National Technical University of Athens
School of Electrical and Computer Engineering
MSc Data Science & Machine Learning

Network Team Growth Dynamics

Diploma Thesis

of

NIKOS THEOLOGIS

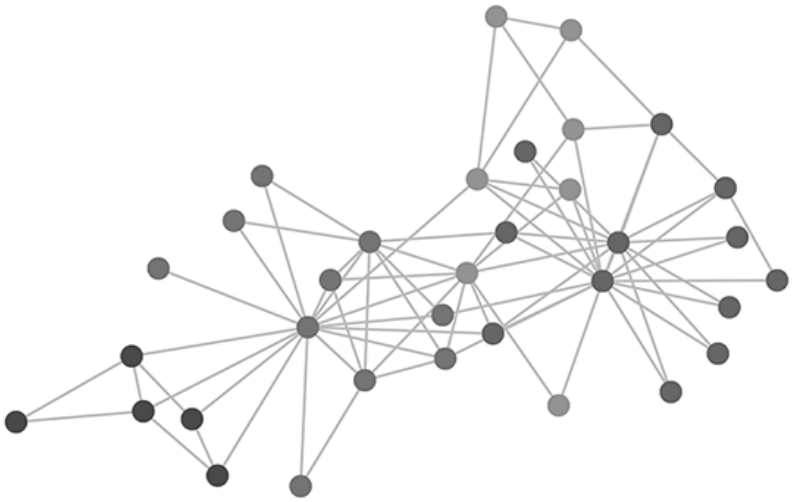

Supervisors: Evimaria Terzi, Dimitris Fotakis
Professors: BU, NTUA

Athens, June 2024

National Technical University of Athens
School of Electrical and Computer Engineering
MSc Data Science & Machine Learning

Network Team Growth Dynamics

Diploma Thesis

of

NIKOS THEOLOGIS

Supervisors: Evimaria Terzi, Dimitris Fotakis
Professors: BU, NTUA

Approved by the examination committee on 25th June, 2024.

(Signature)	*(Signature)*	*(Signature)*
.....	.....	.....
Dimitris Fotakis	Evimaria Terzi	Aris Pagourtzis
Professor, NTUA	Professor, BU	Professor, NTUA

Athens, June 2024

National Technical University of Athens
School of Electrical and Computer Engineering
MSc Data Science & Machine Learning

DISCLAIMER ON ACADEMIC ETHICS AND INTELLECTUAL PROPERTY RIGHTS

Being fully aware of the implications of copyright laws, I expressly state that this diploma thesis, as well as the electronic files and source codes developed or modified in the course of this thesis, are solely the product of my personal work and do not infringe any rights of intellectual property, personality and personal data of third parties, do not contain work / contributions of third parties for which the permission of the authors / beneficiaries is required and are not a product of partial or complete plagiarism, while the sources used are limited to the bibliographic references only and meet the rules of scientific citing. The points where I have used ideas, text, files and / or sources of other authors are clearly mentioned in the text with the appropriate citation and the relevant complete reference is included in the bibliographic references section. I fully, individually and personally undertake all legal and administrative consequences that may arise in the event that it is proven, in the course of time, that this thesis or part of it does not belong to me because it is a product of plagiarism.

(Signature)

.....
Nikos Theologis

25th June, 2024

Abstract

Previous research has convincingly demonstrated that in organizational settings, teams characterized by a diverse range of information and perspectives tend to outperform their homogeneous counterparts. Despite this evidence, why do we frequently observe predominantly homogeneous teams in practice? One prevailing explanation posits that the advantages of informational diversity are in tension with affinity bias. To delve deeper into the implications of this conflict on team composition, we study a sequential model of team formation. In this model, individuals prioritize their team's performance, as measured by its ability to accurately predict future outcomes based on various features, while also considering the potential costs associated with interacting with teammates who employ different approaches to the prediction task. Our work extends this initial team formation model by adding an underlying graph structure that changes how both the accuracy of the team and the disagreement between team members are calculated. We study two different graph structures. The first is a random undirected graph for which we have the freedom of changing and adjusting the edges in order to reach the optimal cost, while the second is a rigid hierarchical pyramid structure in which the edges are fixed in place, allowing us only the freedom to optimally position the agents within the pyramid. These extensions keep the tension between informational diversity and affinity bias, which we aim to optimize either by ensuring the optimal connections within the team or by strategically positioning the team members.

Περίληψη

Προηγούμενες έρευνες έχουν δείξει ότι εντός οργανωτικών πλαισίων, οι ομάδες που διαθέτουν ποικιλία πληροφοριών και οπτικών είναι πιο αποτελεσματικές από τις ομάδες που δεν κατέχουν αντίστοιχο εύρος. Εάν αυτό το είδος πληροφοριακής ποικιλομορφίας (informational diversity) προσδίδει πλεονεκτήματα απόδοσης, γιατί βλέπουμε συχνά στην πράξη έντονα ομοιογενείς ομάδες; Ένα επιχείρημα είναι ότι τα οφέλη της πληροφοριακής ποικιλομορφίας βρίσκονται σε εναντίωση με την μεροληψία απέναντι σε ομοιομορφία απόψεων (affinity bias). Για να κατανοήσουμε τον αντίκτυπο αυτής της τριβής στη σύνθεση των ομάδων, μελετάμε ένα διαδοχικό μοντέλο σχηματισμού ομάδας στο οποίο τα μέλη νοιάζονται για την απόδοση της ομάδας τους, αλλά υπάρχει επίσης ένα κόστος όταν αλληλεπιδρούν με συμπαίκτες διαφορετικών απόψεων. Σε αυτη την διπλωματική εργασία θα επεκτείνουμε αυτό το αρχικό μοντέλο σχηματισμού ομάδας προσθέτοντας μια υποκείμενη δομή γράφου η οποία αλλάζει τον τρόπο που υπολογίζεται η συνολική άποψη, όπως και η διαφωνία μεταξύ των μελών της ομάδας. Μελετάμε δύο διαφορετικές δομές γραφημάτων. Η μια είναι ένα τυχαίο μη κατευθυνόμενο γράφημα για το οποίο έχουμε την ελευθερία της προσαρμογής των ακμών προκειμένου να επιτευχθεί το βέλτιστο κόστος, ενώ το δεύτερο είναι μια προκαθορισμένη ιεραρχική δομή πυραμίδας στην οποία οι ακμές είναι σταθερές, επιτρέποντας έτσι μόνο την ελευθερία να τοποθετούμε τους παίκτες με τον καλύτερο δυνατό τρόπο πάνω στην πυραμίδα. Αυτές οι επεκτάσεις διατηρούν την τριβή μεταξύ της πληροφοριακής ποικιλομορφίας (informational diversity) και της μεροληψίας απέναντι σε ομοιομορφία απόψεων (affinity bias), την οποία τριβή στοχεύουμε να ελαχιστοποιήσουμε είτε διασφαλίζοντας τις βέλτιστες συνδέσεις εντός του γράφου είτε τοποθετώντας στις βέλτιστες θέσεις τα μέλη της ομάδας.

our goal should be the right questions, not the answers

Acknowledgements

At first, I would like to express my sincere gratitude to Prof. Dimitris Fotakis for facilitating my collaboration with Prof. Evimaria Terzi, whose supervision was crucial for the development of this thesis. I am particularly thankful to Evimaria Terzi for her constant support since day one and for steering me in the right direction whenever she thought I needed it. Her guidance has been instrumental in the shaping of this thesis. I should also thank PhD candidate Iasonas Nikolaou for his constant willingness to help and for putting up with my never-ending questions for the past few months. Lastly, I want to thank my family for their love and encouragement and for always believing in me not only throughout my studies, but my whole life. Without their support, I wouldn't be where I am today.

Athens, June 2024

Nikos Theologis

Contents

List of Figures

Chapter 1

Introduction

1.1 English

Extensive research in the social sciences has consistently highlighted the benefits of perspective diversity within organizational contexts. Groups comprising individuals with varied perspectives tend to outperform largely homogeneous groups. This diversity fosters the availability of a broader range of insights, facilitating constructive synergies among these diverse viewpoints. As a result, team performance is enhanced [1] [2]. In the literature, this form of diversity is occasionally labeled as cognitive diversity [3]. However, we prefer the term informational diversity to underscore the idea that team members contribute novel informational resources to the problem-solving efforts. In addition to empirical observations of this phenomenon in real-world scenarios, a series of mathematical models have attempted to formalize these informational advantages. These models operate within abstract contexts where groups of agents collaborate on collective problem-solving tasks [4].

If informational diversity indeed provides performance benefits to organizational teams, why do we frequently observe predominantly homogeneous teams in practice? A prevailing argument suggests that the advantages of informational diversity clash with affinity bias, a human behavioral tendency wherein individuals gravitate towards interacting with others who share similar perspectives. This inclination is extensively documented in prior research within organizational psychology [5] [6]. Affinity bias is an aggregate effect that can result from various underlying mechanisms. For instance, individuals may exhibit a natural preference for those who share similar perspectives, struggle to evaluate those with differing viewpoints, or favor teams with fewer disagreements or whose overall stance aligns closely with their own. Each of these scenarios manifests as a form of affinity bias. In our discussion, we will concentrate on the observable outcomes of these mechanisms, encapsulated in the concept of affinity bias, without confining ourselves to a particular underlying mechanism.

The tension between informational diversity and affinity bias underlies several empirical findings, which demonstrate that teams characterized by informational diversity can produce both higher-quality solutions and lower group cohesion simultaneously [7] [8]. These findings underscore the challenge of building informationally diverse teams: while restructuring a team to include members with diverse perspectives has the potential to

boost performance, it may also decrease subjective satisfaction among participants due to affinity bias. Thus, the question arises: what is the optimal team structure? We are interested in answering this question by understanding the fundamental phenomena that emerges from this conflict between informational diversity and affinity bias.

In this thesis, we further build upon a model proposed for team formation in the presence of both informational diversity and affinity bias [9]. In particular for this model, agents forming a team are faced with a prediction task: they see instances of a prediction problem encoded by features, and they must make a prediction about some future outcome for each instance. Consider various teams engaged in different prediction tasks such as policymakers aiming to forecast policy outcomes, investors seeking to identify promising start-ups, or doctors grappling with intricate medical diagnoses. These scenarios fall under the scope of our framework. Each team member operates with an objective function comprising two key components: one evaluates the team's accuracy, while the other measures their divergence from fellow members. The balance between these components is controlled by a single parameter, enabling examination of scenarios where emphasis is placed either on team performance or team cohesion.

While the original work [9] mainly studied the process by which teams grow over time, as they decide sequentially which new members to add, we choose to focus more on the optimal connectivity of such teams by introducing an underlying graph structure. We study two different graph structures. The first is a random undirected graph for which we have the freedom of changing and adjusting the edges in order to reach the optimal cost, while the second is a rigid hierarchical pyramid structure in which the edges are fixed in place, allowing us only the freedom to optimally position the players within the pyramid. These extensions keep the tension between informational diversity and affinity bias, which we aim to optimize either by ensuring the optimal connections within the team or by strategically positioning the team members.

1.2 Ελληνικά

Εκτεταμένη έρευνα στις κοινωνικές επιστήμες έχει αναδείξει επανειλημμένα τα οφέλη της διαφορετικότητας των οπτικών εντός οργανωτικών πλαισίων. Ομάδες που περιλαμβάνουν άτομα με ποικίλες οπτικές και απόψεις τείνουν να ξεπερνούν σε μεγάλο βαθμό τις ομοιογενείς ομάδες. Αυτή η ποικιλομορφία φαίνεται να ενισχύει το ευρύτερο φάσμα γνώσεων, διευκολύνοντας έτσι στην συνέργεια μεταξύ αυτών των διαφορετικών απόψεων. Ως αποτέλεσμα, η απόδοση της ομάδας βελτιώνεται [1] [2]. Στη βιβλιογραφία, αυτή η μορφή διαφορετικότητας περιστασιακά χαρακτηρίζεται ως ποικιλομορφία αντιλήψεων (cognitive diversity) [3]. Ωστόσο, προτιμούμε τον όρο πληροφοριακή ποικιλομορφία (informational diversity) για να τονίσουμε την ιδέα ότι τα μέλη της ομάδας συνεισφέρουν νέους πληροφοριακούς πόρους και οπτικές στις προσπάθειες επίλυσης προβλημάτων. Εκτός από τις εμπειρικές παρατηρήσεις αυτού του φαινομένου σε σενάρια του πραγματικού κόσμου, μια σειρά από μαθηματικά μοντέλα έχουν προσπαθήσει να επισημοποιήσουν αυτά τα πλεονεκτήματα. Τα μοντέλα που έχουν προταθεί λειτουργούν μέσα σε αφηρημένα περιβάλλοντα όπου ομάδες παικτών συνεργάζονται σε συλλογικές εργασίες επίλυσης προβλημάτων [4].

Εάν η πληροφοριακή ποικιλομορφία (informational diversity) παρέχει πράγματι οφέλη απόδοσης στις οργανωτικές ομάδες, γιατί παρατηρούμε συχνά στην πράξη κυρίως ομοιογενείς ομάδες; Το επικρατέστερο επιχείρημα υποδηλώνει ότι τα πλεονεκτήματα της πληροφοριακής ποικιλομορφίας είναι σε σύγκρουση με την μεροληψίας απέναντι σε ομοιομορφία απόψεων (affinity bias), μια ανθρώπινη τάση συμπεριφοράς όπου τα άτομα έλκονται προς την συναναστροφή με ανθρώπους που μοιράζονται παρόμοιες απόψεις με τους ίδιους. Αυτή η τάση έχει τεκμηριωθεί εκτενώς σε προηγούμενες έρευνες στην οργανωτική ψυχολογία [5] [6]. Η μεροληψία απέναντι σε ομοιομορφία απόψεων είναι ένα γενικό αποτέλεσμα που μπορεί να προκύψει από διάφορους υποκείμενους μηχανισμούς. Για παράδειγμα, τα άτομα μπορεί να επιδεικνύουν φυσική προτίμηση για εκείνους που μοιράζονται παρόμοιες απόψεις, να δυσκολεύονται να αξιολογήσουν άτομα με διαφορετικές απόψεις, να προτιμούν ομάδες με λιγότερες διαφωνίες ή ομάδες των οποίων η συνολική στάση ευθυγραμμίζεται στενά με τη δική τους. Κάθε ένα από αυτά τα σενάρια εκδηλώνεται ως μια μορφή αυτης της μεροληψίας. Σε αυτη την εργασία, θα επικεντρωθούμε στα παρατηρήσιμα αποτελέσματα αυτών των μηχανισμών, που ενσωματώνονται στην έννοια της μεροληψίας απέναντι σε ομοιομορφία απόψεων (affinity bias), χωρίς να περιοριστούμε σε έναν συγκεκριμένο υποκείμενο μηχανισμό.

Η τριβή μεταξύ της πληροφοριακής ποικιλομορφίας (informational diversity) και της μεροληψίας απέναντι σε ομοιομορφία απόψεων (affinity bias) βασίζεται σε πολλά εμπειρικά ευρήματα, τα οποία δείχνουν ότι οι ομάδες που χαρακτηρίζονται από πληροφοριακή ποικιλομορφία μπορούν να παράγουν λύσεις υψηλού επιπέδου ενώ όμως η συνοχή της ομάδας μειώνεται [7] [8]. Αυτά τα ευρήματα υπογραμμίζουν την πρόκληση της δημιουργίας πληροφοριακά ποικιλόμορφων ομάδων: ενώ η αναδιάρθρωση μιας ομάδας ώστε να περιλαμβάνει μέλη με διαφορετικές οπτικές έχει τη δυνατότητα να ενισχύσει την απόδοση της, μπορεί επίσης να μειώσει την ομαδική συνοχή μεταξύ των συμμετεχόντων λόγω της μεροληψίας απέναντι σε ομοιομορφία απόψεων. Έτσι, τίθεται το ερώτημα: ποια είναι η βέλτιστη δομή της ομάδας; Μας ενδιαφέρει να απαντήσουμε σε αυτό το ερώτημα κατανοώντας τα θεμελιώδη φαινόμενα που προκύπτουν από αυτή τη σύγκρουση μεταξύ της πληροφοριακής ποικιλομορφίας και της μεροληψίας απέναντι σε ομοιομορφία απόψεων.

Σε αυτή τη διατριβή, επεκτείνουμε περαιτέρω ένα μοντέλο που έχει προταθεί για το σχηματισμό ομάδας παρουσία τόσο πληροφοριακής ποικιλομορφίας όσο και μεροληψίας απέναντι σε ομοιομορφία απόψεων [9]. Ειδικότερα για αυτό το μοντέλο, οι παίκτες που σχηματίζουν την ομάδα πρέπει να εκτελέσουν μια πρόβλεψη: βλέπουν περιπτώσεις ενός προβλήματος που κωδικοποιείται από χαρακτηριστικά και πρέπει να κάνουν μια πρόβλεψη για κάποιο μελλοντικό αποτέλεσμα. Για παράδειγμα μπορούμε να θεωρήσουμε, μια ομάδα υπεύθυνων μιας πολιτικής που προσπαθεί να προβλέψει το αποτέλεσμα των παρεμβάσεων τους, μια ομάδα επενδυτών που προσπαθεί να προβλέψει ποιες νεοσύστατες εταιρείες θα είναι επιτυχημένες ή μια ομάδα γιατρών που αντιμετωπίζουν μια περίπλοκη ιατρική διάγνωση. Όλα αυτά είναι σενάρια που καταγράφονται μέσα σε αυτό το πλαίσιο. Κάθε παίκτης έχει μια αντικειμενική συνάρτηση που αποτελείται από το άθροισμα δύο όρων: ο ένας όρος είναι το ποσοστό σφάλματος της ομάδας και ο άλλος όρος είναι το επίπεδο ομοιότητάς του με τα άλλα μέλη της ομάδας. Επιπλέον μια μονοδιάστατη παράμετρος ελέγχει το σχετικό βάρος αυτών των δύο όρων στην αντικειμενική συνάρτηση. Αυτή η γενική μορφή για την αντικειμενική συνάρτηση μας επιτρέπει να μελετήσουμε τα άκρα στα οποία οι παίκτες ενδιαφέρονται είτε

κυρίως για την απόδοση της ομάδας είτε κυρίως για την ομοιογένεια της ομάδας.

Ενώ η αρχική εργασία [9] μελέτησε κυρίως τη διαδικασία με την οποία αυτες οι ομάδες αναπτύσσονται με την πάροδο του χρόνου, καθώς αποφασίζουν διαδοχικά ποια νέα μέλη θα προσθέσουν, εμείς επιλέγουμε να εστιάσουμε περισσότερο στη βέλτιστη σύνθεση τέτοιων ομάδων εισάγοντας μια υποκείμενη δομή γράφου. Διαλέγουμε να μελετήσουμε δύο διαφορετικές δομές γράφου. Η πρώτη είναι ένας τυχαίος μη κατευθυνόμενος γράφος για τον οποίο έχουμε την ελευθερία να αλλάξουμε και να προσαρμόσουμε τις ακμές προκειμένου να φτάσουμε στο βέλτιστο κόστος, ενώ η δεύτερη είναι μια σταθερή ιεραρχική δομή πυραμίδας στην οποία οι ακμές είναι σταθερές στη θέση τους, επιτρέποντάς μας μόνο την ελευθερία ως προς την τοποθέτηση των παικτών εντός της πυραμίδας. Αυτές οι επεκτάσεις διατηρούν την τριβή μεταξύ της πληροφοριακής ποικιλομορφίας (informational diversity) και της μεροληψίας απέναντι σε ομοιομορφία απόψεων (affinity bias), την οποία στοχεύουμε να βελτιστοποιήσουμε είτε διασφαλίζοντας τις βέλτιστες συνδέσεις εντός της ομάδας είτε τοποθετώντας στρατηγικά κάθε μέλος της ομάδας.

Chapter 2

Background

2.1 Base Model

The ideas presented in this thesis are extensions of a model proposed for team formation in the presence of both informational diversity and affinity bias [9]. In this section, we will introduce the foundational model along with some key results. While we will not reproduce the proofs, which can be found in the original paper, we will highlight the most relevant findings that will serve as crucial building blocks in the subsequent chapters of this thesis.

Let $\mathcal{X}$ denote the set of all possible states of the world distributed according to a probability distribution $\mathcal{P}$. We assume each state of the world is described by a feature vector $\mathbf{x} = (x_1, \ldots, x_r) \in \mathcal{X}$, consisting of uncorrelated attributes $x_1, \ldots, x_r$, i.e., $\text{cov}(x_i, x_j) = 0$ for all $j \neq i$. Each state of the world, $\mathbf{x}$, leads to an outcome $y \in \mathcal{Y}$. We assume there exists a true outcome function f^*, such that for any $\mathbf{x} \in \mathcal{X}$, $y = f^*(\mathbf{x})$ is the true outcome of the world sate $\mathbf{x}$.

Imagine a group of agents, each equipped with the ability to forecast the actual outcome given the state of the world.. An agent i has a fixed predictive model of the world, denoted by $f_i : \mathcal{X} \to \mathcal{Y}$, which maps each possible state of the world, $\mathbf{x} \in \mathcal{X}$, to a predicted outcome, $\hat{y}_i = f_i(\mathbf{x})$. We will use L_i to denote the accuracy loss of agent i's predictions using the loss function $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$. Also for any two predictive models f_i, f_j, we define the level of disagreement $d_{i,j}$ between them through a distance metric, $\delta : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}^+$.

$$L_i = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\hat{y}_i, y)] = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\hat{y}_i - y)^2].$$
$$d_{i,j} = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\delta(f_i(x), f_j(x))] = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(f_i(\mathbf{x}) - f_j(\mathbf{x}))^2].$$

We assume the true function f^* is linear in the feature vector $\mathbf{x}$. Therefore, it can be decomposed into two components, corresponding to the two types' feature sets. In particular, $\forall \mathbf{x} = (x_1, \ldots, x_r) \in \mathcal{X}$:

$$f^*(\mathbf{x}) = \partial^* \cdot \mathbf{x} = \partial^*_1 x_1 + \ldots + \partial^*_k x_k + \partial_{k+1} x_{k+1} + \ldots + \partial^*_r x_r \tag{2.1}$$
$$= \partial^*_{1,\ldots,k} \cdot (x_1, \ldots, x_k) + \partial^*_{k+1,\ldots,r} \cdot (x_{k+1}, \ldots, x_r) \tag{2.2}$$

We will refer to $\partial^*_{1,\ldots,k}$ as ∂^A (since it's the accuracy-optimal weights on A's features) and

$\partial^*_{k+1,\dots,r}$ as ∂^B, so that: $f^*(\mathbf{x}) = \partial^A \cdot \mathbf{x} + \partial^B \cdot \mathbf{x}$. Additionally, for simplicity, we assume $\mathbb{E}[\mathbf{x}] = 0$. This may seem arbitrary but we can ensure this condition by standardizing all features. As such given an instance $\mathbf{x}$, individuals of each type can produce a noisy prediction f^A, f^B given ϵ_A, ϵ_B are i.i.d. noise sampled from a mean-zero Gaussian distribution with variance σ_A, σ_A.

$$f^A(\mathbf{x}) = \partial^A \cdot \mathbf{x} + \epsilon_A = \partial^A(\mathbf{x}) + \epsilon_A \tag{2.3}$$

$$f^B(\mathbf{x}) = \partial^B \cdot \mathbf{x} + \epsilon_B = \partial^B(\mathbf{x}) + \epsilon_B \tag{2.4}$$

We will also use L^A, L^B to refer to the (noise-less) accuracy loss of each type's predictive model (i.e., $L^A = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\partial^A \cdot \mathbf{x}, y)]$ and $L^B = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\partial^B \cdot \mathbf{x}, y)]$). A key formula, which is not shown or proven in the original paper [9], that connects the noise-less accuracy losses with the ∂ values is the following:

Lemma 1 (Noise-Less Losses: L^A, L^B and ∂^A, ∂^B). By initializing the parameters ∂^A and ∂^B, we gain the ability to tailor the values of L^A and L^B to align with our specific requirements, since:

$$\begin{aligned} L^A &= (\partial^B)^2 \\ L^B &= (\partial^A)^2 \end{aligned} \tag{2.5}$$

Proof. We will write the proof for L^B and the same holds for L^A:

$$\begin{aligned} L^B &= \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\partial^B \cdot \mathbf{x}, y)] = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\partial^B \cdot \mathbf{x} - \partial^* \cdot \mathbf{x})^2] = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\partial^A \cdot \mathbf{x})^2] \\ &= \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\partial^*_1 x_1 + \dots + \partial^*_k x_k)^2] \quad \text{(Using 2.2)} \\ &= \mathbb{E}[(\partial^*_1 x_1)^2] + \dots + \mathbb{E}[(\partial^*_k x_k)^2] + 2\mathbb{E}[(\cancelto{0}{\sum_{i=1}^{k-1}\sum_{j=i+1}^{k}(\partial^*_i x_i)(\partial^*_j x_j)}] \quad \text{(Using } \mathbb{E}[\mathbf{x}] = 0,\ \mathrm{cov}(x_i, x_j) = 0) \\ &= (\partial^*_1)^2 \cancelto{1}{\mathbb{E}[x_1^2]} + \dots + (\partial^*_k)^2 \cancelto{1}{\mathbb{E}[x_k^2]} = \partial^A \cdot \partial^A = (\partial^A)^2 \quad \text{(Using } \mathcal{P} = \mathcal{N}(0,1)) \end{aligned}$$

□

A team, denoted as T, consists of a collection of agents who merge their predictions using a specified aggregation function, represented by $\mathcal{G}_T$. For any $\mathbf{x} \in \mathcal{X}$, the aggregation function $\mathcal{G}_T$ receives the predictions made for input $\mathbf{x}$ by all members of T, and outputs a team prediction for $\mathbf{x}$. Given the team T, we use $\mathcal{G}_T(\mathbf{x})$ to refer to the aggregated prediction of team members for state $\mathbf{x}$. In the original work [9] a general class of aggregation functions inspired by Tullock's contest success function [10] [11] is used, defined as follows: Given a team consisting of n_A individuals of type A and n_B individuals of type B, we define the following parametric class of aggregation functions:

$$\forall a \in [0, \infty): \quad \mathcal{G}^a_{n_A,n_B}(\mathbf{x}) = \left(\frac{n_A^a}{n_A^a + n_B^a}\right) f^A(\mathbf{x}) + \left(\frac{n_B^a}{n_A^a + n_B^a}\right) f^B(\mathbf{x}) \tag{2.6}$$

Note that when $a = 1$, the expression simplifies to a straightforward average. Moreover, as $a \to \infty$, then $\mathcal{G}^a_{n_A,n_B}(\mathbf{x})$ tends towards resembling the median.

Adding individual agents to a team can help alleviate the team's overall disutility or cost. The cost an agent i incurs as a member of team T is defined as a combination of (a) their level of disagreement with other team members, and (b) the team's overall accuracy loss. More precisely,

$$c_i(T) = \lambda \times \frac{1}{|T|^\beta} \sum_{j \in T} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[\delta(f_j(\mathbf{x}), f_i(\mathbf{x}))] + (1 - \lambda) \times \mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[\ell(\mathcal{G}_T(\mathbf{x}), y)]$$

The parameter $\beta \in [0, 1]$ specifies how individual i's perception of disagreement is influenced by the team's size. Specifically, it indicates the extent to which perceptions of disagreement are influenced by the absolute versus relative size of opposing viewpoints within the team. To elucidate, consider a hypothetical scenario where $\mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[\delta(f_j(\mathbf{x}), f_i(\mathbf{x}))]$ is held constant at a value δ for all $j \neq i$. When $\beta = 0$, i perceives disagreement with team members as $(|T-1|\delta)$, and this perception scales linearly with team size $|T|$. In other words, a larger team amplifies i's sense of discord with teammates. Conversely, when $\beta = 1$, i's perception of disagreement level remains approximately constant at $(|T - 1|\delta/|T|)$.

We assume a team T would be willing to accept a new member if it reduces the team's average cost/disutility across its current members:

$$\lambda \times \frac{1}{|T|^{1+\beta}} \sum_{i,j \in T} \mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[\delta(f_j(\mathbf{x}), f_i(\mathbf{x}))] + (1 - \lambda) \times \mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[\ell(\mathcal{G}_T(\mathbf{x}), y)] \tag{2.7}$$

The team growth dynamics unfold in the following manner: Initially, Team T comprises n_A individuals of type A and n_B individuals of type B. New potential team members arrive sequentially over steps $t = 1, 2, \ldots$. Let's denote the t'th individual as i_t, with their type indicated by $s_t \in A, B$. The team incorporates i_t only if it diminishes the current team's cost as per Equation (2.7).

2.1.1 Minimum λ value ($\lambda = 0$)

When $\lambda = 0$, the team adds new members if and only if the new member reduces the team's mean squared error without considering the team's disagreement. We will now write the main results that can be derived in this case from [9].

Lemma 2 (Team's Error Decomposition). Consider a team with a composition of n_A type A members and of n_B type B members where the expectation is with respect to $(\mathbf{x}, y) \sim P$ and $\epsilon_c \sim \mathcal{N}(0, \sigma_c^2)$ for $c \in \{A, B\}$. Then:

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2] = \frac{n_A^{2a}}{(n_A^a + n_B^a)^2} L^A + \frac{n_B^{2a}}{(n_A^a + n_B^a)^2} L^B + \frac{n_A^{2a}\sigma_A^2 + n_B^{2a}\sigma_B^2}{(n_A^a + n_B^a)^2} \tag{2.8}$$

Proposition 1 (Accuracy-optimal composition). Consider a team with an initial composition of $n_A > 0$ members of type A and no member of type B. The optimal number of type

B members whose addition minimizes the team's accuracy loss in (2.7) is equal to:

$$n_B^* = n_A \left(\frac{L^A + \sigma_A^2}{L^B + \sigma_B^2} \right)^{1/a} \tag{2.9}$$

Remark 1 Note that due to the symmetry of Equation (2.7) in A and B, the partial derivatives of the team's accuracy with respect to n_A and n_B always have opposing signs, therefore, at any $n_A, n_B \geq 0$, it is either beneficial to add a new member of type A or a new member of type B, but never both. These trends remain unchanged even if the agents predictions are noisy or not.

2.1.2 Maximum β value ($\beta = 1$)

When $\beta = 1$, the team adds new members if and only if the new member reduces the team's disagreement without considering the team's overall accuracy. We will now write the main results that can be derived in this case from [9].

Lemma 3 (Team's Disagreement Decomposition). Consider a team with a composition of n_A type A members and of n_B type B members, $|T| = n_A + n_B$, where the expectation is with respect to $(\mathbf{x}, y) \sim P$ and $\epsilon_c \sim \mathcal{N}(0, \sigma_c^2)$ for $c \in \{A, B\}$. Then:

$$\frac{1}{|T|^{1+\beta}} \mathbb{E}\left[\sum_{i,j \in T} d_{i,j} \right] = \frac{2}{|T|^{1+\beta}} \left(n_A n_B (L^A + L^B + \sigma_A^2 + \sigma_B^2) + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2 \right) \tag{2.10}$$

Proposition 2 Consider a team with an initial composition of $n_A > 0$ members of type A and n_B members of type B. Suppose $\sigma_B > 0$ and $\beta < 1$. Then there exist $n_B^{lower}, n_B^{upper} \in \mathbb{R}$ such that adding a type B member reduces the team's disagreement in (2.10) if and only if $n_B^{lower} \leq n_B \leq n_B^{upper}$.

2.2 DeGroot Learning

The DeGroot learning model, stated in its general form by the American statistician Morris H. DeGroot [12], is a simple and important model of how people in a network update their opinions over time and eventually reach a group consensus. Antecedents of this model were articulated by John R. P. French [13] and Frank Harary [14]. The model has been extensively used in physics, computer science and most widely in the theory of social networks.

In the DeGroot model links between agents and the weight they put on each other's opinions is represented by a trust matrix T where T_{ij} is the weight that agent i puts on agent j's opinion. The trust matrix is thus in a one-to-one relationship with a weighted, directed graph where there is an edge between i and j if and only if $T_{ij} > 0$. The trust matrix is stochastic, meaning that its rows consists of non-negative real numbers, with each row summing up to 1. Each agent starts with an initial belief $b_i(0)$ and all the beliefs

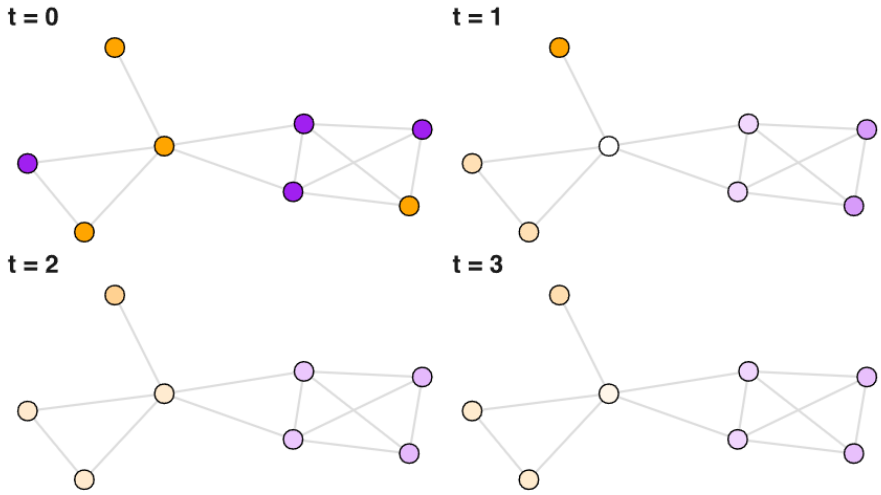

Figure 2.1. *Example of DeGroot Learning. Orange and Purple nodes represent the initial beliefs and as the rounds progress their opinions converge.*

are being updated as:

$$b_i(t) = \sum_j T_{ij} b_j(t-1)$$

An important question is whether beliefs converge to a limit and to each other in the long run. As the trust matrix is stochastic, standard results in Markov chain theory [15] can be used to state conditions under which the limit exists for any initial beliefs.

Theorem 1 (DeGroot Network Convergence). If the social network graph (represented by the trust matrix T) is strongly connected (every node is reachable from every other node), convergence of beliefs is equivalent to each of the following properties:

- T is Convergent $\Leftrightarrow$ T is aperiodic.
- T is Convergent $\Leftrightarrow \lim_{t\to\infty} T^t = [1,\dots,1]^T\mathbf{s}$

Where $\mathbf{s}$ is the unique left eigenvector of T with eigenvalue 1 whose entries sum up to 1. Moreover aperiodicity is easy to satisfy by simply adding a value in the diagonal of matrix T so that even one agent takes into account his own opinion during updates.

Proof. Suppose matrix T is strongly connected.

Definition: Matrix A is primitive if and only if $\exists t_0 : A^t_{ij} > 0, \forall t \geq t_0$.

$\Leftarrow$

Given T is aperiodic, using [16] we know that if matrix A is strongly connected and stochastic (true for matrix T) then aperiodicity $\Leftrightarrow$ primitive. Moreover from [17] we know that if matrix A is strongly connected and primitive then: $\lim_{t\to\infty} T^t = [1,\dots,1]^T\mathbf{s}$, where $\mathbf{s}$ is the unique left hand side eigenvector with eigenvalue 1 and all positive entries, giving us this way convergence.

$\Rightarrow$
Let $S = \lim_{t\to\infty} T^t$ by convergence. Then $ST = \lim_{t\to\infty} T^t T = S$. From the Perron-Frobenius theorem: An eigenvector of an irreducible (i.e. strongly connected) non-negative matrix is strictly positive if and only if it is associated with its largest eigenvalue. This vector is unique if the matrix is primitive. So since S is all positive $\Rightarrow$ T is primitive which from [16] gives us aperiodicity.

□

Lemma 4 (DeGroot Network Consensus). With a strongly connected and aperiodic network the whole group reaches a consensus. This means that each agent's final belief will be the same. Defining $\mathbf{b}$ as the vector of all the initial opinions $b_i(0)$, the final beliefs of all the agents in the network will be equal to:

$$\lim_{t\to\infty} b_i(t) = \mathbf{s}\cdot\mathbf{b} \tag{2.11}$$

A very simple version of the DeGroot model that we will implement makes the following assumptions. Firstly, to make sure that the network converges we make sure that all agents listen to themselves while updating, ensuring aperiodicity by making the diagonal values of the transition matrix $T_{ii} = \epsilon$. This hyperparameter ϵ can be though of as the stubbornness of the agents. Also we assume that $T_{ij} > 0$ if and only if $T_{ji} > 0$. This makes sense in the frame of mutual friendships or workplace relationships. Note here that T_{ij} does not have to be equal to T_{ji} thus the matrix T is not necessarily symmetric. Now also suppose that the agents equally weight their connections so that $T_{ij} = (1-\epsilon)/d_i$, where d_i is the agents outward degree. Finally lets define the total number of outward degrees as: $D = \sum_k d_k$. Assuming all of the above we can claim the following.

Theorem 2 (DeGroot Network Influence). Each agent's influence, defined as their corresponding element in the eigenvector $\mathbf{s}$, for the above simplified version, is just proportional to their degree:

$$s_i = d_i/D \tag{2.12}$$

Proof. Note that in the DeGroot Model because of convergence: $\mathbf{s}\cdot\mathbf{b} = (\mathbf{s}T)\cdot\mathbf{b} \Rightarrow \mathbf{s} = \mathbf{s}T$. Now lets verify this for our version of the model using the fact that $T_{ij} > 0 \Leftrightarrow T_{ji} > 0$, $T_{ii} = \epsilon$ and $T_{ij} = (1-\epsilon)/d_i$:

$$s_i = \sum_j T_{ji}s_j = s_i\epsilon + \sum_{i\neq j} \frac{1-\epsilon}{\cancel{d_j}}\cdot\frac{\cancel{d_j}}{D} = \frac{d_i}{D}\cdot\epsilon + \frac{d_i}{D}(1-\epsilon) = \frac{d_i}{D}$$

□

Chapter 3

Graph Models

In this chapter, we expand upon the findings of the base model by incorporating an underlying graph structure into the team dynamics. Each team member is conceptualized as a node within a graph, with the edges representing the interconnected network of workplace relationships and communication channels. The motivation behind this extension is to generalize the team's opinion formation process beyond the simple Tullock's aggregate function (2.6) to a more realistic framework under the DeGroot learning model.

3.1 Undirected Graph

The DeGroot model, as discussed in Chapter 2, serves as the foundation for our analysis. We implement a version of the DeGroot model outlined in the previous chapter, allowing us to apply the result stated in Equation (2.12), which asserts that the influence of a team member is proportionate to the corresponding node's degree. To facilitate this analysis, we construct a random undirected graph denoted as $G(n_A, n_B, p, q)$, where edges occur independently with probabilities $0 < p < 1$ between nodes of the same type and $0 < q < 1$ between nodes of different types. This enables us to adjust the connectivity between agents of the same and different types. Each node or agent initially holds an opinion $b_i(0) = f^{A,B}(\mathbf{x})$ based on their type, and they converge to a final consensus opinion, as determined by Equation (2.11). Moving forward it will be helpful to analyze the two parts of the total Cost function (2.7) separately as:

$$Cost^T = \lambda \times Cost^D + (1 - \lambda) \times Cost^A \tag{3.1}$$

where the Accuracy Cost is represented by: $Cost^A = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\mathcal{G}_T(\mathbf{x}), y)]$ and the Disagreement Cost by: $Cost^D = \frac{1}{|T|^{1+\beta}} \sum_{i,j\in T} \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\delta(f_j(\mathbf{x}), f_i(\mathbf{x}))]$.

3.1.1 Accuracy

Let's begin by analyzing how the Accuracy Cost changes with the addition of the underlying graph. In order to calculate the Accuracy Cost: $Cost^A = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\mathcal{G}_T(\mathbf{x}), y)]$ we must first calculate how the aggregate opinion function changes compared to the initial Tullock aggregation function (2.6).

Lemma 5 (Undirected Graph Aggregation Function). Consider a team of n_A members of type A and n_B members of type B connected on a specific graph $G(n_A, n_B)$. Then:

$$\mathcal{G}_G^{n_A,n_B}(\mathbf{x}) = \left(\frac{d_A}{D}\right) f^A(\mathbf{x}) + \left(\frac{d_B}{D}\right) f^B(\mathbf{x}) \tag{3.2}$$

Where $d_A = \sum_{i \in N_A} d_i$ and $d_B = \sum_{i \in N_B} d_i$, with N_A, N_B being the sets of agents A, B.

Proof. All the agents will converge to the same final opinion according to (2.11) and as such we define this final opinion to be the aggregate:

$$\begin{aligned}
\mathcal{G}_G^{n_A,n_B}(\mathbf{x}) &= \mathbf{s} \cdot \mathbf{b} \\
&= [\frac{d_1}{D}, \ldots, \frac{d_{n_A}}{D}, \frac{d_{n_A+1}}{D}, \ldots, \frac{d_{n_A+n_B}}{D}] \cdot [f^A(\mathbf{x}), \ldots, f^A(\mathbf{x}), f^B(\mathbf{x}), \ldots, f^B(\mathbf{x})] \quad \text{(Using 2.12)} \\
&= \left(\frac{\sum_{i \in N_A} d_i}{D}\right) f^A(\mathbf{x}) + \left(\frac{\sum_{i \in N_B} d_i}{D}\right) f^B(\mathbf{x}) = \left(\frac{d_A}{D}\right) f^A(\mathbf{x}) + \left(\frac{d_B}{D}\right) f^B(\mathbf{x})
\end{aligned}$$

□

This new Aggregation function has a lot of resemblance with the original Tullock Aggregate (2.6) used in [9] with the changes being: $n_A^a \to d_A = \sum_{i \in N_A} d_i$, $n_B^a \to d_B = \sum_{i \in N_B} d_i$, $n_A^a + n_B^a \to D$. Using this Aggregation function now lets see how the Accuracy Cost changes compared to its initial form (2.8).

Lemma 6 (Undirected Graph Accuracy Cost). Consider a team of n_A members of type A and n_B members of type B connected on a specific graph $G(n_A, n_B)$. The Accuracy Cost of the network according to (2.7) will be:

$$Cost^A = \mathbb{E}_{\mathbf{x} \sim \mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2] = \left(\frac{d_A}{D}\right)^2 L^A + \left(\frac{d_B}{D}\right)^2 L^B + \frac{(d_A \sigma_A)^2 + (d_B \sigma_B)^2}{D^2} \tag{3.3}$$

Proof. We can write:

$$\begin{aligned}
\mathbb{E}[(\mathcal{G}_T(\mathbf{x}) - y)^2] &= \mathbb{E}[(\mathcal{G}_T(\mathbf{x}) - f^*(\mathbf{x}))^2] \\
&= \mathbb{E}\left[\left(\frac{d_A}{D}\partial^A(\mathbf{x}) + \frac{d_B}{D}\partial^B(\mathbf{x}) + \frac{d_A\epsilon_A + d_B\epsilon_B}{D} - f^*(\mathbf{x})\right)^2\right] \quad \text{(Using 2.3, 2.4)} \\
&= \mathbb{E}\left[\left(\frac{-d_B}{D}\partial^A(\mathbf{x}) + \frac{-d_A}{D}\partial^B(\mathbf{x}) + \frac{d_A\epsilon_A + d_B\epsilon_B}{D}\right)^2\right] \quad \text{(Using 2.1)} \\
&= \left(\frac{d_B}{D}\right)^2 \mathbb{E}\left[\partial^A(\mathbf{x})^2\right] + \left(\frac{d_A}{D}\right)^2 \mathbb{E}\left[\partial^B(\mathbf{x})^2\right] \\
&+ \left(\frac{2d_A d_B}{D^2}\right) \mathbb{E}\left[\partial^A(\mathbf{x})\partial^B(\mathbf{x})\right] + \left(\frac{1}{D^2}\right) \mathbb{E}\left[(d_A\epsilon_A + d_B\epsilon_B)^2\right]
\end{aligned}$$

Next, using the fact that the two types don't have access to common features, and the fact that $\text{cov}(x_i, x_j) = 0, \forall j \neq i$, we know $\mathbb{E}[\partial_A(x)\partial_B(x)] = 0$. Additionally, $\mathbb{E}[\epsilon_A \epsilon_B] = 0$ due to the assumption of independent noises. Therefore, the above equation can be simplified to:

$$
\begin{aligned}
&= \left(\frac{d_B}{D}\right)^2 \mathbb{E}\left[\left(f^*(\mathbf{x}) - \partial^B(\mathbf{x})\right)^2\right] + \left(\frac{d_A}{D}\right)^2 \mathbb{E}\left[\left(f^*(\mathbf{x}) - \partial^A(\mathbf{x})\right)^2\right] \\
&+ \left(\frac{d_A}{D}\right)^2 \mathbb{E}\left[\epsilon_A^2\right] + \left(\frac{d_B}{D}\right)^2 \mathbb{E}\left[\epsilon_B^2\right] \\
&= \left(\frac{d_B}{D}\right)^2 L^B + \left(\frac{d_A}{D}\right)^2 L^A + \frac{(d_A\sigma_A)^2 + (d_B\sigma_B)^2}{D^2}
\end{aligned}
$$

□

Where the expectation is with respect to $(\mathbf{x}, y) \sim P$ and $\epsilon_c \sim \mathcal{N}(0, \sigma_c^2)$ for $c \in \{A, B\}$. Again we can see that this new Accuracy Cost function has a lot of resemblance with the decomposed one in the case of $\lambda = 0$ (2.8) used in [9] with the changes again being: $n_A^a \to d_A = \sum_{i \in N_A} d_i$, $n_B^a \to d_B = \sum_{i \in N_B} d_i$, $n_A^a + n_B^a \to D$.

Now let's examine the behavior of the network $G(n_A, n_B, p, q)$ in its two limit cases and compare them to the original results presented in Proposition 1 (2.9) and Remark 1. It's important to note that even for random graphs G with identical values of n_A, n_B, p, and q, the randomness associated with the creation of edges can result in different values for d_A, d_B, and D. Consequently, this variability affects the values of the Aggregate function (3.2) and of the Accuracy Cost (3.3). Explicit calculation of the Aggregate function and the Accuracy Cost is feasible only in two limit cases: $p = q = 1$ and $p = 1$, $q = 0$.

Lemma 7 (Complete Undirected Graph Accuracy). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$ with $p = q = 1$. The Aggregate function of this network as well as the Cost function will be the same with the Tullock Aggregate (2.6) for $a = 1$, and it's corresponding Cost (2.8).

Proof. Since the Graph is Complete, which means that every node is connected with every other node it is true that: $d_i = n_A + n_B - 1, \forall i$. As such $d_A = \sum_{i \in N_A} d_i = n_A(n_A + n_B - 1)$, $d_B = \sum_{i \in N_B} d_i = n_B(n_A + n_B - 1)$ and $D = (n_A + n_B)(n_A + n_B - 1)$. So we can write:

$$
\begin{aligned}
\mathcal{G}^{n_A, n_B}_{G(p=q=1)}(\mathbf{x}) &= \frac{n_A(n_A + n_B - 1)}{(n_A + n_B)(n_A + n_B - 1)} f^A(\mathbf{x}) + \frac{n_B(n_A + n_B - 1)}{(n_A + n_B)(n_A + n_B - 1)} f^B(\mathbf{x}) \\
&= \frac{n_A}{n_A + n_B} f^A(\mathbf{x}) + \frac{n_B}{n_A + n_B} f^B(\mathbf{x}) = \mathcal{G}^{a=1}_{n_A, n_B}(\mathbf{x}) \quad \text{(Using 2.6)}
\end{aligned}
$$

□

Using the Tullock's Aggregate function with $a = 1$, which is reproduced in the case of the Complete graph, we can show that the results of Lemma 2 (2.8), Proposition 1 (2.9) and Remark 1 still hold in the case of the complete undirected Graph. This result is expected since in the base model proposed in [9] there is no underlying graph and as such every team member is connected with every other team member. Let us now study the second limit case of the undirected graph, where type A and type B nodes are fully connected among themselves but do not listen to nodes of the different type.

Lemma 8 (Disconnected Undirected Graph Accuracy). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$ with $p = 1, q = 0$. The Aggregate function as well as the Accuracy Cost function will be:

$$\mathcal{G}^{n_A,n_B}_{G(p=1,q=0)}(\mathbf{x}) = \frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)} f^A(\mathbf{x}) + \frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)} f^B(\mathbf{x})$$

$$\mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}^{n_A,n_B}_{G(p=1,q=0)}(\mathbf{x}) - y)^2] = \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^A + \sigma_A^2) + \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^B + \sigma_B^2)$$

Proof. Since $p = 1$ and $q = 0$, the nodes that are connected are only of the same type. An important note here is that in this case, the DeGroot model does not converge because the graph is not strongly connected, as is required in Theorem 1. We can solve this by adding just one edge between two different types of nodes, changing this way the Trust matrix T as well as the final opinion (2.11) by a small margin. This change, however, is minuscule, and for large networks, it adds up almost to zero while saving the convergence of the DeGroot network. As such, for $p = 1$ and $q \approx 0$, we now have that: $d_i = n_A - 1, \forall i \in N_A$ and $d_i = n_B - 1, \forall i \in N_B$. As such: $d_A = \sum_{i\in N_A} d_i = n_A(n_A-1)$, $d_B = \sum_{i\in N_B} d_i = n_B(n_B-1)$, and $D = n_A(n_A-1) + n_B(n_B-1)$. So we can calculate the value for the Aggregate function as:

$$\begin{aligned}\mathcal{G}^{n_A,n_B}_{G(p=1,q=0)}(\mathbf{x}) &= \left(\frac{d_A}{D}\right) f^A(\mathbf{x}) + \left(\frac{d_B}{D}\right) f^B(\mathbf{x}) \\ &= \frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)} f^A(\mathbf{x}) + \frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)} f^B(\mathbf{x})\end{aligned}$$

Using now this new Aggregate function we can calculate the Accuracy Cost function following similar steps as before:

$$\begin{aligned}&\mathbb{E}[\mathcal{G}^{n_A,n_B}_{G(p=1,q=0)}(\mathbf{x}) - y)^2] = \mathbb{E}[(\mathcal{G}_T(\mathbf{x}) - f^*(\mathbf{x}))^2] = \\ &= \mathbb{E}\left[\left(\frac{-n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\partial^A(\mathbf{x}) + \frac{-n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\partial^B(\mathbf{x}) + \frac{n_A(n_A-1)\epsilon_A+n_B(n_B-1)\epsilon_B}{n_A(n_A-1)+n_B(n_B-1)}\right)^2\right] \quad \text{(Using 2.1, 2.3, 2.4)} \\ &= \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 \mathbb{E}\left[\partial^A(\mathbf{x})^2\right] + \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 \mathbb{E}\left[\partial^B(\mathbf{x})^2\right] + \\ &+ \frac{2n_A(n_A-1)n_B(n_B-1)}{(n_A(n_A-1)+n_B(n_B-1))^2}\mathbb{E}\left[\partial^A(\mathbf{x})\partial^B(\mathbf{x})\right] + \\ &+ \frac{1}{(n_A(n_A-1)+n_B(n_B-1))^2}\mathbb{E}\left[(n_A(n_A-1)\epsilon_A + n_B(n_B-1)\epsilon_B)^2\right]\end{aligned}$$

Next, using the fact that the two types don't have access to common features, and the fact that $\text{cov}(x_i, x_j) = 0, \forall j \neq i$, we know $\mathbb{E}[\partial_A(x)\partial_B(x)] = 0$. Additionally, $\mathbb{E}[\epsilon_A\epsilon_B] = 0$ due to the assumption of independent noises. Therefore, the above equation can be simplified to:

$$\begin{aligned}&= \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 L^B + \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 L^A + \\ &+ \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 \mathbb{E}[\epsilon_B^2] + \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 \mathbb{E}[\epsilon_A^2] \\ &= \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^A + \sigma_A^2) + \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^B + \sigma_B^2)\end{aligned}$$

Where the expectation is with respect to $(\mathbf{x}, y) \sim P$ and $\epsilon_c \sim \mathcal{N}(0, \sigma_c^2)$ for $c \in \{A, B\}$. Again we can see that this new Cost function has a lot of resemblance with the one in the case of $\lambda = 0$ (2.8) used in [9] with the changes this time being: $n_A^a \to n_A(n_A - 1)$, $n_B^a \to n_B(n_B - 1)$. □

Proposition 3 (Disconnected Undirected Graph Accuracy-optimal composition). Consider a Disconnected team with an initial composition of $n_A > 0$ members of type A and no member of type B. The optimal number of type B members whose addition minimizes the team's Accuracy Cost is equal to:

$$n_B^* = \frac{1}{2} + \frac{1}{2}\sqrt{1 + 4n_A(n_A - 1)\frac{L^A + \sigma_A^2}{L^B + \sigma_B^2}} \tag{3.4}$$

Proof. According to Lemma 8 the Disconnected team's accuracy can be written as:

$$\mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}^{n_A,n_B}_{G(p=1,q=0)}(\mathbf{x}) - y)^2] = \left(\frac{n_A(n_A-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^A + \sigma_A^2) + \left(\frac{n_B(n_B-1)}{n_A(n_A-1)+n_B(n_B-1)}\right)^2 (L^B + \sigma_B^2)$$

Taking the derivative of the right hand side with respect to n_B, we obtain:

$$\frac{2n_A^2(n_A-1)^2(2n_B-1)}{(n_A(n_A-1)+n_B(n_B-1))^3}(L^A + \sigma_A^2) - \frac{2n_An_B(n_B-1)(2n_B-1)(n_A-1)}{(n_A(n_A-1)+n_B(n_B-1))^3}(L^B + \sigma_B^2) =$$
$$= \frac{2n_A(n_A-1)(2n_B-1)}{(n_A(n_A-1)+n_B(n_B-1))^3}\Big(n_A(n_A-1)(L^A + \sigma_A^2) - n_B(n_B-1)(L^B + \sigma_B^2)\Big)$$

To obtain the zero of the derivative, we can write:

$$n_B^2 - n_B - n_A(n_A-1)\frac{L^A + \sigma_A^2}{L^B + \sigma_B^2} = 0 \Leftrightarrow n_B = \frac{1}{2} + \frac{1}{2}\sqrt{1 + 4n_A(n_A-1)\frac{L^A + \sigma_A^2}{L^B + \sigma_B^2}}$$

□

We must note here that this value of n_B^* from *Proposition* 3 (3.4) is the same as the value of n_B^* from *Proposition* 1 (2.9) only in the case of: $L^A + \sigma_A^2 = L^B + \sigma_B^2$, where for the complete and for the disconnected graph the optimal team composition is: $n_B^* = n_A$. It is also interesting to point out the different team growth dynamics for these two n_B^* values. We will make this comparison along with some simulated results in Chapter 4.

3.1.2 Disagreement

Let's continue by analyzing how the Disagreement Cost changes with the addition of the underlying graph. By using the DeGroot learning process all agents reach the same final opinion and that can be thought of as canceling out the teams' disagreement. However we want to preserve the tension between Accuracy and Disagreement Costs just as in the base model [9]. One approach to accomplish this is by stipulating that, in our extended model, disagreement occurs solely between agents directly connected in the underlying graph—meaning their nodes are directly linked by an edge. By adopting this criterion, we can reintroduce the Disagreement Cost term:

Lemma 9 (Undirected Graph Disagreement Cost). Consider a team of n_A members of type A and n_B members of type B connected on a specific graph $G(n_A, n_B, p, q)$. Then:

$$Cost^D = \frac{1}{|T|^{1+\beta}} \sum_{i,j \in T} \mathbb{E}[(f_j(\mathbf{x}) - f_i(\mathbf{x}))^2] = \frac{2}{|T|^{1+\beta}}(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2) \quad (3.5)$$

Where $d_A = d_{AA} + d_{AB}$ and $d_B = d_{BB} + d_{AB}$ and d_{AA}, d_{BB} represent the edges between nodes of the same type (i.e. A to A or B to B) and d_{AB} represent the edges between nodes of different type (i.e. A to B). The expectation is with respect to $(\mathbf{x}, y) \sim \mathcal{P}$ and $\epsilon_c, \epsilon_c' \sim \mathcal{N}(0, \sigma_c^2)$ for $c \in \{A, B\}$

Proof. We can write the left hand side of (3.5) as follows:

$$\begin{aligned}
&\frac{1}{|T|^{1+\beta}} \sum_{i,j \in T} \mathbb{E}[(f_j(\mathbf{x}) - f_i(\mathbf{x}))^2] = \frac{1}{|T|^{1+\beta}} \sum_{i,j \in T} \mathbb{E}[(\partial_i(\mathbf{x}) + \epsilon_i - \partial_j(\mathbf{x}) - \epsilon_j)^2] \\
&= \frac{1}{|T|^{1+\beta}} \left(2d_{AB}\mathbb{E}[(\partial^A(\mathbf{x}) + \epsilon^A - \partial^B(\mathbf{x}) - \epsilon^B)^2] + d_{AA}\mathbb{E}[(\epsilon_A - \epsilon_A')^2] + d_{BB}\mathbb{E}[(\epsilon_B - \epsilon_B')^2]\right) \\
&= \frac{1}{|T|^{1+\beta}}(2d_{AB}\mathbb{E}[(\partial^A(\mathbf{x}) - \partial^B(\mathbf{x}))^2] + 2d_{AB}\mathbb{E}[(\epsilon^A - \epsilon^B)^2] \\
&\qquad + d_{AA}(\mathbb{E}[\epsilon_A^2] + \mathbb{E}[\epsilon_A'^2]) + d_{BB}(\mathbb{E}[\epsilon_B^2] + \mathbb{E}[\epsilon_B'^2])) \\
&= \frac{1}{|T|^{1+\beta}}(2d_{AB}(L^A + L^B) + 2d_{AB}(\mathbb{E}[\epsilon_A^2] + \mathbb{E}[\epsilon_B^2]) + 2d_{AA}\sigma_A^2 + 2d_{BB}\sigma_B^2) \\
&= \frac{2}{|T|^{1+\beta}}(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2)
\end{aligned}$$

□

Again we can see that this new Disagreement Cost function has a lot of resemblance with the base model's decomposed one in the case of $\lambda = 1$ (2.10) used in [9] with the changes being: $n_A n_B \to d_{AB}$, $n_A(n_A - 1) \to d_{AA}$, $n_B(n_B - 1) \to d_{BB}$.

Now lets see how the two limit cases of the network $G(n_A, n_B, p, q)$ behave in the case of the Disagreement Cost and also compare them to the original results of [9]. We must note again that even for random graphs G with the same n_A, n_B, p, q values because of the randomness associated with the creation of the edges in the network the d_{AA}, d_{BB} and d_{AB} values will be different and as such the values of the Disagreement Cost (3.5) will also be different. The only 2 limit cases where we can explicitly calculate the Disagreement Cost are for $p = q = 1$ and for $p = 1, q = 0$.

Lemma 10 (Complete Undirected Graph Disagreement). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$ with $p = q = 1$. The teams Disagreement Cost (2.10) as well as Proposition 2 still hold just as in the base model.

Proof. Since the Graph is Complete, which means that every agent/node is connected with every other agent/node it is true that: $d_i = n_A + n_B - 1, \forall i$. As such $d_{AA} = n_A(n_A - 1)$, $d_{BB} = n_B(n_B - 1)$ and $d_{AB} = n_A n_B$. So (3.5) becomes:

$$Cost^D_{(p=q=1)} = \frac{2}{(n_A + n_B)^{1+\beta}} \left(n_A n_B(L^A + L^B + \sigma_A^2 + \sigma_B^2) + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2\right)$$

□

Lemma 11 (Disconnected Undirected Graph Disagreement). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$ with $p = 1$, $q \approx 0$, where $q \approx 0$ implies that we let only 1 edge to connect type A and type B agents in order to keep the graph strongly connected. The teams Disagreement Cost decomposition for this limit case becomes:

$$Cost^D_{(p=1,q\approx 0)} = \frac{2}{(n_A + n_B)^{1+\beta}} \left(L^A + L^B + \sigma_A^2 + \sigma_B^2 + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2 \right)$$

Proof. Since the Graph is Disconnected we now have that: $d_{AA} = n_A(n_A - 1)$, $d_{BB} = n_B(n_B - 1)$, and $d_{AB} = 1$, for the 1 edge that keeps the graph strongly connected. Replacing these values to (3.5) we get the above formula.

□

Proposition 4 Consider a team with an initial composition of $n_A > 0$ members of type A and n_B member of type B. Suppose $\sigma_B > 0$ and $\beta < 1$. Then there exist $n_B^{lower}, n_B^{upper} \in \mathbb{R}$ such that adding a type B member reduces the team's Disagreement Cost (3.5) if and only if $n_B^{lower} \leq n_B \leq n_B^{upper}$.

Proof. Taking the derivative with respect to n_B of the Disconnected Disagreement Cost of Lemma 11 we obtain:

$$\frac{\partial}{\partial n_B}\left(Cost^D_{(p=1,q=0)}\right) = \frac{2(1+\beta)}{(n_A+n_B)^{1+\beta}} \left(\frac{-(1+\beta)}{n_A+n_B}(L^A + L^B + \sigma_A^2 + \sigma_B^2 + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2) + (2n_B - 1)\sigma_B^2 \right)$$

Setting the derivative to zero, is equivalent to solving the roots of the following equation:

$$n_B^2(1 - \beta)\sigma_B^2 + n_B(\beta\sigma_B^2 + 2n_A\sigma_B^2) - const = 0$$

where:

$$const = (1 + \beta)(L^A + L^B + \sigma_A^2 + \sigma_B^2 + n_A(n_A - 1)\sigma_A^2) + n_A\sigma_B^2$$

Note that since $\sigma_B > 0$ and $(1 - \beta) > 0$, the above is a quadratic polynomial in n_B with a positive leading coefficient (i.e., $(1 - \beta)\sigma_B^2$). Let n_B^{lower}, n_B^{upper} denote the roots of this polynomial. Since the leading coefficient is positive, for any $n_B \in [n_B^{lower}, n_B^{upper}]$, the derivative of the disagreement term is negative, indicating that adding new members of type B will reduce the disagreement. Similarly, outside this range, the derivative is positive indicating that new type B members will only worsen the team's disagreement. This is the same mathematical behaviour with the complete graph case just with different n_B^{lower}, n_B^{upper} values.

□

A discussion must be made here for the normalization factor $|T|^{1+\beta}$ in the new Disagreement Cost function (3.5). Unlike the base model [9], where it is intuitive to normalize the Disagreement Cost by the total number of agents to the power of $1 + \beta$ ($\beta \in [0, 1]$), our generalized framework now bases the Disagreement Cost on the connections between directly linked agents (d_{AA}, d_{BB}, d_{AB}). In this context, one might propose normalizing by

the sum of these connections, represented by $D = d_A + d_B = d_{AA} + d_{BB} + 2d_{AB}$. However, this approach poses a challenge. While replacing $|T|^{1+\beta}$ with D restores the base model Disagreement Cost in the complete network, where $p = q = 1$, it does not hold universally across all values of β. Specifically, Lemma 10 remains applicable only for a specific value of β: $\beta = \frac{ln(n_A+n_B-1)}{ln(n_A+n_B)}$. This discrepancy is notable, particularly as it pertains to the Disagreement Cost. While the Accuracy Cost consistently reverts to the base model's cost in the complete network scenario (as demonstrated in Lemma 7), the Disagreement Cost does not exhibit such consistency. Given these considerations, we opt to retain the original normalization factor $|T|^{1+\beta}$ rather than transitioning to D. This decision ensures the best consistency and coherence within our extended model.

3.2 Pyramid Graph

In the previous section we applied a very simple version of DeGroot learning that extended the base model's Cost function (2.7) into including the degree's (d_{AA}, d_{BB}, d_{AB}) of the underlying graph. In that case the team member with the most influence was simply the one with the highest number of connection as shown in Theorem 2 (2.12). In this section we want to make this difference in influence more explicit by introducing a rigid hierarchical structure in which our team operates. Our graph now will be defined not by the p, q connection probability parameters but instead with the parameters k, ℓ. The parameter k defines the maximum number of "subordinates" that each node in the pyramid has, while ℓ the number of layers of the pyramid. In Figure 3.2 is an example.

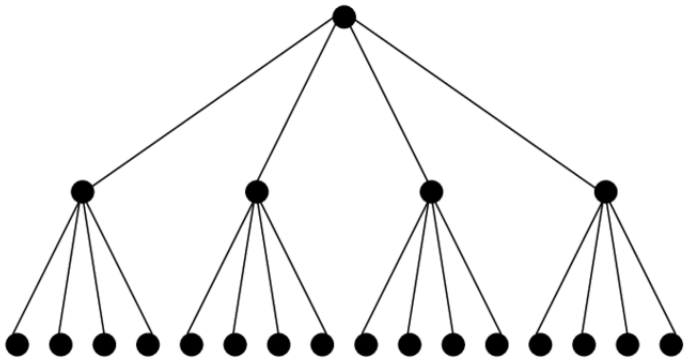

Figure 3.1. *Example of Pyramid graph with ℓ=3 and k=4*

To incorporate the hierarchical influence dynamics within the pyramid structure, we introduce the hyperparameter γ. Unlike the previous undirected graph, where each node equally considered inputs from all connected nodes, the hierarchical framework assigns varying degrees of influence based on the node's position within the pyramid. So for example in Figure 3.2 the nodes of the 2nd level will listen to their 4 subordinates equally, to their top node γ times higher and to themselfs in order to maintain aperiodicity. Formally, this hierarchical influence structure is encoded within a Trust matrix T, from which the corresponding influence vector $\mathbf{s}$ can be derived.

Theorem 3 (Pyramid Graph Influence). In a pyramid hierarchical structure with ℓ layers and k subordinates each agent's influence, defined as their corresponding element in the eigenvector $\mathbf{s}$, remains constant for agents within the same layer. However, the influence increases progressively as we ascend to higher layers of the pyramid following the formula:

$$\begin{aligned}\mathbf{s} = [\ &\underbrace{k\gamma^{\ell-2}}_{\text{Top Node}} \quad \underbrace{\gamma^{\ell-3}(\gamma+k)}_{2^{nd}\text{ Layer } k \text{ nodes}} \quad \underbrace{\gamma^{\ell-4}(\gamma+k)}_{3^{rd}\text{ Layer } k^2 \text{ nodes}} \quad \dots \quad \underbrace{\gamma(\gamma+k)}_{k^{\ell-3}\text{ nodes}} \quad \underbrace{(\gamma+k)}_{k^{\ell-2}\text{ nodes}} \quad \underbrace{1}_{k^{\ell-1}\text{ nodes}}\ ]s \\ = [\ &s_1 \qquad s_2 \qquad s_3 \qquad \dots \qquad s_{\ell-2} \qquad s_{\ell-1} \qquad s_\ell\ ]\end{aligned} \tag{3.6}$$

where $s = 1/\sum_i s_i$ is a normalization factor so that the sum of all the terms inside the vector $\mathbf{s}$ adds up to 1.

Proof. Note that in the DeGroot Model because of convergence: $\mathbf{s}\cdot\mathbf{b} = (\mathbf{s}T)\cdot\mathbf{b} \Rightarrow \mathbf{s} = \mathbf{s}T$. For our pyramid graph structure the conditions for convergence hold since the graph is strongly connected and also aperiodic. Now let us verify $\mathbf{s} = \mathbf{s}T \Rightarrow s_i = \sum_j T_{ij}s_j$. For this proof to be straightforward we must have a clear image of the matrix T. The diagonal values of the matrix are $T_{ii} = \epsilon$ and under the diagonal there are alternating columns of k length, filed with the values $T_b = \frac{\gamma(1-\epsilon)}{\gamma+k}$. For the last half of the matrix, which corresponds to the nodes of the base of the pyramid, the values become $T_{bb} = 1-\epsilon$. Above the diagonal there are alternating rows of k length with values $T_t = \frac{1-\epsilon}{\gamma+k}$. Only the 1st row which corresponds to the top node has k values of $T_{tt} = \frac{1-\epsilon}{k}$. Lets write a simple example of a full pyramid with $\ell = 3$ layers and $k = 2$. In this case we have 7 nodes and the 7 by 7 Trust Matrix is:

$$\mathrm{T} = \begin{bmatrix} \epsilon & \frac{1-\epsilon}{2} & \frac{1-\epsilon}{2} & 0 & 0 & 0 & 0 \\ \frac{\gamma(1-\epsilon)}{\gamma+2} & \epsilon & 0 & \frac{(1-\epsilon)}{\gamma+2} & \frac{(1-\epsilon)}{\gamma+2} & 0 & 0 \\ \frac{\gamma(1-\epsilon)}{\gamma+2} & 0 & \epsilon & 0 & 0 & \frac{(1-\epsilon)}{\gamma+2} & \frac{(1-\epsilon)}{\gamma+2} \\ 0 & 1-\epsilon & 0 & \epsilon & 0 & 0 & 0 \\ 0 & 1-\epsilon & 0 & 0 & \epsilon & 0 & 0 \\ 0 & 0 & 1-\epsilon & 0 & 0 & \epsilon & 0 \\ 0 & 0 & 1-\epsilon & 0 & 0 & 0 & \epsilon \end{bmatrix}$$

So in the general case that the pyramid has ℓ layers and k subordinates we can write for the top node that:

$$\begin{aligned} s_1 &= \sum_j T_{j1}s_j = T_{ii}s_1 + kT_b s_2 \\ &= \left(\epsilon k\gamma^{\ell-2} + k\frac{\gamma(1-\epsilon)}{\gamma+k}\gamma^{\ell-3}(\gamma+k)\right)s \\ &= \left(\epsilon k\gamma^{\ell-2} + k\gamma^{\ell-2}(1-\epsilon)\right)s = k\gamma^{\ell-2}s = s_1 \end{aligned}$$

For the nodes in the second layer we have that:

$$\begin{aligned} s_2 &= \sum_j T_{j2} s_j = T_{tt} s_1 + T_{ii} s_2 + k T_b s_3 \\ &= \left(\frac{1-\epsilon}{k} k\gamma^{\ell-2} + \epsilon\gamma^{\ell-3}(k+\gamma) + k\frac{\gamma(1-\epsilon)}{\gamma+k}\gamma^{\ell-4}(\gamma+k) \right) s \\ &= \left((1-\epsilon)\gamma^{\ell-2} + \epsilon\gamma^{\ell-3}(k+\gamma) + k(1-\epsilon)\gamma^{\ell-3} \right) s \\ &= \gamma^{\ell-3}\left((1-\epsilon)\gamma + \epsilon(k+\gamma) + (1-\epsilon)k \right) s = \gamma^{\ell-3}(\gamma+k)s = s_2 \end{aligned}$$

Similar proofs hold for the rest of the middle layers of the pyramid that have agents both above and bellow them. For the 3rd layer for example we have that:

$$\begin{aligned} s_3 &= \sum_j T_{j3} s_j = T_t s_2 + T_{ii} s_3 + k T_b s_4 \\ &= \left(\frac{1-\epsilon}{\gamma+k}\gamma^{\ell-3}(\gamma+k) + \epsilon\gamma^{\ell-4}(k+\gamma) + k\frac{\gamma(1-\epsilon)}{\gamma+k}\gamma^{\ell-5}(\gamma+k) \right) s \\ &= \left((1-\epsilon)\gamma^{\ell-3} + \epsilon\gamma^{\ell-4}(k+\gamma) + k(1-\epsilon)\gamma^{\ell-4} \right) s \\ &= \gamma^{\ell-4}\left((1-\epsilon)\gamma + \epsilon(k+\gamma) + (1-\epsilon)k \right) s = \gamma^{\ell-4}(\gamma+k)s = s_3 \end{aligned}$$

We can keep moving down the layers of the pyramid using the same method until we reach the semi-final layer for which we can write that:

$$\begin{aligned} s_{\ell-1} &= \sum_j T_{j(\ell-1)} s_j = T_t s_{\ell-2} + T_{ii} s_{\ell-1} + k T_{bb} s_\ell \\ &= \left(\frac{1-\epsilon}{\gamma+k}\gamma(\gamma+k) + \epsilon(k+\gamma) + k(1-\epsilon) \right) s \\ &= \left((1-\epsilon)\gamma + \epsilon\gamma(k+\gamma) + k(1-\epsilon) \right) s \\ &= (\gamma+k)s = s_{\ell-1} \end{aligned}$$

And for the nodes at the base of the pyramid it is easy to show that:

$$\begin{aligned} s_\ell &= \sum_j T_{j\ell} s_j = T_t s_{\ell-1} + T_{ii} s_\ell \\ &= \left(\frac{1-\epsilon}{k+\gamma}(k+\gamma) + \epsilon \right) s \\ &= (1-\epsilon+\epsilon)) s = s = s_\ell \end{aligned}$$

□

So now based on the result of Theorem 3 (3.6), for the pyramid graph which node has the most influence? Defining for the parameters that: $\gamma > 1$ and for the integer k that: $k \geq 2$ it easy to see that the top node has the most influence followed by the nodes of the second layer, then the third layer and so forth until we reach the nodes on the base of the pyramid which have the least influence. Also very interesting is to define the influence that each layer has as a whole.

Definition (Pyramid Graph Layer Influence). Consider a team of members connected on a pyramid graph with parameters k, ℓ. We define the influence of layer $n \in \mathbb{N}$ in this pyramid as:

$$infl(n) = k^{n-1} s_n \tag{3.7}$$

Where k^{n-1} is the number of nodes in the layer and s_n their corresponding influence. For example: $infl(1) = s_1$, $infl(2) = ks_2$, $infl(3) = k^2 s_3$, ... , $infl(\ell) = k^{\ell-1} s_\ell$.

Lemma 12 (Top and Base Layer Influence). Consider a team of members connected on a pyramid graph with parameters k, ℓ. It is true for their layer influence that:

$$infl(1) < infl(2)$$
$$infl(\ell) < infl(\ell - 1)$$

Proof. Solving the inequalities by using the results of Theorem 3 (3.6) we get:

$$infl(1) < infl(2) \Leftrightarrow k\gamma^{\ell-2} < k\gamma^{\ell-3}(\gamma + k) \Leftrightarrow \gamma < \gamma + k \Leftrightarrow k > 0$$
$$infl(\ell) < infl(\ell - 1) \Leftrightarrow k^{\ell-1} < k^{\ell-2}(\gamma + k) \Leftrightarrow k < \gamma + k \Leftrightarrow \gamma > 0$$

□

Lemma 13 (Middle Layer Influence). Consider a team of members connected on a pyramid graph with parameters k, ℓ. For their layer influence in layer $n \in [2, \ell - 2]$:

$$\gamma > k \Leftrightarrow infl(n) > infl(n + 1)$$
$$\gamma < k \Leftrightarrow infl(n) < infl(n + 1)$$

Proof. Solving the first inequality by using the results of Theorem 3 (3.6) we get:

$$infl(n) > infl(n + 1) \Leftrightarrow k^{n-1}\gamma^{\ell-n-1}(\gamma + k) > k^n\gamma^{\ell-n-2}(\gamma + k) \Leftrightarrow \gamma > k$$

□

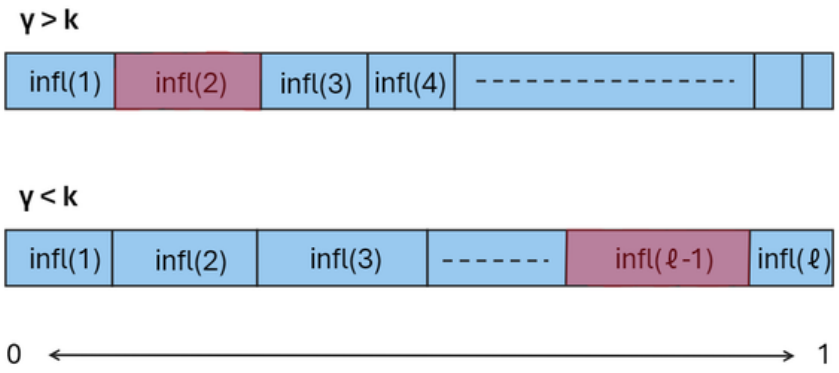

Figure 3.2. *The results of Lemma 12 and 13 expressed visually. The total influence is fixed and the layer with the most influence is underlined with red.*

3.2.1 Accuracy

Let's analyze now how the Accuracy Cost changes in the frame of the pyramid graph. In order to calculate the Accuracy Cost: $Cost^A = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[\ell(\mathcal{G}_T(\mathbf{x}), y)]$ we must first calculate how the aggregate opinion function changes compared to the undirected graph (3.2).

Lemma 14 (Pyramid Graph Aggregation Function). Consider a team of n_A members of type A and n_B members of type B connected on a pyramid graph. Then:

$$\begin{aligned}\mathcal{G}(\mathbf{x}) = (i_1 s_1 + i_2 s_2 + \cdots + i_{\ell-1} s_{\ell-1} + i_\ell s_\ell) f^A(\mathbf{x}) \\ + \left(i'_1 s_1 + i'_2 s_2 + \cdots + i'_{\ell-1} s_{\ell-1} + i'_\ell s_\ell\right) f^B(\mathbf{x})\end{aligned} \tag{3.8}$$

Where $i_n, i'_n \in \mathbb{N}$ with $n \in \mathbb{N}$ representing the layer. It must hold that: $i_1 + i'_1 = 1$, $i_2 + i'_2 = k$, $i_3 + i'_3 = k^2, \ldots, i_{\ell-1} + i'_{\ell-1} = k^{\ell-2}$, $i_\ell + i'_\ell = k^{\ell-1}$. Where the parameters i_n show the number of type A agents in each layer and i'_n the number of type B agents.

Proof. All the agents will converge to the same final opinion according to (2.11) and as such we define this final opinion to be the aggregate:

$$\begin{aligned}\mathcal{G}(\mathbf{x}) &= \mathbf{s}\cdot\mathbf{b} \\ &= [\ \underbrace{k\gamma^{\ell-2}}_{\text{Top Node}}\quad \underbrace{\gamma^{\ell-3}(\gamma+k)}_{2^{nd}\text{ Layer }k\text{ nodes}}\quad \underbrace{\gamma^{\ell-4}(\gamma+k)}_{3^{rd}\text{ Layer }k^2\text{ nodes}}\quad \ldots\quad \underbrace{\gamma(\gamma+k)}_{k^{\ell-3}\text{ nodes}}\quad \underbrace{(\gamma+k)}_{k^{\ell-2}\text{ nodes}}\quad \underbrace{1}_{k^{\ell-1}\text{ nodes}}\ ]s\cdot\mathbf{b} \\ &= [\quad s_1 \qquad s_2 \qquad s_3 \qquad \ldots \qquad s_{\ell-2} \qquad s_{\ell-1} \qquad s_\ell\quad]\cdot\mathbf{b}\end{aligned}$$

Where $\mathbf{b}$ is the vector of the initial beliefs:

$$\mathbf{b} = [f^A(\mathbf{x}), f^B(\mathbf{x}), \ldots, f^A(\mathbf{x}), f^B(\mathbf{x})]$$

As such based on the positioning of $f^A(\mathbf{x})$ and $f^B(\mathbf{x})$ in the vector of the initial beliefs $\mathbf{b}$ aka the positioning of the agents of type A and type B at the pyramid we get the values of $i_n, i'_n \in \mathbb{N}$ in order to calculate $\mathcal{G}(\mathbf{x})$.

□

Lemma 15 (Pyramid Network Accuracy Cost). Consider a team of n_A members of type A and n_B members of type B connected on a full pyramid graph. The Accuracy cost of the network will be:

$$Cost^A = (i_1 s_1 + \cdots + i_\ell s_\ell)^2 (L^A + \sigma_A^2) + (i'_1 s_1 + \cdots + i'_\ell s_\ell)^2 (L^B + \sigma_B^2) \tag{3.9}$$

Proof. We can write:

$$\mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2] = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}\left[\left((\partial^A(\mathbf{x}) + \epsilon_A)\sum_{n=1}^{\ell} i_n s_n + (\partial^B(\mathbf{x}) + \epsilon_B)\sum_{n=1}^{\ell} i'_n s_n - f^*(\mathbf{x})\right)^2\right]$$

$$= \mathbb{E}\left[\left(-\partial^B(\mathbf{x})\sum_{n=1}^{\ell} i_n s_n - \partial^A(\mathbf{x})\sum_{n=1}^{\ell} i'_n s_n + \epsilon_A\sum_{n=1}^{\ell} i_n s_n + \epsilon_B\sum_{n=1}^{\ell} i'_n s_n\right)^2\right] \quad \text{(Using 2.1)}$$

$$= \left(\mathbb{E}\left[\partial^B(\mathbf{x})^2\right] + \mathbb{E}\left[\epsilon_A^2\right]\right)\left(\sum_{n=1}^{\ell} i_n s_n\right)^2 + \left(\mathbb{E}\left[\partial^A(\mathbf{x})^2\right] + \mathbb{E}\left[\epsilon_B^2\right]\right)\left(\sum_{n=1}^{\ell} i'_n s_n\right)^2$$

$$= (L^A + \sigma_A^2)(i_1 s_1 + \cdots + i_\ell s_\ell)^2 + (L^B + \sigma_B^2)(i'_1 s_1 + \cdots + i'_\ell s_\ell)^2$$

□

3.2.2 Disagreement

Let's continue by analyzing how we can introduce the Disagreement Cost with the addition of the pyramid graph. By using the DeGroot learning process all agents reach the same final opinion and that can be thought of as canceling out the teams' disagreement. However again we want to preserve the tension between Accuracy and Disagreement Costs just as in the base model [9]. We will again apply the same approach as in the undirected graph in order to accomplish this. As such in our extended model, disagreement occurs solely between agents directly connected in the underlying graph—meaning their nodes are directly linked by an edge. By adopting this criterion for the pyramid graph the results of Lemma 9 still hold and we can write that the Disagreement Cost again is:

$$Cost^D = \frac{1}{|T|^{1+\beta}}\sum_{i,j\in T}\mathbb{E}[(f_j(\mathbf{x}) - f_i(\mathbf{x}))^2] = \frac{2}{|T|^{1+\beta}}(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2)$$

We must note that, unlike the undirected graph analyzed in the previous section, the pyramid graph does not allow for the definition of complete and disconnected cases. This is due to the fixed edge connections inherent in the structure of the pyramid graph. Therefore, our scope for adjustments is limited to the positioning of agent types A or B within the network. Understanding the optimal placement of team members is crucial for maximizing the efficiency and effectiveness of the team dynamics within this hierarchical structure. In the following chapter, Chapter 4, we delve deeper into the optimal positioning of team members within the pyramid graph and present simulated results.

Chapter 4

Results and Analysis

In this chapter, we present the outcomes of simulations and algorithmic analyses aimed at elucidating the dynamics of our extended model proposed in previous chapters. We explore the interplay between team structures and decision-making processes, focusing on two fundamental aspects: the connectivity of the network and the positioning of team members within it. Through a series of experiments, we investigate the implications of the complete versus the disconnected undirected graph, examining how different network configurations influence the team's performance and decision outcomes. Additionally, we delve into the optimization of edge connections, exploring various approaches to minimize the cost function for both the undirected and the pyramid graph. Furthermore, we scrutinize the optimal positioning of team members within the pyramid graph, shedding light on strategies to enhance team cohesion and efficiency. By synthesizing simulation results and algorithmic insights, this chapter offers valuable insights into the dynamics of team decision-making in complex network structures.

4.1 Undirected Graph: Complete vs Disconnected

Let's begin by confirming the results of Propositions 1 (2.9) and 3 (3.4) through team growth simulations. The results presented in Figures 4.1 and 4.2 were created using the following conditions: Each world state $\mathbf{x}$ comprises 10 features x_i, each drawn from a standard normal distribution with mean = 0 and standard deviation = 0.1, so that $\text{cov}(x_i, x_j) = 0$ for all $j \neq i$. With a total of 10,000 world states and 10 features each, this yields a 10 by 10,000 matrix describing the world. After creating the world states, in order to reach the desirable L^A and L^B we generate appropriate values of ∂^A and ∂^B each targeting half of the features, following the equations of Lemma 1 (2.5). As such by concatenating the created ∂_A and ∂_B we get the vector ∂^*.

With the foundational elements of our simulation in place, we proceed to construct the undirected graph. Leveraging the capabilities of the NetworkX package [18], we create an undirected graph with n_A nodes of type A and n_B nodes of type B connected with edges based on the probabilities p and q. It is paramount to ensure that the graph remains strongly connected for any possible p and q value, guaranteeing this way convergence for the DeGroot learning process, as stated in Theorem 1. From the constructed graph we derive the corresponding trust matrix T, enabling us to compute the influence vector

s (2.12), the dot product of which with the initial opinions **b** results into the networks Aggregate opinion $\mathcal{G}_T(\mathbf{x})$.

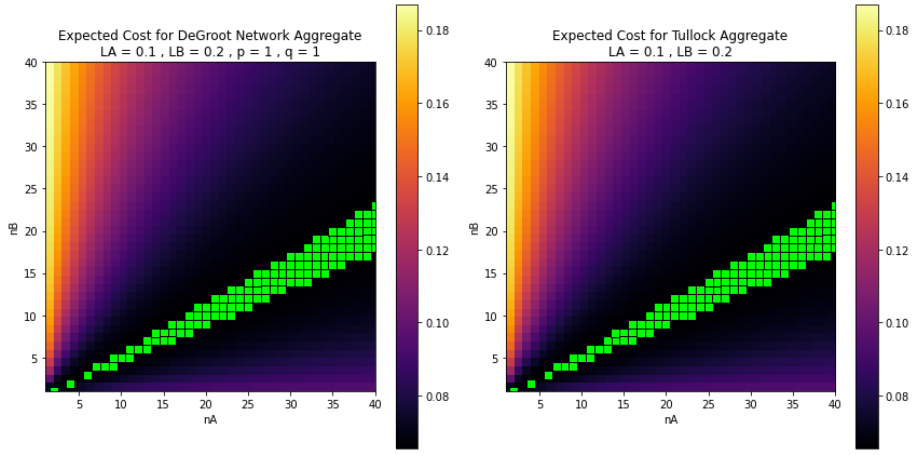

Figure 4.1. *Simulation of the Team Growth Dynamics with the Base model Tullock Aggregate on the right and the Undirected graph with DeGroot Aggregate on the left. Here the graph is Complete, $p = q = 1$ and the green cells represent the 50 (n_A, n_B) points with the lowest values.*

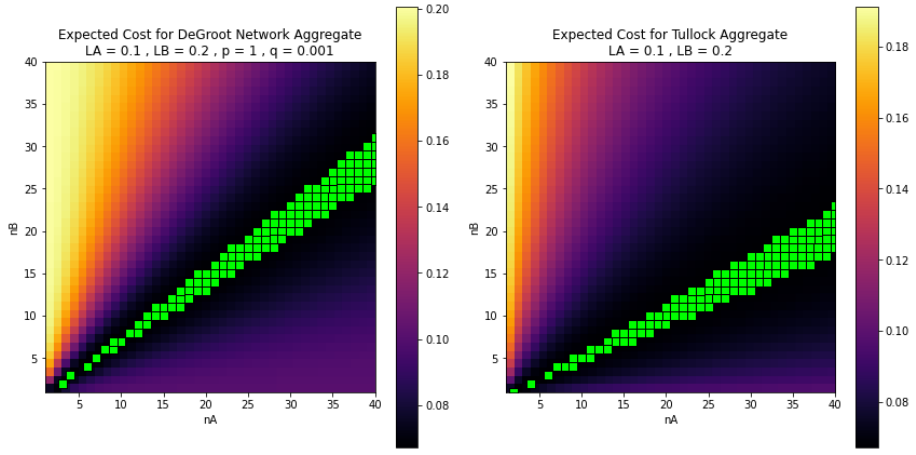

Figure 4.2. *Simulation of the Team Growth Dynamics with the Base model Tullock Aggregate on the right and the Undirected graph with DeGroot Aggregate on the left. Here the graph is Disconnected, $p = 1$, $q \approx 0$ and the green cells represent the 50 (n_A, n_B) points with the lowest values.*

Thus, by calculating the Accuracy Cost, which represents the expected squared difference between the Aggregate $\mathcal{G}_T(\mathbf{x})$ and the true outcome function $f^*(\mathbf{x})$ across all world states $\mathbf{x}$ in a given network $G(n_A, n_B, p, q)$, we can validate the results of Propositions 1 and 3. We compute Accuracy Cost values for various combinations of n_A and n_B (ranging from 1 to 40) and visualize them in heat maps shown in Figures 4.1 and 4.2. Additionally, we plot the Accuracy Cost values that would result from using the Base Model's Tullock Aggregate function instead of the DeGroot learning Aggregate, providing further insight. As depicted in Figure 4.1, the complete network maintains the results of the base model,

while the Disconnected network in Figure 4.2 deviates and follows the Accuracy-optimal composition outlined in Equation (3.4). Furthermore, for intermediate values of $q \in (0, 1)$, we observe the minimum line shifting between these two extremes, indicating how the value of q influences the teams optimal composition. Despite this observation however, we are unable to derive closed-form equations for $q \in (0, 1)$, like the ones for the Complete (Equation (2.9)) or Disconnected (Equation (3.4)) cases due to the randomness of the graph $G(n_A, n_B, p, q)$.

It's noteworthy to consider the theoretical disparity between the n_B^* values for complete and disconnected graphs. Proposition 1 establishes an Accuracy-optimal Composition for n_B^* with a clear linear relationship with n_A, characterized by the slope $\frac{L^A+\sigma_A^2}{L^B+\sigma_B^2}$. Similarly, Proposition 2 indicates a nearly linear relationship between n_B^* and n_A. However, upon comparison across various values of $\frac{L^A+\sigma_A^2}{L^B+\sigma_B^2}$, we observe that the disconnected graph tends to adhere more closely to the line $n_B^* = n_A$ than the complete graph. This tendency of the disconnected graph can be attributed to the absence of across the aisle d_{AB} edges, necessitating either more 'bad' agents (with higher Loss) or fewer 'good' agents (with lower Loss) to reach the Cost minimum. This is because information propagation within the disconnected network is suboptimal.

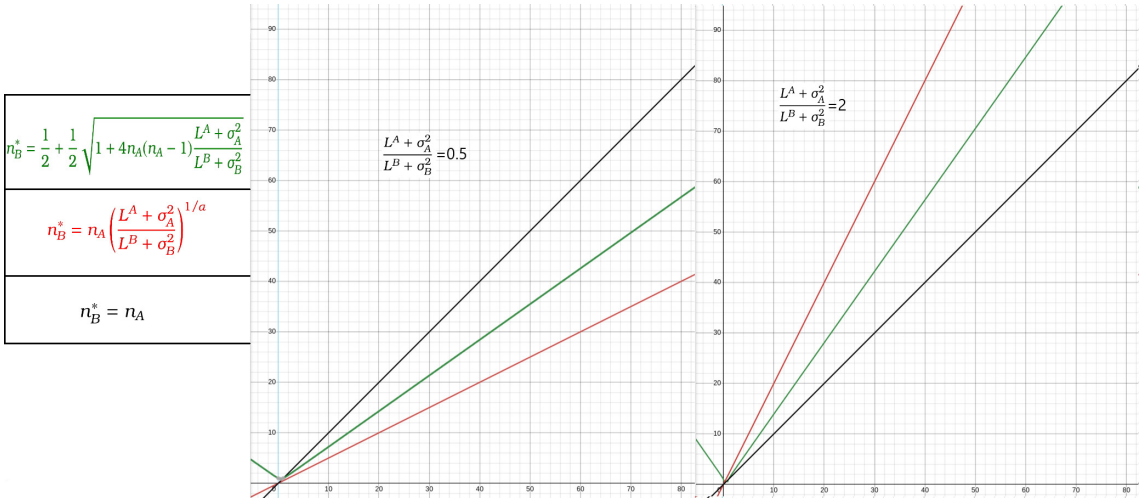

Figure 4.3. *Theoretical Accuracy-optimal Team Compositions for different values of* $\frac{L^A+\sigma_A^2}{L^B+\sigma_B^2}$. *Green represents the optimal composition for the Disconnected network, Red for the Complete and Black is the composition with the same number of agents in each type.*

Let's analyze, for example, Figure 4.3, considering the case where $\frac{L^A+\sigma_A^2}{L^B+\sigma_B^2} = \frac{1}{2}$. Here, agents of type B experience twice the Loss compared to agents of type A. In the Complete model's Accuracy-optimal composition, this leads to the requirement for double the number of type A agents compared to type B agents to minimize the Cost. For instance, in Figure 4.3, if $n_A = 40$, the minimum Cost occurs when $n_B = 20$ (red line). However, in the Disconnected network, achieving the minimum Cost necessitates more type B agents. With $n_A = 40$, we find that $n_B = 29$ (green line). Therefore, in this case, the Disconnected graph requires more type B agents, which are the agents with the higher Loss.

4.2 Undirected Graph: Optimal Edges

Let's now consider a new problem that arises from the introduction of the undirected graph in the base model. As mentioned before, for random graphs G with the same n_A, n_B, p, q values because of the randomness associated with the creation of the edges the d_A, d_B and D values can vary and as such the values of the Aggregate function (3.2) and of the Accuracy Cost (3.3) can also vary. So a natural question to ask here is: given a network G of n_A and n_B agents how can we change the edge connections between them in order to minimize the Accuracy Cost? Let's start answering this question by first finding when the Accuracy Cost is minimized and under which condition.

Lemma 16 (Undirected Graph Minimum Accuracy Cost). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$. The Accuracy Cost function will be minimized when $d_A(L^A + \sigma_A^2) = d_B(L^B + \sigma_B^2)$ and will be equal to:

$$Cost^* = \frac{(L^A + \sigma_A^2)(L^B + \sigma_B^2)}{(L^A + \sigma_A^2) + (L^B + \sigma_B^2)} \tag{4.1}$$

Proof. According to Lemma 6 (3.3) the team's Accuracy Cost can be written as:

$$\mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2] = \left(\frac{d_A}{D}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B}{D}\right)^2 (L^B + \sigma_B^2)$$

Taking the derivative of the right hand side with respect to d_B, we obtain:

$$\frac{\partial}{\partial d_B}\left(\mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2]\right) = \frac{2d_A d_B}{(d_A + d_B)^3}(L^A + \sigma_A^2) - \frac{2d_B^2}{(d_A + d_B)^3}(L^B + \sigma_B^2)$$

To obtain the zero of the derivative, we can write:

$$\frac{2d_A d_B}{(d_A + d_B)^3}(L^A + \sigma_A^2) - \frac{2d_B^2}{(d_A + d_B)^3}(L^B + \sigma_B^2) = 0 \Leftrightarrow d_A(L^A + \sigma_A^2) = d_B(L^B + \sigma_B^2)$$

Now replacing the value of d_B with $d_B^* = d_A\frac{L^A+\sigma_A^2}{L^B+\sigma_B^2}$ we can see than $\frac{d_B}{D} = \frac{L^A+\sigma_A^2}{(L^A+\sigma_A^2)+(L^B+\sigma_B^2)}$ and $\frac{d_A}{D} = \frac{L^B+\sigma_B^2}{(L^A+\sigma_A^2)+(L^B+\sigma_B^2)}$. Applying these values to (3.3) we can find the minimum Cost value as:

$$Cost^* = \frac{(L^A + \sigma_A^2)(L^B + \sigma_B^2)}{(L^A + \sigma_A^2) + (L^B + \sigma_B^2)}$$

□

We can use now the fact that the minimum for any given graph G is reached when $d_A(L^A + \sigma_A^2) = d_B(L^B + \sigma_B^2)$ to find which "moves" can be made in order to minimize the Accuracy Cost. All the possible types of edge additions and removals in a graph are a lot but in our case every possible move can be boiled down to only 6 categories each affecting the Accuracy Cost in the same way: Add/Remove Edge between A types, Add/Remove Edge between B types, Add/Remove Edge between A and B types. Because of the way

the undirected network calculates the Accuracy Cost (3.3) each of these 6 "moves" affects the network in the same way independently of which type A or type B nodes are affected. An important note here is that these moves must be made with respect to the **Boundary Conditions** which are:

- The Graph G must be kept strongly connected (applied during removals)
- No new edges can be added when the category is full (applied during additions)

Respecting these boundary conditions we can see that the minimum is reached when:

$$d_A = d_B \frac{L^B + \sigma_B^2}{L^A + \sigma_A^2} = d_B C$$

Where $C = \frac{L^B + \sigma_B^2}{L^A + \sigma_A^2}$ is a constant since $L^A, L^B, \sigma_A, \sigma_B$ are all fixed values defined at the start of the simulation process. As such for any initial d_A, d_B values after $k \in \mathbb{N}$ "moves" of adding or subtracting edges from the graph we are trying to reach the value:

$$C = \frac{d_A + 2(k_1 - k_2) + (k_5 - k_6)}{d_B + 2(k_3 - k_4) + (k_5 - k_6)} \tag{4.2}$$

Where $k = k_1 + k_2 + k_3 + k_4 + k_5 + k_6$ with $k_{1,2}$ = Add/Remove Edge between A types, $k_{3,4}$ = Add/Remove Edge between B types, $k_{5,6}$ = Add/Remove Edge between A and B types. The above expression can be simplified after the observation that we can use only half of the k_i moves and still reach the same value, since the moves $k_{1,2}$, $k_{3,4}$, $k_{5,6}$ mutually exclude each other.

Lemma 17 (Adding/Removing same Types of Edges). Through the process of trying to reach the value C through the addition or removal of a particular type of edge, each move is counteracted when its opposite action is taken. This means that (4.2) can be simplified to:

$$C = \frac{d_A + 2\mu + \xi}{d_B + 2\nu + \xi} \tag{4.3}$$

Where $\mu, \nu, \xi \in \mathbb{Z}$ and when $\mu, \nu, \xi > 0$ only the moves k_1, k_3, k_5 are used and when $\mu, \nu, \xi < 0$ only the moves k_2, k_4, k_6 are used, with all $2^3 = 8$ combinations between them.

Proof. Lets write how k_1 and k_2 cancel each other out and the same holds for $k_{3,4}$ and $k_{5,6}$. For any given graph initially we start with:

$$Cost^{init} = \left(\frac{d_A}{D}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B}{D}\right)^2 (L^B + \sigma_B^2)$$

Let's start with an example: After $k = 3$ steps lets assume $k_1 \to k_3 \to k_6$ the Cost will be:

$$\begin{aligned} Cost^{k=3}_{k_1 \to k_3 \to k_6} &= \left(\frac{d_A + 2 - 1}{D + 2 + 2 - 2}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B + 2 - 1}{D + 2 + 2 - 2}\right)^2 (L^B + \sigma_B^2) \\ &= \left(\frac{d_A + 1}{D + 2}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B + 1}{D + 2}\right)^2 (L^B + \sigma_B^2) \end{aligned}$$

Now making the k_2 move the Cost will become:

$$\begin{aligned} Cost^{k=4}_{k_1\to k_3\to k_6\to k_2} &= \left(\frac{d_A+1-2}{D+2-2}\right)^2 (L^A+\sigma_A^2) + \left(\frac{d_B+1}{D+2-2}\right)^2 (L^B+\sigma_B^2) \\ &= \left(\frac{d_A-1}{D}\right)^2 (L^A+\sigma_A^2) + \left(\frac{d_B+1}{D}\right)^2 (L^B+\sigma_B^2) \end{aligned}$$

And we can see that by making the move k_2 the move k_1 has been canceled since:

$$Cost^{k=4}_{k_1\to k_3\to k_6\to k_2} = Cost^{k=2}_{k_3\to k_6}$$

This can be generalized for any possible number of k steps thus the moves k_1 cancel the effect of the moves k_2 and vice versa. As a result at the end only k_1 or k_2 moves remain and can be expressed through the integer $\mu = k_1 - k_2$ which is positive when $k_1 > k_2$ and negative when $k_1 < k_2$. Following the same logic we can introduce the integers $\upsilon = k_3 - k_4$ and $\xi = k_5 - k_6$ for the moves k_3, k_4 and k_5, k_6 respectively.

□

Now a fair question to ask is if the ratio in (4.3) can produce any possible rational number C. Lets start answering this question by first setting aside the consideration of the Boundary Conditions.

Lemma 18 (Ratio $\frac{d_A+2\mu+\xi}{d_B+2\upsilon+\xi} \in \mathbb{Q}^+$). For any initial $d_A, d_B \in \mathbb{N}$ with both d_A, d_B being *even* or both d_A, d_B being *odd*, the ratio:

$$\frac{d_A+2\mu+\xi}{d_B+2\upsilon+\xi}$$

Can reach any possible $C \in \mathbb{Q}^+$, where $\mu, \upsilon, \xi \in \mathbb{Z}$.

Proof. Lets start by writing down that since $C \in \mathbb{Q}^+ \Rightarrow C = p/q$, where $p, q \in \mathbb{N}$. As such whether p, q are odd or even they can be both turned even by simply multiplying the nominator and the denominator with an even number for example lets say 2. As such $C = p/q = 2p/2q$, where both $2p$ and $2q$ are even numbers. Now lets think about the two cases: First case is that both d_A and d_B are even. In that case we can use the fact that every even number is reachable from every other even number by adding or subtracting 2 so that:

$$\frac{d_A}{d_B} \to \frac{d_A+2\mu}{d_B+2\upsilon} = \frac{2p}{2q} = C$$

Now for the second case that both d_A and d_B are odd we can start by turning d_A and d_B even by adding an odd ξ at the nominator and the denominator. This will give us as before a ratio of even numbers and now as before we can reach C:

$$\frac{d_A}{d_B} \to \frac{d_A+\xi}{d_B+\xi} \to \frac{d_A+\xi+2\mu}{d_B+\xi+2\upsilon} = \frac{2p}{2q} = C$$

□

At this point a good question to ask is how can we find the values of μ, υ, ξ. The first

and simplest solution is an exhaustive search of all the possible $\mu, \nu, \xi \in \mathbb{Z}$ combinations based on the Boundary Conditions of the graph. In order to do that we must "break down" again the d_A and d_B values:

$$d_A = d_{AA} + d_{AB} \quad , \quad d_B = d_{BB} + d_{AB}$$

As mentioned in Chapter 3 d_{AA} and d_{BB} represent the edges between nodes of the same type (i.e. A to A or B to B) and d_{AB} represent the edges between nodes of different type (i.e. A to B). By doing that we can now write the **Boundary Conditions** as:

- $1 \leq d_{AB}^{init} + \xi \leq n_A n_B$
- $2(n_A - 1) \leq d_{AA}^{init} + 2\mu \leq n_A(n_A - 1)$
- $2(n_B - 1) \leq d_{BB}^{init} + 2\nu \leq n_B(n_B - 1)$

A comment needs to be made here about the lower limits of these inequalities. Firstly the upper limits are applied while adding edges to the graph. It is easy to see the upper limits of d_{AB}, d_{AA} and d_{BB} as the values they would have if the graph was complete (aka every node is connected with every other node). It is also easy to find the lowest value for d_{AB} as equal to 1 since we want the graph to be strongly connected. The same is not so obvious however for the lowest values of d_{BB} and d_{AA}. It is theoretically possible for large enough graphs to get d_{AA} and d_{BB} as low as 0 while still being strongly connected. This however requires a large enough value of d_{AB}. So in order to keep the boundary conditions as simple as possible we choose to keep the lowest values of d_{AA} and d_{BB} at $2(n_A - 1)$ and $2(n_B - 1)$. These lower bounds for d_{AA} and d_{BB} make sure that the nodes of the same type are all reachable among themselves ensuring this way a level of connectivity between the agents of the same type. As such by using the 3 boundary condition inequalities for any initial graph $G(n_A, n_B, q, p)$ we can find all the possible values of $\mu, \nu, \xi \in \mathbb{Z}$ that respect the Boundary Conditions and try their combinations in order to reach C.

However this can be very computationally expensive especially for large graphs. Can we apply another approach? One other approach can be a greedy algorithm which given a graph $G(n_A, n_B, q, p)$ tries out iteratively the 6 possible "moves" of addition/removal of edges and applies the one which lowers the Cost the most. This greedy algorithm will terminate when no "move" can be applied which lowers the Cost any further. A first observation that we can make is that the Greedy Algorithm is not Optimal since it can be trapped close to unreachable solutions due to the boundary conditions.

Lemma 19 (Greedy Search Algorithm is Not Optimal). In order to find the steps μ, ν, ξ that minimize the Accuracy Cost (3.3) the Greedy search Algorithm is not Optimal since it can be trapped close to an optimal solution without ever reaching it due to the boundary conditions of the problem.

Proof. Lets prove this with an example: Lets assume we have a network of $n_A = n_B = 5$ type A and type B agents with each having Losses: $L^A = 0.1$ and $L^B = 0.2$ with no noise $\sigma_A = \sigma_B = 0$. The optimal ratio that needs to be reached in order to minimize the Cost is:

$C = d_A/d_B = L_B/L_A = 2$. In the case that the network is complete ($p = q = 1$) aka every node is connected with every other node directly then the steps that the Greedy Algorithm takes are the following:

$$\frac{d_A^{init}}{d_B^{init}} = \frac{(n_A + n_B)(n_A + n_B - 1)/2}{(n_A + n_B)(n_A + n_B - 1)/2} = \frac{45}{45} \xrightarrow{10k_4} \frac{45}{25} \xrightarrow{5k_6} \frac{40}{20} = 2 = C$$

In the case however that the network is disconnected ($p = 1, q \approx 0$) aka every node of the same type is connected but the two different types of nodes are connected only with 1 edge so that the network is strongly connected, the steps that the Greedy Algorithm takes are the following:

$$\frac{d_A^{init}}{d_B^{init}} = \frac{((n_A)(n_A - 1) + 1)/2}{((n_B)(n_B - 1) + 1)/2} = \frac{21}{21} \xrightarrow{5k_4} \frac{21}{11} \neq 2 \neq C$$

As can be seen in this case the Greedy Algorithm is trying to reach the ratio 20/10. However when it reaches the ratio 21/11 it cannot do the move k_6 because removing the connection between the type A and the type B agents will disconnect the network (boundary condition). As a result it terminates in a sub optimal solution since from the ratio 21/11 no moves are beneficial, from it's myopic point of view, in lowering the Cost. However if the algorithm was able to add edges in the graph it would be able to reach the ratio 40/20 which, as we saw from the case that the initial graph was complete, is an optimal solution.

□

For small graphs where the boundary conditions are very restrictive or for irrational values of C (might happen if we choose irrational $L^A, L^B, \sigma_A, \sigma_B$) the minimum value cannot be reached. However this does not pose a problem because we can show that the closer we are at the minimum (4.1) the less is the values of the Accuracy Cost.

Lemma 20 (Accuracy Cost Proximity). For the Accuracy Cost function of the Undirected Graph (3.3) with $d_A, d_B > 0$ as the Euclidean distance of a point (d_A, d_B) to the line $d_A(L^A + \sigma_A^2) = d_B(L^B + \sigma_B^2)$ decreases, the value of the Cost function $Cost^A$ also decreases.

Proof. The Accuracy Cost of the network can be written as:

$$Cost^A = \mathbb{E}_{\mathbf{x}\sim\mathcal{P}}[(\mathcal{G}_T(\mathbf{x}) - y)^2] = \left(\frac{d_A}{D}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B}{D}\right)^2 (L^B + \sigma_B^2)$$

Given now that $d_A, d_B > 0$ and also that $\sigma_A, \sigma_B > 0$, by considering d_A, d_B as continues and not integer values we introduce the Euclidean distance d from a point (d_A, d_B) to the line $d_A(L^A + \sigma_A^2) = d_A(L^B + \sigma_B^2)$:

$$d = \frac{\left|d_A(L^A + \sigma_A^2) - d_B(L^B + \sigma_B^2)\right|}{\sqrt{L^A + \sigma_A^2 + L^B + \sigma_B^2}}$$

As d decreases, it indicates that the point (d_A, d_B) is getting closer to the line $d_A(L^A + \sigma_A^2) = d_A(L^B + \sigma_B^2)$. When $d = 0$ we have shown in Lemma 16 that the Accuracy Cost function

is minimized. As such when d decreases, the numerator $\left|d_A(L^A + \sigma_A^2) - d_B(L^B + \sigma_B^2)\right|$ decreases. Consequently, $Cost^A$ decreases because the ratios d_A/D and d_B/D deviate less form their optimal values, which are attained at $d_A(L^A + \sigma_A^2) = d_A(L^B + \sigma_B^2)$.

□

Based now on Lemma 20 we can posit that the closer we are to the C value the closer we are at the minimum. Because of this even when the C value cannot be reached we know that the Accuracy Cost minimum corresponds to the ratio d_A/d_B that is the closest to C. This can be seen visually by plotting the graph of the Accuracy Cost function in Figure 4.4. Moreover from our simulation results we can see that for reasonably large graphs most of the time even the Greedy search Algorithm can reach values very close to the minimum, while the method of exhaustive search can tell us if the minimum is achievable and exactly under which steps.

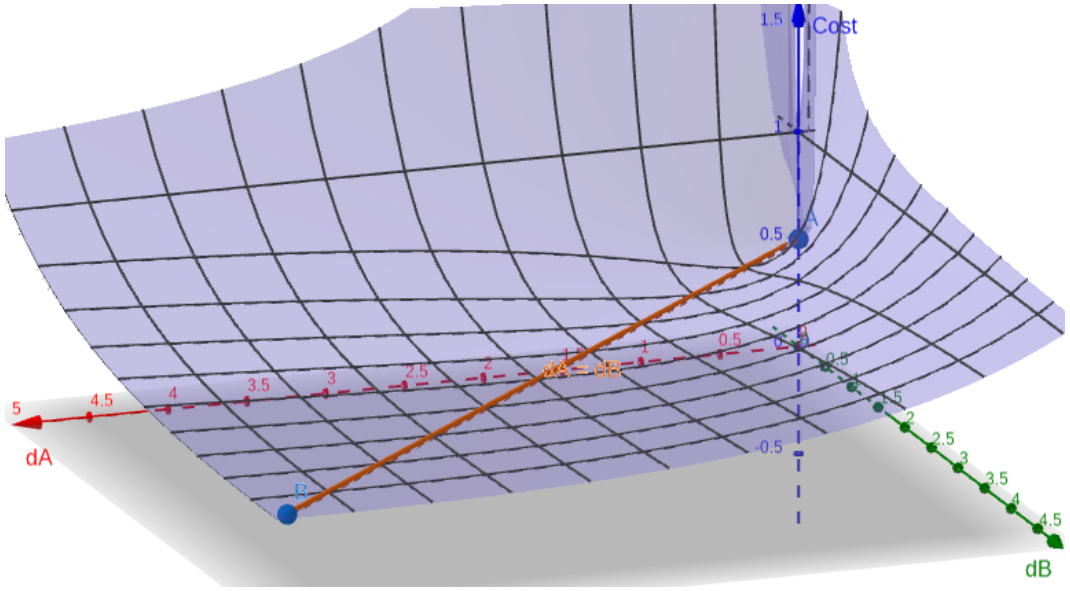

Figure 4.4. *Accuracy Cost function for the undirected graph with* $L^A = L^B = 1$ *and* $\sigma_A = \sigma_B = 0$. *We see that the minimum values are in the line* $d_A = d_B$ *since:* $C = L^A/L^B = 1$

We can also answer this question from the point of view of the Disagreement Cost (3.5). So now given a network G of n_A and n_B agents how can we change the edge connections between the nodes in order to minimize the Disagreement Cost? The answer turns out to be much simpler compared to the Accuracy:

Lemma 21 (Undirected Graph Minimum Disagreement Cost). Consider a team of n_A members of type A and n_B members of type B connected based on the graph $G(n_A, n_B, q, p)$. The Disagreement Cost is minimized by removing as many of the graph's edges as possible.

Proof. According to Lemma 9 the team's Disagreement Cost can be written as:

$$Cost^D = \frac{2}{|T|^{1+\beta}}\left(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2\right)$$

Taking the derivatives with respect to d_{AA}, d_{BB} and d_{AB}, we obtain:

$$\frac{\partial}{\partial d_{AA}}\left(Cost^D\right) = \frac{2\sigma_A^2}{(n_A+n_B)^{1+\beta}} > 0$$
$$\frac{\partial}{\partial d_{BB}}\left(Cost^D\right) = \frac{2\sigma_B^2}{(n_A+n_B)^{1+\beta}} > 0$$
$$\frac{\partial}{\partial d_{AB}}\left(Cost^D\right) = \frac{2(L^A+L^B+\sigma_A^2+\sigma_B^2)}{(n_A+n_B)^{1+\beta}} > 0$$

As such removing edges of any type always reduces the value of the Disagreement Cost.

□

As expected adding edge connection in the graph, based on how we defined the Disagreement Cost, will make $Cost^D$ get higher. We must however keep in mind that for our analysis to work the graph must be kept strongly connected. So there is a lower limit to the amount of edges that can be removed in order to reach the minimum value of the Disagreement Cost.

So what occurs when we combine the Disagreement with the Accuracy Cost? The total Cost, denoted as $Cost^T$ can be written as:

$$\begin{aligned} Cost^T &= \lambda \times Cost^D + (1-\lambda) \times Cost^A \\ &= \lambda \times \left(\frac{2}{|T|^{1+\beta}}\left(d_{AB}(L^A+L^B+\sigma_A^2+\sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2\right)\right) \\ &+ (1-\lambda) \times \left(\left(\frac{d_{AA}+d_{AB}}{d_{AA}+d_{AA}+2d_{AB}}\right)^2 (L^A+\sigma_A^2) + \left(\frac{d_{BB}+d_{AB}}{d_{AA}+d_{AA}+2d_{AB}}\right)^2 (L^B+\sigma_B^2)\right) \end{aligned}$$

For this more complicated case we choose to study the model without any noise on the outcome functions ($\sigma_A = \sigma_B = 0$). As such from now on we will study $Cost^T$ in it's simplified form as:

$$Cost^T = \lambda \times \left(\frac{2(L^A+L^B)}{|T|^{1+\beta}} d_{AB}\right) + (1-\lambda) \times \left(L^A\left(\frac{d_{AA}+d_{AB}}{d_{AA}+d_{AA}+2d_{AB}}\right)^2 + L^B\left(\frac{d_{BB}+d_{AB}}{d_{AA}+d_{AA}+2d_{AB}}\right)^2\right)$$

Let's begin by taking the derivatives of $Cost^T$ with respect to d_{AA}, d_{BB} and d_{AB}:

$$\frac{\partial}{\partial d_{AA}}\left(Cost^T\right) = (1-\lambda) \times \frac{2(d_{BB}+d_{AB})}{(d_{AA}+d_{BB}+2d_{AB})^3}\left(L^A(d_{AA}+d_{AB}) - L^B(d_{BB}+d_{AB})\right)$$
$$\frac{\partial}{\partial d_{BB}}\left(Cost^T\right) = (1-\lambda) \times \frac{2(d_{AA}+d_{AB})}{(d_{AA}+d_{BB}+2d_{AB})^3}\left(L^B(d_{BB}+d_{AB}) - L^A(d_{AA}+d_{AB})\right)$$
$$\frac{\partial}{\partial d_{AB}}\left(Cost^T\right) = (1-\lambda) \times \frac{2(d_{BB}-d_{AA})}{(d_{AA}+d_{BB}+2d_{AB})^3}\left(L^A(d_{AA}+d_{AB}) - L^B(d_{BB}+d_{AB})\right) + \lambda \times \frac{2(L^A+L^B)}{|T|^{1+\beta}}$$

We can observe from these results that for non-limit values of λ, i.e. $\lambda \in (0,1)$, the above three derivatives cannot be equal to zero at the same time. This it true because the derivatives of d_{AA} and d_{BB} are zeroed out only when: $L^A(d_{AA}+d_{AB}) = L^B(d_{BB}+d_{AB})$. However when this happens the derivative of d_{AB} is strictly positive since it is equal to: $\lambda \times \frac{2(L^A+L^B)}{|T|^{1+\beta}} > 0$. We remind the fact that: $\sigma^A, \sigma^B, |T|, d_{AA}, d_{BB}, d_{AB} > 0$.

Based on this observation we can deduce that the minimum for the Total Cost function will be at the bounds of the d_{AA}, d_{BB} and d_{AB} values. In order to make clear where those minimums are, lets run some simulations. We will start on Figure 4.5 by plotting on 3D space with (x, y, z) coordinates being the values of (d_{AA}, d_{BB}, d_{AB}), the 150 points where $Cost^T$ has the lowest value. Let's begin with the case where $\lambda = 0$, i.e. there is no Disagreement Cost and also: $L^A = 0.1, L^B = 0.2, |T| = 14, \beta = 1$.

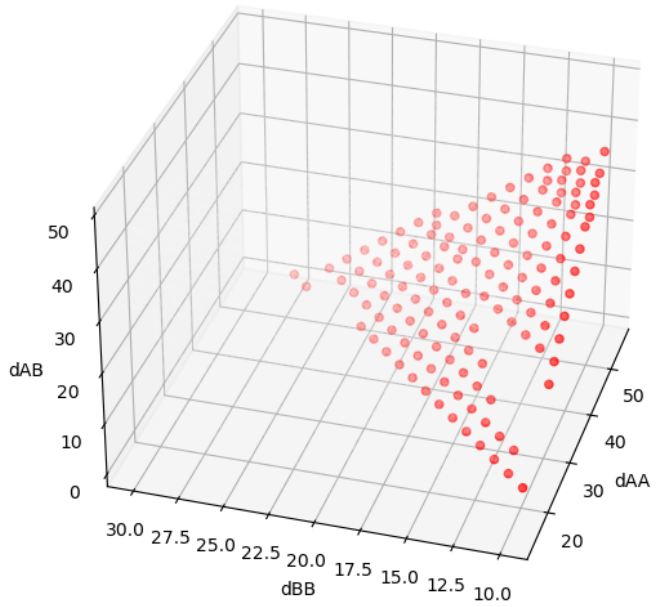

Figure 4.5. *$Cost^T(d_{AA}, d_{BB}, d_{AB})$ plot with red colored points where the Total Cost function is minimized. $L^A = 0.1, L^B = 0.2, |T| = 14, \beta = 1$ and $\lambda = 0$.*

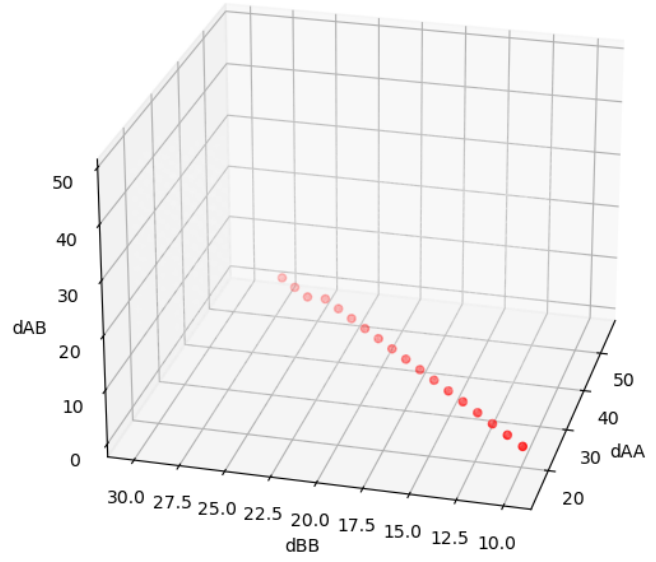

Figure 4.6. *$Cost^T(d_{AA}, d_{BB}, d_{AB})$ plot with red colored points where the Total Cost function is minimized. $L^A = 0.1, L^B = 0.2, |T| = 14, \beta = 1$ and $\lambda = 0.001$.*

From Figure 4.5 we can clearly see that all the minimum points of just the Accuracy Cost function fall on the plane $L^A(d_{AA} + d_{AB}) = L^B(d_{BB} + d_{AB})$. This plane is equivalent with Figures 4.4 line: $(L^A + \sigma_A^2)d_A = (L^B + \sigma_B^2)d_B$. Now what will occur if we introduce Disagreement in the model, i.e. a small $\lambda > 0$ value? While keeping $L^A, L^B, |T|, \beta$ the same and changing λ from: $\lambda = 0$ to $\lambda = 0.001$, we can see in Figure 4.6 that the 18 points where $Cost^T$ has the lowest value stay on the plain $L^A(d_{AA} + d_{AB}) = L^B(d_{BB} + d_{AB})$ but move to where the value of $d_{AB} = 1$. We can show that this result holds theoretically using the Karush–Kuhn–Tucker conditions:

Definition (KKT Conditions). Consider a nonlinear optimization problem with differentiable objective function and functional constraints in the form:

$$\begin{aligned} &\min_x f(x) \\ &\text{s.t.} \quad \begin{aligned} h_i(x) &= 0, \quad i \in \{1, \dots, r\} \\ g_j(x) &\le 0, \quad j \in \{1, \dots, s\} \end{aligned} \end{aligned}$$

The KKT conditions at point x are given by:

$$-\nabla f(x) = \sum_{i=1}^{r} \lambda_i \nabla h_i(x) + \sum_{j=1}^{s} \mu_j \nabla g_j(x) \qquad \text{("Stationarity")}$$

$$\lambda_i \in \mathbb{R}\ ,\ \mu_j \ge 0,\ \forall i, j \qquad \text{("Dual feasibility")}$$

$$\mu_j \cdot g_j(x) = 0, \forall j \qquad \text{("Complementary slackness")}$$

We can draw an intuative understanding of how the KKT conditions generalize the method of Lagrange multipliers by allowing for inequality constraints from Figure (4.7).

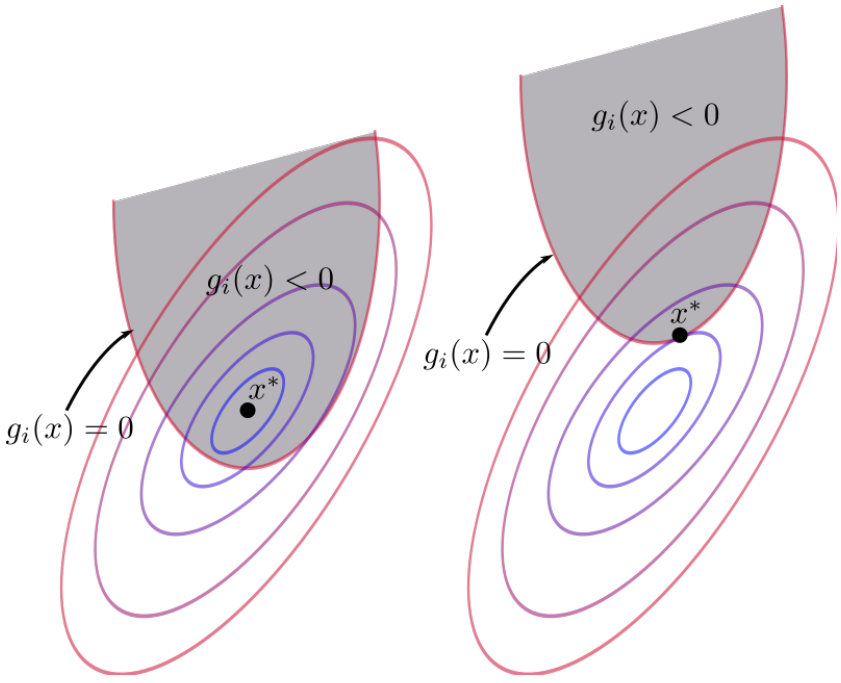

Figure 4.7. *KKT Condition Inequality constraint diagram for optimization problems*

Applying the KKT condition in our case we have that:

$$\begin{aligned} min_{\mathbf{x}} f(\mathbf{x}) &= Cost^T(d_{AA}, d_{BB}, d_{AB}) \\ g_1 &= 1 - d_{AB} \leq 0 \\ g_2 &= d_{AB} - n_A n_B \leq 0 \\ g_3 &= 2(n_A - 1) - d_{AA} \leq 0 \\ g_4 &= d_{AA} - n_A(n_A - 1) \leq 0 \\ g_5 &= 2(n_B - 1) - d_{BB} \leq 0 \\ g_6 &= d_{BB} - n_B(n_B - 1) \leq 0 \end{aligned} \quad \text{s.t.}$$

From Stationarity we get the following 3 equations:

$$\begin{bmatrix} \frac{\partial}{\partial d_{AA}}\left(Cost^T\right) \\ \frac{\partial}{\partial d_{BB}}\left(Cost^T\right) \\ \frac{\partial}{\partial d_{AB}}\left(Cost^T\right) \end{bmatrix} = \begin{bmatrix} \mu_3 - \mu_4 \\ \mu_5 - \mu_6 \\ \mu_1 - \mu_2 \end{bmatrix}$$

Also from Dual feasibility we know that all μ_j values are non-negative while from the 6 Complementary slackness equations, we can choose whether μ_j or g_j equals zero. We can make two observations using the Complementary slackness equations. Firstly all the Lagrange multipliers μ_j cannot equal to zero since in that case we get the equation: $\nabla Cost^T = 0$, which as we have discussed before does not have a solution. As such at least one μ_j must stay positive. Furthermore the pairs of Lagrange multipliers: (μ_1, μ_2) or (μ_3, μ_4) or (μ_5, μ_6) cannot both be equal to zero at the same time. This is true because for example when: $\mu_3 = \mu_4 = 0$ then from the Complementary slackness equations we get that: $d_{AA} = 2(n_A - 1)$ and $d_{AA} = n_A(n_A - 1)$. In order for this to hold n_A must be equal to 2 which restrains a parameter we want to be freely adjustable (the number of type A or B agents). Similar restraining results hold for the other two pairs (μ_1, μ_2) and (μ_5, μ_6), which shows that these pairs of Lagrange multipliers cannot both be positive at the same time. This now leaves us with a specific number of μ_j combinations that are applicable, one of which is:

$$\mu_3 = \mu_4 = \mu_5 = \mu_6 = 0 \Rightarrow L^A(d_{AA} + d_{AB}) = L^B(d_{BB} + d_{AB})$$
$$\frac{\partial}{\partial d_{AB}}\left(Cost^T\right) = \lambda \times \frac{2(L^A + L^B)}{|T|^{1+\beta}} = \mu_1 - \mu_2$$

where the value $\lambda \times \frac{2(L^A+L^B)}{|T|^{1+\beta}}$ is positive and since both μ_1, μ_2 are non negative and at least one of the two Lagrange multipliers must be equal to zero then:

$$\mu_2 = 0 \Rightarrow \mu_1 > 0 \Rightarrow d_{AB} = 1$$

As such for this case of the Lagrange multipliers we get the result hinted by the simulation of Figure (4.6) where the minimum is on the plane $L^A(d_{AA} + d_{AB}) = L^B(d_{BB} + d_{AB})$ where d_{AB} is minimized i.e. $d_{AB} = 1$. We can check one by one the other combinations of

μ_j and find more points (d_{AA}, d_{BB}, d_{AB}) that satisfy the KKT conditions, however this is not necessary since we know that this combination is the one that minimizes the $Cost^T$ function. We can conclude this because: $Cost^T = \lambda \times Cost^D + (1-\lambda) \times Cost^A$ and it is true that: $min(Cost^D) \Rightarrow d_{AB} = 1$ and $min(Cost^A) \Rightarrow L^A(d_{AA} + d_{AB}) = L^B(d_{BB} + d_{AB})$. Thus while respecting the KKT condition this result gives us the true minimum.

4.3 Pyramid Graph: Optimal Positioning

Let's now consider a new problem that arises from the introduction of the pyramid graph in the base model. Unlike the undirected graph analyzed in the previous section, the pyramid graph does not allow the altering of the edges due to it's rigid structure. Therefore, our scope for adjustments is limited to the positioning of agent types A or B within the network. So a natural question to ask is: given a pyramid graph G of n_A and n_B agents what is the optimal way to position them in the pyramid in order to minimize the Accuracy Cost? Let's start answering this question by first finding when the Pyramid graph Accuracy Cost is minimized and under which condition.

Lemma 22 (Pyramid Graph Minimum Accuracy Cost). Consider a team of n_A members of type A and n_B members of type B connected on a full pyramid graph. The Cost function will be minimized when $\sum_{n=1}^{\ell} i_n s_n = (L^B+\sigma_B^2)/(L^A+\sigma_A^2+L^B+\sigma_B^2)$ or equivalently $\sum_{n=1}^{\ell} i'_n s_n = (L^A+\sigma_A^2)/(L^A+\sigma_A^2+L^B+\sigma_B^2)$ and will be equal to:

$$Cost^* = \frac{(L^A+\sigma_A^2)(L^B+\sigma_B^2)}{(L^A+\sigma_A^2)+(L^B+\sigma_B^2)} \tag{4.4}$$

Proof. According to Lemma 15 (3.9) the team's Accuracy Cost can be rewritten as:

$$\begin{aligned} Cost^A &= \left(\sum_{n=1}^{\ell} i_n s_n\right)^2 (L^A+\sigma_A^2) + \left(\sum_{n=1}^{\ell} i'_n s_n\right)^2 (L^B+\sigma_B^2) \\ &= \left(\sum_{n=1}^{\ell} i_n s_n\right)^2 (L^A+\sigma_A^2) + \left(1-\sum_{n=1}^{\ell} i_n s_n\right)^2 (L^B+\sigma_B^2) \\ &= \left(\sum_{n=1}^{\ell} i_n s_n\right)^2 (L^A+\sigma_A^2+L^B+\sigma_B^2) - 2\left(\sum_{n=1}^{\ell} i_n s_n\right)(L^B+\sigma_B^2) + (L_B+\sigma_B^2) \end{aligned}$$

Note that since $L^A, L^B, \sigma_A, \sigma_B > 0$, the above is an always positive quadratic polynomial in $\sum_{n=1}^{\ell} i_n s_n$ with a positive leading coefficient. Taking the derivative of the right hand side with respect to $\sum_{n=1}^{\ell} i_n s_n$, we can find the minimum:

$$\frac{\partial\left(Cost^A\right)}{\partial\left(\sum_{n=1}^{\ell} i_n s_n\right)} = 2\left(\sum_{n=1}^{\ell} i_n s_n\right)(L^A+\sigma_A^2+L^B+\sigma_B^2) - 2(L^B+\sigma_B^2)$$

To obtain the zero of the derivative, we can write:

$$2\left(\sum_{n=1}^{\ell} i_n s_n\right)(L^A + \sigma_A^2 + L^B + \sigma_B^2) - 2(L^B + \sigma_B^2) = 0$$

$$\sum_{n=1}^{\ell} i_n s_n = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$$

Following the same steps but solving for $\sum_{n=1}^{\ell} i'_n s_n$ instead of $\sum_{n=1}^{\ell} i_n s_n$ we can find that the minimum is reached when: $\sum_{n=1}^{\ell} i'_n s_n = (L^A + \sigma_A^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$. Now replacing these values to (3.9) we can find the minimum Cost as:

$$Cost^* = \frac{(L^A + \sigma_A^2)(L^B + \sigma_B^2)}{(L^A + \sigma_A^2) + (L^B + \sigma_B^2)}$$

□

Lets now consider what this result tells us for the optimal distribution of type A and type B agents into our hierarchical structure from an Accuracy point of view. Lemma 22 shows that in order to reach the Accuracy Cost minimum we must solve the equation:

$$i_1 s_1 + i_2 s_2 + \cdots + i_{\ell-1} s_{\ell-1} + i_\ell s_\ell = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$$

where $L^A, L^B, \sigma_A, \sigma_B$ are known constants initialized at the start of the simulation and the values of s_n can be found using Theorem 3 based on the k, ℓ values of the pyramid and the parameter γ. As such this equation has ℓ unknowns which are the values of $i_n \in \mathbb{N}$. The i_n are bounded and can take the values: $i_1 = \{0, 1\}$, $i_2 = \{0, 1, \ldots, k\}$, $i_3 = \{0, 1, \ldots, k^2\}$, $\ldots$, $i_{\ell-1} = \{0, 1, \ldots, k^{\ell-2}\}$, $i_\ell = \{0, 1, \ldots, k^{\ell-1}\}$. Also it must hold that: $i_1 + i_2 + \cdots + i_\ell = n_A$ as well as that: $i'_1 + i'_2 + \cdots + i'_\ell = n_B$. Because of these constraints this equation can have many or no solutions. There can be no solutions in the case that the parameters $L^A, L^B, \sigma_A, \sigma_B$ and s_n cannot produce any combinations of i_n which can reach the target, as well as many solutions in cases where different combinations of i_n can solve the equation. For example in the case of a pyramid with hyperparameters: $k = 3, \ell = 3, \gamma = 2$ and Losses: $L^A = L^B$, while $\sigma_A = \sigma_B = 0$, there are 4 combinations of i_n values that fulfill the requirements of equation: $\sum_{n=1}^{\ell} i_n s_n = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$ and these combinations are: $(i_1 = 0, i_2 = 2, i_3 = 5)$, $(i_1 = 0, i_2 = 3, i_3 = 0)$, $(i_1 = 1, i_2 = 0, i_3 = 9)$, $(i_1 = 1, i_2 = 1, i_3 = 4)$.

Now a good question to ask is how can we find these i_n optimal values that minimize the Accuracy Cost. To help us answer this question we must first observe that for the pyramid graph occurs something similar with the result of Lemma 20 of the undirected graph. That is the fact that the closer the $\sum_{n=1}^{\ell} i_n s_n$ gets to the new target value: $C = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$, the closer we are at the Accuracy Cost minimum. This can be easily seen from Lemma 22 (4.4) where the form of $Cost^A$ can be written as an always positive parabola. As such even if the target C value cannot be reached from the sum: $\sum_{n=1}^{\ell} i_n s_n$, we know that the best combination of i_n values is the one that reaches the closest to it.

Based on this observation we can again propose, as in the previous subsection, one

exhaustive and one greedy algorithm in order to find the i_n values. The exhaustive search algorithm simply checks all the possible combinations of i_n, which are bounded and can take the values: $i_1 = \{0, 1\}$, $i_2 = \{0, 1, \ldots, k\}$, $i_3 = \{0, 1, \ldots, k^2\}$, $\ldots$, $i_{\ell-1} = \{0, 1, \ldots, k^{\ell-2}\}$, $i_\ell = \{0, 1, \ldots, k^{\ell-1}\}$. By knowing the s_n values from the pyramid's hyperparameters we can check all the i_n combinations in the sum: $\sum_{n=1}^{\ell} i_n s_n$, and simply choose the one that is the closest to the target C. Again however we can see that this method is very computationally expensive especially in the case of pyramids with large values of ℓ, because for the last layers the values of i_n get exponentially large.

To address this we can propose a greedy algorithm which is different from the one used in the previous subsection. This time we can use the fact that the influence of the agents in the top layers is bigger than the influence of the agents in the bottom layers:

$$s_1 > s_2 > s_3 > \cdots > s_{\ell-2} > s_{\ell-1} > s_\ell$$

The greedy algorithm designed for approximating the target sum C works by iteratively selecting integer coefficients i_n for a series of known values s_n such that $\sum_{n=1}^{\ell} i_n s_n$ is as close to C as possible without exceeding it. The algorithm begins by initializing the sum S to zero. It then processes each coefficient i_n starting from i_1 and ending with i_ℓ. For each i_n, it increments i_n from zero, adding it's corresponding s_n to S as long as the resulting sum does not exceed C and i_n remains within its bounded range. If the sum S surpasses C after an increment, the algorithm decrements i_n by one and adjusts S accordingly. This process ensures that the sum S approaches C as closely as possible without overshooting. Finally, the algorithm returns the set of coefficients $i_1, i_2, \ldots, i_\ell$ which represent the closest possible approximation to the target sum C using the given values s_n. This approach leverages the properties of greedy algorithms by making locally optimal choices at each step. We can think of this approach as beginning with large steps when using s_1, s_2 and as we get closer to the target C, we use smaller and smaller steps.

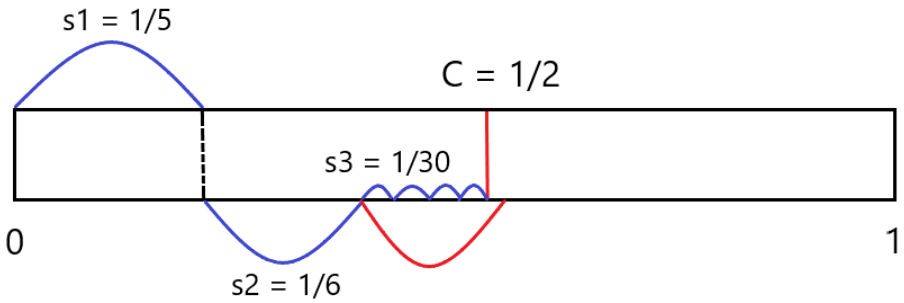

Figure 4.8. *Example of applying the Greedy Algorithm in the case of the pyramid graph with hyperparameters: $k = \ell = 3, \gamma = 2$ and $L^A = L^B, \sigma_A = \sigma_B = 0$. In this case we can see that greedy does find one of the four optimal solutions*

Conjecture 1 (Greedy Algorithm is Optimal). In order to find the steps i_n that minimize the Pyramid graph's Accuracy Cost (3.9) the proposed Greedy Algorithm is Optimal and does find one correct set of i_n values if and only if C is reachable. C is defined as reachable when there exists at least one combination of i_n that satisfy: $\sum_{n=1}^{\ell} i_n s_n = C$.

Intuition It is easy see from Theorem 3 (3.6) that the step sizes: $s_2, s_3 \dots s_{\ell-1}, s_\ell$ are perfect divisors of one another. To be exact: $s_{\ell-1}$ contains $\gamma + k$ steps of size s_ℓ, $s_{\ell-2}$ contains γ steps of size $s_{\ell-1}$, $s_{\ell-3}$ contains γ steps of size $s_{\ell-2}$. This pattern continues up to: s_2 which contains γ steps of size s_3. The only step which breaks this pattern is s_1, which can easily be written as an linear combination of the next $\ell - 1$ steps, i.e. $s_1 = as_2 + bs_3 + \dots + cs_\ell$. These observations give us an intuitive understanding of why when we choose the target to be reached with a combination of s_2, s_3 and s_5 steps, the greedy algorithm is able to reach it with a combination of s_1, s_2 and s_6 steps.

The following pseudocode describes the greedy algorithm for approximating the target sum C using the provided values s_n and integer coefficients i_n. Running many simulations with many different pyramid hyperparameters k, ℓ, γ and many different reachable targets C we have consistently confirmed Conjecture 1.

```
Input: target C, values s_n, parameter k, length ℓ
Output: coefficients i_1, i_2, ..., i_ℓ or None if no solution is found
Initialize solution ← [0, 0, ..., 0] (length ℓ)
Initialize current_sum ← 0
Initialize ε ← 1 × 10^−9
for i = 0 to ℓ − 1 do
    for j = 0 to k^i − 1 do
        if current_sum + s_i ≤ C + ε then
            solution[i] ← solution[i] + 1
            current_sum ← current_sum + s_i
            if |current_sum − C| < ε then
                return solution
            end if
        end if
    end for
end for
if |current_sum − C| < ε then
    return solution
else
    return None
end if
```

We can also answer this question from the point of view of the Disagreement Cost (3.5). Again we keep the definition that in our network model, Disagreement happens only between the agents that are directly connected in the graph i.e. they have an edge that connects them directly. As such the results of Lemma 9 hold for the pyramid network and the Disagreement Cost can be written as:

$$Cost^D = \frac{2}{|T|^{1+\beta}} \left(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2 \right)$$

Now in the pyramid network structure we must remind again that we do not have the freedom to directly change the edges by connecting or disconnecting specific nodes since

the network structure is rigid and predetermined. We have the power however to choose in which positions agents of type A and type B go to. As such by optimally choosing the positioning of the agents in the pyramid we can adjust the values of d_{AA}, d_{BB} and d_{AB}. Let's think of the case where there is no noise in the predictions of the agents and as such: $\sigma_A = \sigma_B = 0$. In this simplified case only connections between nodes of type A and type B create Disagreement Cost and as such we must minimize the value of d_{AB}. The obvious solution that minimizes the Disagreement Cost in this case is when there are only type A or only type B agents in the pyramid, however what happens when we must position $n_A > 0$ type A and $n_B > 0$ type B nodes in the pyramid? The answer is that in order to keep the Disagreement Cost at a minimum we must position the agents of the same type at sub-pyramids, thus minimizing the value of d_{AB}. This can be seen visually for different values of n_A and n_B in Figure 4.9

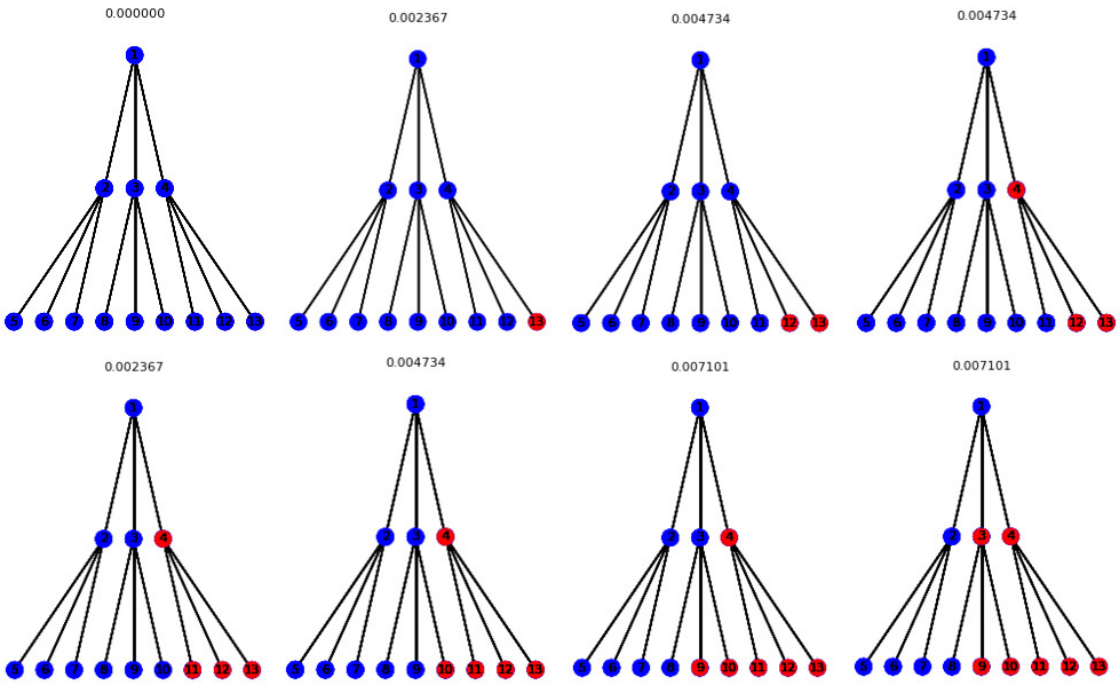

Figure 4.9. *Optimal Positioning of type A, B agents in k=ℓ=3 pyramid with only Disagreement Cost ($\lambda = 1$). New n_A agents position themselves inside a sub-pyramid*

So now for the pyramid graph what occurs when we combine the Disagreement with the Accuracy Cost? The total Cost, denoted as $Cost^T$ can be written as:

$$\begin{aligned} Cost^T &= \lambda \times Cost^D + (1-\lambda) \times Cost^A \\ &= \lambda \times \left(\frac{2}{|T|^{1+\beta}} \left(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2 \right) \right) \\ &+ (1-\lambda) \times \left(\left(\sum_{n=1}^{\ell} i_n s_n \right)^2 (L^A + \sigma_A^2) + \left(\sum_{n=1}^{\ell} i'_n s_n \right)^2 (L^B + \sigma_B^2) \right) \end{aligned}$$

For this more complicated case again we choose to study the model without any noise on the outcome functions ($\sigma_A = \sigma_B = 0$). As such from now on we will study $Cost^T$ in it's

simplified form as:

$$Cost^T = \lambda \times \left(\frac{2(L^A + L^B)}{|T|^{1+\beta}} d_{AB} \right) + (1 - \lambda) \times \left(L^A \left(\sum_{n=1}^{\ell} i_n s_n \right)^2 + L^B \left(\sum_{n=1}^{\ell} i'_n s_n \right)^2 \right)$$

Right away it is clear that for the pyramid graph we cannot try to find a theoretical solution as we did in the previous sub-section for the undirected graph. This is because we are not able to take derivatives of $Cost^T$ with respect to d_{AA}, d_{BB}, d_{AB} since now the Accuracy Cost term only considers the layer positioning of the agent types (i_n) and not their edge connections. So we choose to deploy a pure algorithmic approach for the pyramid graph's Total Cost. We will deploy 3 algorithms: an Exhaustive, a Greedy and a Local Search.

The Exhaustive algorithm explores all possible 2^n type combinations, where n is the number of nodes in the k, ℓ pyramid. For each combination, it calculates the corresponding $Cost^T$ and selects the one with the lowest value. As expected, this approach becomes very computationally expensive as the pyramid graph grows larger. The Greedy algorithm is the least computationally demanding. It starts by assigning the same type to all nodes in the pyramid. Then, from top to bottom, it attempts to change the type of each node one by one. Each change is kept only if it reduces the $Cost^T$. It is important to note that this greedy algorithm differs from the one described for the Accuracy Cost and shown in Figure (4.8). In this case, we do not have a target value C to indicate when we overshoot, so this greedy algorithm is not optimal, even when $\lambda = 0$ (i.e., when only the Accuracy Cost is considered). The Local Search/Swaps algorithm begins by randomly assigning types to all nodes in the pyramid graph. It then iteratively improves the assignment by considering pairs of nodes with different types. For each pair, the algorithm swaps their types and calculates the new cost. If the swap results in a lower cost, the change is kept; otherwise, it is reverted. This process continues until no further improvements can be found. The swaps algorithm is more computationally expensive compared to the greedy algorithm but less so compared to the exhaustive search.

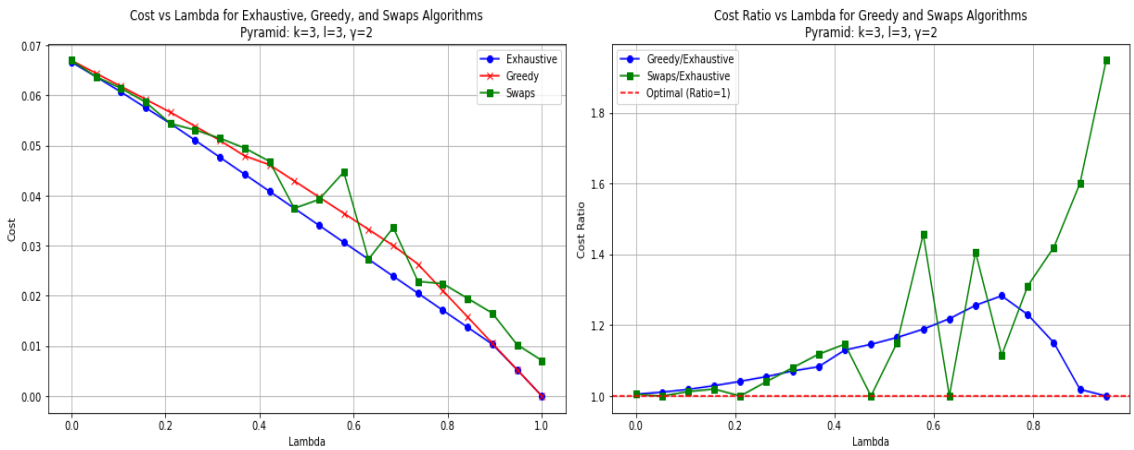

Figure 4.10. *Cost comparison and cost ratio for Exhaustive, Greedy and Swaps algorithms for different values of λ. Pyramid parameters: $k = 3$, $\ell = 3$, $\gamma = 2$*

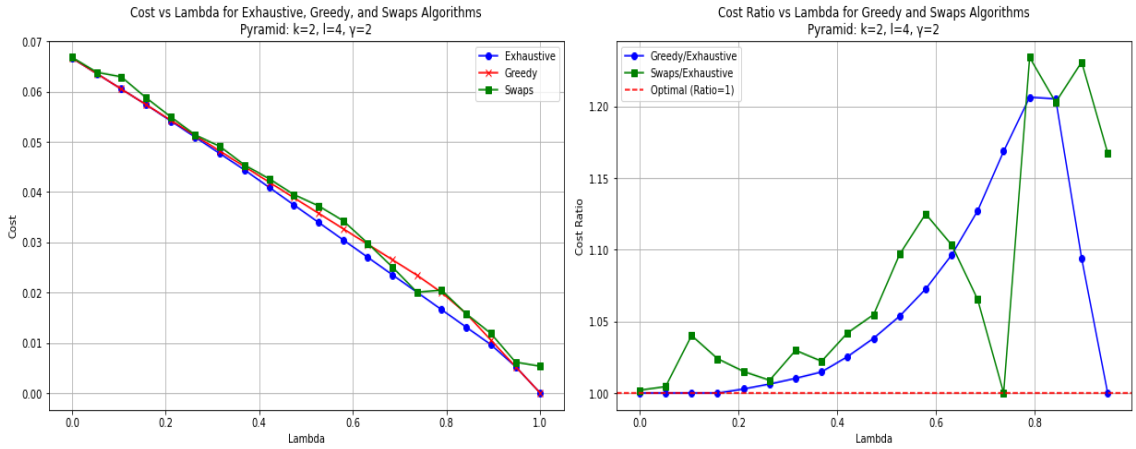

Figure 4.11. *Cost comparison and cost ratio for Exhaustive, Greedy and Swaps algorithms for different values of λ. Pyramid parameters: $k = 2$, $\ell = 4$, $\gamma = 2$*

The performance of the three algorithms is illustrated in Figures (4.10) and (4.11). These figures compare the cost values obtained by the Exhaustive, Greedy, and Swaps algorithms across a range of λ values. In both figures, the Greedy algorithm demonstrates its effectiveness particularly for λ values close to 0 or 1, where it closely approximates the results of the Exhaustive algorithm. This suggests that the Greedy algorithm is well-suited for scenarios where either the Accuracy Cost or the Disagreement Cost predominates. Conversely, the Swaps algorithm often outperforms the Greedy algorithm at intermediate λ values, showcasing its strength in balancing both costs. However, it is crucial to note that the performance of the Swaps algorithm is highly sensitive to the initial random assignment of node types in the pyramid. Different random initializations can lead to varying results, which highlights the algorithm's dependence on its starting configuration.

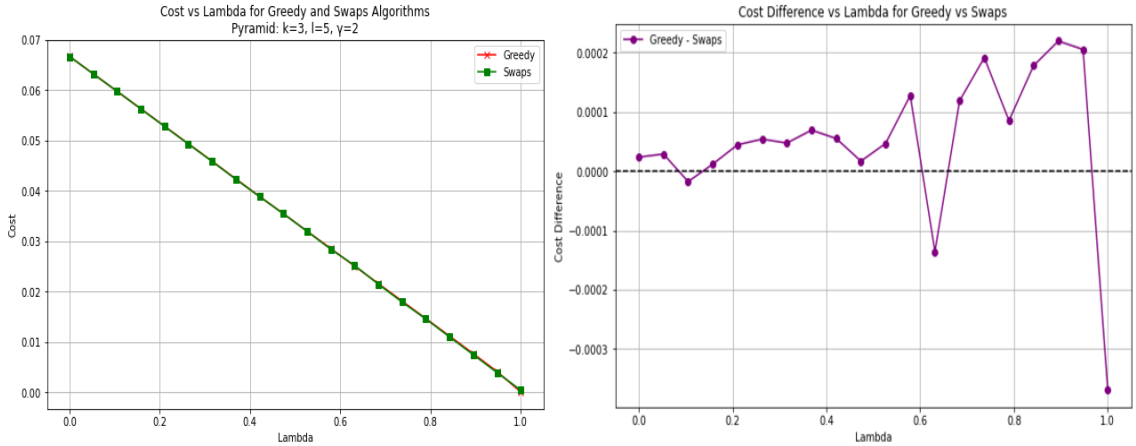

Figure 4.12. *Cost comparison and cost difference for Greedy and Swaps algorithms for different values of λ. Pyramid parameters: $k = 3$, $\ell = 5$, $\gamma = 2$*

Figure (4.12) focuses on a larger pyramid structure with $k = 3$, $\ell = 5$, resulting in 121 nodes. Due to the significant computational demands, the Exhaustive algorithm was not

feasible for this case, and thus, only the results from the Greedy and Swaps algorithms are presented. Remarkably, both algorithms yield very similar results, as evidenced by the left plot where their lines almost overlap. The right plot, which shows the cost difference between the two algorithms, indicates that their differences are on the order of 10^{-4}. This negligible difference underscores the efficiency of both algorithms in approximating the minimal $Cost^T$. For example, at $\lambda = 1$, where only the Disagreement Cost is considered, both algorithms achieve values close to zero. The Greedy algorithm, in particular, reaches exactly $Cost^T = 0$ since it starts with all nodes assigned the same type, reflecting its effectiveness in this limit case.

Chapter 5

Discussion

5.1 Future Work & Limitations

Building on the research presented in this thesis, several avenues for future work can be suggested to enhance and expand these findings. Firstly, the algorithms proposed, although generally effective for large graphs, can perform poorly for specific hyperparameter values, deviating significantly from the optimal value. Future work should focus on developing more sophisticated algorithms to achieve optimal positioning in the pyramid and optimal edge selection in the undirected graph. Additionally, proving Conjecture 1 remains a critical task for future research.

Moreover, our current models use fixed values for the weights in the edge connections between nodes, with the transition matrix T being constant in both cases. Future research should explore integrating learning algorithms that allow the weights to adjust dynamically as rounds of DeGroot or other learning processes occur. This approach would result in a more dynamic graph, better reflecting the evolving relationships and interaction dynamics within teams. Beyond the simplistic DeGroot learning process, where all nodes converge to the same final opinion, future work could investigate other learning algorithms. For instance, applying the Friedkin-Johnsen (FJ) model, where nodes converge to different opinions, could provide a more generalized approach to the Disagreement Cost term in the Total Cost function compared to our current model. This would enable a richer understanding of opinion dynamics within teams.

Despite the advancements and findings presented in this thesis, several limitations should be acknowledged. One critical limitation is the fundamental assumptions of the base model presented in [9]. The base model assumes distinct features of a prediction task that are uncorrelated, so that each type of agent (A, B, or potentially more types) considers a specific subset of these features. However, in reality, it is more plausible that there is some mixture of knowledge between different types of agents. This raises questions about the realism of the model, as it may not fully capture the complexity and interdependence of features in actual prediction tasks.

Another significant limitation lies in the intrinsic nature of modeling social environments. Unlike physical or biological sciences, which are deterministic, social sciences involve human behavior, making them more stochastic and probabilistic. This introduces a fundamental question: how do we know if teams actually operate in the manner

our model suggests? Proving the validity of our model requires empirical experiments, which present their own set of challenges. Determining which teams to study, the size of these teams, and the appropriate methods for experimentation are complex issues that must be addressed to validate the mathematical models beyond their theoretical elegance.

In summary, this thesis has provided significant insights into network team growth dynamics through the development and analysis of various algorithms and models. However, the limitations discussed highlight the need for caution in interpreting the results and underscore the importance of further empirical validation. Future work should aim to address these limitations by developing more robust algorithms, incorporating dynamic and adaptive modeling techniques, and exploring a broader range of learning processes. Only through such continued efforts can we hope to fully understand the complex and nuanced nature of team dynamics and opinion formation. By advancing the research in these directions, we can move closer to models that not only exhibit mathematical rigor but also capture the intricate realities of social interactions and team behaviors. This ongoing quest for understanding will undoubtedly contribute to more effective team management and optimization strategies in various organizational settings.

Κεφάλαιο 6

Εκτεταμένη Περίληψη

Σε αυτό το κεφάλαιο θα παρουσιάσουμε στα ελληνικά τα βασικά αποτελέσματα των κεφαλαίων 3 και 4, που αποτελούν τα καινοτόμα αποτελέσματα αυτής της διπλωματικής εργασίας. Στόχος μας είναι να συνοψίσουμε τα κύρια ευρήματα, λήμματα και θεωρήματα που αναπτύχθηκαν και αποδείχθηκαν στα προηγούμενα κεφάλαια, χωρίς να επαναλάβουμε το βασικό μοντέλο [9] το οποίο παρουσιάσαμε σύντομα στο Κεφάλαιο 2. Παρόλο που θα αναφερθούμε στα βασικά σημεία και στα αποτελέσματα των μαθηματικών μας προσεγγίσεων, δεν θα ξαναγράψουμε τις αποδείξεις, καθώς αυτές είναι ήδη καταγεγραμμένες αναλυτικά στα προηγούμενα κεφάλαια. Με αυτόν τον τρόπο, προσφέρουμε μια συμπυκνωμένη επισκόπηση των νέων γνώσεων στα ελληνικά, διατηρώντας την πλήρη επιστημονική τεκμηρίωση στο κύριο σώμα της διατριβής.

6.1 Θεωρητικές Επεκτάσεις

Στην βάση των θεωρητικών μας επεκτάσεων εξετάσαμε και αναπτύξαμε δύο διαφορετικές δομές γραφημάτων, τον τυχαίο μη κατευθυνόμενο γράφο και τον πυραμιδικό γράφο. Εισάγοντας αυτά τα γραφήματα στο αρχικό μοντέλο σχηματισμού ομάδας, το οποίο αρχικά δεν περιλάμβανε καμία τέτοια δομή, μπορέσαμε να γενικεύσουμε τη διαδικασία σχηματισμού άποψης μέσω της εκμάθησης DeGroot [12]. Επιπλέον η εισαγωγή αυτών των υποκείμενων γραφημάτων στο βασικό μοντέλο αλλάζει τον τρόπο υπολογισμού της διαφωνίας μεταξύ των παικτών της ομάδας, όπως και την συνολική άποψη της ομάδας. Έτσι το συνολικό κόστος ($Cost^T$) μιας ομάδας γενικεύετε.

Το συνολικό κόστος ($Cost^T$) αποτελείται από δύο βασικά μέρη: το κόστος διαφωνίας ($Cost^D$) και το κόστος ακρίβειας ($Cost^A$). Αυτά τα δύο μέρη συνδυάζονται με έναν παράγοντα λ, ο οποίος κυμαίνεται μεταξύ 0 και 1, και ορίζει τη σχετική βαρύτητα κάθε κόστους ως προς το συνολικό. Ο συνολικός τύπος του κόστους δίνεται από τη σχέση:

$$Cost^T = \lambda \times Cost^D + (1 - \lambda) \times Cost^A$$

Ας μελετήσουμε τώρα πως αλλάζει κάθε επιμέρους όρος για τις περιπτώσεις των δύο γραφημάτων που προτείναμε, ξεκινόντας από τον τυχαίο μη κατευθυνόμενο γράφο.

Ο τυχαίος μη κατευθυνόμενος γράφος, που θα τον συμβολίσουμε ως $G(n_A, n_B, p, q)$, δημιουργεί τις ακμές μεταξύ των κόμβων του ανεξάρτητα με πιθανότητα $0 < p < 1$ για κόμβους ίδιου τύπου και $0 < q < 1$ για κόμβους διαφορετικού τύπου. Αλλάζοντας τις τιμές των p και q μπορούμε να προσαρμόσουμε τη συνδεσιμότητα μεταξύ κόμβων ίδιου και διαφορετικου τύπων. Κάθε κόμβος αρχικά έχει μια άποψη $b_i(0) = f^{A,B}(\mathbf{x})$ βάση του τύπου του και μέσω της διαδικασίας εκμάθησης DeGroot όλοι οι κόμβοι συγκλίνουν σε μια τελική κοινή άποψη που θα ορίσουμε ως τη συγκεντρωτική άποψη $\mathcal{G}_G^{n_A,n_B}(\mathbf{x})$. Κάνοντας χρήση μιας συγκεκριμένης εκδοχής εκμάθησης DeGroot μπορούμε να αποδείξουμε ότι η συγκεντρωτική άποψη της ομάδας για τον μη κατευθυνόμενο γράφο είναι:

$$\mathcal{G}_G^{n_A,n_B}(\mathbf{x}) = \left(\frac{d_A}{D}\right) f^A(\mathbf{x}) + \left(\frac{d_B}{D}\right) f^B(\mathbf{x})$$

όπου $d_A = \sum_{i \in N_A} d_i$, $d_B = \sum_{i \in N_B} d_i$ και $D = d_A + d_B$, με N_A, N_B να είναι τα σύνολα των κόμβων τύπου A,B και d_i οι βαθμοί των κόμβων αυτών. Απευθείας μπορούμε να παρατηρήσουμε πως η επέκτασή μας αλλάζει τον απλό υπολογισμό συγκεντρωτικής γνώμης που χρησιμοποιεί το βασικό μοντέλο. Υπενθυμίζουμε εδώ πως στο βασικό μοντέλο [9] η συγκεντρωτική γνώμη της ομάδας υπολογιζεται βάση του Tullock Aggregate [10] και δίνεται από την σχέση:

$$\mathcal{G}_{n_A,n_B}^{a}(\mathbf{x}) = \left(\frac{n_A^a}{n_A^a + n_B^a}\right) f^A(\mathbf{x}) + \left(\frac{n_B^a}{n_A^a + n_B^a}\right) f^B(\mathbf{x})$$

όπου n_A, n_B ο αριθμός των κόμβων τύπου A και B και $a \in [0, \infty)$. Στην περίπτωσή του τυχαίου μη κατευθυνόμενου γράφου λοιπόν έχουν αντικατασταθεί τα n_A, n_B με τους βαθμούς των κόμβων d_A, d_B. Χρησιμοποιώντας αυτήν τη συνάρτηση υπολογισμού της συγκεντρωτικής γνώμης, μπορούμε τώρα να υπολογίσουμε το νέο Κόστος Ακρίβειας το οποίο είναι:

$$Cost^A = \left(\frac{d_A}{D}\right)^2 (L^A + \sigma_A^2) + \left(\frac{d_B}{D}\right)^2 (L^B + \sigma_B^2)$$

Ας συνεχίσουμε αναλύοντας πώς αλλάζει το κόστος Διαφωνίας με την εισαγωγή του υποκείμενου γράφου. Χρησιμοποιώντας τη διαδικασία εκμάθησης DeGroot, όλοι οι παίκτες φτάνουν στην ίδια τελική άποψη και αυτό αρχικά μπορεί να θεωρηθεί πως ακυρώνει τη διαφωνία εντός της ομάδας. Ωστόσο, θέλουμε να διατηρήσουμε την τριβή μεταξύ του κόστους Ακρίβειας και του κόστους Διαφωνίας όπως και στο βασικό μοντέλο [9]. Μια προσέγγιση για να επιτευχθεί αυτό είναι να ορίσουμε ότι, για το εκτεταμένο μοντέλο μας, η διαφωνία εμφανίζεται μόνο μεταξύ παικτών που επικοινωνούν άμεσα, το οποίο στον υποκείμενο γράφο σημαίνει ότι οι κόμβοι τους συνδέονται απευθείας με μια ακμή. Υιοθετώντας αυτό το κριτήριο, μπορούμε να εισάγουμε τον όρο του κόστους Διαφωνίας ως:

$$Cost^D = \frac{2}{|T|^{1+\beta}} (d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2)$$

όπου $d_A = d_{AA} + d_{AB}$ και $d_B = d_{BB} + d_{AB}$ με d_{AA}, d_{BB} να συμβολίζουν τις ακμές μεταξύ κόμβων ίδιου τύπου (δηλαδή A με A ή B με B) και d_{AB} να συμβολίζει τις ακμές μεταξύ κόμβων διαφορετικού τύπου (δηλαδή από A σε B ή αντιστρόφως).

Ας εξετάσουμε τώρα τη συμπεριφορά του δικτύου $G(n_A, n_B, p, q)$ ως προς το $Cost^A$ και το $Cost^D$ στις δύο οριακές του περιπτώσεις. Αυτές είναι για: $p = q = 1$ (πλήρης γράφος) όπου όλοι οι κόμβοι είναι συνδεδεμένοι με όλους και $p = 1$, $q \approx 0$ (αποσυνδεδεμένος γράφος) όπου οι κόμβοι ίδιου τύπου είναι πλήρως συνδεδεμένοι μεταξύ τους αλλά οι κόμβοι διαφορετικού τύπου συνδέονται μεταξύ τους μόνο με μια ακμή ($d_{AB} = 1$).

Αρχικά στην περίπτωση του πλήρες συνδεδεμένου γράφου παρατηρούμε ότι το κόστος ακρίβειας όπως και το κόστος διαφονείας επανέρχονται στην μορφή του βασικού μοντέλου. Συγκεκριμένα όταν $p = q = 1$ τα δύο κόστη είναι:

$$Cost^A_{(p=1,q=1)} = \left(\frac{n_A}{n_A + n_B}\right)^2 (L^A + \sigma_A^2) + \left(\frac{n_B}{n_A + n_B}\right)^2 (L^B + \sigma_B^2)$$

$$Cost^D_{(p=1,q=1)} = \frac{2}{|T|^{1+\beta}} \left(n_A n_B (L^A + L^B + \sigma_A^2 + \sigma_B^2) + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2\right)$$

Αυτό αποδεικνύει ότι η γενίκευσή μας για τον τυχαίο μη κατευθυνόμενο γράφο ενσωματώνει όλα τα αποτελέσματα του βασικού μοντέλου στην περίπτωση του πλήρους γράφου [9]. Επιπλέον, με την εναλλαγή των τιμών των παραμέτρων p και q, μπορούμε να επεκτείνουμε το βασικό μοντέλο, ανοίγοντας τον δρόμο για νέες ερευνητικές δυνατότητες.

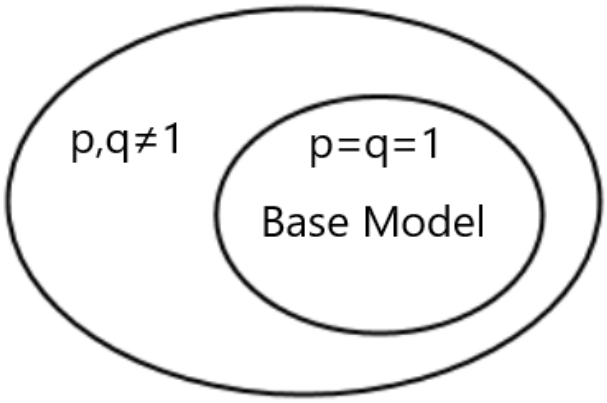

Εικόνα 6.1. *Διάγραμμα Ven που συμβολίζει ότι το βασικό μουτέλο εμπεριέχετε στον τυχαίο μη κατευθυνόμενο γράφο στην περίπτωση $p = q = 1$*

Συνεχίζοντας στην περίπτωση του αποσυνδεδεμένου γράφου παρατηρούμε ότι το κόστος ακρίβειας όπως και το κόστος διαφονείας αλλάζουν. Συγκεκριμένα όταν $p = 1$ και $q \approx 0$ τα δύο κόστη γίνονται:

$$Cost^A_{(p=1,q\approx 0)} = \left(\frac{n_A(n_A - 1)}{n_A(n_A - 1) + n_B(n_B - 1)}\right)^2 (L^A + \sigma_A^2) + \left(\frac{n_B(n_B - 1)}{n_A(n_A - 1) + n_B(n_B - 1)}\right)^2 (L^B + \sigma_B^2)$$

$$Cost^D_{(p=1,q\approx 0)} = \frac{2}{(n_A + n_B)^{1+\beta}} \left(L^A + L^B + \sigma_A^2 + \sigma_B^2 + n_A(n_A - 1)\sigma_A^2 + n_B(n_B - 1)\sigma_B^2\right)$$

Αυτά τα νέα αποτελέσματα δείχνουν ότι η γενίκευσή μας για τον τυχαίο μη κατευθυνόμενο γράφο επεκτείνει το βασικού μοντέλου, παρέχοντας σημαντικά διαφορετικές εξισώσεις τόσο για το κόστος Ακρίβειας όσο και για το κόστος Διαφωνίας.

Συνεχίζοντας με την εισαγωγή του πυραμιδικού γράφου, θέλουμε να εισαγάγουμε μια ιεραρχική δομή εντός της οποίας να λειτουργεί η ομάδα μας. Προηγουμένως, στον τυχαίο μη κατευθυνόμενο γράφο, τη μεγαλύτερη επιρροή είχε ο κόμβος με το μεγαλύτερο βαθμό. Τώρα το γράφημά μας θα ορίζεται όχι με τις παραμέτρους πιθανότητας p και q, αλλά με τις παραμέτρους k, ℓ και με τον γράφο να συμβολίζετε ως $G(n_A, n_B, k, \ell, \gamma)$. Η παράμετρος k ορίζει τον αριθμό των υφισταμένων που θα έχει κάθε κόμβος από κάτω του, ενώ το ℓ τον αριθμό των επιπέδων της πυραμίδας. Στην παρακάτω εικόνα δίνετε ένα παράδειγμα.

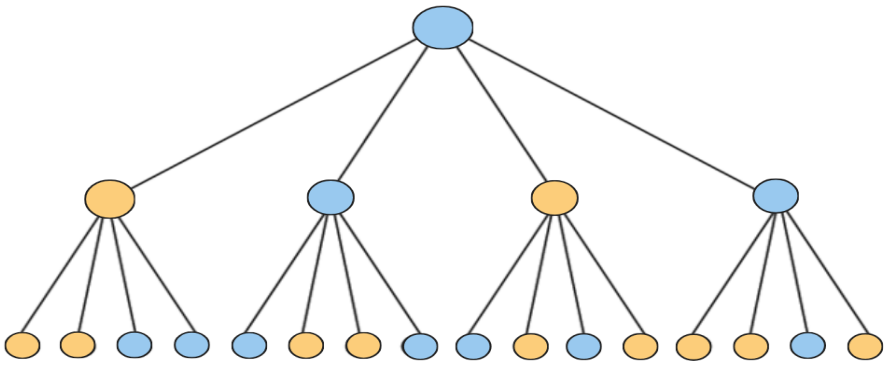

Εικόνα 6.2. *Παράδειγμα πυραμιδικού γράφου με ℓ=3 και k=4*

Για να ενσωματώσουμε τη δυναμική της ιεραρχικής επιρροής μέσα στη δομή της πυραμίδας, εισάγουμε την υπερπαράμετρο γ. Σε αντίθεση με το τυχαίο μη κατευθυνόμενο γράφημα, όπου κάθε κόμβος ακούει ισάξια την άποψη κάθε κόμβου με τον οποίο συνδέετε, τώρα η επιρροή κάθε κόμβου εξαρτάται από τη θέση του στην ιεραρχία. Έτσι για παράδειγμα στην παραπάνω εικόνα οι κόμβοι του 2ου επιπέδου θα ακούσουν εξίσου τους 4 υφισταμένους τους, τον προϊστάμενό τους κόμβο γ φορές περισσότερο, όπως επίσης και τον εαυτό τους κατά ϵ προκειμένου να διατηρηθεί η απεριοδικότητα του γράφου. Βάση αυτών των παραμέτρων μπορεί να υπολογιστεί το νέο διάνυσμα επιρροής $\mathbf{s}$ που δίνεται από την σχέση:

$$\begin{aligned}\mathbf{s} = [\ & \underbrace{k\gamma^{\ell-2}}_{\text{1ος Κόμβος}} \quad \underbrace{\gamma^{\ell-3}(\gamma+k)}_{\text{2ο Επίπεδο } k \text{ κόμβοι}} \quad \underbrace{\gamma^{\ell-4}(\gamma+k)}_{\text{3ο Επίπεδο } k^2 \text{ κόμβοι}} \quad \dots \quad \underbrace{\gamma(\gamma+k)}_{k^{\ell-3} \text{ κόμβοι}} \quad \underbrace{(\gamma+k)}_{k^{\ell-2} \text{ κόμβοι}} \quad \underbrace{1}_{k^{\ell-1} \text{ κόμβοι}} \ ]\mathbf{s} \\ = [\ & s_1 \qquad s_2 \qquad s_3 \qquad \dots \qquad s_{\ell-2} \qquad s_{\ell-1} \qquad s_\ell \ ]\end{aligned}$$

Ας αναλύσουμε τώρα το πώς αλλάζει το Κόστος Ακρίβειας στο πλαίσιο του πυραμιδικού γράφου. Για να υπολογίσουμε το Κόστος Ακρίβειας ($Cost^A$) πρέπει πρώτα να υπολογίσουμε την συγκεντρωτική άποψη των κόμβων της πυραμίδας η οποία δίνεται από την σχέση:

$$\begin{aligned}\mathcal{G}(\mathbf{x}) = (i_1 s_1 + i_2 s_2 + \cdots + i_{\ell-1} s_{\ell-1} + i_\ell s_\ell) f^A(\mathbf{x}) \\ + \left(i'_1 s_1 + i'_2 s_2 + \cdots + i'_{\ell-1} s_{\ell-1} + i'_\ell s_\ell\right) f^B(\mathbf{x})\end{aligned}$$

Όπου $i_n, i'_n \in \mathbb{N}$ με το $n \in \mathbb{N}$ να αντιπροσωπεύει το επίπεδο. Πρέπει να ισχύει ότι: $i_1 + i'_1 = 1$, $i_2 + i'_2 = k$, $i_3 + i'_3 = k^2$, $\dots$, $i_{\ell-1} + i'_{\ell-1} = k^{\ell-2}$, $i_\ell + i'_\ell = k^{\ell-1}$. Όπου οι παράμετροι i_n δείχνουν τον αριθμό των τύπου Α κόμβων σε κάθε επίπεδο και i'_n τον αριθμό των τύπου Β κόμβων.

Βάση αυτης της συγκεντρωτικής άποψης λοιπόν μπορούμε να υπολογίσουμε το νέο

κόστος Ακρίβειας για τον πυραμίδικό γράφο που δίνεται από την σχέση:

$$Cost^A = (i_1 s_1 + \cdots + i_\ell s_\ell)^2 (L^A + \sigma_A^2) + (i'_1 s_1 + \cdots + i'_\ell s_\ell)^2 (L^B + \sigma_B^2)$$

Επιπλέον το κόστος Διαφωνίας πάλι ορίζεται όπως και στον τυχαίο μη κατευθυνόμενο γράφο. Δηλαδή ορίζουμε ξανά πως το φαινόμενο της διαφωνίας εμφανίζεται μόνο μεταξύ παικτών που επικοινωνούν άμεσα, το οποίο στον υποκείμενο γράφο σημαίνει ότι οι κόμβοι τους συνδέονται απευθείας με μια ακμή. Υιοθετώντας αυτό το κριτήριο, μπορούμε να εισάγουμε ξανά τον όρο του κόστους Διαφωνίας ως:

$$Cost^D = \frac{2}{|T|^{1+\beta}} (d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2)$$

Πρέπει να σημειώσουμε εδώ ότι, σε αντίθεση με τον μη κατευθυνόμενο γράφο, ο πυραμιδικός δεν επιτρέπει τον ορισμό πλήρες συνδεδεμένων και αποσυνδεδεμένων περιπτώσεων. Αυτό οφείλεται στη σταθερή δομή και τις προκαθορισμένες ακμές που είναι εγγενείς στη δομή του πυραμιδικού γράφου. Επομένως, το μόνο που μπορούμε να προσαρμώσουμε σε αυτη την περίπτωση είναι η τοποθέτηση των παικτών τύπου Α ή Β εντός της δομής του δικτύου. Η κατανόηση της βέλτιστης τοποθέτησης των μελών της ομάδας εντός της ιεραρχικής δομής είναι ένα πολύ ενδιαφέρον πρόβλημα βελτιστοποίησης που θα αναλυθεί στην επόμενη ενότητα.

6.2 Αποτελέσματα Προσομοιώσεων

Σε αυτή την τελευταία ενότητα θα παρουσιάσουμε τα αποτελέσματα των προσομοιώσεων και των αλγοριθμικών αναλύσεων που εφαρμόσαμε με στόχο να αποσαφηνιστεί η δυναμική του εκτεταμένου μοντέλου που προτείναμε. Θα διερευνήσουμε την αλληλεπίδραση μεταξύ των δομών της ομάδας και των διαδικασιών λήψης αποφάσεων, εστιάζοντας σε τρείς πτυχές: τη βέλτιστη σύνθεση της ομάδας, τη συνδεσιμότητα του δικτύου της ομάδας και την τοποθέτηση των μελών τύπου Α ή Β εντός της ομάδας.

Ας ξεκινήσουμε αναλύοντας τη βέλτιστη σύνθεση της ομάδας ως προς το κόστος Ακρίβειας για τον βασικό όπως και για τον επεκτεταμένο μοντέλο στην περίπτωση του τυχαίου μη κατευθυνόμενου γράφου. Ως προς το βασικό μοντέλο γνωρίζουμε πως το κόστος Ακρίβειας ελαχιστοποιείται όταν ισχύει η σχέση:

$$n_B^* = n_A \left(\frac{L^A + \sigma_A^2}{L^B + \sigma_B^2} \right)^{1/a}$$

Για την επέκταση του τυχαίου μη κατευθυνόμενου γράφου γνωρίζουμε πως η ίδια σχέση ισχύει και για τον πλήρες γράφου ($p = q = 1$) μιας και σε αυτη την περίπτωση διατηρείται το βασικό μοντέλο. Επιπλέον έχει θεωρητικό ενδιαφέρον να δούμε πως αλλάζει αυτη η σχέση για την περίπτωση του αποσυνδεδεμένου γράφου ($p = 1, q \approx 0$), η οποία γίνετε:

$$n_B^* = \frac{1}{2} + \frac{1}{2} \sqrt{1 + 4n_A(n_A - 1) \frac{L^A + \sigma_A^2}{L^B + \sigma_B^2}}$$

Με στόχο την επιβεβαίωση και ανάλυση αυτών των δύο θεωρητικών αποτελεσμάτων θα εκτελέσουμε προσομοιώσεις ανάπτυξης ομάδας. Τα αποτελέσματα που παρουσιάζονται στις Εικόνες 6.3 και 6.4 δημιουργήθηκαν χρησιμοποιώντας τις ακόλουθες συνθήκες: Κάθε κατάσταση $\mathbf{x}$ περιλαμβάνει 10 χαρακτηριστικά x_i, το καθένα από μια τυπική κανονική κατανομή με μέση τιμή 0 και τυπική απόκλιση 0.1 έτσι ώστε $cov(x_i, x_j) = 0$ για όλα τα i, j. Συνολικά δημιουργούμε 10.000 καταστάσεις $\mathbf{x}$ από 10 χαρακτηριστικά η κάθε μια που μας δίνει έναν πίνακα 10 επί 10.000. Επιπλέον, για να έχουμε τα επιθυμητά L^A και L^B δημιουργούμε κατάλληλες τιμές των ∂^A και ∂^B με το κάθε ένα να επιδρά στα μισά χαρακτηριστικά.

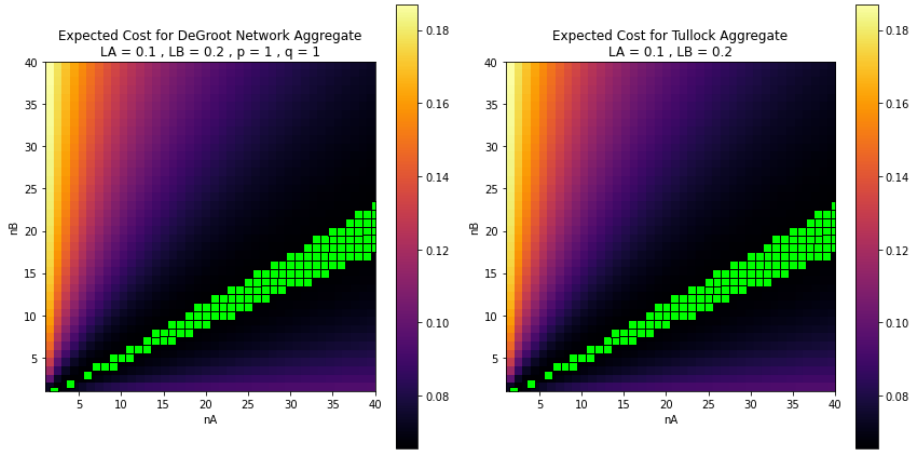

Εικόνα 6.3. *Προσομοίωση της βέλτιστης σύνθεσης της ομάδας με το βασικό μοντέλο στα δεξιά και του μη κατευθυνόμενο γράφο στα αριστερά. Εδώ το γράφημα είναι πλήρες, $p = q = 1$ και τα πράσινα κελιά αντιπροσωπεύουν τα 50 (n_A, n_B) σημεία με τις χαμηλότερες τιμές.*

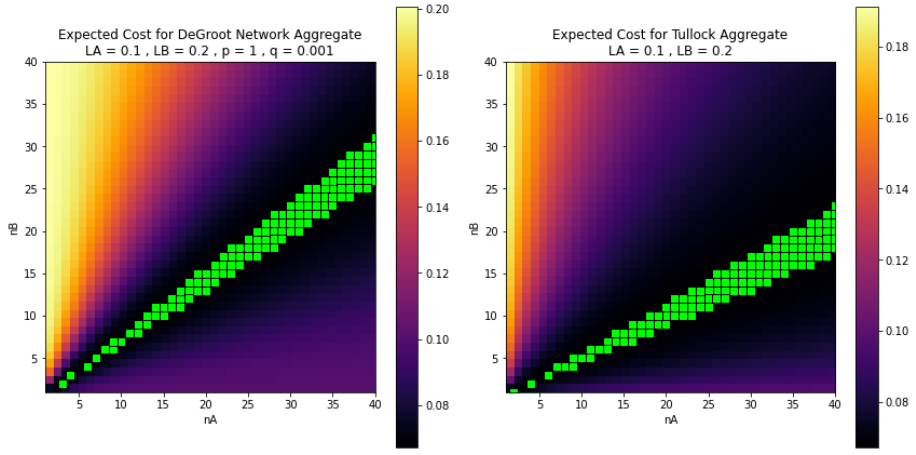

Εικόνα 6.4. *Προσομοίωση βέλτιστης σύνθεσης ομάδας με το βασικό μοντέλο στα δεξιά και του μη κατευθυνόμενο γράφο στα αριστερά. Εδώ το γράφημα είναι αποσυνδεδεμένο, $p = 1, q \approx 0$ και τα πράσινα κελιά αντιπροσωπεύουν τα 50 (n_A, n_B) σημεία με τις χαμηλότερες τιμές.*

Έχοντας ορίσει τα θεμελιώδη στοιχεία της προσομοίωσης, προχωράμε στην κατασκευή του μη κατευθυνόμενου γράφου. Χρησιμοποιώντας το πακέτο NetworkX [18], δημιουργούμε ένα μη κατευθυνόμενο γράφημα με n_A κόμβους τύπου A και n_B κόμβους τύπου B, οι οποίοι

συνδέονται με ακμές βάσει των πιθανοτήτων p και q. Είναι κρίσιμη η διασφάλιση ότι το γράφημα παραμένει ισχυρά συνδεδεμένο για οποιαδήποτε πιθανή τιμή των p και q, ώστε να εξασφαλιστεί η σύγκλιση της διαδικασίας εκμάθησης DeGroot.

Υπολογίζουμε τις τιμές του κόστους ακρίβειας για διάφορους συνδυασμούς των n_A και n_B (που κυμαίνονται από 1 έως 40) και τις οπτικοποιούμε στους θερμικούς χάρτες των Εικόνων 6.3 και 6.4. Επιπλέον, σε κάθε εικόνα, παρουσιάζουμε και τον θερμικό χάρτη των τιμών κόστους ακρίβειας που θα προέκυπταν από το βασικό μοντέλο. Σύμφωνα με την Εικόνα 6.3, το πλήρες δίκτυο διατηρεί τα αποτελέσματα του βασικού μοντέλου, ενώ το αποσυνδεδεμένο δίκτυο στην Εικόνα 6.4 φαίνεται να αποκλίνει. Επιπλέον, για ενδιάμεσες τιμές του $q \in (0, 1)$, παρατηρούμε τη μετατόπιση της πράσινης γραμμής του ελαχίστου μεταξύ αυτών των δύο ακραίων σημείων, υποδεικνύοντας πώς η τιμή του q επηρεάζει και αυτη τη βέλτιστη σύνθεση των ομάδων.

Είναι αξιοσημείωτο να εξεταστεί η θεωρητική διαφορά μεταξύ των τιμών n_B^* για πλήρη και αποσυνδεδεμένα γραφήματα. Γνωρίζουμε πως για το βασικό μοντέλο η βέλτιστη ακρίβεια σύνθεσης ως προς το n_B^* δίνεται μέσω μιας γραμμικής σχέσης ως προς το n_A, με κλίση $\frac{L_A+\sigma_A^2}{L_B+\sigma_B^2}$. Ομοίως, για τον αποσυνδεδεμένο γράφο έχουμε μια σχεδόν γραμμική σχέση μεταξύ n_B^* και n_A. Κατά τη σύγκριση όμως μεταξύ διαφορετικών τιμών της αναλογίας $\frac{L_A+\sigma_A^2}{L_B+\sigma_B^2}$, παρατηρούμε ότι το αποσυνδεδεμένο γράφημα τείνει περισσότερο προς τη γραμμή $n_B^* = n_A$ σε σύγκριση με το πλήρες γράφημα. Αυτή η τάση του αποσυνδεδεμένου γράφου μπορεί να αποδοθεί στην απουσία ακμών d_{AB}, απαιτώντας έτσι είτε περισσότερους «κακούς» παίκτες (με μεγαλύτερο L) είτε λιγότερους «καλούς» παίκτες (με χαμηλότερο L) για να επιτευχθεί το ελάχιστο κόστος. Αυτό συμβαίνει επειδή η διάδοση πληροφοριών εντός του αποσυνδεδεμένου δικτύου δεν είναι βέλτιστη.

Ας προχωρήσουμε τώρα στην περίπτωση του πυραμιδικού γράφου και συγκεκριμένα στην βέλτιστη τοποθέτηση των μελών τύπου A ή B εντός της ομάδας. Ο πυραμιδικός γράφος δεν επιτρέπει την αλλαγή των ακμών λόγω της άκαμπτης δομής του. Επομένως, το πεδίο των ρυθμίσεων μας περιορίζεται στη θέση των παικτών τύπου A ή B εντός του δικτύου. Έτσι, ένα φυσικό ερώτημα που προκύπτει είναι: δεδομένου ενός γραφήματος πυραμίδας $G(n_A, n_B, k, \ell, \gamma)$, ποιος είναι ο βέλτιστος τρόπος να τοποθετηθήσουμε τους παίκτες τύπου A και B εντός της πυραμίδας προκειμένου να ελαχιστοποιηθεί το Κόστος. Ας αρχίσουμε να απαντάμε σε αυτό το ερώτημα δίνοτας την συνθήκη κάτω απο την οποία το Κόστος Ακρίβειας του γραφήματος πυραμίδας ελαχιστοποιείται:

$$\sum_{n=1}^{\ell} i_n s_n = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$$

Η εξίσωση αυτή έχει ℓ αγνώστους, οι οποίες είναι οι τιμές των $i_n \in \mathbb{N}$. Υπενθυμίζουμε εδώ πως τα i_n είναι φυσικοί αριθμοί και μπορούν να λάβουν τις εξής τιμές: $i_1 = \{0, 1\}$, $i_2 = \{0, 1, \ldots, k\}$, $i_3 = \{0, 1, \ldots, k^2\}$, $\ldots$, $i_{\ell-1} = \{0, 1, \ldots, k^{\ell-2}\}$, $i_\ell = \{0, 1, \ldots, k^{\ell-1}\}$. Επίσης, πρέπει να ισχύει ότι: $i_1 + i_2 + \cdots + i_\ell = n_A$ καθώς και ότι: $i'_1 + i'_2 + \cdots + i'_\ell = n_B$. Λόγω αυτών των περιορισμών, η εξίσωση μπορεί να έχει πολλές ή καμία λύση. Μπορεί να μην υπάρχουν λύσεις στην περίπτωση που οι παράμετροι $L^A, L^B, \sigma_A, \sigma_B$ και s_n δεν μπορούν να παράγουν συνδυασμούς των i_n που να φτάνουν τον στόχο: $C = (L^B + \sigma_B^2)/(L^A + \sigma_A^2 + L^B + \sigma_B^2)$, όπως

επίσης μπορεί να υπάρχει μεγάλος αριθμός λύσεων όταν αρκετοί διαφορετικοί συνδυασμοί των i_n ικανοποιούν την εξίσωση.

Τώρα, ένα καλό ερώτημα που προκύπτει είναι το πώς μπορούμε να βρούμε αυτές τις βέλτιστες τιμές i_n που ελαχιστοποιούν το Κόστος Ακρίβειας. Για να μας βοηθήσει να απαντήσουμε σε αυτό το ερώτημα, πρέπει πρώτα να παρατηρήσουμε ότι στο γράφημα πυραμίδας συμβαίνει εξής: Όσο πιο κοντά πλησιάζει το $\sum_{n=1}^{\ell} i_n s_n$ στην τιμή στόχου: $C = \frac{L^B+\sigma_B^2}{L^A+\sigma_A^2+L^B+\sigma_B^2}$, τόσο πιο κοντά βρισκόμαστε στο ελάχιστο του Κόστους Ακρίβειας. Έτσι, ακόμη και αν η τιμή στόχος C δεν μπορεί να επιτευχθεί από το άθροισμα: $\sum_{n=1}^{\ell} i_n s_n$, γνωρίζουμε ότι ο καλύτερος συνδυασμός των τιμών i_n είναι αυτός που πλησιάζει περισσότερο σε αυτήν.

Βάσει αυτών των παρατηρήσεων, μπορούμε να προτείνουμε αρχικά έναν εξαντλητικό αλγόριθμο προκειμένου να βρούμε τις βέλτιστες τιμές των i_n. Ο εξαντλητικός αλγόριθμος ελέγχει όλους τους δυνατούς συνδυασμούς των i_n, οι οποίοι είναι περιορισμένοι και μπορούν να λάβουν τις εξής τιμές: $i_1 = \{0, 1\}$, $i_2 = \{0, 1, \ldots, k\}$, $i_3 = \{0, 1, \ldots, k^2\}$, $\ldots$, $i_{\ell-1} = \{0, 1, \ldots, k^{\ell-2}\}$, $i_\ell = \{0, 1, \ldots, k^{\ell-1}\}$. Γνωρίζοντας τις τιμές των s_n από τις υπερπαραμέτρους της πυραμίδας, μπορούμε να ελέγξουμε όλους τους συνδυασμούς των i_n στο άθροισμα: $\sum_{n=1}^{\ell} i_n s_n$, και απλώς να επιλέξουμε αυτόν που είναι πιο κοντά στην τιμή στόχου C. Ωστόσο, μπορούμε να δούμε πως αυτή η μέθοδος είναι υπολογιστικά ακριβή, ειδικά στην περίπτωση πυραμίδων με μεγάλες τιμές του ℓ.

Για να αντιμετωπίσουμε αυτό το πρόβλημα, μπορούμε να προτείνουμε έναν άπληστο αλγόριθμο. Μπορούμε να χρησιμοποιήσουμε το γεγονός ότι η επιρροή των πρακτόρων στα ανώτερα επίπεδα είναι μεγαλύτερη από την επιρροή των πρακτόρων στα κατώτερα επίπεδα:

$$s_1 > s_2 > s_3 > \cdots > s_{\ell-2} > s_{\ell-1} > s_\ell$$

Ο άπληστος αλγόριθμος C λειτουργεί επιλέγοντας επαναληπτικά ακέραιους συντελεστές i_n για μια σειρά γνωστών τιμών s_n έτσι ώστε $\sum_{n=1}^{\ell} i_n s_n$ να είναι όσο το δυνατόν πιο κοντά στο C χωρίς να το υπερβαίνει. Ο αλγόριθμος ξεκινά με την αρχικοποίηση του αθροίσματος S στο μηδέν. Στη συνέχεια, επεξεργάζεται κάθε συντελεστή i_n ξεκινώντας από το i_1 και καταλήγοντας στο i_ℓ. Για κάθε i_n, αυξάνει το i_n ξεκινόντας από το μηδέν, προσθέτοντας την αντίστοιχη τιμή s_n στο S εφόσον το προκύπτον άθροισμα δεν υπερβαίνει το C και το i_n παραμένει εντός των ορίων του. Αν το άθροισμα S υπερβεί το C μετά από μια αύξηση, ο αλγόριθμος μειώνει το i_n κατά ένα και προσαρμόζει το S αναλόγως. Αυτή η διαδικασία διασφαλίζει ότι το άθροισμα S πλησιάζει όσο το δυνατόν πιο κοντά στο C χωρίς να το υπερβαίνει. Τελικά, ο αλγόριθμος επιστρέφει το σύνολο των συντελεστών $i_1, i_2, \ldots, i_\ell$ που αντιπροσωπεύουν την πλησιέστερη δυνατή προσέγγιση στο άθροισμα στόχο C χρησιμοποιώντας τις δεδομένες τιμές s_n. Αυτή η προσέγγιση εκμεταλλεύεται τις ιδιότητες των άπληστων αλγορίθμων κάνοντας τοπικά μυωπικές βέλτιστες επιλογές σε κάθε βήμα. Μπορούμε να σκεφτούμε αυτή την προσέγγιση σαν να ξεκινά με μεγάλα βήματα όταν χρησιμοποιούμε τα s_1, s_2 και καθώς πλησιάζουμε στον στόχο C, χρησιμοποιούμε όλο και μικρότερα βήματα.

Εικασία Για να βρούμε τα i_n που ελαχιστοποιούν το Κόστος Ακρίβειας του γράφου Πυραμίδας, ο προτεινόμενος Άπληστος Αλγόριθμος είναι Βέλτιστος και βρίσκει παντα ένα σωστό σύνολο i_n αν και μόνο αν το C είναι επιτεύξιμο. Το C ορίζεται ως επιτεύξιμο όταν υπάρχει τουλάχιστον ένας συνδυασμός των i_n που ικανοποιεί: $\sum_{n=1}^{\ell} i_n s_n = C$.

Μπορούμε επίσης να απαντήσουμε αυτό το ερώτημα κοιτάζοντας μόνο το Κόστους Διαφωνίας. Ορίζουμε στο μοντέλο του δικτύου μας πως η Διαφωνία συμβαίνει μόνο μεταξύ των πρακτόρων που είναι άμεσα συνδεδεμένοι στο γράφημα, δηλαδή έχουν μια ακμή που τους συνδέει άμεσα. Έτσι, για το δίκτυο πυραμίδας το Κόστος Διαφωνίας μπορεί να γραφτεί ως:

$$Cost^D = \frac{2}{|T|^{1+\beta}} \left(d_{AB}(L^A + L^B + \sigma_A^2 + \sigma_B^2) + d_{AA}\sigma_A^2 + d_{BB}\sigma_B^2 \right)$$

Με την κατάλληλη επιλογή της τοποθέτησης των παικτών τύπου A ή B στην πυραμίδα, μπορούμε να προσαρμόσουμε τις τιμές των d_{AA}, d_{BB} και d_{AB}. Ας εξετάσουμε την περίπτωση όπου δεν υπάρχει θόρυβος στις προβλέψεις των πρακτόρων: $\sigma_A = \sigma_B = 0$. Σε αυτήν την απλουστευμένη περίπτωση, μόνο οι συνδέσεις μεταξύ κόμβων τύπου A και τύπου B δημιουργούν Κόστος Διαφωνίας και έτσι πρέπει να ελαχιστοποιήσουμε την τιμή του d_{AB}. Η προφανής λύση που ελαχιστοποιεί το Κόστος Διαφωνίας σε αυτή την περίπτωση είναι όταν υπάρχουν μόνο πράκτορες τύπου A ή μόνο πράκτορες τύπου B στην πυραμίδα. Ωστόσο, τι συμβαίνει όταν πρέπει να τοποθετήσουμε $n_A > 0$ κόμβους τύπου A και $n_B > 0$ κόμβους τύπου B στην ίδια πυραμίδα· Η απάντηση είναι ότι για να διατηρήσουμε το Κόστος Διαφωνίας στο ελάχιστο, πρέπει να τοποθετήσουμε τους πράκτορες του ίδιου τύπου σε υποπυραμίδες, ελαχιστοποιώντας έτσι την τιμή του d_{AB}. Αυτό μπορεί και φανεί και οπτικά για διάφορες τιμές των n_A και n_B στην Εικόνα 6.5.

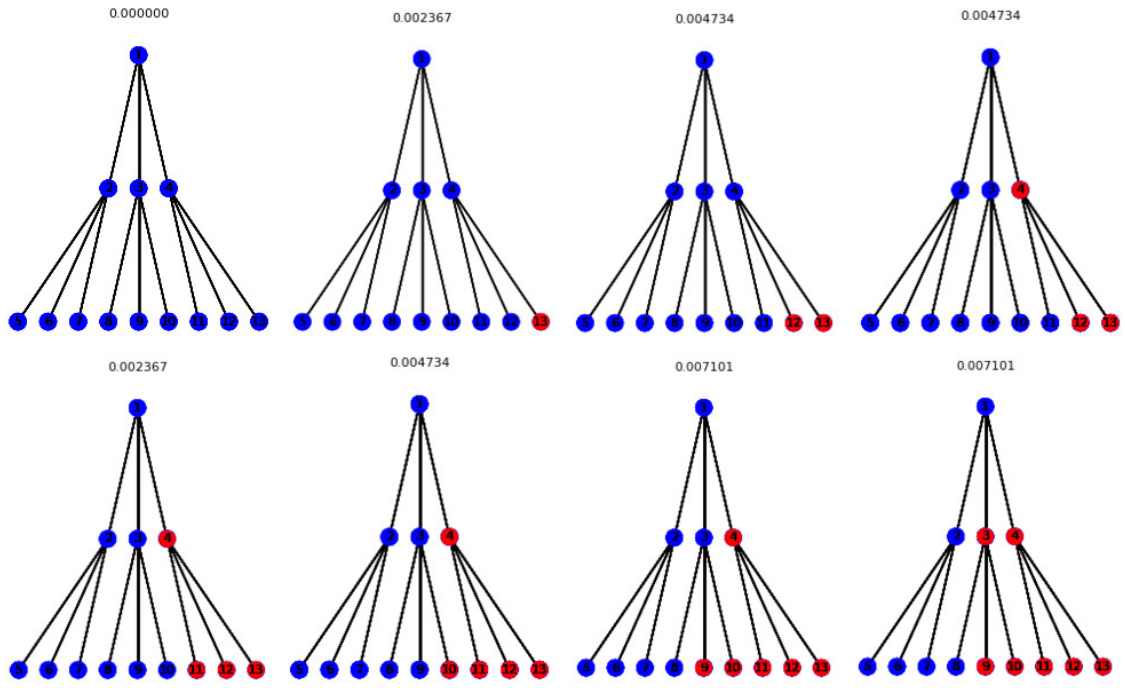

Εικόνα 6.5. *Βέλτιστη τοποθέτηση παικτών τύπου A και B σε πυραμίδα με $k = \ell = 3$ όπου υπάρχει μόνο Κόστος Διαφωνίας ($\beta = 1$). Οι νέοι παίκτες n_A τοποθετούνται εντός υποπυραμίδας.*

Τι συμβαίνει λοιπόν στον πυραμιδικό γράφο όταν συνδυάζουμε το Κόστος Διαφωνίας με το Κόστος Ακρίβειας· Για αυτήν την πιο περίπλοκη περίπτωση επιλέγουμε να μελετήσουμε το μοντέλο χωρίς κανένα θόρυβο ($\sigma_A = \sigma_B = 0$). Ως εκ τούτου, από εδώ και στο εξής θα

μελετάμε το $Cost^T$ στην απλουστευμένη του μορφή ως:

$$Cost^T = \lambda \times \left(\frac{2(L^A + L^B)}{|T|^{1+\beta}} d_{AB} \right) + (1 - \lambda) \times \left(L^A \left(\sum_{n=1}^{\ell} i_n s_n \right)^2 + L^B \left(\sum_{n=1}^{\ell} i'_n s_n \right)^2 \right)$$

Αμέσως γίνεται σαφές ότι για την συγκεκριμένη συνάρτηση κόστους δεν είναι εφικτό να βρούμε μια θεωρητική λύση. Αυτό συμβαίνει διότι δεν είμαστε ικανοί να πάρουμε παραγώγους του $Cost^T$ ως προς τα d_{AA}, d_{BB}, d_{AB} καθώς τώρα ο όρος Κόστος Ακρίβειας λαμβάνει υπόψη μόνο την τοποθέτηση του τύπου παικτών στα επίπεδα i_n και όχι τις μεταξύ τους συνδέσεις. Έτσι, επιλέγουμε να αναπτύξουμε μια καθαρά αλγοριθμική προσέγγιση για το Συνολικό Κόστος του πυραμίδικού γράφου. Θα αναπτύξουμε 3 αλγόριθμους: έναν Εξαντλητικό, έναν Άπληστο και έναν Τοπικής Αναζήτησης.

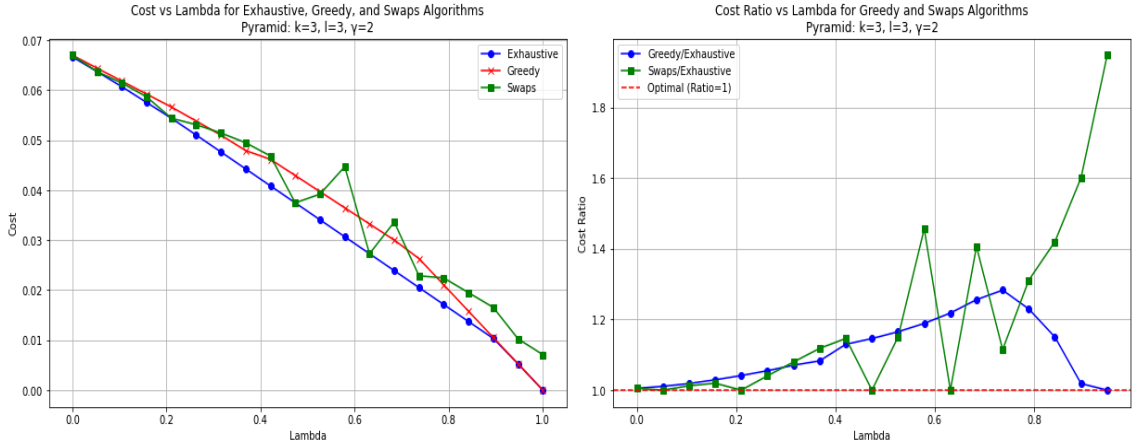

Εικόνα 6.6. *Σύγκριση κόστους και λόγου κόστους για τους αλγόριθμους: Εξαντλητικής, Άπληστης και Τοπικής αναζήτησης για διαφορετικές τιμές του λ. Παράμετροι πυραμίδας:* $k = 3$, $\ell = 3$, $\gamma = 2$.

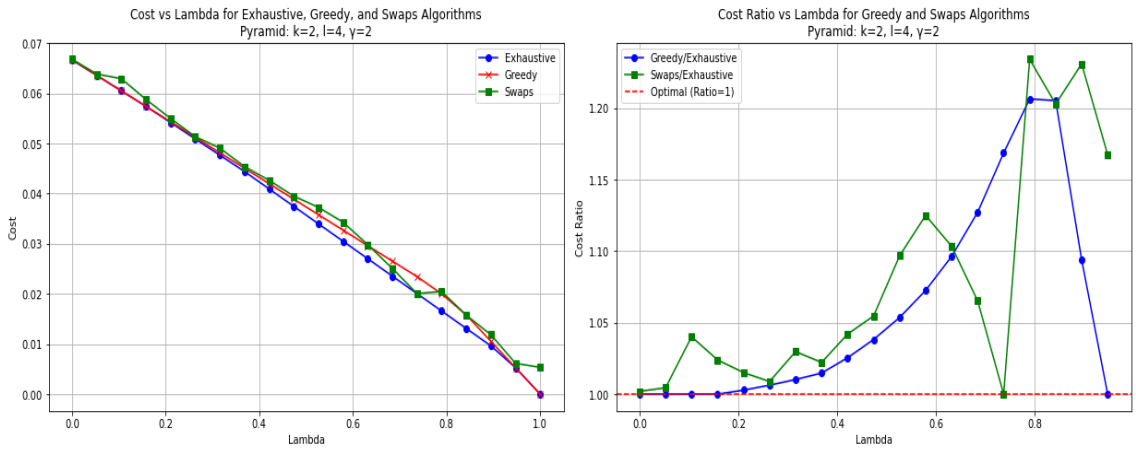

Εικόνα 6.7. *Σύγκριση κόστους και λόγου κόστους για τους αλγόριθμους: Εξαντλητικής, Άπληστης και Τοπικής αναζήτησης για διαφορετικές τιμές του λ. Παράμετροι πυραμίδας:* $k = 2$, $\ell = 4$, $\gamma = 2$.

Ο Εξαντλητικός αλγόριθμος εξετάζει όλους τους 2^n δυνατούς συνδυασμούς τύπου Α ή Β εντός της πυραμίδας, όπου n είναι ο αριθμός των κόμβων στην $G(n_A, n_B, k, \ell, \gamma)$ πυραμίδα. Για κάθε συνδυασμό, υπολογίζει το αντίστοιχο $Cost^T$ και επιστρέφει στο τέλεος το συνδιασμό με το χαμηλότερο κόστος. Όπως αναμένεται, αυτή η προσέγγιση γίνεται αρκετά υπολογιστικά δαπανηρή καθώς ο γράφος της πυραμίδας μεγαλώνει. Ο Άπληστος αλγόριθμος από την άλλη είναι ο λιγότερο απαιτητικός υπολογιστικά. Ξεκινά αναθέτοντας τον ίδιο τύπο (Α ή Β) σε όλους τους κόμβους της πυραμίδας. Στη συνέχεια, από την κορυφή προς τα κάτω, προσπαθεί να αλλάξει τον τύπο κάθε κόμβου έναν προς έναν. Κάθε αλλαγή διατηρείται μόνο αν μειώνει το $Cost^T$. Είναι σημαντικό να σημειωθεί ότι αυτός ο άπληστος αλγόριθμος διαφέρει από αυτόν που περιγράφηκε προηγουμένως αποκλειστικά για το Κόστος Ακρίβειας. Αυτο συμβαίνει διότι για τη γενική περίπτωση, δεν έχουμε τιμή στόχου C για να υποδείξει πότε υπερβαίνουμε το όριο, οπότε αυτός ο άπληστος αλγόριθμος δεν είναι βέλτιστος, ακόμη και όταν $\lambda = 0$ (δηλαδή, όταν λαμβάνεται υπόψη μόνο το Κόστος Ακρίβειας). Ο αλγόριθμος Τοπικής Αναζήτησης ξεκινά αναθέτοντας τυχαία τύπους σε όλους τους κόμβους στον γράφο της πυραμίδας. Στη συνέχεια, βελτιώνει επαναληπτικά την ανάθεση εξετάζοντας ζεύγη κόμβων με διαφορετικούς τύπους. Για κάθε ζεύγος, ο αλγόριθμος ανταλλάσσει τους τύπους των δύο κόμβων και υπολογίζει το νέο κόστος. Αν η ανταλλαγή έχει ως αποτέλεσμα χαμηλότερο κόστος, η αλλαγή διατηρείται. Διαφορετικά, επανέρχεται στην αρχική κατάσταση. Αυτή η διαδικασία συνεχίζεται μέχρι να μην μπορούν να βρεθούν περαιτέρω βελτιώσεις. Ο αλγόριθμος τοπικής αναζήτησης είναι πιο υπολογιστικά ακριβός σε σύγκριση με τον άπληστο αλγόριθμο, αλλά λιγότερο ακριβός σε σύγκριση με την εξαντλητική αναζήτηση.

Η απόδοση των τριών αλγορίθμων απεικονίζεται στις Εικόνες 6.6. και 6.7. Σε αυτά τα σχήματα συγκρίνονται οι τιμές κόστους που λαμβάνονται απο τους 3 αλγόριθμους για ένα εύρος τιμών του λ. Και στα δύο σχήματα, ο άπληστος αλγόριθμος δείχνει την αποτελεσματικότητά του ιδιαίτερα για τιμές του λ κοντά στο 0 ή 1, όπου προσεγγίζει αρκετά τα αποτελέσματα του αλγορίθμου Εξαντλητικής Αναζήτησης. Αυτό υποδηλώνει ότι ο άπληστος αλγόριθμος είναι κατάλληλος για σενάρια όπου είτε το Κόστος Ακρίβειας είτε το Κόστος Διαφωνίας επικρατεί. Αντίθετα, ο αλγόριθμος Τοπικής αναζήτησης συχνά υπερέχει του άπληστου αλγορίθμου για ενδιάμεσες τιμές του λ, επιδεικνύοντας τη δύναμή του στην εξισορρόπηση και των δύο κόστων. Ωστόσο, είναι σημαντικό να σημειωθεί ότι η απόδοση του αλγορίθμου Τοπικής αναζήτησης είναι αρκετά ευαίσθητη στην αρχική τυχαία κατανομή του τύπου των κόμβων της πυραμίδας. Διαφορετικές τυχαίες αρχικοποιήσεις μπορούν να οδηγήσουν σε ποικίλα αποτελέσματα, γεγονός που αναδεικνύει την εξάρτηση του αλγορίθμου από την αρχική διαμόρφωση της πυραμίδας.

Η Εικόνα 6.8. επικεντρώνεται σε μια μεγαλύτερη πυραμίδα με $k = 3$, $\ell = 5$, η οποία έχει 121 κόμβους. Λόγω των σημαντικών υπολογιστικών απαιτήσεων, ο αλγόριθμος Εξαντλητικής Αναζήτησης δεν είναι εφικτός για αυτή την περίπτωση, και έτσι παρουσιάζονται μόνο τα αποτελέσματα των αλγορίθμων Άπληστης και Τοπικής αναζήτησης. Είναι αξιοσημείωτο πως και οι δύο αλγόριθμοι παράγουν παρόμοια αποτελέσματα, όπως φαίνεται από το αριστερό διάγραμμα όπου οι γραμμές τους σχεδόν επικαλύπτονται. Το δεξιό διάγραμμα, που δείχνει τη διαφορά κόστους μεταξύ των δύο αλγορίθμων, υποδεικνύει ότι οι διαφορές τους είναι της τάξης του 10^{-4}. Αυτή η αμελητέα διαφορά υπογραμμίζει την αποτελεσματικότητα και των δύο αλγορίθμων στην προσέγγιση του ελάχιστου $Cost^T$. Για παράδειγμα, όταν $\lambda = 1$, όπου μόνο

το Κόστος Διαφωνίας λαμβάνεται υπόψη, και οι δύο αλγόριθμοι επιτυγχάνουν τιμές κοντά στο μηδέν. Ο Άπληστος αλγόριθμος, ειδικότερα, φτάνει ακριβώς στο $Cost^T = 0$ καθώς ξεκινά με όλους τους κόμβους να έχουν τον ίδιο τύπο, αντανακλώντας την αποτελεσματικότητά του για αυτή την οριακή περίπτωση.

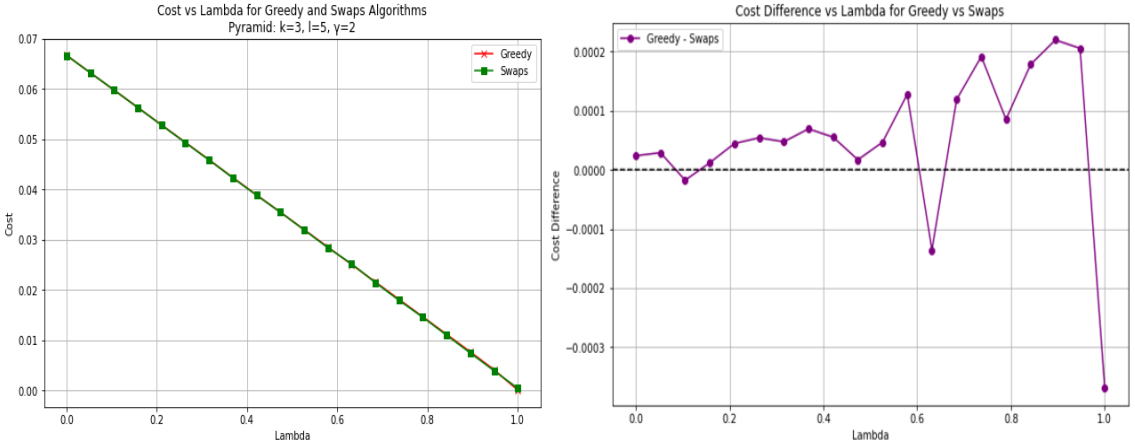

Εικόνα 6.8. *Σύγκριση κόστους και διαφοράς κόστους για τους αλγορίθμους Άπληστης και Τοπικής αναζήτησης για διάφορες τιμές του λ. Παράμετροι πυραμίδας:* $k = 3$, $\ell = 5$, $\gamma = 2$

Bibliography

[1] S. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. The William G. Bowen Memorial Series in Higher Education Series, Princeton University Press, 2007.

[2] R. Burt, "Structural holes and good ideas," *American Journal of Sociology*, vol. 110, no. 2, pp. 349–399, 2004.

[3] S. Page, N. Cantor, and K. Phillips, *The Diversity Bonus: How Great Teams Pay Off in the Knowledge Economy*. Our Compelling Interests, Princeton University Press, 2019.

[4] L. Hong and S. Page, "Groups of diverse problem solvers can outperform groups of high-ability problem solvers," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, pp. 16385–16389, 12 2004.

[5] H. McCormick, "The real effects of unconscious bias in the workplace," *UNC Executive Development, Kenan-Flagler Business School. DIRECCIÓN*, pp. 2–12, 2015.

[6] H. Oberai and I. M. Anand, "Unconscious bias: thinking without thinking," *Human Resource Management International Digest*, vol. 26, no. 6, pp. 14–17, 2018.

[7] F. J. Milliken and L. L. Martins, "Searching for common threads: Understanding the multiple effects of diversity in organizational groups," *Academy of management review*, vol. 21, no. 2, pp. 402–433, 1996.

[8] K. W. Phillips, K. A. Liljenquist, and M. A. Neale, "Is the pain worth the gain? the advantages and liabilities of agreeing with socially distinct newcomers," *Personality and Social Psychology Bulletin*, vol. 35, no. 3, pp. 336–350, 2009.

[9] H. Heidari, S. Barocas, J. Kleinberg, and K. Levy, "Informational diversity and affinity bias in team growth dynamics," in *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pp. 1–10, 2023.

[10] H. Jia, S. Skaperdas, and S. Vaidya, "Contest functions: Theoretical foundations and issues in estimation," *International Journal of Industrial Organization*, vol. 31, no. 3, pp. 211–222, 2013.

[11] S. Skaperdas, "Contest success functions," *Economic Theory*, vol. 7, no. 2, pp. 283–290, 1996.

[12] M. H. Degroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[13] J. R. French Jr, "A formal theory of social power.," *Psychological review*, vol. 63, no. 3, p. 181, 1956.

[14] F. Harary, "A criterion for unanimity in french's theory of social power.," 1959.

[15] B. Golub and M. O. Jackson, "Naïve learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, pp. 112–49, February 2010.

[16] P. Perkins, "A theorem on regular matrices," *Pacific Journal of Mathematics*, vol. 11, no. 4, pp. 1529–33, 1961.

[17] C. D. Meyer, *Matrix analysis and applied linear algebra*. Philadelphia: Society for Industrial and Applied Mathematics (SIAM), 2000.

[18] A. Hagberg, P. Swart, and D. S Chult, "Exploring network structure, dynamics, and function using networkx," tech. rep., Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.