9, 0.84, 0.83, 0.2, 0.18)
- normalized independent weights with reciprocal generating functions, i.e.:

$$f_i(d) = \frac{1}{d + 0.0001}$$

Furthermore, we consider the following two points of the opinion space:

$$x^1 = (0.84, 0.89, 0.8, 0.49, 0.3)$$

 $x^2 = (0.83, 0.87, 0.8, 0.49, 0.3)$

We start by computing the weight matrix $Q(x) = \{q_{ij}(x)\}_{1 \leq i,j \leq N}$ on both of these points, and from that calculate the pseudogradient of the game on these points, which we write in matrix form as:

$$v(x) = 2(x - \operatorname{diag}\{1 - \rho_1, \dots, 1 - \rho_N\}Q(x)x - \operatorname{diag}\{\rho_1, \dots, \rho_N\}s)$$

So, we have:

$$Q(x^1) = \begin{pmatrix} 0 & 0.8753825 & 0.08832531 & 0.02155904 & 0.01473315 \\ 0.997044213 & 0 & 0.00248639455 & 0.000284788407 & 0.000184603631 \\ 0.82182545 & 0.13811043 & 0 & 0.02440587 & 0.01565825 \\ 0.1065098 & 0.09879281 & 0.11360927 & 0 & 0.68108812 \\ 0.19520722 & 0.18146217 & 0.20779921 & 0.4155314 & 0 \end{pmatrix}$$

and

$$Q(x^2) = \begin{pmatrix} 0 & 0.702121785 & 0.211127530 & 0.0515334449 & 0.0352172400 \\ 0.769574434 & 0 & 0.193832962 & 0.0222013761 & 0.0143912272 \\ 0.997032393 & 0.00248636507 & 0 & 0.000293158598 & 0.000188083832 \\ 0.107700542 & 0.101271618 & 0.113084671 & 0 & 0.677943169 \\ 0.196585808 & 0.185191181 & 0.206096491 & 0.412126520 & 0 \end{pmatrix}$$

from which we get:

$$Q(x^1)x^1 = \begin{bmatrix} 0.86473455 \\ 0.83970118 \\ 0.82990801 \\ 0.47260768 \\ 0.69532515 \end{bmatrix} \text{ and } Q(x^2)x^2 = \begin{bmatrix} 0.81556454 \\ 0.80900919 \\ 0.8299001 \\ 0.47134844 \\ 0.69110174 \end{bmatrix}$$

so, finally, with $\rho_i = \frac{1}{2}$ for all i, we get:

$$v(x^{1}) = 2 \begin{pmatrix} \begin{bmatrix} 0.84 \\ 0.89 \\ 0.8 \\ 0.49 \\ 0.3 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 0.86473455 \\ 0.83970118 \\ 0.82990801 \\ 0.47260768 \\ 0.69532515 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 0.9 \\ 0.84 \\ 0.83 \\ 0.2 \\ 0.18 \end{bmatrix} \right) = \begin{bmatrix} -0.08473455 \\ 0.10029882 \\ -0.05990801 \\ 0.30739232 \\ -0.27532515 \end{bmatrix}$$

and

$$v(x^{2}) = 2 \begin{pmatrix} \begin{bmatrix} 0.83 \\ 0.87 \\ 0.8 \\ 0.49 \\ 0.3 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 0.81556454 \\ 0.80900919 \\ 0.8299001 \\ 0.47134844 \\ 0.69110174 \end{bmatrix} - \frac{1}{2} \begin{bmatrix} 0.9 \\ 0.84 \\ 0.83 \\ 0.2 \\ 0.18 \end{bmatrix} \right) = \begin{bmatrix} -0.05556454 \\ 0.090999081 \\ -0.0599001 \\ 0.30865156 \\ -0.27110174 \end{bmatrix}$$

which can now use to determine if diagonal strict concavity holds:

$$\langle v(x^1) - v(x^2), x^1 - x^2 \rangle = \left\langle \begin{bmatrix} -0.0291700099 \\ 0.00930801053 \\ -0.00000791657539 \\ -0.00125923916 \\ -0.00422341428 \end{bmatrix}, \begin{bmatrix} 0.01 \\ 0.02 \\ 0 \\ 0 \\ 0 \end{bmatrix} \right\rangle = -0.00010554 < 0$$

which means that the definition of diagonal strict concavity does not hold for those two opinion profiles.

Παράρτημα Β


Simulation Plots


B.1 Convergence to Equilibrium for Large Numbers of Players

In this section, we present all graphical results from the simulations we conducted on games with up to several thousands of players, in order to keep the main body of the text decluttered, since we simulated a lot of games.


We have gathered here both the plots that focus on the convergence of the algorithm and those that focus on the speed of convergence.

Specifically, for each weight function we tried, we first show the distance to minimum costs at the end of the algorithm and the whole time evolution of that distance, for a select subset of the games run. Then, we show both the number of steps and the running length that the algorithm needed before reaching an ε -neighbourhood of the equilibrium, as a function of $1/\varepsilon$; again, for a select subset of games, of course.







Σχήμα B.1: Convergence for simple exponential weights




Σχήμα B.2: Simple exponential weights, convergence speed




Σχήμα B.3: Convergence for adjusted exponential weights




Σχήμα B.4: Adjusted exponential weights, convergence speed




Σχήμα B.5: Convergence for linear weights




Σχήμα B.6: Linear weights, convergence speed




Σχήμα B.7: Convergence for logarithmic weights




Σχήμα B.8: Logarithmic weights, convergence speed




Σχήμα B.9: Convergence for adjusted logarithmic weights




Σχήμα B.10: Adjusted logarithmic weights, convergence speed




Σχήμα B.11: Convergence for logistic weights




Σχήμα B.12: Logistic weights, convergence speed




Σχήμα B.13: Convergence for reciprocal weights




Σχήμα B.14: Reciprocal weights, convergence speed




Σχήμα B.15: Convergence for adjusted reciprocal weights



Σχήμα B.16: Adjusted reciprocal weights, convergence speed



Σχήμα B.17: Convergence for square root weights



Σχήμα B.18: Square root weights, convergence speed