Doctoral Dissertation

Doctoral Program in Pure and Applied Mathematics

Intervention Design on Super-Modular Network Games

By

Sebastiano Messina

******

Supervisor(s):

Prof. Giacomo Como, Supervisor

Prof. Fabio Fagnani, Co-Supervisor

Politecnico di Torino

2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and do not compromise in any way the rights of third parties, including those relating to the security of personal data.

Sebastiano Messina

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

Abstract

The design of optimal intervention policies by a central regulator in networked multi-agent systems is investigated in this dissertation. Modeled as games exhibiting strategic complementarities, these systems are characterized by interactions wherein an increase in the effort of one agent incentivizes corresponding increases among peers. While this specific class of interactions, formally defined as super-modular games, guarantees the existence of Nash equilibria, it is frequently observed to exhibit equilibrium multiplicity. Such multiplicity typically induces severe coordination failures and systemic inefficiencies. To address these inefficiencies, mechanism design principles are applied to compute minimum-cost interventions. These interventions are structured to guide decentralized systems deterministically toward socially optimal, maximal equilibria, without directly commanding the individual actions of the agents.

For finite, non-binary unimodal games, interventions are modeled as lower-bound restrictions on action sets, effectively raising the minimum effort required by agents. A distributed algorithmic approach leveraging time-reversible Markov chains is proposed, establishing that minimum sufficient interventions can be identified by iteratively tracing weak improvement paths from the maximal strategy profile. Furthermore, the scalability limitations inherent to large-scale binary network games are systematically addressed. In these environments, optimization problems such as Optimal Seeding and Least-Cost Influence are strictly NP-hard, and complete topological data is rarely accessible to a central planner. By applying a directed configuration model and Local Mean-Field approximations, these computationally intractable network optimization problems are rigorously reformulated as dimension-independent linear programs. This analytical transformation enables the design of near-optimal intervention strategies reliant exclusively on aggregate network statistics and degree distributions, thereby circumventing the requirement for exact topological knowledge.

Contents

Chapter 1

Introduction and Methodological Overview

The analysis of strategic interactions among agents connected through complex networks constitutes a central focus in modern game theory, economics, and applied mathematics. In numerous economic, social, and technological systems, decisions are executed by individual agents whose payoffs depend not solely on their own actions, but significantly on the actions of interacting peers.

Network games are employed to model populations of agents whose decisions are rendered interdependent by a graph-structured interaction, capturing empirical patterns of influence, information exchange, or physical connectivity. This dissertation primarily focuses on games exhibiting **strategic complementarities**, formally designated as **super-modular games**. Within these strategic environments, heightened action levels by one agent strictly increase the marginal incentives for adjacent peers to similarly elevate their actions. This specific class of games has been extensively analyzed due to its highly structured equilibrium properties, monotone comparative statics, and broad applicability. Crucially, supermodularity facilitates the application of powerful mathematical frameworks, such as lattice-theoretic methods, which enable sharp equilibrium characterizations even within highly complex or large-scale systems.

From a theoretical perspective, network games characterized by complementarities introduce fundamental challenges regarding equilibrium selection, coordination, and systemic robustness.

Although the existence of Nash equilibria is mathematically guaranteed in supermodular games, the system remains highly susceptible to equilibrium multiplicity. It is empirically and theoretically observed that this multiplicity acts as a primary catalyst for coordination failures, systemic inefficiencies, and stochastically unpredictable outcomes. The reduction of equilibrium multiplicity constitutes a primary objective in the mechanism design of multi-agent systems, particularly within economic and financial networks where uncoordinated decision-making frequently yields severe inefficiencies. Without regulatory intervention, these systems are susceptible to suboptimal equilibria.

To address these inefficiencies, regulatory mechanisms are introduced to constrain the strategic environment. By restricting the domains of suboptimal equilibria, targeted interventions are utilized to guarantee deterministic convergence toward the maximal, Pareto-optimal equilibrium. Consequently, the formulation of coordination protocols designed to structurally shape the equilibrium landscape and maximize aggregate social welfare is established as the central thesis of this analysis. From an applied perspective, a multitude of real-world systems are regulated indirectly via policy instruments rather than through direct, authoritative commands. Examples include:

- **Policy Design:** Interventions are frequently aimed at influencing local interactions to maximize social welfare. Network effects make direct prediction difficult, highlighting the importance of network-aware interventions.
- **Decentralized System Coordination:** In engineered networked systems, such as smart grids or autonomous vehicle ecosystems, actions are executed based strictly on local information. Incentive mechanisms are required to ensure stability and desirable collective behavior.
- **Platform Design:** Online market platforms implement design choices, such as pricing rules, matching algorithms, and information disclosure policies, to influence user behavior and market outcomes.

The main objective of this thesis is to rigorously investigate the mechanisms through which interventions on network games with strategic complementarities can be optimally designed by a central regulatory authority. It is strictly assumed that the individual strategies of the agents cannot be directly dictated by the planner. Instead, a regulatory policy, such as the modification of localized incentive structures or the explicit imposition of action-set restrictions, is deployed.

This policy is mathematically calibrated to ensure the induced game satisfies specific topological and strategic properties, such as equilibrium uniqueness, dynamic stability, or the implementation of socially optimal outcomes, while strictly preserving the decentralized autonomy of the decision-making process. This formulation raises the following fundamental questions:

- Under the constraint that agents retain autonomy over their individual strategies and that intervention budgets are finitely bounded, how can regulatory policies be designed to reduce equilibrium multiplicity?
- Can convergence to stable, maximal equilibria be guaranteed, particularly in time-critical systems (i.e., dynamic environments where delayed coordination results in severe financial losses, cascading defaults, or catastrophic infrastructure failures)?
- How can *strategic complementarity effects* be leveraged to promote coordination that is efficient with respect to the maximization of aggregate social welfare and the minimization of intervention costs?

1.1 Context and Motivation

A central problem in the analysis of multi-agent systems is the design of regulatory interventions intended to direct network dynamics toward socially optimal states (Riehl et al., 2018b). Given that individual actions in decentralized environments are generally unobservable, direct control by a central planner is precluded. Consequently, the strategic environment must be regulated through indirect mechanisms. Coordination is thereby achieved strictly through the targeted modification of systemic parameters, such as incentive structures, information topologies, or network connectivities (Sakhaei et al., 2023).

Diffusion and cascading processes constitute a widely investigated class of network dynamics. These processes mathematically describe the local propagation of behaviors, information, or structural failures strictly through peer interactions. Examples include:

- In *social sciences*, viral marketing campaigns are studied to promote the diffusion of social trends and/or products through word-of-mouth effects; these mechanisms are formalized in foundational literature concerning influence maximization (Kempe et al., 2003, Richardson and Domingos, 2002)
- In *economics*, the focus is placed on understanding cascading effects that may lead to the bankruptcy of financial institutions; such dynamics are rigorously analyzed through network-theoretic models of systemic risk (Acemoglu et al., 2015; Elliott et al., 2014).
- Cascading failures in brain networks are investigated within *neuroscience*; these vulnerabilities are typically illustrated through the graph-theoretical analysis of complex neural topologies (Bullmore and Sporns, 2009).
- Failure propagation in critical infrastructure systems is addressed by *infrastructure engineering*; these phenomena are formally modeled through the mathematical study of load-based attacks and interdependent spatial networks (Buldyrev et al., 2010; Motter and Lai, 2002).
- The propagation of financial distress within interbank lending systems is analyzed through *financial networks* models; the characterization of default probabilities and interbank liquidity contagion is provided by classic analytical frameworks (Allen and Gale, 2000; Gai and Kapadia, 2010).

These phenomena are characterized by information cascades, where agent behavior is highly correlated and driven by peer decisions. Although global cascades are statistically infrequent, they can be initiated by minor exogenous shocks. A priori, shocks leading to large cascades cannot be distinguished from those causing localized effects. Consequently, the propagation of influence or failure is strictly determined by the underlying network topology. Understanding diffusion therefore requires analyzing the network embedding of individual components and their susceptibility to neighborhood actions.

From a dynamical perspective, a central question concerns how initial perturbations propagate through the network and potentially disrupt systems that were previously stable. In many models, the state of each agent evolves according to local interaction rules depending on the states of its neighbors. For monotone dynamics, a stationary configuration is eventually reached by the system, raising the question of whether the diffusion process results in a global cascade affecting a large fraction of the population or if it remains confined to a small subset of nodes.

The study of influence propagation in social networks has attracted significant attention in recent years. These models capture situations in which individuals choose among alternatives or levels of engagement, with payoffs that depend on the behavior of others. In many social and economic environments, the observed behavior of peers is relied upon because individuals possess limited information or constrained cognitive capacities. It is observed that equilibrium multiplicity frequently leads to coordination failures, systemic inefficiencies, or unpredictable outcomes (Ramazi et al., 2016). From an applied viewpoint, many real-world systems are regulated indirectly through policy instruments rather than direct commands. Furthermore, in complex decision-making scenarios, such as the evaluation of novel technologies, risky financial assets, or prospective job candidates, insufficient information frequently exists. In such environments, an inherent incentive is generated for individual decision-makers to monitor and emulate the decisions of others, thereby leading to the emergence of information cascades.

In economic environments characterized by bounded rationality or asymmetric information, agents rely upon the observable actions of connected peers, thereby generating decision problems with externalities. Threshold-based diffusion frameworks provide a mathematically tractable formalization of these externalities.
A canonical example is the **Linear Threshold Model** (LTM), introduced by Granovetter, 1978. In this framework, the network is represented by a directed graph, where the vertices represent individuals and the edges represent influence relationships. Each agent can be in one of two states, typically referred to as inactive or active (mapping directly to the binary actions 0 and 1, respectively).

Each node is endowed with a threshold, and each incoming edge carries a weight measuring the strength of influence from a neighbor. At each time step, an inactive agent becomes active (i.e., switches to action 1) if the sum of incoming influence from active neighbors exceeds its threshold. Once activated, nodes remain active, and the process continues until no further activations are possible.

A central optimization problem in this framework is **Influence Maximization**: identifying a small but effective set of initially active nodes whose activation leads to a large diffusion throughout the network. This problem is motivated by numerous applications, including viral marketing, targeted advertising, and the protection of critical infrastructure components. More generally, one can consider optimization problems such as selecting a fixed number of initial adopters in order to maximize the final number of active agents, or determining the minimal set of initial adopters required to trigger a cascade that eventually activates the entire population. In practice, the overall dynamics of the system can be controlled at a relatively low cost through the identification of an appropriate seed set. For instance, a product may be promoted by a firm by targeting a small number of influential individuals in a social network, whereas a limited number of structurally important components may be strengthened or protected by policy-makers to prevent cascading failures in infrastructure systems. Similar considerations arise in systemic-risk analysis, where identifying highly sensitive or highly contagious nodes may help prevent large-scale financial crises. Despite its conceptual simplicity, the problem of optimizing the trajectory of a diffusion process over a network remains computationally challenging. Interventions affecting a single individual may have complex indirect effects on the entire network, and these effects are typically highly nonlinear. Existing theoretical results are often limited to deterministic infinite graphs or rely on local structural conditions that must be verified node by node (Morris, 2000). In networks characterized by significant heterogeneity among nodes, such approaches become difficult to apply in practice. These dynamics are described by network games exhibiting **strategic complementarities**. This connection provides a bridge between the literature on diffusion processes and the theory of super-modular games.

1.2 Literature Review

The problems addressed herein lie at the intersection of **Economic Theory** (super-modular, network games), **Operations Research, Control Theory** (network optimization, equilibrium selection and mechanism design), and **Applied Mathematics** (lattice methods, learning dynamics). Each of these fields provides complementary tools for studying coordination, strategic interaction, and policy design in networked systems. The most relevant contributions are subsequently reviewed to position the present work within this interdisciplinary literature.

Strategic Complementarities and Network Games

Games with strategic complementarities have been extensively studied in economic theory because of their ability to capture coordination phenomena and the emergence of multiple equilibria. The concept of super-modularity was introduced and lattice-theoretic methods for analyzing equilibrium behavior were established in the foundational work of Topkis, 1978, 1998. Building on these ideas, the role of increasing differences and monotone comparative statics was formalized by Milgrom and Roberts, 1990, providing powerful tools to characterize how equilibrium outcomes respond to parameter changes. These techniques were subsequently applied by Vives, 1990 to several economic environments including oligopolistic competition, technology adoption, and macroeconomic coordination. A central insight from this literature is that complementarities often generate equilibrium multiplicity, making coordination failures a central concern in many economic settings. At the same time, the existence of extremal equilibria in super-modular games allows for partial equilibrium selection through monotone parameter changes. This framework is extended by network games by explicitly modeling the interaction structure among agents.

The economic importance of network structures in shaping strategic outcomes was highlighted in early contributions such as Jackson and Wolinsky, 1996. A particularly influential contribution is found in Ballester et al., 2006, wherein linear-quadratic games on networks were studied and the dependence of equilibrium actions on network centrality was demonstrated. Subsequent work has further explored how network topology affects equilibrium behavior and welfare outcomes. In games with complementarities, network structure plays a crucial role in determining the diffusion of behaviors and stability of equilibria. These models have been widely applied to social interactions, peer effects and technology adoption.

Optimization and Control of Networked Strategic Systems

Operations research and control theory provide complementary perspectives by focusing on the design of interventions and regulatory mechanisms in networked systems. Classical operations research models analyze resource allocation, flows, and congestion in networks, often under decentralized decision-making. The interaction between individual optimization and system-level performance has motivated the study of equilibrium concepts such as Wardrop equilibria and Nash equilibria in networks. More recently, this literature has increasingly incorporated strategic interactions and network externalities.

In many settings, the regulator's problem can be formulated as a bilevel optimization problem, where the upper level represents the regulator's intervention and the lower level captures agents' equilibrium responses. These models naturally arise in applications such as transportation networks, telecommunications, and energy markets. Control theory offers a systems-oriented framework for studying the regulation of such decentralized environments. When agents are strategic, regulatory policies can be interpreted as control inputs that modify the incentives faced by agents while preserving decentralized decision-making. This perspective has led to the development of research areas such as distributed control, mechanism design for control systems, and mean-field control.

Super-modular games are considered particularly attractive from a control perspective because their monotone best-response structure often leads to stable adjustment dynamics and predictable equilibrium selection. Regulatory interventions can therefore exploit these structural properties to guide the system toward desired outcomes. Applications to energy systems are particularly prominent. Electricity markets involve decentralized decisions by producers and consumers, whose incentives are dependent upon network constraints, prices, and the behavior of others. Complementarities arise through congestion effects, demand response, and technology adoption (for example, renewable integration). Regulatory instruments such as pricing rules, subsidies, and capacity mechanisms can be interpreted as parameters shaping the strategic environment.

Targeted interventions in networked strategic environments have also been studied in the economic and control literature, including Acemoglu et al., 2015 and Galeotti et al., 2020, analyzing how interventions directed at specific agents or nodes can influence aggregate outcomes. Related work has further examined optimal intervention in network games through centrality-based policies and optimal taxation or subsidies in network interactions (Elliott et al., 2014, Bloch, 2016).

Mathematical Foundations and Learning Dynamics

From the perspective of applied mathematics, the analysis of super-modular games relies heavily on lattice theory, fixed-point results, and monotone operator methods. A central tool is Tarski's fixed-point theorem (Tarski, 1955), which guarantees the existence of equilibria in games with complementarities and provides a constructive basis for equilibrium computation via monotone iterative procedures. Order-theoretic methods have been widely used to study convergence of iterative algorithms and equilibrium computation.
In many network games with finite or partially ordered action spaces, these tools allow for the characterization of the structure of equilibrium sets and to analyze comparative statics with respect to parameters.

These mathematical properties are closely connected to the study of learning dynamics. In many practical applications, Nash equilibrium strategies are not instantaneously played by agents; rather, actions are adjusted over time based on observations and beliefs, a process that frequently leads to the eventual emergence of an equilibrium. As a consequence, adaptive processes such as best-response dynamics, fictitious play, or reinforcement learning frequently converge to extremal equilibria (Ramazi and Cao, 2017). Furthermore, recent work on mean-field approximations demonstrates that discrete fluctuations in such binary best-response dynamics almost surely vanish as the population size increases, resulting in predictable continuous dynamics (Aghaeeyan and Ramazi, 2025). In games with strategic complementarities, learning dynamics are often particularly well behaved because best responses are increasing in the actions of others.

Such dynamics are especially relevant in social networks, where behavior spreads through imitation, reinforcement, and local interaction. The study of diffusion processes in networks has therefore become an important bridge between economic theory, algorithmic network science, and applied mathematics.

A major contribution in this direction is the work of Kempe et al., 2003, who introduced the *Influence Maximization Problem* (IMP). In this framework, a planner selects a limited number of initial "seed" agents with the objective of maximizing the eventual spread of adoption under the Linear Threshold Model (LTM). Although the Influence Maximization Problem is NP-hard, Kempe et al., 2003 showed that when individual thresholds are drawn independently from a uniform distribution over $[0, 1]$, the expected spread becomes a submodular function. This property allows the use of greedy algorithms, achieving a $\left(1 - \frac{1}{e}\right)$ approximation guarantee, where e denotes Euler's number.

Additional control methodologies for asynchronous best-response dynamics on binary decision networks have been explored to design optimal incentive-based interventions (Riehl et al., 2018a). Subsequent research was directed toward improving algorithmic scalability and refining the seed-selection process (see, Goyal et al., 2011). Parallel work examined cost-minimization variants of the problem. The Target Set Selection (TSS) problem, introduced by N. Chen, 2009, seeks the smallest set of initial adopters capable of activating the entire network.

Surveys and further developments can be found in W. Chen et al., 2013 and Como et al., 2021. Weighted extensions such as the WTSS problem allow heterogeneous activation costs Cordasco et al., 2015; Raghavan and Zhang, 2019. A particularly relevant extension for policy design is the *Least-Cost Influence Problem* (LCIP), introduced in Gunnec, 2012 and further studied in Günneç et al., 2020a, 2020b. In this framework, the planner can assign fractional incentives that effectively reduce individuals' adoption thresholds. This allows the modeling of realistic interventions such as targeted subsidies or discounts, where the objective is to minimize intervention costs while ensuring large-scale adoption.

Related fractional formulations have been considered by Demaine et al., 2014, who show that allowing fractional allocations can substantially alter diffusion outcomes without increasing computational complexity. Similar problems appear under different names, such as the Targeting with Partial Incentives (TPI) problem studied by Cordasco et al., 2015.

The scalability of these approaches remains a major challenge in large social networks. To address this issue, analytical tools have been developed within the study of diffusion processes on random networks, based on configuration models and branching process approximations. Mean-field techniques enable the characterization of cascade behavior in terms of degree distributions and threshold statistics, thereby enabling tractable analysis of large populations without explicitly enumerating all nodes (see Rossi et al., 2017 and Ravazzi et al., 2023).

Additional work has connected diffusion models with game-theoretic learning dynamics and coordination processes in networks (e.g., Morris, 2000, Acemoglu and Ozdaglar, 2011). Building on these mathematical and algorithmic foundations, the present dissertation investigates how regulatory interventions affect both equilibrium outcomes and the dynamic paths leading to them. By focusing on subclasses of super-modular network games, it identifies conditions under which interventions preserve monotonicity and ensure convergence toward socially desirable equilibria in decentralized systems.

1.3 Contributions and Outline

Several contributions to the theory of network games and mechanism design under complementarities are established herein.
Specifically, a comprehensive framework for **Intervention Design on Super-Modular Network Games** is introduced. By explicitly modeling the choice of payoff parameters by the regulator, equilibrium theory is connected with design problems. By leveraging methods from lattice theory and monotone comparative statics, specific conditions are identified under which interventions can induce equilibrium uniqueness or dominance solvability.

Furthermore, subclasses of network games are isolated where coordination can be achieved through minimal or localized interventions, advancing the literature on equilibrium selection. Finally, a unified mathematical treatment is provided, capable of addressing both finite, deterministic action spaces and large-scale networks.

The structural organization of the manuscript reflects these methodological advances. The thesis is structured as follows:

- **Theoretical Foundations: Game Theory and Network Dynamics.**
  This chapter establishes the mathematical and theoretical foundations for the dissertation. It introduces the fundamental notation and core concepts of game theory, with a focus on network games characterized by supermodular payoffs. A key component of this chapter is the formalization of policy interventions, describing the mechanisms through which a central regulator can influence incentives and system parameters to steer the system toward social optima. Additionally, it presents the learning dynamics which provide the computational framework for analyzing the emergence of equilibria in decentralized environments.

- **Optimal Intervention in Non-Binary Super-Modular Games.**
  This chapter examines mechanism designs aimed at steering decentralized systems toward their maximal Nash equilibrium. Focusing on finite unimodal games (a concept formally defined in Chapter 3, Definition 3.1.1), interventions are modeled as lower-bound restrictions on action sets, effectively preventing agents from adopting strategies below a designated threshold, thereby pruning suboptimal branches of the decision space. A distributed algorithmic approach based on time-reversible Markov chains is proposed to compute minimum-cost sufficient interventions by tracing weak improvement paths.

- **Optimal Seeding in Large-Scale Super-Modular Network Games.** Shifting the focus to binary-action environments, this chapter tackles the scalability and information limitations of the Optimal Seeding Problem in large populations. To overcome the inherent NP-hardness of exact solutions, the analysis leverages a directed configuration model and Local Mean-Field approximations. This methodology rigorously reformulates the seeding problem into a dimension-independent linear program that relies exclusively on aggregate network statistics rather than exact topological data.

- **Optimal Intervention in Large-Scale Super-Modular Network Games.** Building upon the statistical framework of the previous chapter, this section generalizes the intervention scope to address the Least-Cost Influence Problem (LCIP). It expands beyond binary seed selection by allowing the central planner to assign fractional incentives (e.g., targeted subsidies) that continuously lower individual adoption thresholds. This complex optimization problem is successfully recast as a highly scalable linear program.

- **Conclusions and Future Research.** The final chapter synthesizes the main theoretical and algorithmic contributions of the work. It also outlines potential avenues for future research, including the extension of these frameworks to environments with mixed externalities and dynamically evolving (endogenous) network topologies.

Chapter 2

Theoretical Foundations: Game Theory and Network Dynamics

In this chapter, the notation, theoretical concepts, and mathematical tools utilized throughout this dissertation are introduced.

Notation

Initially, the fundamental notation employed throughout this work is established. Vectors are denoted by bold lowercase letters, matrices and random variables by uppercase letters, and sets (as well as set-valued functions) by calligraphic symbols. All vectors are assumed to be column vectors. The transpose of a vector $\mathbf{x}$ is denoted by $\mathbf{x}^\top$; the same notation applies to matrices. The vector of all ones is denoted by the symbol $\mathbf{1}$, while the identity matrix is represented by $\mathbb{1}$, with the respective dimensions being dictated by the mathematical context. Throughout the dissertation, the standard componentwise partial order on $\mathbb{R}^n$ is considered. In particular, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, the relation $\mathbf{x} \leq \mathbf{y}$ is understood to hold if and only if $x_i \leq y_i$ for all $i = 1, 2, \ldots, n$. Moreover, $\mathbf{x} < \mathbf{y}$ indicates that $\mathbf{x} \leq \mathbf{y}$ and that the inequality is strict for at least one component. Additional notations will be introduced throughout this work and explained when needed.

2.1 Strategic Interactions on Networks

In this section, the foundational concepts of game theory, a mathematical framework utilized for the analysis of competitive environments characterized by interacting decision-makers, are established. These decision-makers are formally designated as agents (or players). A defined set of feasible strategies is available to each agent, and a specific strategy is selected with the strictly rational objective of maximizing a utility function. This utility is contingent upon both the agent's individual strategy and the strategy profile adopted by the complementary set of agents. Formally, a (strategic form) game is defined as a tuple $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ where

- $\mathcal{V} = \{1, \dots, n\}$ is the finite set of agents (or players),

while, for each agent $i \in \mathcal{V}$

- $\mathcal{A}_i$ is the set of strategies (or actions). Throughout most cases examined in this dissertation, the set $\mathcal{A}_i$ is assumed to be finite subset of $\mathbb{R}$.
- $u_i : \mathcal{A}_1 \times \cdots \times \mathcal{A}_n \to \mathbb{R}$ is the utility function (reward or payoff function).

The strategy executed by an agent $i \in \mathcal{V}$ is denoted by $x_i \in \mathcal{A}_i$ and the strategies of all agents are gathered in a vector $\mathbf{x} \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ which is formally referred to as a strategy profile or configuration. The configuration space is uniformly denoted by $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. The utility function $u_i : \mathcal{A} \to \mathbb{R}$ returns the payoff $u_i(\mathbf{x})$ received by agent i when a strategy profile $\mathbf{x} \in \mathcal{A}$ is collectively realized, meaning each agent j executes strategy $x_j \in \mathcal{A}_j$. Let $\mathbf{x}_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ be the vector obtained from the strategy profile $\mathbf{x}$ by removing its i-th entry, such that $\mathbf{x}_{-i} \in \mathcal{A}_{-i} = \prod_{j\neq i} \mathcal{A}_j$. Then, with a slight abuse of notation, it can be written

$$u_i(x_i, \mathbf{x}_{-i}) = u_i(\mathbf{x})$$

for the utility received by agent i when strategy x_i is selected by agent i, and strategy profile $\mathbf{x}_{-i}$ is executed by the remaining agents. A foundational assumption of non-cooperative game theory dictates that every agent $i \in \mathcal{V}$ operates rationally, selecting a strategy x_i to maximize their individual utility $u_i(x_i, \mathbf{x}_{-i})$.

Given the dependence of the utility on the strategies of the opposing agents $\mathbf{x}_{-i}$, the (set-valued) *Best Response* correspondence is introduced as:

$$\mathcal{B}_i(\mathbf{x}_{-i}) = \underset{x_i \in \mathcal{A}_i}{\operatorname{argmax}}\, u_i(x_i, \mathbf{x}_{-i}) \tag{2.1}$$

Under the assumption that agent $i \in \mathcal{V}$ possesses complete information regarding the static strategies executed by the remaining agents, the best-response correspondence identifies the mathematically rational choice; specifically, it identifies the strategy (or strategies) $x_i \in \mathcal{A}_i$ that maximizes the utility function, conditional upon the opponents playing $\mathbf{x}_{-i}$. Notice that $\mathcal{B}_i(\mathbf{x}_{-i})$ can be an empty set when $\mathcal{A}_i$ is not finite.
The definition of a (pure strategy) Nash equilibrium is thus formalized as follows.

Definition 2.1.1. *Given a game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$, a (pure strategy)* ***Nash equilibrium*** *is a strategy configuration $\mathbf{x}^* \in \mathcal{A}$ such that*

$$x_i^* \in \mathcal{B}_i(\mathbf{x}_{-i}^*), \qquad \forall i \in \mathcal{V}. \tag{2.2}$$

The Nash equilibrium is strict if, moreover, $|\mathcal{B}_i(\mathbf{x}_{-i}^*)| = 1$ for every agent i.

Consequently, a Nash equilibrium constitutes a strategy profile wherein no agent possesses an incentive to unilaterally deviate from the selected strategy; the utility derived is optimized relative to the strategies fixed by the other agents. It should be noted that a Nash equilibrium does not preclude the possibility that coordinated deviations by multiple agents could yield higher collective or individual utilities.
Throughout, the notation $\mathcal{A}^*$ will denote the set of Nash equilibria of a game, which can be empty or include one or more elements. Existence and uniqueness of Nash equilibria are not guaranteed in general.

Network Games

Within this subsection, a class of games particularly relevant to the study of interconnected systems, designated as network games, is examined. In these theoretical models, players are structurally represented as nodes within a graph, and their corresponding utility functions are strictly contingent upon their own strategic choices as well as the actions executed by their connected peers. Prior to the formal introduction of network games, foundational graph-theoretic concepts must be established; comprehensive notation and extended structural details are provided in Appendix A. Specifically, for a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where the node set $\mathcal{V}$ corresponds to the individual agents, the out-neighborhood of an agent i, formally expressed as $\mathcal{N}_i = \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\}$ is defined as the subset of nodes to which agent i directs a connection. Furthermore, the out-degree, denoted by $k_i = |\mathcal{N}_i|$ is utilized to quantify the total number of these outgoing connections, effectively capturing the scope of the agent's interactions. In weighted configurations, k_i is generalized as the weighted out-degree of agent i, given by $k_i = \sum_{j \in \mathcal{V}} W_{i,j}$. These topological constructs are strictly required to formalize the localized dependencies of the utility functions within the network.

Definition 2.1.2. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph. A game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i \in \mathcal{V}}, \{u_i\}_{i \in \mathcal{V}})$ is called a* ***network game*** *over $\mathcal{G}$ (or simply a $\mathcal{G}$-game) if, for every player $i \in \mathcal{V}$, it holds that*

$$u_i(\mathbf{x}) = u_i(\mathbf{y})$$

for any configurations $\mathbf{x}, \mathbf{y} \in \mathcal{A}$ such that

$$x_j = y_j, \qquad \textit{for all } j \in \mathcal{N}_i \cup \{i\}$$

Naturally, every game is a network game with respect to the complete graph.

The foundational premise of a network game dictates that interactions among players are structurally constrained by the underlying graph topology. The utility derived by any individual player is evaluated strictly based on their chosen strategy and the specific action configuration adopted by their out-neighborhood. When individual decisions are restricted to binary spaces, strategic preferences for one action over another frequently emerge based on the proportion of peers selecting the identical action. For instance, in models concerning the diffusion of innovations and beliefs, the utility derived from adopting a specific behavior increases as a function of the number of neighbors adopting that same behavior, a dynamic that is formally illustrated in the subsequent example.

Example 2.1.1. *A **Network Coordination Game** on a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ is defined as a game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ characterized by binary action sets $\mathcal{A}_i = \{0, 1\},\ \forall i \in \mathcal{V}$. The utility functions are formulated as:*

$$u_i(\mathbf{x}) = \sum_{j\in\mathcal{V}} W_{i,j}((1 - x_i)(1 - x_j) + x_i x_j) + b_i x_i \tag{2.3}$$

where the term $W_{i,j} \geq 0$ denotes the non-negative weight of the directed connection. Furthermore, a constant parameter b_i in $[-k_i, k_i]$ is introduced to model a possible bias of player i toward action 0 (if $b_i < 0$) or action 1 (if $b_i > 0$).

Given an agent $i \in \mathcal{V}$ and any choice of actions $\mathbf{x}_{-i} \in \mathcal{A}_{-i}$, action 1 is determined to be a Best Response, in accordance with Equation (2.1)*, if and only if:*

$$u_i(1, \mathbf{x}_{-i}) \geq u_i(0, \mathbf{x}_{-i}) \iff \sum_{j\in\mathcal{V}} W_{i,j} x_j + b_i \geq \sum_{j\in\mathcal{V}} W_{i,j}(1 - x_j).$$

This condition implies that the best response evaluates to 1 if and only if the cumulative weighted sum of active neighbors strictly exceeds a defined threshold. Given that $k_i = \sum_{j\in\mathcal{V}} W_{i,j}$, the inequality is algebraically reformulated as:

$$2\sum_{j\in\mathcal{V}} W_{i,j} x_j \geq k_i - b_i$$

Therefore, the best response correspondence for an agent $i \in \mathcal{V}$ is given by

$$\mathcal{B}_i(\mathbf{x}_{-i}) = \begin{cases} \{1\} & \textit{if } \sum\limits_{j\in\mathcal{V}} W_{i,j} x_j > r_i \\ \{0,1\} & \textit{if } \sum\limits_{j\in\mathcal{V}} W_{i,j} x_j = r_i \\ \{0\} & \textit{if } \sum\limits_{j\in\mathcal{V}} W_{i,j} x_j < r_i \end{cases} \tag{2.4}$$

where

$$r_i = \frac{k_i - b_i}{2} \tag{2.5}$$

denotes the threshold of player $i \in \mathcal{V}$.

While the intrinsic bias b_i is naturally constrained within the interval $[-k_i, k_i]$ for agents susceptible to peer influence, relaxing these bounds provides a rigorous theoretical mechanism to model stubborn agents, individuals whose strategic choices remain invariant to network effects.

Specifically, an agent i is classified as stubborn if their baseline bias strictly dominates the aggregate network influence:

- *If $b_i > k_i$ (yielding a threshold $r_i < 0$), action 0 is strictly dominated, and the agent's best response is invariably $\mathcal{B}_i(\mathbf{x}_{-i}) = \{1\}$ for any feasible configuration of opponents $\mathbf{x}_{-i}$.*
- *Conversely, if $b_i < -k_i$ (yielding $r_i > k_i$), action 1 is strictly dominated, fixing the best response at $\mathcal{B}_i(\mathbf{x}_{-i}) = \{0\}$ regardless of peer actions.*

In the absence of such stubborn agents, that is, when $-k_i \le b_i \le k_i$ for all $i \in \mathcal{V}$, strictly dominated actions do not exist, allowing the system to sustain multiple equilibria, notably including the consensus configurations $\mathbf{x} = \mathbf{0}$ and $\mathbf{x} = \mathbf{1}$.

Example 2.1.2. ***Quadratic Games*** *constitute a family of games that has garnered significant attention recently and appeared in various applicative scenarios within the social and economic sciences.*

In the general form, quadratic games possess continuous strategy sets, i.e., $\mathcal{A}_i \subseteq \mathbb{R}$, for all i, and the utilities are given by

$$u_i(\mathbf{x}) = \rho_i(x_i) + \delta x_i \sum_{j \in \mathcal{V}} W_{i,j} x_j \tag{2.6}$$

where $\rho_i : \mathcal{A}_i \to \mathbb{R}$.

These games model situations wherein agents are engaged in some common activity: x_i is the level of activity of agent i and $\rho_i(x_i)$ represents the utility in the absence of social interactions. Social interactions are captured by $\delta W_{i,j}$, where δ parameterizes the importance of social interaction and $W_{i,j}$ quantifies the strength of the influence of j on i.

In the special case when each $\rho_i(x_i) = b_i x_i - \frac{1}{2}x_i^2$, i.e., is quadratic in x_i, the canonical quadratic game is obtained, with utilities formulated as, for $i \in \mathcal{V}$

$$u_i(x_i, \mathbf{x}_{-i}) = \underbrace{\left(b_i + \delta \sum_{j \in \mathcal{N}} W_{i,j} x_j\right) x_i}_{\textit{returns from own strategy}} - \underbrace{\frac{1}{2} x_i^2}_{\textit{private costs of own strategy}} \tag{2.7}$$

where $x_i \in \mathbb{R}_+$ is the activity level of agent i, b_i is the marginal individual reward from activity level x_i, $\frac{1}{2}x_i^2$ is the cost for providing activity level x_i and $W_{i,j} \ge 0$ is the benefit from peer interaction.

It is noted that the best response of agent i is always a singleton and is linear in x_j for all $j \in \mathcal{V} \setminus \{i\}$. More precisely, it is given by

$$\mathcal{B}_i(\mathbf{x}_{-i}) = b_i + \delta \sum_{j \in \mathcal{V}} W_{i,j} x_j. \tag{2.8}$$

Then, if they exist, Nash equilibria are the solutions of the system

$$\mathbf{x} = \mathbf{b} + \delta W \mathbf{x}. \tag{2.9}$$

If $\delta\rho(W) < 1$, the game possesses a unique Nash equilibrium given by

$$\mathbf{x}^* = (\mathbb{1} - \delta W)^{-1} \mathbf{b}$$

Remark 1. *The unique Nash equilibrium is fundamentally linked to the Katz-Bonacich centrality (Ballester et al., 2006; Bonacich, 1987). While traditional network centrality measures are frequently formulated to evaluate prestige via incoming links (as detailed in Appendix A.1), in the context of network games with strategic complementarities, an agent's equilibrium action is governed by the influence exerted by their observed peers.*

It is observed that, under the assumption $\delta\rho(W) < 1$ the equilibrium profile $\mathbf{x}^*$ corresponds precisely to the weighted Katz-Bonacich centrality vector, which accounts for the number and discounted length of all outgoing walks. The centrality of agent i can thus be explicitly expressed as:

$$x_i^* = \sum_{j \in \mathcal{V}} \sum_{k=0}^{\infty} \left(\delta^k W^k\right)_{i,j} b_j,$$

where b_j represents the intrinsic baseline utility of agent j, and $W_{i,j}^k$ is the weight of walks from i to j of length k.

2.2 Mechanism Design and Policy Interventions

The concept of interventions in network games is introduced in this section. A game-theoretic framework is developed for the analysis of these games and the design of computationally efficient interventions implementable by a central planner, both in deterministic networks and in large-scale networks generated by stochastic formation models when exact network data are unavailable.

In networked systems, agent decisions are governed jointly by direct incentives and strategic interactions. Due to the unobservability of individual actions, indirect interventions are required to shape equilibrium behavior and facilitate coordination. Because agent responses to these interventions are contingent upon both individual utility maximization and the strategic adjustments of peers, policy design must explicitly account for these interdependencies. This is accomplished by evaluating how parameter modifications alter the equilibrium landscape to maximize the aggregate action.

Intervention Mechanism

In non-cooperative environments, suboptimal collective behavior are frequently generated by the uncoordinated maximization of individual payoffs. To mitigate these systemic inefficiencies, targeted interventions are implemented by the central regulator to steer the system toward socially optimal configurations.
From the perspective of a central planner (CP), the previously developed characterization of equilibria is frequently exploited to determine the optimal intervention required to maximize a system-level objective. An intervention parameter

$$\mathbf{h} \in \mathcal{H},$$

is selected by the central regulator, where $\mathcal{H}$ is a nonempty design parameter space (finite-dimensional vector space, function space, or abstract parameter set).

The game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ is modified by an intervention by altering payoffs, action sets, information, or feasible outcomes. This is modeled through a game transformation mapping,

$$\mathcal{T} \,:\, \mathcal{H} \to \mathfrak{G},$$

where $\mathfrak{G}$ is a class of strategic-form games with the same player set.

For each $\mathbf{h} \in \mathcal{H}$,

$$\Gamma(\mathbf{h}) := \mathcal{T}(\mathbf{h}) = \Big(\mathcal{V}, \{\mathcal{A}_i(\mathbf{h})\}_{i\in\mathcal{V}}, \{u_i^{\mathbf{h}}\}_{i\in\mathcal{V}}\Big)$$

Common special cases include:

- **Payoff Intervention**

  Action sets remain unchanged and utilities are modified:

  $$\mathcal{A}_i(\mathbf{h}) = \mathcal{A}_i, \quad u_i^{\mathbf{h}}(x_i, \mathbf{x}_{-i}) = u_i(x_i, \mathbf{x}_{-i}) + f_i(x_i, \mathbf{h}).$$

  The term $f_i(x_i, \mathbf{h})$ represents the intervention-induced perturbation to the marginal utility of agent i. It is quantified as the net deviation in standalone returns resulting directly from the implemented regulatory policy.

- **Action-Set Intervention**

  The regulator restricts feasible actions: $\mathcal{A}_i(\mathbf{h}) \subseteq \mathcal{A}_i$.

  Under this intervention mechanism, the underlying utility functions of the agents are preserved, while the set of permissible actions is strictly constrained by the central regulator. Rather than altering the intrinsic payoff structures, the strategic environment is modified by imposing direct limitations on the feasible choices available to individual players.

  Practically, this is frequently modeled as the imposition of a lower bound on the action space, thereby ensuring that a minimum level of effort or engagement is exerted by the targeted agents. This condition is formally expressed by restricting the action set to

  $$\mathcal{A}_i(\mathbf{h}) = \{x_i \in \mathcal{A}_i \mid x_i \geq h_i\}, \qquad \forall\, i \in \mathcal{V}.$$

  Furthermore, this mathematical formulation naturally encompasses specific deployment scenarios, such as the seeding problem, wherein selected agents are compelled to adopt a predetermined action, effectively reducing their viable strategy set to a single mandatory element. By directly restricting the configuration space, decentralized systems can be effectively steered toward socially optimal equilibria.

Figure 2.1 illustrates a sample network of agents, and how the regulator's intervention h_i influences their payoff. This graph highlights the local interactions that determine best response updates.

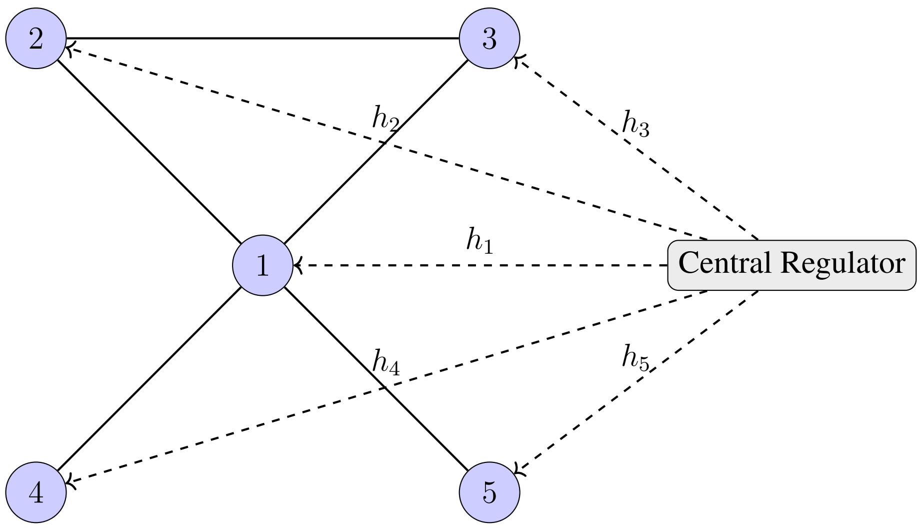

Local strategic complementarities

Figure 2.1: Network interaction graph. Players interact locally through the network, while a central regulator influences incentives via parameters h_i.

Target Class of Games. The game structure is reshaped by the central planner through interventions, ensuring its membership within a subclass of games characterized by specific properties. Let $\Gamma^* \subseteq \mathfrak{G}$ denote a **target class of games** defined by structural or equilibrium properties. Examples include:

- Games with unique Nash equilibrium
- Games whose equilibria satisfy a system-level constraint
- Games implementing a desired social choice rule

Planner's Objective. To every intervention $\mathbf{h}$ is associated a cost

$$C \;:\; \mathcal{H} \to \mathbb{R}_+. \tag{2.10}$$

Intervention mechanisms $\mathbf{h} \in \mathcal{H}$ must be designed such that the induced game

$$\Gamma(\mathbf{h}) = \Big(\mathcal{V}, \{\mathcal{A}_i(\mathbf{h})\}_{i\in\mathcal{V}}, \{u_i^{\mathbf{h}}\}_{i\in\mathcal{V}}\Big)$$

resides within a target class Γ^*.

The regulator, therefore, aims to solve the following problem

$$\begin{array}{ll} \underset{\mathbf{h} \in \mathcal{H}}{\text{minimize}} & C(\mathbf{h}) \\ \text{subject to} & \Gamma(\mathbf{h}) \in \Gamma^* \end{array} \tag{2.11}$$

Remark 2. *Note that while the constraints on the action set are local, namely, $x_i \in \mathcal{A}_i$, the interventions constraint set $\mathcal{H}$ allows both local, e.g., $\mathcal{H} = \mathbb{R}_+^{\mathcal{V}}$, and coupled constraints, e.g., $\mathcal{H} = \{\mathbf{h} \in \mathbb{R}^{\mathcal{V}} : \|\mathbf{h}\| \leq c\}$, for some $c > 0$. Another notable example is given by*

$$\mathcal{H} = \{\mathbf{h} \in \mathbb{R}^{\mathcal{V}} \ : \ h_i \in \mathbb{R}, \ \forall i \in \mathcal{I} \text{ and } h_i = 0, \ \forall i \in \mathcal{V} \setminus \mathcal{I}\},$$

where $\mathcal{I} \subseteq \mathcal{V}$ denotes a specific subset of targeted players subject to intervention.

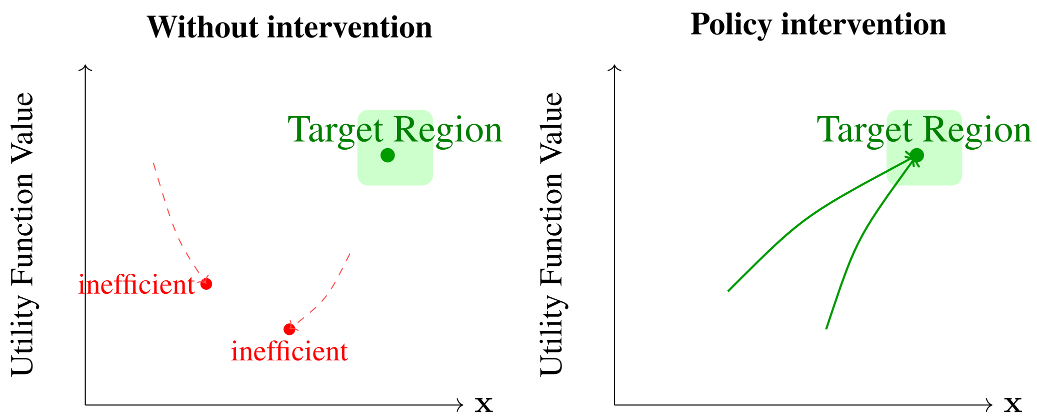

Figure 2.2: Policy intervention in a network game. Inefficient equilibria (red) arise before intervention. By redesigning game parameters, the planner eliminates these equilibria and ensures convergence toward the target region (green).

The effectiveness of the designed interventions is illustrated in Figure 2.2. It is shown how the modification of game parameters geometrically eliminates sub-optimal equilibrium regions, thereby steering the network toward a target space characterized exclusively by efficient equilibria.

Example 2.2.1. *Consider the canonical quadratic game with utility functions specified by Equation* (2.7)*, where the central planner intervenes by modifying the individual marginal reward via the parameter $h_i, \ \forall i \in \mathcal{V}$. The induced game*

$$\Gamma(\mathbf{h}) = \left(\mathcal{V}, \{\mathcal{A}_i(\mathbf{h})\}_{i \in \mathcal{V}}, \{u_i^{\mathbf{h}}\}_{i \in \mathcal{V}}\right)$$

without imposing restrictions on the action sets (i.e., $\mathcal{A}_i(\mathbf{h}) = \mathcal{A}_i$, for all $i \in \mathcal{V}$), and with the modified utility function:

$$
\begin{aligned}
u_i^{\mathbf{h}}(x_i, \mathbf{x}_{-i}) &= b_i x_i + h_i x_i - \frac{1}{2}x_i^2 + \delta x_i \sum_{j \in \mathcal{V}} W_{i,j} x_j \\
&= (b_i + h_i)x_i - \frac{1}{2}x_i^2 + \delta x_i \sum_{j \in \mathcal{V}} W_{i,j} x_j
\end{aligned}
$$

Prior to targeted interventions, the existence and uniqueness of an equilibrium configuration, $\mathbf{x}^*$,cannot be guaranteed. To establish well-posedness, strict spectral bounds, $\delta\rho(W) < 1$, must be imposed. Provided this criterion is satisfied, the resulting Nash equilibrium is analytically derived as $\mathbf{x}^* = (\mathbb{1} - \delta W)^{-1}(\mathbf{b} + \mathbf{h})$. The planner seeks to minimize the intervention cost while ensuring the aggregate effort exceeds a target threshold α. This yields the following convex program:

$$
\begin{aligned}
\underset{\mathbf{h} \in \mathbb{R}^{\mathcal{V}}}{\text{minimize}} \quad & \frac{1}{2}\|\mathbf{h}\|^2 \\
\text{subject to} \quad & \mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}(\mathbf{b} + \mathbf{h}) \geq \alpha
\end{aligned}
$$

The associated Lagrangian function is defined as

$$
\mathcal{L}(\mathbf{h}, \mu) = \frac{1}{2}\mathbf{h}^\top \mathbf{h} - \mu \left(\mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}(\mathbf{b} + \mathbf{h}) - \alpha \right)
$$

The optimal solution must satisfy the Karush-Kuhn-Tucker (KKT, Kuhn and Tucker, 2013) conditions:

- *Stationarity:*

$$
\begin{aligned}
\nabla_{\mathbf{h}} \mathcal{L} =& \mathbf{h} - \mu \left((\mathbb{1} - \delta W)^{-1} \right)^\top \mathbf{1} = 0 \\
\Longrightarrow& \mathbf{h} = \mu \left((\mathbb{1} - \delta W)^{-1} \right)^\top \mathbf{1}
\end{aligned}
$$

- *Primal feasibility:*

$$
\begin{aligned}
& \mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}(\mathbf{b} + \mathbf{h}) \geq \alpha \\
\Longrightarrow & \mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}\mathbf{h} \geq \alpha - \mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}\mathbf{b}
\end{aligned}
$$

- *Dual feasibility:* $\mu \geq 0$.
- *Complementary slackness:* $\mu \left(\mathbf{1}^\top (\mathbb{1} - \delta W)^{-1}(\mathbf{b} + \mathbf{h}) - \alpha \right) = 0$.

If the constraint is assumed to be inactive, then it holds that

$$\mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{h} > \alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b}.$$

complementary slackness forces $\mu = 0$. *From stationarity, this implies* $\mathbf{h} = 0$. *the condition* $\mathbf{1}^T(\mathbb{1} - \delta W)^{-1}\mathbf{h} = \mathbf{0}$ *is mathematically consistent if and only if the original inequality is strictly satisfied in the absence of intervention, meaning:*

$$\alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b} \leq 0.$$

If the constraint is assumed to be active, $\mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{h} = \alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b}$ *and complementary slackness implies* $\mu > 0$.

From stationarity: $\mathbf{h} = \mu\left((\mathbb{1} - \delta W)^{-1}\right)^\top \mathbf{1}$, *and imposing the active constraint:*

$$\mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{h} = \mu\mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\left((\mathbb{1} - \delta W)^{-1}\right)^\top \mathbf{1} = \alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b}$$

$$\implies \mu = \frac{\alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b}}{\left\|\left((\mathbb{1} - \delta W)^{-1}\right)^\top \mathbf{1}\right\|^2} > 0$$

Thus, the optimal solution is

$$\mathbf{h}^* = \begin{cases} 0 & \text{if } \alpha \leq \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b} \\ \frac{\alpha - \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b}}{\left\|\left((\mathbb{1} - \delta W)^{-1}\right)^\top \mathbf{1}\right\|^2}\left((\mathbb{1} - \delta W)^{-1}\right)^\top \mathbf{1} & \text{if } \alpha > \mathbf{1}^\top(\mathbb{1} - \delta W)^{-1}\mathbf{b} \end{cases}$$

The derived optimal policy dictates that a non-zero intervention ($\mathbf{h}^* > 0$*) is implemented if and only if the target aggregate effort* α *strictly exceeds the autonomous equilibrium level* $\sum_{i \in \mathcal{V}} x_i^*$. *Otherwise, the optimal intervention is null.*

Example 2.2.2 (Seeding Problem)**.** *Consider the coordination game of Example* (2.3)*. Consider a setting where initially all agents have selected strategy 0, except for an initial set* $\mathcal{S}(0)$ *of agents playing strategy 1. The set* $\mathcal{S}(0)$ *is called the seed set and represents early adopters (e.g.,first adopters of a new product or innovation, initial bank failures, etc.). At every iterative step* $t + 1$ *each agent takes a best response action according to* (2.4)*:*

$$x_i(t+1) = \begin{cases} \{1\} & \text{if } \sum\limits_{j \in \mathcal{V}} W_{i,j}x_j(t) \geq r_i \\ \{0\} & \text{if } \sum\limits_{j \in \mathcal{V}} W_{i,j}x_j(t) < r_i \end{cases}$$

and therefore selects strategy 1 *if at least* r_i *of their neighbors have selected strategy* 1 *at step* t *and* 0 *otherwise.*

Algorithm 1 Diffusion Dynamics for the Coordination Games

Set: $k = 0$, Initialize $x_i = 1$ for all $i \in \mathcal{S}(0)$ and $x_i = 0$, otherwise.
Iterate:

$$\mathcal{S}(t+1) = \left\{ i \in \mathcal{V} \ : \ \sum_{j \in \mathcal{S}(t)} W_{i,j} \geq r_i \right\}$$

The process defines a monotone update, such that

$$\mathcal{S}(t) \subseteq \mathcal{S}(t+1), \qquad \textit{for all } t \geq 0.$$

Since the network is finite (of order $n = |\mathcal{V}|$), the sequence stabilizes in at most n iterations. The set of final adopters, given the initial seed set $\mathcal{S}^*(\mathcal{S}(0))$, is therefore defined as

$$\mathcal{S}^*(\mathcal{S}(0)) = \bigcup_{t=0}^{\infty} \mathcal{S}(t) = \mathcal{S}(n).$$

By construction, this terminal set constitutes a stable configuration, inherently satisfying the fixed-point condition wherein no further activations occur.
Intervention: *The central planner's objective is to select an optimal initial seed set $\mathcal{S}(0) \subseteq \mathcal{V}$ to maximize the cardinality of the final adopting population, subject to a strict budget constraint k on the number of initially active agents. This intervention strategy is formalized by the following optimization problem:*

$$\begin{aligned} \underset{\mathcal{S}(0) \subseteq \mathcal{V}}{\text{maximize}} \quad & |\mathcal{S}^*(\mathcal{S}(0))| \\ \text{subject to} \quad & |\mathcal{S}(0)| = k \end{aligned}$$

The identification of targeted individuals to maximize information propagation, known as the seeding problem, represents a well-established paradigm in the study of viral marketing, financial contagion, and technology adoption. Various network-centrality-based heuristics have been proposed in the literature to select seeds in a manner conducive to boosting diffusion. The determination of optimal 'seeds' for maximizing information spread is maintained as a critical policy question in contexts like brand awareness (Richardson and Domingos, 2002), microfinance (Banerjee et al., 2013), and technology adoption (Beaman et al., 2021).

This optimization problem is not computationally tractable for large populations. When agent i threshold is selected from $[0, \ldots k_i]$ uniformly at random, (and one aims at maximizing the expected size of the final set given the randomness on thresholds) the problem in NP-hard, as shown by Kempe et al., 2003.

2.3 Strategic Complementarities and Super-Modular Structures

Super-modular games constitute an essential class of games (Milgrom and Roberts, 1990), utilized to model "strategic complementarities" (Jackson and Zenou, 2015), wherein an increase in strategy by one agent incentivizes others to respond similarly. Complementarities are expressed both in terms of constraints and payoff functions. In terms of constraints, two activities are complementary if an increase in one activity does not reduce the feasible domain of the other. This is mathematically captured by *lattices*. In terms of payoffs, complementarity between two activities is established if an increase in one enhances the marginal benefit of the other. These games possess significant properties in terms of existence and structure of Nash equilibria.

In the general setting, the definition of super-modular games requires that the strategy sets $\mathcal{A}_i$ be partially ordered sets satisfying the properties of a lattice. The set $\mathcal{A}_i$ is assumed to be a finite (or countable) compact subset of $\mathbb{R}$ for all $i \in \mathcal{V}$. Then, a component-wise partial order $\leq$ is considered on the strategy configuration space $\mathcal{A}$, formally defined by $\mathbf{x} \leq \mathbf{y}$ if and only if $x_i \leq y_i$ for all $i \in \mathcal{V}$. The fundamental feature of super-modular games is the increasing difference property, defined as follows.

Definition 2.3.1. *A utility function u_i defined on a lattice $\mathcal{A}$ satisfies the* ***increasing differences property*** *if for all $y_i \geq x_i$ and $\mathbf{y}_{-i} \geq \mathbf{x}_{-i}$, it holds that:*

$$u_i(y_i, \mathbf{y}_{-i}) - u_i(x_i, \mathbf{y}_{-i}) \geq u_i(y_i, \mathbf{x}_{-i}) - u_i(x_i, \mathbf{x}_{-i}) \tag{2.12}$$

This condition can also be written as

$$u_i(y_i, \mathbf{y}_{-i}) - u_i(y_i, \mathbf{x}_{-i}) \geq u_i(x_i, \mathbf{y}_{-i}) - u_i(x_i, \mathbf{x}_{-i}).$$

Consequently, the variation in a player's utility resulting from an increased action is monotonically non-decreasing with respect to the actions executed by other players. The class of games with the increasing difference property models the so-called *strategic complements effect* in economics: an increase in the effort exerted by opponents incentivizes a corresponding increase in an individual player's effort. Such effects are frequently observed in the modeling of phenomena such as technology adoption or social behaviors.

The increasing differences condition in Equation (2.12) is structurally equivalent to the defining properties of A-coordinating games, formalized by Sakhaei et al., 2023. Within this framework, strategic complementarities induce a monotonic alignment of marginal utilities. Consequently, unilateral upward deviations by any subset of agents monotonically increase the best-response thresholds across the network. This property facilitates the iterative elimination of strictly dominated strategies, thereby ensuring the existence of extremal equilibria. The formal characterization of these extremal equilibria is required to establish the predictability necessary for robust mechanism design. Specifically, the deterministic convergence to maximal equilibria, guaranteed by these coordinating properties, enables the rigorous formulation of optimal intervention mechanisms.

This theoretical progression is directly exemplified by Como et al., 2021, wherein it is demonstrated that optimal targeting strategies relies upon exploiting such underlying strategic complementarities to predictably steer network dynamics.

A notable example of games with the increasing difference property is the network coordination game discussed in the previous section.

Example 2.3.1. *Network coordination games posses the increasing difference property. From expression* (2.4)*, the Best Response for player i is monotonically non-decreasing function in any of the components of* $\mathbf{x}_{-i}$.

The increasing difference property is fundamental to many important properties of a game. For games with non-finite action sets, additional regularity assumptions are required, which motivates the following definition.

The definition of super-modular games is formally established as follows.

Definition 2.3.2. *A game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ is* ***super-modular*** *if:*

- $\mathcal{A}_i$ *is a compact subset of* $\mathbb{R}$*;*
- $u_i(x_i, \mathbf{x}_{-i})$ *is upper semi-continuous in x_i and continuous in* $\mathbf{x}_{-i}$*;*
- $u_i(x_i, \mathbf{x}_{-i})$ *satisfy the increasing difference property* (2.12)*,*

for all $i \in \mathcal{V}$.

Within continuous action spaces, the verification of the increasing differences property is rendered analytically tractable by restricting the payoff mappings to $\mathcal{C}^2(\mathcal{A}, \mathbb{R})$, the class of real-valued functions over a domain $\mathcal{A}$ endowed with continuous second-order partial derivatives, thereby constraining the game's differential structure sufficiently to enable local characterizations of global phenomena.

Lemma 1. *Let a strategic game* $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ *be considered. It is assumed that the global configuration space is given by* $\mathcal{A} = \prod_{j\in\mathcal{V}} \mathcal{A}_j$*, where each individual action set* $\mathcal{A}_j \subset \mathbb{R}$ *is defined as a compact interval. Furthermore, it is assumed that the utility functions* $u_i : \mathcal{A} \to \mathbb{R}$ *belong to the space* $\mathcal{C}^2(\mathcal{A}, \mathbb{R})$ *for all* $i \in \mathcal{V}$*. Under these conditions, it is established that the game exhibits the increasing differences property if and only if the cross-partial derivatives of the utility functions are everywhere non-negative; that is:*

$$\frac{\partial^2 u_i}{\partial x_i \partial x_j}(\mathbf{x}) \geq 0, \quad \forall \mathbf{x} \in \mathcal{A}, \quad \forall\, j \neq i. \tag{2.13}$$

Proof. The necessity of the differential condition is established by assuming the utility function u_i possesses the increasing differences property. For any $y_i > x_i$, the incremental difference $u_i(y_i, \mathbf{x}_{-i}) - u_i(x_i, \mathbf{x}_{-i})$ is observed to be monotonically non-decreasing with respect to the complementary profile $\mathbf{x}_{-i}$. By dividing this difference by the strictly positive scalar $(y_i - x_i)$ and evaluating the limit as $y_i \to x_i$, the first-order partial derivative $\frac{\partial u_i}{\partial x_i}(\mathbf{x})$ is obtained. Since the pointwise limit of a sequence of non-decreasing functions strictly preserves monotonicity, the marginal utility $\frac{\partial u_i}{\partial x_i}(\mathbf{x})$ is guaranteed to be non-decreasing with respect to any opposing variable x_j, for $j \neq i$. Consequently, when the differential operator $\frac{\partial}{\partial x_j}$ is applied to this marginal utility, the established monotonicity dictates that the ensuing second-order cross-partial derivative is non-negative. Thus, the condition $\frac{\partial^2 u_i}{\partial x_i \partial x_j}(\mathbf{x}) \geq 0$ is obtained.

On the other hand, for every $a_i, b_i \in \mathcal{A}_i$ with $b_i \geq a_i$ and every $\mathbf{x}_{-i} \in \mathcal{A}_{-i}$ the following can be written

$$u_i(b_i, \mathbf{x}_{-i}) - u_i(a_i, \mathbf{x}_{-i}) = \frac{\partial u_i}{\partial x_j}(c_i, \mathbf{x}_{-i})$$

for some $c_i \in [a_i, b_i]$ by virtue of the mean value theorem.

If relation (2.13) holds, this last expression is non-decreasing in x_j for every $j \neq i$. This yields the increasing difference property (2.12). ■

Example 2.3.2. *Quadratic games with $\delta W_{i,j} \geq 0$ for every i, j qualify as the type of super-modular network games defined over positively weighted graphs that are studied in this thesis.*

The following fundamental statement summarizes the key properties that characterize these games features.

Theorem 2.3.1. *For a super-modular game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ the following facts hold:*

- *the set of pure Nash equilibria $\mathcal{A}^*$ is non-empty;*
- *there exist two extremal (i.e., maximum and minimum) pure Nash equilibria, denoted as the least equilibrium $\underline{\mathbf{x}}^*$ and the greatest equilibrium $\overline{\mathbf{x}}^*$, such that every pure Nash equilibrium $\mathbf{x}^*$ satisfies $\underline{\mathbf{x}}^* \leq \mathbf{x}^* \leq \overline{\mathbf{x}}^*$.*

Proof. Consider a super-modular game where the strategy sets $\mathcal{A}_i$ are compact subsets of $\mathbb{R}$. The configuration space $\mathcal{A}$ is thus a complete lattice under the component-wise partial order. Under the assumption of strategic complementarities (increasing differences), the best response correspondence $\mathcal{B}(\mathbf{x})$ maps the lattice $\mathcal{A}$ into the collection of its sub-lattices. By Tarski's Fixed Point Theorem, the set of Nash equilibria $\mathcal{A}^*$ is non-empty. Furthermore, the set $\mathcal{A}^*$ itself constitutes a complete lattice, which implies the existence of two extremal pure Nash equilibria, $\underline{\mathbf{x}}^*$ and $\overline{\mathbf{x}}^*$ such that every Nash equilibrium $\mathbf{x}^*$ satisfies $\underline{\mathbf{x}}^* \leq \mathbf{x}^* \leq \overline{\mathbf{x}}^*$. For a complete treatment, reference is made to Topkis, 1998. ■

Comparative Statics

In this subsection, the variation of equilibrium behavior with respect to exogenous parameters is analyzed. To avoid notational conflict with the policy interventions $\mathbf{h} \in \mathcal{H}$ defined previously, the parameter is modeled as a vector $\theta \in \Theta$, where Θ is a partially ordered set. Although comparative statics results are applicable to more abstract settings, focus is maintained on $\Theta \subseteq \mathbb{R}^{\mathcal{V}}$, which is a partially ordered set with the usual componentwise order. The mathematical formalization of parameterized super-modular structures and their associated comparative statics relies on foundational framework elements established by Milgrom and Roberts, 1990 and Topkis, 1998.

Definition 2.3.3. *Let Θ, the parameter space, be a partially ordered set.*
A ***parameterized super-modular game*** $\Gamma(\theta) = \left(\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i^\theta\}_{i\in\mathcal{V}}\right)$ *is defined as a family of super-modular games with payoff functions parameterized by θ in Θ, such that for each $i \in \mathcal{V}$:*

- *$\mathcal{A}_i$ is a compact subset of $\mathbb{R}$;*
- *$u_i^\theta(x_i, \mathbf{x}_{-i})$ is upper semi-continuous in x_i, for a fixed strategy profile $\mathbf{x}_{-i}$ and parameter θ.*
- *$u_i^\theta(x_i, \mathbf{x}_{-i})$ satisfies the increasing differences property with respect to the parameter θ:*

$$u_i^{\theta'}(y_i, \mathbf{y}_{-i}) - u_i^{\theta'}(x_i, \mathbf{y}_{-i}) \geq u_i^\theta(y_i, \mathbf{x}_{-i}) - u_i^\theta(x_i, \mathbf{x}_{-i}) \tag{2.14}$$

for all $y_i \geq x_i,\ \mathbf{y}_{-i} \geq \mathbf{x}_{-i}$ and $\theta' \geq \theta$

Theorem 2.3.2. *Assuming $\Gamma(\theta) = \left(\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i^\theta\}_{i\in\mathcal{V}}\right)$ is a parametrized super-modular game, the maximum and minimum Nash equilibria, denoted respectively by $\overline{\mathbf{x}}^*(\theta)$ and $\underline{\mathbf{x}}^*(\theta)$, are both non-decreasing functions, for each $\theta \in \Theta$.*

Proof. Let $\Gamma(\theta)$ be a parameterized super-modular game where the design parameter space Θ is a partially ordered set. The utility functions u_i^θ satisfy the increasing difference property not only with respect to the strategies $(x_i, \mathbf{x}_{-i})$ but also between the strategies and the parameter θ. This structural property guarantees that the best response correspondence shifts monotonically upward as θ increases. By applying the theory of monotone comparative statics to the equilibrium lattice, it is established that the maximum Nash equilibrium $\overline{\mathbf{x}}^*(\theta)$ and the minimum Nash equilibrium $\underline{\mathbf{x}}^*(\theta)$ are non-decreasing functions of θ. For a complete treatment, reference is made to Milgrom and Roberts, 1990. ∎

2.4 Learning and Best-Response Dynamics

Equilibrium analysis based on Nash equilibrium forms the foundation of most economic theory. In strategic environments characterized by equilibrium multiplicity, classical rationality and common knowledge assumptions are frequently insufficient for equilibrium selection. Consequently, the emergence of an equilibrium is typically modeled as the asymptotic limit of an iterative learning process, wherein strategies are sequentially adjusted over time. In contexts where the game is used as a modeling tool to solve multi-agent decision and optimization problems, such evolutionary dynamics can be interpreted as distributed algorithms.

The individual payoff function is maximized by each player, given the aggregated actions of their neighbors and the current value of the intervention signal. To capture this, it is assumed that the action of each player $i \in \mathcal{V}$ evolves over time according to the following dynamic

Maximal Best Response Dynamics

In this subsection, a game-theoretic learning process is introduced: the *Maximal Best Response Dynamics*. To ensure that such dynamics are well-defined, sufficient topological regularity is imposed on the strategic environment.

Assumption 1. *For a strategic-form game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$, it is assumed that for every player $i \in \mathcal{V}$, the action set $\mathcal{A}_i$ is a non-empty compact subset of $\mathbb{R}$, and the utility function $u_i(x_i, \mathbf{x}_{-i})$ is upper semi-continuous with respect to x_i for any fixed opponents' profile $\mathbf{x}_{-i}$.*

By Assumption 1, the extreme value theorem ensures that the utility function attains its maximum on the action space. Consequently, the set of maximizers (the best-response set $\mathcal{B}_i(\mathbf{x}_{-i})$) is both non-empty and compact.

Consequently, for each player $i \in \mathcal{V}$ and opponents' profile $\mathbf{x}_{-i}$, the maximal element of the best-response set is always uniquely determined. Formally:

$$\overline{\mathcal{B}}_i(\mathbf{x}_{-i}) = \max\{\mathcal{B}_i(\mathbf{x}_{-i})\} = \max\left\{\underset{x_i\in\mathcal{A}_i}{\operatorname{argmax}}\, u_i(x_i, \mathbf{x}_{-i})\right\} \tag{2.15}$$

Thus $\overline{\mathcal{B}}_i(\mathbf{x}_{-i})$ is inherently a single-valued map.

The *maximal best-response operator* is defined as $\overline{\mathcal{B}} : \mathcal{A} \to \mathcal{A}$, by

$$\overline{\mathcal{B}}(\mathbf{x}) = \big(\overline{\mathcal{B}}_1(\mathbf{x}_{-1}), \ldots, \overline{\mathcal{B}}_n(\mathbf{x}_{-n})\big)$$

Thus, the discrete-time synchronous maximal best response dynamic is the sequence $\mathbf{x}(t)$ with state space $\mathcal{A}$,

$$\mathbf{x}(t+1) = \overline{\mathcal{B}}(\mathbf{x}(t)), \qquad t \geq 0. \tag{2.16}$$

Equivalently,

$$x_i(t+1) = \overline{\mathcal{B}}_i(\mathbf{x}_{-i}(t)), \qquad t \geq 0.$$

wherein updates are executed simultaneously by every agent at each step, according to the maximal best response (2.15). It is noted that, because the choice is unique, this constitutes a deterministic dynamic.

It is observed that Assumption 1 constitutes a strictly weaker topological requirement than the formal definition of super-modularity (Definition 2.3.2).

While Assumption 1 merely demands the compactness of the action sets and the upper semi-continuity of the utility functions with respect to a player's own action, super-modularity necessitates these fundamental properties alongside continuity with respect to opponents' actions and the increasing differences property.
Consequently, the regularity conditions established in Assumption 1 are inherently satisfied by any super-modular game.

Given that super-modular network games represent the primary focus of this dissertation, this structural assumption is naturally subsumed and implicitly maintained throughout the remainder of the analysis.

The majority of the properties inherent to super-modular games are derived from the following proposition.

Proposition 2.4.1. *Let $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ be a super-modular game, then*

- *For every player $i \in \mathcal{V}$ and $\mathbf{x}_{-i}$, the best response set $\mathcal{B}_i(\mathbf{x}_{-i})$ is nonempty and has a maximum element $\overline{\mathcal{B}}_i(\mathbf{x}_{-i})$,* (2.15).
- *$\overline{\mathcal{B}}_i(\mathbf{x}_{-i})$ is monotone non-decreasing in $\mathbf{x}_{-i}$, i.e.*

$$\mathbf{x}'_{-i} \geq \mathbf{x}_{-i} \Longrightarrow \overline{\mathcal{B}}_i(\mathbf{x}'_{-i}) \geq \overline{\mathcal{B}}_i(\mathbf{x}_{-i}).$$

Proof.

- The non-emptiness of the best response set $\mathcal{B}_i(\mathbf{x}_{-i})$ and the existence of its maximum element $\overline{\mathcal{B}}_i(\mathbf{x}_{-i})$ are directly guaranteed by the regularity conditions established in Assumption 1. Specifically, since the action set $\mathcal{A}_i$ is compact and the utility function u_i is upper semi-continuous in x_i, the existence of a maximum is ensured by the extreme value theorem.

- The monotonicity of the operator $\overline{\mathcal{B}}_i(\mathbf{x}_{-i})$ is established by leveraging the increasing differences property inherent to super-modular functions. Let two strategy profiles of the opponents be considered, $\mathbf{x}'_{-i}$ and $\mathbf{x}_{-i}$, such that $\mathbf{x}'_{-i} \geq \mathbf{x}_{-i}$. Let $y_i = \overline{\mathcal{B}}_i(\mathbf{x}'_{-i})$ and $x_i = \overline{\mathcal{B}}_i(\mathbf{x}_{-i})$. It must be proven that $y_i \geq x_i$. It is assumed, for the sake of contradiction, that $y_i < x_i$.

  Since x_i is best response to $\mathbf{x}_{-i}$, the utility derived from playing x_i is at least as great as the utility derived from playing any other feasible action, including y_i. Thus, it is written:

  $$u_i\left(x_i, \mathbf{x}_{-i}\right) \geq u_i\left(y_i, \mathbf{x}_{-i}\right)$$

  which implies that the marginal difference is non-negative:

  $$u_i\left(x_i, \mathbf{x}_{-i}\right) - u_i\left(y_i, \mathbf{x}_{-i}\right) \geq 0$$

  By the property of increasing differences (Definition 2.12), since $x_i > y_i$, and $\mathbf{x}'_{-i} \geq \mathbf{x}_{-i}$, it follows that:

  $$u_i\left(x_i, \mathbf{x}'_{-i}\right) - u_i\left(y_i, \mathbf{x}'_{-i}\right) \geq u_i\left(x_i, \mathbf{x}_{-i}\right) - u_i\left(y_i, \mathbf{x}_{-i}\right)$$

  Substituting the non-negative difference from the best response condition, it is obtained that:

  $$u_i\left(x_i, \mathbf{x}'_{-i}\right) - u_i\left(y_i, \mathbf{x}'_{-i}\right) \geq 0$$

  which yields:

  $$u_i\left(x_i, \mathbf{x}'_{-i}\right) \geq u_i\left(y_i, \mathbf{x}'_{-i}\right).$$

  This result implies that, the action x_i is also an element of the best response set $\mathcal{B}_i(\mathbf{x}'_{-i})$. However, y_i is explicitly defined as the maximum element of $\mathcal{B}_i(\mathbf{x}'_{-i})$. The condition $x_i > y_i$, with $x_i \in \mathcal{B}_i(\mathbf{x}'_{-i})$ constitutes a direct contradiction to the maximality of y_i. Therefore, the initial assumption must be false, and it is concluded that $y_i \geq x_i$, which proves the monotonicity of the maximal best response operator. ■

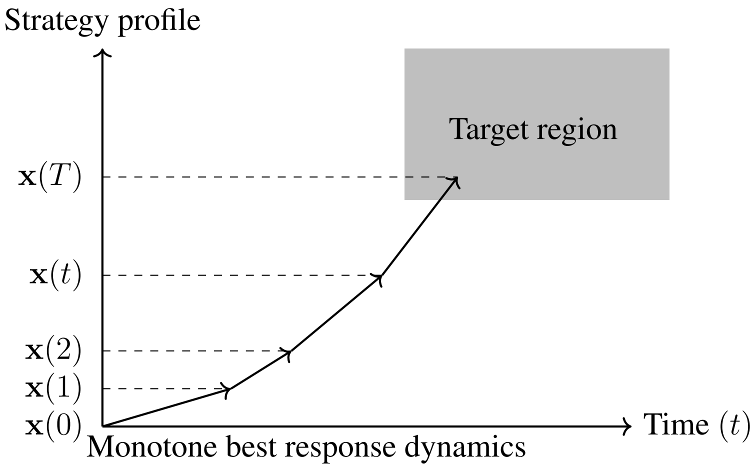

Figure 2.3: Best response dynamics in a super-modular network game. Strategic complementarities induce monotone trajectories that reach a Target region.

A fundamental question subsequently arises regarding the assumptions under which the dynamics (2.16), initiated from a specific state $\mathbf{x}_0 \in \mathcal{A}$, are guaranteed to converge to the set of Nash equilibria $\mathcal{A}^*$.

Best response dynamics are visualized in Figure 2.3, where strategies are iteratively updated by agents over time. Convergence toward a **target region** is achieved, as discussed previously with Figure 2.2, with arrows indicating the monotone trajectory of updates. In super-modular games, this convergence is guaranteed under monotone dynamics, and interventions can **accelerate convergence to finite time**.

Example 2.4.1 (Best response quadratic games)**.** *Consider a quadratic game* $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$*, with utility functions defined in equation* (2.7)*. The Best Response in* (2.8) *is unique, and therefore coincides with the maximal Best Response* (2.15)*. If the strategies are updated by the agents simultaneously to their current best response, the following dynamical system is obtained:*

$$\mathbf{x}(t+1) = \mathbf{b} + \delta W \mathbf{x}(t).$$

If $\delta\rho(W) < 1$*, the evolution* $\mathbf{x}(t)$ *converges to the unique Nash equilibrium* $\mathbf{x}^* = \mathbf{b} + \delta W \mathbf{x}^*$*. Notice that this is also an algorithm to compute the Nash equilibrium.*

Transition graphs

This subsection introduces several graphs in the configuration space of a game that play a key role in the analysis of learning systems.
Let $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ be strategic form games. For every player $i \in \mathcal{V}$, consider $\mathcal{A}_i$ to be a finite subset of $\mathbb{R}$ and denote by

$$m_i = \min\{\mathcal{A}_i\} \quad \text{and} \quad M_i = \max\{\mathcal{A}_i\}$$

the minimal and maximal actions of player i, respectively. Then, $\mathbf{m} = (m_1, \ldots, m_n)$ and $\mathbf{M} = (M_1, \ldots, M_n)$ are the least and the greatest strategy profiles.

Definition 2.4.1. *Consider a finite game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$. The following directed graphs are defined on the configuration space $\mathcal{A} = \prod_{i\in\mathcal{V}} \mathcal{A}_i$*

- *The hypercube graph is defined as $\mathcal{G}_h = (\mathcal{A}, \mathcal{E}_h)$ where*

$$\mathcal{E}_h = \{(\mathbf{x}, \mathbf{y}) \mid \exists\, i \in \mathcal{V}, \mathbf{x}_{-i} = \mathbf{y}_{-i},\ x_i \neq y_i\}$$

- *The Improvement graph (I-graph) is defined as $\mathcal{G}_I = (\mathcal{A}, \mathcal{E}_I)$ where*

$$\mathcal{E}_I = \{(\mathbf{x}, \mathbf{y}) \mid \exists\, i \in \mathcal{V}, \mathbf{x}_{-i} = \mathbf{y}_{-i},\ x_i \neq y_i,\ u_i(\mathbf{x}) \leq u_i(\mathbf{y})\}$$

- *The Best Response graph ($\mathcal{B}$-graph) is defined as $\mathcal{G}_{\mathcal{B}} = (\mathcal{A}, \mathcal{E}_{\mathcal{B}})$ where*

$$\mathcal{E}_{\mathcal{B}} = \{(\mathbf{x}, \mathbf{y}) \mid \exists\, i \in \mathcal{V}, \mathbf{x}_{-i} = \mathbf{y}_{-i},\ y_i \in \mathcal{B}_i(\mathbf{x}_{-i})\}$$

Remark 3. *Notice that the $\mathcal{G}_{\mathcal{B}} \subseteq \mathcal{G}_I \subseteq \mathcal{G}_h$. In the special case when all $\mathcal{A}_i$ are binary sets, the I-graph and the $\mathcal{B}$-graph coincide.*

Definition 2.4.2. *Let $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ be a finite game, then*

*(i) A **$\mathcal{B}$-path** from $\mathbf{x}$ to $\mathbf{z}$ is a path in $\mathcal{G}_{\mathcal{B}} = (\mathcal{A}, \mathcal{E}_{\mathcal{B}})$, i.e., a tuple of strategy profiles $\left(\mathbf{x}^{(k)}\right)_{0\leq k\leq l}$ such that $\mathbf{x}^{(0)} = \mathbf{x}$, $\mathbf{x}^{(l)} = \mathbf{z}$, and for every $0 \leq k < l$ there exists a player i_k in $\mathcal{V}$ such that*

$$x_{i_k}^{(k+1)} \in \mathcal{B}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right), \qquad \mathbf{x}_{-i_k}^{(k+1)} = \mathbf{x}_{-i_k}^{(k)};$$

*(ii) a **$\overline{\mathcal{B}}$-path** is a $\mathcal{B}$-path such that*

$$x_{i_k}^{(k+1)} \in \overline{\mathcal{B}}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right), \qquad 0 \leq k < l;$$

Remark 4. *Analytically, a Best Response path ($\mathcal{B}$-path) represents a sequence of strategic updates wherein, at each discrete step, a single agent unilaterally deviates to any strategy that is an element of its best-response set, conditional on the fixed strategies of the opponents. In contrast, a $\overline{\mathcal{B}}$-path imposes a strictly tighter algebraic constraint: the updating agent is required to systematically select the maximal element from the set of available best responses. In strategic environments characterized by complementarities, where best-response correspondences are frequently multi-valued, the $\overline{\mathcal{B}}$-path deterministically tracks the highest attainable trajectory of rational deviations, structurally guaranteeing monotonic convergence toward the maximal, extremal equilibrium.*

Proposition 2.4.2. *Consider a finite super-modular game. For a strategy profile* $\mathbf{x}$*, the following conditions are equivalent:*

(i) there exists a $\mathcal{B}$-path from $\mathbf{x}$ to $\mathbf{M}$;

(ii) there exists a $\overline{\mathcal{B}}$-path from $\mathbf{x}$ to $\mathbf{M}$.

Proof. For a strategy profile $\mathbf{x} \in \mathcal{A}$

(ii) $\Rightarrow$ (i) By definition

(i) $\Rightarrow$ (ii) Consider a $\mathcal{B}$-path $\left(\mathbf{x}^{(k)}\right)_{0\le k\le l}$ from $\mathbf{x}$ to $\mathbf{M}$ and construct a new path $\left(\mathbf{y}^{(k)}\right)_{0\le k\le l}$ with $\mathbf{y}^{(0)} = \mathbf{x}$ and $y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right)$, $\mathbf{y}_{-i_k}^{(k+1)} = \mathbf{y}_{-i_k}^{(k)}$, for every $0 \le k < l$. Clearly, $\left(\mathbf{y}^{(k)}\right)_{0\le k\le l}$ is a $\overline{\mathcal{B}}$-path, leaving only the requirement to demonstrate that $\mathbf{y}^{(l)} = \mathbf{M}$. It is proven by induction that

$$\mathbf{y}^{(k)} \ge \mathbf{x}^{(k)}, \qquad \forall 0 \le k \le l. \tag{2.17}$$

Clearly, $\mathbf{y}^{(0)} = \mathbf{x} = \mathbf{x}^{(0)}$. Moreover, if $\mathbf{y}^{(k)} \ge \mathbf{x}^{(k)}$ for some $0 \le k < l$, then, (2.4.1) and $x_{i_k}^{(k+1)} \in \mathcal{B}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right)$ imply that

$$y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right) \ge \overline{\mathcal{B}}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right) \ge x_{i_k}^{(k+1)},$$

so that $\mathbf{y}^{(k+1)} \ge \mathbf{x}^{(k+1)}$. This proves (2.17). For $k = l$, it follows that $\mathbf{M} \le \mathbf{x}^{(l)} \le \mathbf{y}^{(l)} \le \mathbf{M}$, thus completing the proof. ■

This yields an important consequence on the structure of the set of $\mathbf{M}$-attracted profiles, stated in the following result.

Corollary 2.4.3. *If $\mathbf{x}$ is $\mathbf{M}$-attracted, then every $\mathbf{y} \ge \mathbf{x}$ is $\mathbf{M}$-attracted.*

Proof. By virtue of Proposition 2.4.2, there exists a $\overline{\mathcal{B}}$-path $(\mathbf{x}^{(k)})_{0\le k\le l}$ from $\mathbf{x}$ to $\mathbf{M}$. Consider now the $\overline{\mathcal{B}}$-path

$$\mathbf{y}^{(0)} = \mathbf{y}, \qquad y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right), \quad \forall 0 \le k < l.$$

By induction, utilizing the monotonicity of $\overline{\mathcal{B}}$ as demonstrated previously, it is proven that $\mathbf{y}^{(k)} \ge \mathbf{x}^{(k)}$ for every k. Hence $\mathbf{y}^{(l)} = \mathbf{M}$. ■

The maximal strategy profile $\mathbf{M}$ is defined as being globally reachable from a $\mathcal{B}$-path if, for every initial configuration $\mathbf{x} \in \mathcal{A}$, there exists at least one directed sequence of best-response updates ($\mathcal{B}$-path) originating at $\mathbf{x}$ and terminating strictly at the maximal strategy profile $\mathbf{M}$.

Corollary 2.4.4. $\mathbf{M}$ *is globally reachable from a $\mathcal{B}$-path if and only if* $\mathbf{m}$ *is* $\mathbf{M}$-*attracted*

Proof. For a strategy profile $\mathbf{x} \in \mathcal{A}$

- $\Longrightarrow$ By definition
- $\Longleftarrow$ By Corollary 2.4.3 if $\mathbf{m}$ is $\mathbf{M}$-attracted then, every $\mathbf{x} \ge \mathbf{m}$ is $\mathbf{M}$-attracted. Since $\mathbf{m}$ is the minimal element, then every configuration reaches $\mathbf{M}$ by a $\mathcal{B}$-path. ■

Theorem 2.4.5. *Let* $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ *a finite super-modular game, then* $\mathbf{M}$ *is globally reachable from a $\mathcal{B}$-path if and only if the discrete-time synchronous maximal best response dynamic* (2.16) *converges to* $\mathbf{M}$*, with* $\mathbf{x}(0) = \mathbf{m}$.

Proof.

- $\Longrightarrow$ It is initially assumed that there exists a $\mathcal{B}$-path

$$\mathbf{m} = \mathbf{x}^{(0)} \le \mathbf{x}^{(1)} \le \cdots \le \mathbf{x}^{(l)} = \mathbf{M}$$

The synchronous sequence is defined as

$$\mathbf{y}(0) = \mathbf{m}, \qquad \mathbf{y}(t+1) = \overline{\mathcal{B}}(\mathbf{y}(t))$$

Base case:

$$\mathbf{y}(0) = \mathbf{m} = \mathbf{x}^{(0)}$$

By induction, given that synchronous updating applies all maximal best responses simultaneously, whereas asynchronous updating applies only to a single coordinate

$$\overline{\mathcal{B}}\left(\mathbf{x}^{(k)}\right) \geq \mathbf{x}^{(k+1)}$$

Since $\overline{\mathcal{B}}$ is a monotone non-decreasing map and $\mathbf{y}(k) \geq \mathbf{x}^{(k)}$

$$\mathbf{y}(k+1) = \overline{\mathcal{B}}(\mathbf{y}(k)) \geq \overline{\mathcal{B}}\left(\mathbf{x}^{(k)}\right) \geq \mathbf{x}^{(k+1)},$$

thus

$$\mathbf{y}(l) \geq \mathbf{x}^{(l)} = \mathbf{M}.$$

But $\mathbf{M}$ is the maximal element of the configuration space, hence

$$\mathbf{y}(l) = \mathbf{M}.$$

Therefore synchronous dynamics reaches $\mathbf{M}$.

$\Longleftarrow$ Conversely, it is assumed that

$$\mathbf{y}(0) = \mathbf{m}, \qquad \mathbf{y}(t+1) = \overline{\mathcal{B}}(\mathbf{y}(t)),$$

and there exists some T, such that

$$\mathbf{y}(T) = \mathbf{M}.$$

An intermediate time step is fixed $0 < t < T$; the configuration

$$\mathbf{y}(t+1) = \overline{\mathcal{B}}(\mathbf{y}(t))$$

differs from $\mathbf{y}(t)$ in some coordinates, $\emptyset \neq \mathcal{I}(t) \subseteq \mathcal{V}$.

A path is defined $\left(\mathbf{x}^{(k)}\right)_{0\leq k\leq l}$ with $\mathbf{x}^{(0)} = \mathbf{y}(0)$, wherein players in $\mathcal{I}(t)$ are updated sequentially, in any order, by interpolating the sequence of profiles from $\mathbf{y}(t)$ to $\mathbf{y}(t+1)$ with a subsequence $\mathbf{x}^{(k,s)}$ with $1 \leq s \leq |\mathcal{I}(t)|$ such that

$$x_{i_t}^{(t,s)} = y_i(t+1), \quad \mathbf{x}_{-i_t}^{(t,s)} = \mathbf{x}_{-i_t}^{(t)}, \quad \forall i_t \in \mathcal{I}(t)$$

infact, for any intermediate profile $\mathbf{z}$ such that

$$\mathbf{y}(t) \leq \mathbf{z} \leq \mathbf{y}(t+1),$$

then for any player i,

$$\overline{\mathcal{B}}_i(\mathbf{z}_{-i}) \leq \overline{\mathcal{B}}_i(\mathbf{y}_{-i}(t+1)) = y_i(t+1).$$

But since $y_i(t) \leq y_i(t+1)$, the maximal best response at the intermediate profile moves toward $y_i(t+1)$ ■

Example 2.4.2 (Best response Coordination games)**.** *The network coordination game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ provides a foundational framework for understanding contagion dynamics over networks.*
The best response dynamics of this game (Example 2.1.1) are of particular interest because they naturally yield the well-known linear threshold model *(Granovetter, 1978), where choosing action* 1 *represents the adoption of a new behavior, technology, or innovation. By equation* (2.4)*, the discrete-time synchronous maximal best response dynamic* (2.16)*, is given by*

$$x_i(t+1) = \begin{cases} \{1\} & \textit{if } \sum\limits_{j\in\mathcal{V}} W_{i,j}x_j(t) \geq r_i \\ \{0\} & \textit{if } \sum\limits_{j\in\mathcal{V}} W_{i,j}x_j(t) < r_i \end{cases}$$

This iterative process captures the evolution of the set of adopters $\mathcal{S}(t)$, over time as studied in Algorithm 1.

The theoretical equivalence between the synchronous maximal best-response dynamic and the previously formalized $\mathcal{B}$-path framework must be rigorously established. It is mathematically demonstrated that the synchronous update mechanism formulated herein is topologically isomorphic to a synchronous $\overline{\mathcal{B}}$-path evaluated over a directed configuration graph.

While the Permanent Activation Linear Threshold Model (PALTM) (detailed in Appendix B) is defined as an irreversible process, whereby an agent cannot revert to state 0 *once state* 1 *is adopted, the dynamic described herein allows for reversibility in principle. However, due to the inherent super-modularity of the coordination game, if the network is initialized at the minimal configuration $\mathbf{x}(0) = \mathbf{0}$ (or initialized via a predetermined seed set), the resulting $\overline{\mathcal{B}}$-path is strictly monotonically non-decreasing.*
Consequently, along this specific trajectory, state reversions do not occur, rendering the $\overline{\mathcal{B}}$-path functionally equivalent to the classical irreversible LTM. The LTM can thus be rigorously interpreted as the traversal of a directed $\overline{\mathcal{B}}$-path.

An agent i joins the set $\mathcal{S}(t+1)$ if and only if adoption is the optimal unilateral response to the current state $\mathcal{S}(t)$. The literature on diffusion and coordination games (e.g., Acemoglu et al., 2011), demonstrates that the limiting set of adopters, $\mathcal{S}^*$, depends critically on the structural properties of the network, the distribution of thresholds, and the influence of initially seeded agents.
Once reached, this limiting set represents a Nash equilibrium: a stable configuration where no individual agent has an incentive to deviate.

To systematically analyze these dynamics, it is useful to define the normalized threshold value for each agent $i \in \mathcal{V}$

$$\theta_i = \frac{r_i}{k_i} \in [0, 1].$$

In the complete graph, where the population is fully mixed, the evolution of the fraction of adopters, $Y(t) = \frac{1}{n}\sum_{i \in \mathcal{V}} x_i(t)$*, is governed by a one-dimensional non-linear recursion characterized by the threshold distribution function:*

$$F(\theta) = \frac{1}{n} |\{i \in \mathcal{V} \ : \ \theta_i \leq \theta\}|$$

Then, the sequence of adoptions $Y(t)$ *satisfies*

$$Y(t+1) = F(Y(t)), \quad t \geq 0$$

Since $F(\cdot)$ *is non-decreasing and such that* $F(1) = 1$*, standard dynamical system arguments imply that* $Y(t) \to 1$ *as* t *grows large if and only if*

$$F(Y) > Y, \quad \forall\, Y \in [0, 1)$$

It is observed that the threshold distribution function $F(\cdot)$ *is rightcontinuous and, for a finite population of* n *players as is the present case, it is piecewise constant on* $[0, 1]$*. Consequently, if a fixed point* $F(Y^*) = Y^*$ *exists in* $[0, 1)$*, it must hold that* $F(Y) = Y^* < Y$ *for* Y *in a right neighborhood of* Y^**. This macroscopic formulation elegantly bypasses the need to track a discrete state space that grows exponentially at* 2^n*, resulting in a dramatic reduction in complexity.*

However, this analytical simplicity is lost when the population interacts along an arbitrary, structured graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$*. To evaluate dynamics on complex networks, a standard approach is to assume homogeneous thresholds, assigning identical normalized thresholds* $\theta_i = \theta \in [0, 1]$*,* $\forall i \in \mathcal{V}$ *to all players. In this setting, the fixed points of the threshold model dynamics can be characterized through the graph-theoretic notion of cohesiveness introduced in Morris, 2000.*

Specifically, a subset of nodes $\mathcal{S} \subseteq \mathcal{V}$ *is defined as* θ*-cohesive if the total weight of out-links from any node* $i \in \mathcal{S}$ *towards other nodes within* $\mathcal{S}$ *constitutes at least a fraction* θ *of its total out-degree. Formally, for every* $i \in \mathcal{S}$*,*

$$\frac{1}{k_i} \sum_{j \in \mathcal{S}} W_{i,j} \geq \theta, \quad \forall i \in \mathcal{S}. \tag{2.18}$$

By definition, a θ-cohesive set is also θ'-cohesive for any $\theta' \leq \theta$. In a simple unweighted graph, this implies that every node in $\mathcal{S}$ has at least a proportion θ of its neighbors inside $\mathcal{S}$.

This structural property maps directly to the equilibrium states of the game. A specific configuration of the network, expressed as the vector $x = \mathbf{1}_{\mathcal{S}} - \mathbf{1}_{\mathcal{V}\setminus\mathcal{S}}$ forms a Nash equilibrium if and only if the adopting set $\mathcal{S}$ is θ-cohesive and the non-adopting set $\mathcal{V} \setminus \mathcal{S}$ is $(1-\theta)$-cohesive.

Furthermore, cohesiveness is essential for identifying optimal seeding strategies—the initial nodes that can trigger a global cascade. This requires a stricter boundary condition: a subset is uniformly no more than θ-cohesive if none of its subsets are θ'-cohesive for any $\theta' > \theta$. Building on this, an initial set of adopters $\mathcal{S} \subseteq \mathcal{V}$ functions as a sufficient control set (guaranteeing global adoption) if and only if the remaining network $\mathcal{V} \setminus \mathcal{S}$ is uniformly no more than $(1-\theta)$-cohesive.

While translating the linear threshold model into the language of cohesiveness elegantly links network topology to game-theoretic equilibria, it introduces significant computational hurdles. Moving away from the macroscopic approximation of the complete graph sacrifices analytical tractability. Verifying whether the non-adopting set is uniformly no more than $(1-\theta)$-cohesive iequires evaluating every possible subset of $\mathcal{V} \setminus \mathcal{S}$. Even in the simple (possibly infinite) network, finding θ-cohesive sets with $(1-\theta)$-cohesive complements remains a computationally hard problem, highlighting the inherent trade-off between topological realism and algorithmic efficiency in diffusion models.

The complete characterization of reachable configurations and stationary distributions within fully connected networks under heterogeneous thresholds is provided by Ramazi and Cao, 2020. The structural analysis presented in Example 2.4.2 serves to connect the abstract topological concepts of directed graph paths directly to concrete diffusion processes. Specifically, the deterministic progression of the threshold model corresponds precisely to the synchronous execution of $\overline{\mathcal{B}}$-path over the hypercube of configurations. This demonstrates that topological features like cohesiveness act as the underlying structural conditions that govern the existence and convergence of paths within the graph $\mathcal{G}_{\mathcal{B}}$.

Chapter 3

Optimal Intervention in Non-Binary Super-Modular Games

Within multi-player game theory, the design of optimal intervention policies capable of steering a system toward a desirable configuration remains a fundamental problem. Super-modular games are typically endowed with multiple Nash equilibria that admit a Pareto ordering; consequently, determining the minimal effort required to transition the system from a suboptimal to an optimal equilibrium constitutes a natural and relevant problem across these applied contexts.

In this chapter, intervention problems within finite super-modular games are investigated, with a specific focus directed toward mechanisms designed to drive players to the highest Nash equilibrium. A related contribution is found in Galeotti et al., 2020, wherein an intervention problem concerning marginal individual utilities is considered for quadratic games with continuous action sets, aiming to maximize social welfare subject to a bounded cost. Although these proposed results are also applicable to non super-modular games (e.g., public good games), the mechanism of their intervention is quite specific and the corresponding solution relies on the explicit expression of the Nash equilibrium, which exhibits a linear dependence on marginal utilities.

The intervention on a finite super-modular game is modeled as the imposition of a lower bound on the actions selected by the participating players. The optimization problem of finding the minimum cost intervention is considered (for general cost functions) for which, in such restricted game, there exists a best response path (2.4.2) that leads to the maximal configuration in the action profile.

This model is a natural generalization of the seeding problem considered in the binary case (Como et al., 2021; Cordasco et al., 2015; Günneç et al., 2020b) thereby demonstrating that it can be interpreted within economic applications as an intervention offering a fixed incentive (granted if a player maintains its effort above some prescribed level).

3.1 Finite Super-modular Unimodal Games

Given a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, finite strategic form network games $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ with player set $\mathcal{V}$ are considered. An action x_i is chosen by each player $i \in \mathcal{V}$ from a non-empty finite set $\mathcal{A}_i \subseteq \mathbb{R}$; consequently, each player's set of actions possesses minimal and maximal actions denoted

$$m_i = \min \mathcal{A}_i \quad \text{and} \quad M_i = \max \mathcal{A}_i$$

Then, $\mathbf{m} = (m_1, m_2, \dots, m_n)$ and $\mathbf{M} = (M_1, M_2, \dots, M_n)$ are respectively the least and the greatest strategy profiles. For every non-maximal action $x_i < M_i$ (respectively, non-minimal action $x_i > m_i$), the following are defined

$$x_i^+ = \min\{y \in \mathcal{A}_i \ : \ y > x_i\}, \quad x_i^- = \max\{y \in \mathcal{A}_i \ : \ y < x_i\}$$

Every player i is endowed with a utility function $u_i(\mathbf{x}) \ : \mathcal{A} \to \mathbb{R}$ satisfying the increasing difference property (2.12), hence the game satisfies super-modularity 2.3.2. The analysis is subsequently restricted to finite super-modular games that satisfy the following additional property.

Definition 3.1.1 (Unimodal Games). *A game is* unimodal *if the function $x_i \mapsto u_i(x_i, \mathbf{x}_{-i})$ is unimodal for every player i and strategy profile $\mathbf{x}_{-i}$, i.e., if there exists $r_i(\mathbf{x}_{-i})$ in $\mathbb{R}$ such that $x_i \mapsto u_i(x_i, \mathbf{x}_{-i})$ is non-decreasing for $x_i \leq r_i(\mathbf{x}_{-i})$ and non-increasing for $x_i \geq r_i(\mathbf{x}_{-i})$.*

Example 3.1.1. *For a nonnegative square matrix $W = (W_{i,j})$ in $\mathbb{R}_+^{n\times n}$ and a parameters vector $\mathbf{b} = (b_i)_{i=1,\dots,n}$ in $\mathbb{R}^n$, consider the game where every player $i = 1, \dots, n$ has nonempty finite action set $\mathcal{A}_i \subseteq \mathbb{R}$ and utility*

$$u_i(x_i, \mathbf{x}_{-i}) = b_i x_i - \frac{1}{2}x_i^2 + x_i \sum_{j\neq i} W_{i,j} x_j.$$

Such games can be interpreted as quantized versions of the quadratic games (2.1.2) *with strategic complements.*

Super-modularity is satisfied by this game, given that, for every two actions x_i and y_i in $\mathcal{A}_i$ such that $z_i = x_i - y_i \geq 0$,

$$u_i(x_i, \mathbf{x}_{-i}) - u_i(y_i, \mathbf{x}_{-i}) = b_i z_i - \frac{x_i^2 - y_i^2}{2} + z_i \sum_{j \in \mathcal{V}} W_{i,j} x_j$$

is a nondecreasing function of the strategy profile $\mathbf{x}_{-i}$. This game is also unimodal, since the map $x_i \mapsto u_i(x_i, \mathbf{x}_{-i}) = r_i x_i - x_i^2/2$, where $r_i = b_i + \sum_{j \in \mathcal{V}} W_{i,j} x_i$, is non-decreasing for $x_i \leq r_i$ and non-increasing for $x_i \geq r_i$.

Intervention Optimization

Interventions on finite super-modular unimodal games are formulated as restrictions imposed on the respective action sets. Specifically, for a vector $\mathbf{h}$ in $\mathcal{H} = \mathcal{A}$, referred to as an intervention vector, the restricted game $\Gamma(\mathbf{h})$ is considered, wherein the action set for player i is restricted to

$$\mathcal{A}_i(\mathbf{h}) = \{x_i \in \mathcal{A}_i \mid x_i \geq h_i\}.$$

Hence, in particular, $\Gamma(\mathbf{m}) = \Gamma$ coincides with the original (unrestricted) game. It is observed that the original game Γ is super-modular and unimodal if and only if the restricted game $\Gamma(\mathbf{h})$ maintains these properties for every intervention vector $\mathbf{h}$ in $\mathcal{A}$. A separable cost is associated with every intervention, defined as:

$$C(\mathbf{h}) = \sum_{i \in \mathcal{V}} \gamma_i(h_i), \tag{3.1}$$

where, for every player i in $\mathcal{V}$, $\gamma_i : \mathcal{A}_i \to \mathbb{R}_+$ is a non-decreasing cost function such that $\gamma_i(m_i) = 0$.

The assumption of a separable cost structure, as formulated in Equation 3.1, is required for the subsequent analysis of the minimization problem. Specifically, separability ensures that the variation in the global cost resulting from a unilateral deviation by a single agent depends solely on that agent's individual cost function. This property is strictly necessary for the definition of the transition probabilities of the time-reversible Markov chain introduced in Section 3.3, as it allows the detailed balance equations to be satisfied using purely local cost differentials.

The objective is subsequently established to determine a minimum-cost intervention vector $\mathbf{h}$ such that the maximal profile $\mathbf{M}$ is globally reachable via a $\mathcal{B}$-path within the restricted game $\Gamma(\mathbf{h})$, as formalized below.

Definition 3.1.2. *For an intervention vector* $\mathbf{h}$ *in* $\mathcal{H}$*:*

(i) A strategy profile $\mathbf{x}$ *in* $\mathcal{A}(\mathbf{h})$ *is* $\mathbf{M}$-attracted *in* $\Gamma(\mathbf{h})$ *if there exists a* $\mathcal{B}$*-path in* $\mathcal{G}_{\mathcal{B}} = (\mathcal{A}(\mathbf{h}); \mathcal{E}_{\mathcal{B}})$ *from* $\mathbf{x}$ *to* $\mathbf{M}$*;*

(ii) An intervention $\mathbf{h}$ *is* sufficient *if every* $\mathbf{x}$ *in* $\mathcal{A}(\mathbf{h})$ *is* $\mathbf{M}$*-attracted.*

The optimal intervention problem is subsequently formulated as follows:

$$\underset{\mathbf{h} \in \mathcal{O}}{\text{minimize}} \quad C(\mathbf{h}) = \sum_{i \in \mathcal{V}} \gamma_i(h_i) \tag{3.2}$$

where $\mathcal{O}$ denotes the set of sufficient interventions, and let

$$\mathcal{O}^* = \underset{\mathbf{h} \in \mathcal{O}}{\operatorname{argmin}}\, C(\mathbf{h})$$

denote the set of *minimum cost sufficient interventions*.

In the context of finite super-modular unimodal games, the abstract target class of games Γ^*, initially introduced in Section 2.2, is explicitly characterized by the structural property of $\mathbf{M}$-attractivity. Specifically, Γ^* is defined as the collection of restricted games $\Gamma(\mathbf{h})$ in which the maximal strategy profile $\mathbf{M}$ is globally reachable via a Best-Response path ($\mathcal{B}$-path) from any feasible configuration $\mathbf{x} \in \mathcal{A}$. Consequently, the intervention problem formulated in Problem 3.2 is mathematically equivalent to identifying a parameter vector $\mathbf{h}$ of minimum cost such that the induced game strictly belongs to Γ^*, thereby ensuring that the decentralized system is inexorably steered toward the highest Nash equilibrium.

Proposition 3.1.1. *Assume an asynchronous, discrete-time best-response dynamic, where at each time step a single agent, selected uniformly at random, updates its strategy to a best response within the restricted game* $\Gamma(\mathbf{h})$ *(resolving ties uniformly at random). If an intervention* $\mathbf{h}$ *is sufficient, almost sure convergence to the maximal profile* $\mathbf{M}$ *is guaranteed in finite time from any initial profile* $\mathbf{x} \in \mathcal{A}(\mathbf{h})$.

Proof. Under the assumption of a strictly sufficient intervention $\mathbf{h}$, it is observed that the restricted configuration space $\mathcal{A}(\mathbf{h})$ contains no absorbing states other than the maximal profile $\mathbf{M}$. Due to the finite nature of the discrete state space, the asynchronous dynamic is modeled as a Markov process exhibiting strictly positive transition probabilities along the established $\mathcal{B}$-paths directed toward $\mathbf{M}$.

Consequently, because the dynamic follows a strictly non-decreasing probability sequence along these directed paths, it is mathematically guaranteed that the system reaches the absorbing state $\mathbf{M}$ in finite time with probability one. ■

Example 3.1.2 (Continued). *Let the network coordination game defined in Example 3.1.1 be considered, under the assumption that the minimal action for all agents is exactly $m_i = 0$, for all $i \in \mathcal{V}$.*
Given an intervention $\mathbf{h} \in \mathcal{A}$, it is assumed that a feasible intervention consists of granting each player i an extra utility $\lambda_i \geq 0$ whenever $x_i \geq h_i$. Let a game

$$\tilde{\Gamma}(\lambda, \mathbf{h}) = \left(\mathcal{V}, \{\mathcal{A}_i(\lambda, \mathbf{h})\}_{i\in\mathcal{V}}, \left\{u_i^{\lambda,\mathbf{h}}\right\}_{i\in\mathcal{V}}\right)$$

be considered, with no restriction on the action sets, i.e., $\mathcal{A}_i(\lambda, \mathbf{h}) = \mathcal{A}_i$, for all $i \in \mathcal{V}$ and utility function

$$u_i^{\lambda,\mathbf{h}}(x_i, \mathbf{x}_{-i}) = u_i(\mathbf{x}) + \lambda_i \mathbf{1}_{[h_i,M_i]}(x_i).$$

The objective is subsequently established as finding a pair of vectors $(\lambda, \mathbf{h})$ for which every configuration $\mathbf{x}$ is $\mathbf{M}$-attracted and the cost associated with the intervention, $\sum\limits_{i\in\mathcal{V}} \lambda_i$ is minimized. This problem can be formulated as (3.2).
Initially, the minimum value of $\{\lambda_i\}_{i\in\mathcal{V}}$, for which

$$u_i^{\lambda,\mathbf{h}}(h_i, \mathbf{x}_{-i}) \geq u_i^{\lambda,\mathbf{h}}(x_i, \mathbf{x}_{-i}),$$

holds for every $x_i \leq h_i$ and every $\mathbf{x}_{-i}$ must be determined.

$$\begin{aligned}
u_i^{\lambda,\mathbf{h}}(h_i, \mathbf{x}_{-i}) \geq u_i^{\lambda,\mathbf{h}}(x_i, \mathbf{x}_{-i}) &\iff u_i^{\lambda,\mathbf{h}}(h_i, \mathbf{x}_{-i}) - u_i^{\lambda,\mathbf{h}}(x_i, \mathbf{x}_{-i}) \geq 0 \\
b_i h_i - \frac{1}{2}h_i^2 + h_i \sum_{j\in\mathcal{V}} W_{i,j}x_j + \lambda_i - b_i x_i &+ \frac{1}{2}x_i^2 - x_i \sum_{j\in\mathcal{V}} W_{i,j}x_j \geq 0 \\
\lambda_i \geq \frac{1}{2}(h_i^2 - 2b_i h_i) - \frac{1}{2}(x_i^2 &- 2b_i x_i) - (h_i - x_i) \sum_{j\in\mathcal{V}} W_{i,j}x_j.
\end{aligned}$$

Given that $\sum_{j\in\mathcal{V}} W_{i,j}x_j \geq 0$, the term $\sum_{j\in\mathcal{V}} W_{i,j}x_j$ strictly decreases the overall expression, as it is multiplied by the positive factor $h_i - x_i$ (since $x_i \leq h_i$). Hence, the maximum possible gain from deviating occurs when $\sum_{j\in\mathcal{V}} W_{i,j}x_j = 0$.

The condition $\sum_{j\in\mathcal{V}} W_{i,j}x_j = 0$ is evaluated strictly to isolate the minimum incentive λ_i required to render $x_i \geq h_i$ a dominant strategy. By calibrating the incentive magnitude against this zero-peer-support lower bound, strategic compliance is mathematically guaranteed across all feasible network configurations.

Consequently, it follows that $\lambda_i \geq \frac{1}{2}(h_i^2 - 2b_i h_i) - \frac{1}{2}(x_i^2 - 2b_i x_i)$.

By completing the square, it is obtained that

$$\lambda_i \geq \frac{1}{2}\left((h_i - b_i)^2 - b_i^2\right) - \frac{1}{2}\left((x_i - b_i)^2 - b_i^2\right) = \frac{1}{2}\left((h_i - b_i)^2 - (x_i - b_i)^2\right).$$

Since $(x_i - b_i)^2 \geq 0$ *for all* x_i*,*

$$\lambda_i \geq \frac{1}{2}\left(h_i - b_i\right)^2$$

Equality is achieved when $x_i = b_i$ *(which is feasible whenever* $b_i \leq h_i$*).*
If $b_i > h_i$*, the maximum gain is* 0*. Therefore,*

$$\lambda_i = \gamma_i(h_i) = \frac{1}{2}[h_i - b_i]_+^2$$

The notation $[\cdot]_+$ *is employed to denote the standard positive-part operator. Within this framework, the operator ensures that only strictly non-negative incentive values are allocated by the central regulator. Let*

$$\Gamma(\mathbf{h}) = \left(\mathcal{V}, \{\mathcal{A}_i(\mathbf{h})\}_{i\in\mathcal{V}}, \{u_i^{\mathbf{h}}\}_{i\in\mathcal{V}}\right)$$

be defined as the game with restricted action sets previously introduced.

A comparison is established between $\tilde{\Gamma}(\lambda, \mathbf{h})$ *and the game* $\Gamma(\mathbf{h})$*. It is noted that, for configurations* $\mathbf{x} \geq \mathbf{h}$*, the two games are equivalent, in the sense that best response sets are always equal. This implies that if* $\mathbf{h}$ *is a sufficient intervention vector for* $\Gamma(\mathbf{h})$*, then every configuration* $\mathbf{x}$ *is* $\mathbf{M}$*-attracted also for* $\tilde{\Gamma}(\lambda, \mathbf{h})$*. Indeed, by construction, starting from any initial configuration there exists a* $\mathcal{B}$*-path in* $\tilde{\Gamma}(\lambda, \mathbf{h})$ *leading to a configuration* $\mathbf{x} \geq \mathbf{h}$ *and from* $\mathbf{x}$ *the same* $\mathcal{B}$*-path that leads to* $\mathbf{M}$ *in* $\Gamma(\mathbf{h})$ *also leads to* $\mathbf{M}$ *in* $\tilde{\Gamma}(\lambda, \mathbf{h})$*.*

Conversely, it is assumed that $\mathbf{h}$ *is such that every configuration* $\mathbf{x}$ *is* $\mathbf{M}$*-attracted in* $\tilde{\Gamma}(\lambda, \mathbf{h})$*. If the process is initiated from* $\mathbf{x} \geq \mathbf{h}$*, any* $\mathcal{B}$*-path in* $\tilde{\Gamma}(\lambda, \mathbf{h})$*, from* $\mathbf{x}$ *to* $\mathbf{M}$*, will pass through profiles* $\mathbf{x}^{(k)} \geq \mathbf{h}$*, hence it will also be a* $\mathcal{B}$*-path in* $\Gamma(\mathbf{h})$*. This shows that the optimal intervention problem for this example is equivalent to the minimization problem* (3.2) *with a cost function*

$$C(\mathbf{h}) = \sum_{i\in\mathcal{V}} \gamma_i(h_i) = \sum_{i\in\mathcal{V}} \frac{1}{2}[h_i - b_i]_+^2 \tag{3.3}$$

3.2 Equilibrium Characterization

In this section, an equivalent description of the optimal intervention problem is proposed, allowing for the design of a low-complexity algorithm for its solution. A distinct type of path is defined below, wherein modifications are constrained to be monotonic and of a minimum step, requiring only that utilities do not decrease.

Definition 3.2.1. *For a finite game* $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ *a* weak Improvement path (wI-path) *from the strategy profile* $\mathbf{x}$ *to* $\mathbf{z}$ *is a path in the improvement graph* $\mathcal{G}_I = (\mathcal{A}; \mathcal{E}_I)$*, i.e., a tuple of strategy profiles* $\left(\mathbf{x}^{(k)}\right)_{0\le k\le l}$ *such that* $\mathbf{x}^{(0)} = \mathbf{x}$, $\mathbf{x}^{(l)} = \mathbf{z}$ *and, for every* $0 \le k < l$ *there exists* i_k *in* $\mathcal{V}$ *such that* $\mathbf{x}^{(k+1)}_{-i_k} = \mathbf{x}^{(k)}_{-i_k}$ *and*

$$x^{(k+1)}_{i_k} = \left(x^{(k)}_{i_k}\right)^+, \qquad u_{i_k}\left(\mathbf{x}^{(k+1)}\right) \ge u_{i_k}\left(\mathbf{x}^{(k)}\right). \tag{3.4}$$

A weak improvement path characterizes a sequence of unilateral strategy updates wherein, at each step, a single player increases their action by exactly one atomic increment (to the immediate next available action), provided that such an increment does not result in a strictly lower utility for that player. This mechanism introduces a minimum-step constraint on the action variations, distinguishing it from general best-response paths by strictly limiting transitions to adjacent strategies within the lattice.
It is observed that, unlike $\mathcal{B}$-paths, wI-paths originating from any $\mathbf{x} \ge \mathbf{h}$ remain invariant in the restricted game $\Gamma(\mathbf{h})$ for every intervention vector $\mathbf{h}$.

The following technical result builds upon the unimodality assumption and is instrumental for subsequent derivations.

Lemma 2. *Let a finite super-modular unimodal game and a strategy profile* $\mathbf{x}$ *in* $\mathcal{A}$ *be considered. If there exists a player* i *such that* $x_i < M_i$ *and* $u_i(x_i, \mathbf{x}_{-i}) \le u_i(x_i^+, \mathbf{x}_{-i})$*, then*

$$\overline{\mathcal{B}}_i(\mathbf{x}_{-i}) \ge x_i^+.$$

Proof. Given a strategy profile $\mathbf{x}_{-i}$, since $u_i(x_i, \mathbf{x}_{-i}) \le u_i(x_i^+, \mathbf{x}_{-i})$, the unimodality assumption (Definition (3.1.1)) implies

$$x_i < x_i^+ \le r_i$$

i.e., $x \mapsto u_i(x, \mathbf{x}_{-i})$ is necessarily non-decreasing on the set $\{x \in \mathcal{A}_i \,:\, x \le x_i^+\}$. This implies that the maximal best response must necessarily belong to the set $\{x \in \mathcal{A}_i \,:\, x \ge x_i^+\}$. ■

The following result shows that M-attractivity can be equivalently studied through weak Improvement paths.

Proposition 3.2.1. *For a finite super-modular unimodal game and a strategy profile* $\mathbf{x}$*, the following are equivalent:*

(i) there exists a $\mathcal{B}$-path from $\mathbf{x}$ to M;

(ii) there exists a wI-path from $\mathbf{x}$ to M.

Proof. For a strategy profile $\mathbf{x} \in \mathcal{A}$

(i) $\Rightarrow$ (ii) Assuming the existence of a $\mathcal{B}$-path, due to the super-modularity assumption, Proposition 2.4.2 guarantees the existence of a $\overline{\mathcal{B}}$-path $\left(\mathbf{y}^{(k)}\right)_{0\leq k\leq l}$ from $\mathbf{x}$ to M. A new path $\left(\mathbf{x}^{(k)}\right)_{0\leq k\leq l}$ is defined, with $\mathbf{x}^{(0)} = \mathbf{y}^{(0)}$ and

$$x_{i_k}^{(k+1)} = \max\left\{y_{i_k}^{(k+1)},\ x_{i_k}^{(k)}\right\}, \qquad \mathbf{x}_{-i_k}^{(k+1)} = \mathbf{x}_{-i_k}^{(k)}, \tag{3.5}$$

where i_k denotes the player moving according to $\overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right)$ at time instant k and $\mathbf{y}_{-i_k}^{(k+1)} = \mathbf{y}_{-i_k}^{(k)}$, for $0 \leq k < l$. It should be noted that $\mathbf{x}^{(k)} \geq \mathbf{y}^{(k)}$, for $0 \leq k \leq l$. Furthermore, it is observed that

$$u_{i_k}\left(x_{i_k}^{(k+1)}, \mathbf{y}_{-i_k}^{(k)}\right) \geq u_{i_k}\left(x_{i_k}^{(k)}, \mathbf{y}_{-i_k}^{(k)}\right), \tag{3.6}$$

indeed,

- if $x_{i_k}^{(k+1)} = x_{i_k}^{(k)}$, (3.6) trivially holds true as an equality;
- if $x_{i_k}^{(k+1)} = y_{i_k}^{(k+1)}$, (3.6) holds true since $y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right)$ is a best response to $\mathbf{y}_{-i_k}^{(k)}$.

Equation (3.6) together with $\mathbf{x}^{(k)} \geq \mathbf{y}^{(k)}$ and super-modularity imply

$$\begin{aligned} 0 &\leq u_{i_k}\left(x_{i_k}^{(k+1)}, \mathbf{y}_{-i_k}^{(k)}\right) - u_{i_k}\left(x_{i_k}^{(k)}, \mathbf{y}_{-i_k}^{(k)}\right) \\ &\leq u_{i_k}\left(x_{i_k}^{(k+1)}, \mathbf{x}_{-i_k}^{(k)}\right) - u_{i_k}\left(x_{i_k}^{(k)}, \mathbf{x}_{-i_k}^{(k)}\right), \end{aligned}$$

for $0 \leq k < l$.

This shows that $\left(\mathbf{x}^{(k)}\right)_{0\leq k\leq l}$ satisfies all the requirements for a weak improvement path, except for possibly the first equation in (3.4), as it could still be the case that $x_{i_k}^{(k+1)} > \left(x_{i_k}^{(k)}\right)^+$.

In such cases, the intermediate actions $x_1 = \left(x_{i_k}^{(k)}\right)^+ < x_2 = x_1^+ < \cdots < x_r = x_{i_k}^{(k+1)}$, are considered, interpolating the sequence of profiles from $\mathbf{x}^{(k)}$ to $\mathbf{x}^{(k+1)}$ via a subsequence $\mathbf{x}^{(k,s)}$ (for $1 \leq s \leq r$) such that $x_{i_k}^{(k,s)} = x_s$ and $\mathbf{x}_{-i_k}^{k,s} = \mathbf{x}_{-i_k}^{(k)}$ for all s. It should be noted that, by the unimodality assumption, necessarily $x \mapsto u_{i_k}\left(x, \mathbf{x}_{-i_k}^{(k)}\right)$ is necessarily non-decreasing in the interval $\left[x_{i_k}^{(k)}, x_{i_k}^{(k+1)}\right]$; namely $u_{i_k}(\mathbf{x}^{k,s+1}) \geq u_{i_k}(\mathbf{x}^{k,s})$ for every $1 \leq s < r$. If this interpolation is systematically carried out for every k where $x_{i_k}^{(k+1)} > \left(x_{i_k}^{(k)}\right)^+$, a wI-path leading from $\mathbf{x}$ to $\mathbf{M}$ is ultimately constructed.

(ii) $\Rightarrow$ (i) Let a wI-path $(\mathbf{x}^{(k)})_{0\leq k\leq l}$ from $\mathbf{x}$ to $\mathbf{M}$ be considered, and let a path be defined as $(\mathbf{y}^{(k)})_{0\leq k\leq l}$ with $\mathbf{y}^{(0)} = \mathbf{x}^{(0)}$ and

$$y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right), \qquad \mathbf{y}_{-i_k}^{(k+1)} = \mathbf{y}_{-i_k}^{(k)}, \qquad \forall 0 \leq k < l.$$

By construction, $(\mathbf{y}^{(k)})_{0\leq k\leq l}$ constitutes a $\overline{\mathcal{B}}$-path leaving only the requirement to prove that $\mathbf{y}^{(l)} = \mathbf{M}$. This follows directly from demonstrating that $\mathbf{y}^{(k)} \geq \mathbf{x}^{(k)}$ for every k. This condition is proven by induction, noting first its trivial validity for $k = 0$. Assuming its validity for some arbitrary k, the monotonicity of $\overline{\mathcal{B}}_{i_k}$ implies that

$$y_{i_k}^{(k+1)} = \overline{\mathcal{B}}_{i_k}\left(\mathbf{y}_{-i_k}^{(k)}\right) \geq \overline{\mathcal{B}}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right) \tag{3.7}$$

Since $x_{i_k}^{(k+1)} = \left(x_{i_k}^{(k)}\right)^+$ and $u_{i_k}(\mathbf{x}^{(k+1)}) \geq u_{i_k}(\mathbf{x}^{(k)})$, Lemma 2 implies that $\overline{\mathcal{B}}_{i_k}\left(\mathbf{x}_{-i_k}^{(k)}\right) \geq x_{i_k}^{(k+1)}$. This and (3.7) imply that $\mathbf{y}^{(k+1)} \geq \mathbf{x}^{(k+1)}$, thus proving the claim. ■

The following conclusive result provides a simpler characterization of the set of sufficient interventions, establishing the theoretical basis for the algorithm proposed in the subsequent section.

Corollary 3.2.2. *For a finite super-modular unimodal game Γ, the set of sufficient interventions can be characterized as*

$$\mathcal{O} = \{\mathbf{h} \in \mathcal{A} \mid \exists \textit{ a wI-path from } \mathbf{h} \textit{ to } \mathbf{M}\} \tag{3.8}$$

Proof. To prove that the two sets are equal, a double inclusion is demonstrated.

($\supseteq$) If there exists a wI-path from $\mathbf{h}$ to $\mathbf{M}$, then, by Proposition 3.2.1, there exists a $\mathcal{B}$-path from $\mathbf{h}$ to $\mathbf{M}$ in $\Gamma(\mathbf{h})$. From Corollary 2.4.3, it is deduced that all profiles $\mathbf{x} \in \mathcal{A}(\mathbf{h})$ are $\mathbf{M}$-attracted in $\Gamma(\mathbf{h})$, so that $\mathbf{h} \in \mathcal{O}$.

($\subseteq$) If $\mathbf{h} \in \mathcal{O}$, then there exists a $\mathcal{B}$-path from $\mathbf{h}$ to $\mathbf{M}$ in $\Gamma(\mathbf{h})$. Then, by Proposition 3.2.1, there exists also a wI-path from $\mathbf{h}$ to $\mathbf{M}$. ■

The following example demonstrates that unimodality is necessary for the above result to hold true.

Example 3.2.1. *Consider the* 2*-player* 3*-action game with* $\mathcal{A}_i = \{-1, 0, 1\}$ *for* $i = 1, 2$ *and utilities*

$$u_i(x_i, x_{-i}) = -x_i + 2x_i^2 + x_i x_{-i}, \quad i = 1, 2$$

A direct check shows that the two possible $\mathcal{B}$*-paths are* $(0,0) \to (-1,0) \to (-1,-1)$ *and* $(0,0) \to (0,-1) \to (-1,-1)$*, showing that* $(0,0)$ *is not* $(1,1)$ *attracted. However,* $(0,0) \to (1,0) \to (1,1)$ *is a wI-path leading to* $(1,1)$*.*

3.3 A Markovian Algorithmic Approach

In this section, a simple, provably convergent iterative algorithm is proposed to compute sufficient interventions of minimum cost, based on the characterization of $\mathcal{O}$ in Corollary 3.2.2. This indicates that minimum sufficient interventions can be identified by initiating the search from the maximal strategy profile $\mathbf{M}$ and iteratively tracing a wI-path backwards.
Assuming a fixed finite super-modular unimodal game and non-decreasing cost functions $\gamma_i : \mathcal{A}_i \to \mathbb{R}$, the quantity $d_i(x_i) = \gamma_i\left(x_i^+\right) - \gamma_i(x_i)$ is defined for $x_i < M_i$, with $d_i(M_i) = 0$.

A family of discrete-time Markov chains $(Z_t^\epsilon)_{t\geq 0}$ is introduced on the strategy profile space $\mathcal{A}$, parameterized by a scalar ϵ in $[0, 1]$. It is subsequently proven that, for $0 < \epsilon \leq 1$, the Markov chain $(Z_t^\epsilon)_{t\geq 0}$ is time-reversible and that, as ϵ vanishes, its stationary distribution concentrates on the family of minimum cost sufficient interventions. The dynamics of the Markov chain Z_t^ϵ are described as follows: at every time $t = 0, 1...$, given that $Z_t^\epsilon = \mathbf{x}$, a player i is selected uniformly at random from the entire player set $\mathcal{V}$, and a movement direction (up or down) is selected with probability $1/2$. If the downward direction is chosen, then

- if $u_i\left(x_i^-, \mathbf{x}_{-i}\right) > u_i\left(x_i, \mathbf{x}_{-i}\right)$, then
$$Z_{t+1}^\epsilon = Z_t^\epsilon;$$
- if $u_i\left(x_i^-, \mathbf{x}_{-i}\right) \leq u_i\left(x_i, \mathbf{x}_{-i}\right)$, then
$$\left(Z_{t+1}^\epsilon\right)_i = \left(Z_t^\epsilon\right)_i^-, \qquad \left(Z_{t+1}^\epsilon\right)_{-i} = \left(Z_t^\epsilon\right)_{-i}.$$

Conversely, if the upward direction is chosen, then:

- if $u_i\left(x_i^+, \mathbf{x}_{-i}\right) < u_i\left(x_i, \mathbf{x}_{-i}\right)$, then
$$Z_{t+1}^\epsilon = Z_t^\epsilon;$$
- if $u_i\left(x_i^+, \mathbf{x}_{-i}\right) \geq u_i\left(x_i, \mathbf{x}_{-i}\right)$, then
$$Z_{t+1}^\epsilon = Z_t^\epsilon, \qquad \text{with probability } 1 - \epsilon^{d_i(x_i)},$$
whereas
$$\left(Z_{t+1}^\epsilon\right)_i = \left(Z_t^\epsilon\right)_i^+, \qquad \left(Z_{t+1}^\epsilon\right)_{-i} = \left(Z_t^\epsilon\right)_{-i}, \qquad \text{with probability } \epsilon^{d_i(x_i)}.$$

For a strategy profile $\mathbf{x} \in \mathcal{A}$ the following sets are defined

$$n^+(\mathbf{x}) = \{i \in \mathcal{V} : \; u_i(x_i^+, \mathbf{x}_{-i}) \geq u_i(x_i, \mathbf{x}_{-i})\}$$
$$n^-(\mathbf{x}) = \{i \in \mathcal{V} : \; u_i(x_i, \mathbf{x}_{-i}) \geq u_i(x_i^-, \mathbf{x}_{-i})\}$$

along with the quantity

$$\alpha_\epsilon(\mathbf{x}) = \frac{1}{2n}\left(\sum_{i \in n^+(\mathbf{x})} \epsilon^{d_i(x_i)} + |n^-(\mathbf{x})|\right)$$

Consequently, the transition probabilities of this Markov chain are formally formulated as:

$$P^{(\epsilon)}_{\mathbf{x},\mathbf{y}} = \begin{cases} 1/2n & \text{if } \mathbf{y} = \left(x_i^-, \mathbf{x}_{-i}\right) \text{ and } u_i(\mathbf{y}) \leq u_i(\mathbf{x}) \\ \epsilon^{d_i(x_i)}/2n & \text{if } \mathbf{y} = \left(x_i^+, \mathbf{x}_{-i}\right) \text{ and } u_i(\mathbf{y}) \geq u_i(\mathbf{x}) \\ 1 - \alpha_\epsilon(\mathbf{x}) & \text{if } \mathbf{y} = \mathbf{x} \\ 0 & \text{otherwise}\,. \end{cases} \tag{3.9}$$

The transition probabilities defined in Equation (3.9) reflect a Metropolis-Hastings update mechanism. The parameter ϵ governs the acceptance probability of utility-decreasing transitions. The function $\alpha_\epsilon(\mathbf{x})$ ensures that the transition matrix is stochastic. Consequently, downward transitions (corresponding to the relaxation of action constraints) that do not reduce utility are accepted with a baseline probability of $1/2n$. Conversely, upward transitions, which strictly increase the minimal action and the associated intervention cost, are accepted with a probability penalized by the factor $\epsilon^{d_i(x_i)}$.

Initial considerations are directed toward the special case where $\epsilon = 0$. It is observed that, in this scenario, only transitions from higher to lower actions are permissible; thus, convergence to an absorbing state is guaranteed for the Markov chain $Z_t^{(0)}$. The set of all states reachable by the Markov chain $Z_t^{(0)}$ starting from $\mathbf{M}$ is defined as

$$\mathcal{Z} = \left\{\mathbf{x} \in \mathcal{A} |\; \mathbb{P}\left(\exists t_0 \geq 0 : Z_{t_0}^{(0)} = \mathbf{x} | Z_0^{(0)} = \mathbf{M}\right) \geq 0\right\} \tag{3.10}$$

The following result is established.

Proposition 3.3.1. *For a finite super-modular unimodal game, let $\mathcal{Z}$ be the set defined in* (3.10)*. Then, the equality $\mathcal{O} = \mathcal{Z}$ holds.*

Proof. By definition, $\mathbf{x} \in \mathcal{Z}$ if and only if there exists a tuple of strategy profiles $\left(\mathbf{y}^{(k)}\right)_{0 \leq k \leq l}$, such that $\mathbf{y}^{(0)} = \mathbf{M},\ \mathbf{y}^{(l)} = \mathbf{x}$, and for every $0 \leq k < l$ there exists i_k such that

$$y_{i_k}^{(k+1)} = \left(y_{i_k}^{(k)}\right)^{-}, \qquad \mathbf{y}_{-i_k}^{(k+1)} = \mathbf{y}_{-i_k}^{(k)},$$

and

$$u_{i_k}\left(\mathbf{y}^{(k+1)}\right) \leq u_{i_k}\left(\mathbf{y}^{(k)}\right).$$

It is noted that this condition is equivalent to stating that the reversed path $\left(\mathbf{x}^{(k)}\right)_{0 \leq k \leq l}$ with $\mathbf{x}^{(k)} = \mathbf{y}^{(l-k)}$ for $0 \leq k \leq l$ constitutes a wI-path from $\mathbf{x}$ to $\mathbf{M}$. By Corollary 3.2.2, this property is shown to be equivalent to the condition $\mathbf{x} \in \mathcal{O}$. ■

However, given that the absorbing states of the chain $Z_t^{(0)}$ do not generally belong to $\mathcal{O}^*$, the execution of this Markov chain is not deemed useful for the present objectives. The perturbed Markov chain Z_t^{ϵ} for $\epsilon > 0$ is subsequently analyzed, beginning with the following technical lemma.

Lemma 3. *Let a discrete-time Markov Chain Z_t on $\mathcal{A}$ with transition matrix P be considered, satisfying the following properties:*

(a) if $\mathbf{x} \neq \mathbf{y}$ differ in more than one entry or are such that $\mathbf{y}_{-i} = \mathbf{x}_{-i}$ and $y_i \notin \left\{x_i^-, x_i^+\right\}$, then $P_{\mathbf{x},\mathbf{y}} = 0$;

(b) if $\mathbf{x} \neq \mathbf{y}$ are such that $\mathbf{y}_{-i} = \mathbf{x}_{-i}$ and $y_i = x_i^+$, then

$$P_{\mathbf{x},\mathbf{y}} = P_{\mathbf{y},\mathbf{x}}\epsilon^{d_i(x_i)}.$$

Then, Z_t is reversible with respect to the invariant distribution

$$\mu_{\mathbf{x}} \propto \epsilon^{\sum_i \gamma_i(x_i)}, \qquad \mathbf{x} \in \mathcal{A}.$$

Proof. For $\mathbf{x} \neq \mathbf{y}$ such that $\mathbf{y}_{-i} = \mathbf{x}_{-i}$ and $y_i = x_i^+$,

$$\frac{\mu_{\mathbf{y}}}{\mu_{\mathbf{x}}} = \frac{\epsilon^{\gamma_i(x_i^+)}}{\epsilon^{\gamma_i(x_i)}} = \epsilon^{d_i(x_i)} = \frac{P_{\mathbf{x},\mathbf{y}}}{P_{\mathbf{y},\mathbf{x}}}.$$

For all other $\mathbf{x} \neq \mathbf{y}$, $P_{\mathbf{x},\mathbf{y}} = P_{\mathbf{y},\mathbf{x}} = 0$ holds. Hence, the detailed balance equation $\mu_{\mathbf{x}}P_{\mathbf{x},\mathbf{y}} = \mu_{\mathbf{y}}P_{\mathbf{y},\mathbf{x}}$ is satisfied for all $\mathbf{x}, \mathbf{y}$ in $\mathcal{A}$. ■

Theorem 3.3.2. *Let a finite super-modular unimodal game be considered. For $\epsilon > 0$, the Markov chain Z_t^ϵ with transition probabilities* (3.9)*:*

(i) keeps the set $\mathcal{Z}$ invariant;

(ii) is time-reversible and ergodic on the set $\mathcal{Z}$;

(iiii) has stationary probability distribution

$$\mu_{\mathbf{x}}^{(\epsilon)} = \epsilon^{\sum_i \gamma_i(x_i)}/\mathcal{K}_\epsilon, \qquad \mathbf{x} \in \mathcal{Z} \tag{3.11}$$

where $\mathcal{K}_\epsilon = \sum_{\mathbf{x}\in\mathcal{Z}} \epsilon^{\sum_i \gamma_i(x_i)}$.

Proof.

(i) Let $\mathbf{x}$ in $\mathcal{Z}$ be a strategy profile that is reachable from the profile M by the Markov chain $Z_t^{(0)}$, and let $\mathbf{y}$ in $\mathcal{A}$ be a strategy profile such that $P_{\mathbf{x},\mathbf{y}}^{(\epsilon)} > 0$. It must be proven that $\mathbf{y}$ belongs to $\mathcal{Z}$. If $\mathbf{y} = \left(x_i^-, \mathbf{x}_{-i}\right)$ for some player $i \in \mathcal{V}$ such that $u_i(\mathbf{y}) \leq u_i(\mathbf{x})$, then it follows that $0 \leq P_{\mathbf{x},\mathbf{y}}^{(\epsilon)} = 1/2n$ and $P_{\mathbf{x},\mathbf{y}}^0 = 1/2n > 0$, thus implying that the strategy profile $\mathbf{y}$ belongs to $\mathcal{Z}$. Conversely, if $\mathbf{y} = \left(x_i^+, \mathbf{x}_{-i}\right)$ for some player i in $\mathcal{V}$, the following argument applies. Since $\mathbf{x}$ in $\mathcal{Z}$ is a strategy profile reachable by the Markov chain $Z_t^{(0)}$ from the profile M, a sequence of profiles can be found $\left(\mathbf{x}^{(k)}\right)_{k=0,\dots,l}$ such that $\mathbf{x}^{(0)} = \mathrm{M}$ and $\mathbf{x}^l = \mathbf{x}$ and $P^0_{\mathbf{x}^{(k-1)},\mathbf{x}^{(k)}} > 0$ for $1 \leq k \leq l$. From (3.9), this is equivalent to

$$\mathbf{x}^{(k)} = \left(\left(x_{i_k}^{(k-1)}\right)^-, \mathbf{x}_{-i_k}^{(k-1)}\right) \text{ and } u_{i_k}\left(\mathbf{x}^{(k)}\right) \leq u_{i_k}\left(\mathbf{x}^{(k-1)}\right)$$

Let s in $\{1, ..., l\}$ be such that $i_s = i$ and consider $\left(\mathbf{z}^{(k)}\right)_{0\leq k<1}$ such that $\mathbf{z}^{(k)} = \mathbf{x}^{(k)}$ for $k \leq s-1$ and $\mathbf{z}^{(k)} = \left(\left(x_i^{(k+1)}\right)^+, \mathbf{x}_{-i}^{(k+1)}\right)$ for $k \geq s$. Notice that

$$\begin{aligned}
\mathbf{z}^{(k)} &= \left(\left(x_i^{(k+1)}\right)^+, \mathbf{x}_{-i}^{(k+1)}\right) = \left(\left(x_i^{(k+1)}\right)^+, \left(x_{i_{k+1}}^{(k)}\right)^-, \mathbf{x}_{-i,-i_{k+1}}^{(k)}\right) \\
&= \left(\left(x_i^{(k)}\right)^+, \left(x_{i_{k+1}}^{(k)}\right)^-, \mathbf{x}_{-i,-i_{k+1}}^{(k)}\right) = \left(\left(x_{i_{k+1}}^{(k)}\right)^-, \left(x_i^{(k)}\right)^+, \mathbf{x}_{-i,-i_{k+1}}^{(k)}\right) \\
&= \left(\left(x_{i_{k+1}}^{(k)}\right)^-, \mathbf{z}_{-i_{k+1}}^{(k-1)}\right) = \left(\left(z_{i_{k+1}}^{(k-1)}\right)^-, \mathbf{z}_{-i_{k+1}}^{(k-1)}\right)
\end{aligned}$$

Using this relation and the super-modularity property,

$$0 \leq u_{i_{k+1}}(\mathbf{x}^{(k)}) - u_{i_{k+1}}\left(\mathbf{x}^{(k+1)}\right) \leq u_{i_{k+1}}\left(\mathbf{z}^{(k-1)}\right) - u_{i_{k+1}}\left(\mathbf{z}^{(k)}\right)$$

for every $k \geq s-1$. This implies that $P^0_{\mathbf{z}^{(k-1)},\mathbf{z}^{(k)}} > 0$ for every $1 \leq k < l$. Since $\mathbf{z}^{(l-1)} = \left(\left(x_i^{(l)}\right)^+, \mathbf{x}_{-i}^{(l)}\right) = \mathbf{y}$, this proves that the strategy profile $\mathbf{y}$ belongs to $\mathcal{Z}$.

(ii) It follows from Lemma 3 that

$$\epsilon^{\sum_i \gamma_i(x_i)} P^{(\epsilon)}_{\mathbf{x},\mathbf{y}} = \epsilon^{\sum_i \gamma_i(y_i)} P^{(\epsilon)}_{\mathbf{y},\mathbf{x}} \tag{3.12}$$

for every two strategy profiles $\mathbf{x}$ and $\mathbf{y}$ in $\mathcal{A}$. This implies that Z_t^ϵ is time-reversible with respect to the stationary distribution (3.11). Since the transitions that have positive probability for Z_t^0 have also positive probability for Z_t^ϵ, all profiles in $\mathcal{Z}$ can be reached from the profile $\mathbf{M}$ by the Markov chain Z_t^ϵ. Moreover, (3.12) implies that a transition probability $P^{(\epsilon)}_{\mathbf{x},\mathbf{y}}$ is positive if and only if the reverse transition $P^{(\epsilon)}_{\mathbf{y},\mathbf{x}}$ is positive. This implies that $\mathbf{M}$ is reachable from any other profile in $\mathcal{Z}$ and thus it is concluded that Z_t^ϵ is ergodic on $\mathcal{Z}$.

(iii) Ergodicity and (3.12) imply that, for every $\epsilon > 0$, the unique stationary distribution of Z_t^ϵ on $\mathcal{Z}$ is (3.11). ■

Asymptotic Behavior of the Stationary Distribution

An analytical characterization of the stationary distribution $\mu^{(\epsilon)}$ is required to establish its continuity and asymptotic concentration properties as the perturbation parameter $\epsilon \to 0^+$. By Theorem 3.3.2, for any strictly positive ϵ, the invariant measure over the state space $\mathcal{Z}$ is explicitly defined as: $\mu^{(\epsilon)}_{\mathbf{x}} = \frac{\epsilon^{C(\mathbf{x})}}{\mathcal{K}_\epsilon}$.

To rigorously evaluate the asymptotic limit, the state space $\mathcal{Z}$ is partitioned into the set of minimum-cost sufficient interventions, denoted by $\mathcal{O}^* = \operatorname*{argmin}_{\mathbf{z}\in\mathcal{Z}} C(\mathbf{z})$, and the complementary set of sub-optimal states, $\mathcal{Z} \setminus \mathcal{O}^*$. Let $C^* = \min_{\mathbf{z}\in\mathcal{Z}} C(\mathbf{z})$ denote the globally optimal cost. The partition function $\mathcal{K}_\epsilon$ can be algebraically factored by extracting the dominant exponential term

$$\mathcal{K}_\epsilon = \epsilon^{C^*}\left(|\mathcal{O}^*| + \sum_{\mathbf{z}\in\mathcal{Z}\setminus\mathcal{O}^*} \epsilon^{C(\mathbf{z})-C^*}\right).$$

Consequently, the probability of state $\mathbf{x} \in \mathcal{Z}$ can be analytically rewritten as:

$$\mu_{\mathbf{x}}^{(\epsilon)} = \frac{\epsilon^{C(\mathbf{x})-C^*}}{|\mathcal{O}^*| + \sum_{\mathbf{z} \in \mathcal{Z} \setminus \mathcal{O}^*} \epsilon^{C(\mathbf{z})-C^*}}$$

The limit of this expression as $\epsilon \to 0^+$ is subsequently evaluated. If a state belongs to the optimal set ($\mathbf{x} \in \mathcal{O}^*$), the numerator exponent evaluates to zero, yielding a numerator of 1. For any strictly sub-optimal state $\mathbf{z} \notin \mathcal{O}^*$, the exponent is strictly positive (i.e., $C(\mathbf{z}) - C^* > 0$). Therefore, the exponential term $\epsilon^{C(\mathbf{z})-C^*}$ vanishes asymptotically. As the summation in the denominator vanishes, the probability mass converges strictly to $1/|\mathcal{O}^*|$. Conversely, if $\mathbf{x} \notin \mathcal{O}^*$, the numerator exponent remains strictly positive, dictating that the probability mass converges to 0. This algebraic factorization rigorously proves that the stationary measure exhibits a continuous extension to $\epsilon = 0$, at which point it perfectly degenerates into a uniform distribution supported exclusively on the optimal intervention set $\mathcal{O}^*$.

Corollary 3.3.3. *Let a finite super-modular unimodal game be considered. For any $\epsilon > 0$, let $\mu^{(\epsilon)}$ denote the stationary distribution of the Markov chain Z_t^ϵ governed by the transition probabilities defined in* (3.9)*. Then, the asymptotic behavior of the invariant measure is given by:*

$$\lim_{\epsilon \to 0} \mu_{\mathbf{x}}^{(\epsilon)} = \begin{cases} 1/|\mathcal{O}^*| & \textit{if } \mathbf{x} \in \mathcal{O}^* \\ 0 & \textit{if } \mathbf{x} \notin \mathcal{O}^*. \end{cases}$$

Proof. As the perturbation ϵ vanishes, the stationary distribution, $\mu^{(\epsilon)}$, converges to a uniform distribution supported exclusively on the subset $\operatorname{argmin}_{\mathbf{z} \in \mathcal{Z}} \sum_i \gamma_i(z_i)$. By virtue of Proposition 3.3.1, the reachable state space $\mathcal{Z}$ strictly coincides with the set of sufficient interventions $\mathcal{O}$. Consequently, this set coincides exactly with the set of minimum-cost sufficient interventions, $\mathcal{O}^*$. ■

Operationally, this continuity result allows the stochastic process (Z_t^ϵ), for arbitrarily small $\epsilon > 0$, to be utilized as an approximation for exact combinatorial minimization. The minimal cost intervention is identified by simulating the Markov chain (Z_t^ϵ) initialized at the maximal configuration $\mathbf{M}$. Since the state space $\mathcal{Z}$ strictly coincides with the set of sufficient interventions $\mathcal{O}$, any state visited by the chain constitutes a feasible intervention. The established asymptotic properties ensure that the process concentrates on the optimal set $\mathcal{O}^*$. Consequently, termination of the algorithm is guaranteed to yield a sufficient intervention, although global optimality remains dependent on the mixing time.

3.4 Computational Results

In this section, numerical simulations of the algorithm proposed in Section 3.3 are presented. Specifically, the algorithm is applied to the game presented in Example 3.1.1 with n players and action set $\mathcal{A}_i = \{0, 1, ..., M\}$.

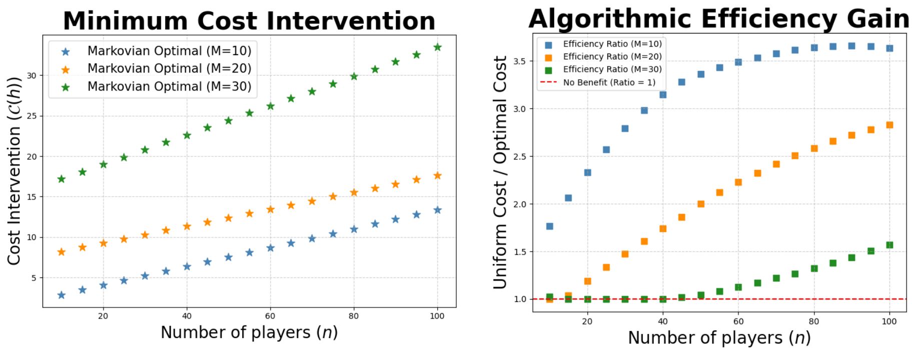

Figure 3.1: Minimum cost of sufficient intervention and algorithmic efficiency gain for Erdős-Rényi random graphs $E(n, p)$ with $p = 4n^{-1} \log n$. The intervention cost, computed via the Markovian algorithm for maximum action bounds $M \in \{10, 20, 30\}$, is illustrated in the left panel. The efficiency ratio relative to the optimal uniform intervention is plotted in the right panel.

The Erdos-Renyi random graph $E(n, p)$ is considered, with n nodes where links between nodes are present with probability p, independently from one another. The regime $p = 4\frac{\log n}{n}$ is considered, leading to a sparse graph that nevertheless remains connected with high probability as n grows large. The intervention cost is evaluated in accordance with Equation (3.3), $C(\mathbf{h}) = \frac{1}{2} \sum_{i \in \mathcal{V}} h_i^2$.
The randomized algorithm Z_t^ϵ is executed with a perturbation parameter $\epsilon = 0.01$ over $T = 100n^2$ steps. The computational outcomes are detailed in Figure 3.1 for maximum action set bounds $M = 2, 5, 10$. As illustrated in Figure 3.1 (left), the minimum cost of sufficient intervention, identified along the Markov chain trajectory, is presented. It is observed that the absolute intervention cost scales monotonically with both the network order n and the maximal strategy profile M. To rigorously evaluate the algorithmic performance, these results are benchmarked against the optimal sufficient homogeneous intervention, wherein an identical lower bound $h_i = h$ is uniformly imposed across all agents.
In sparse network configurations, the variance in nodal degrees and localized structural properties is significantly amplified.

Consequently, homogeneous policy interventions fail to capture the underlying graph topology, inducing economically inefficient over allocations of incentives within highly connected subgraphs. The algorithmic efficiency gain, defined as the ratio of the optimal uniform intervention cost to the minimal targeted cost, is illustrated in 3.1 (right). The efficiency ratio exceeds unity, increasing monotonically with the network order n and decreasing as a function of the maximum action bound $\mathbf{M}$. These computational results empirically verify the economic suboptimality of homogeneous policies within heterogeneous network topologies.

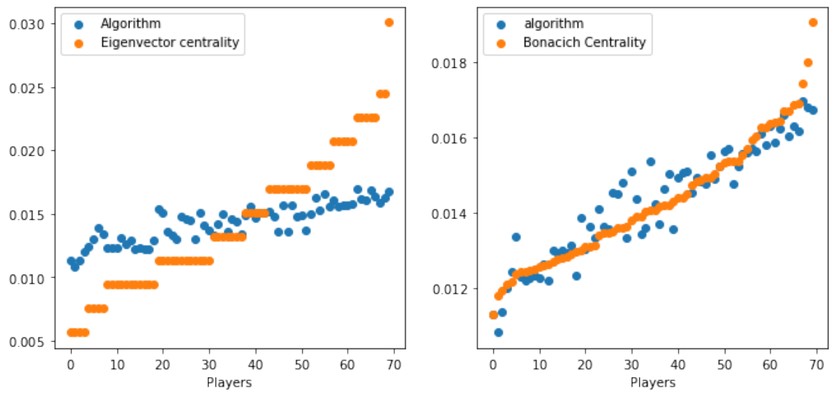

Figure 3.2: For Erdős-Rényi random graphs $E(n,p)$ with $p = 4\frac{\log n}{n}$, comparison between the average cost of the random-node with Eigenvector centrality (left) and Bonacich centrality (right).

The average cost, illustrated in Figure 3.2, is defined as the expected intervention effort required for each individual node, computed over a series of independent stochastic simulations. Mathematically, the average cost for a generic node $i \in \mathcal{V}$, denoted as $\bar{\mathcal{C}}_i$, is evaluated as the arithmetic mean of the interventions $h_i^{(t)}$ applied across T independent iterations. This relationship is formally expressed as:

$$\bar{\mathcal{C}}_i = \frac{1}{T}\sum_{t=1}^{T} h_i^{(t)}$$

For the average cost to be structurally compared with standard probabilistic centrality measures (namely, Eigenvector and Bonacich centralities), it is transformed into a relative metric. This normalization is executed by calculating the ratio of the individual average cost to the aggregate sum of average costs across all nodes within the network, $\bar{\mathcal{C}}_i / \sum_{j\in\mathcal{V}} \bar{\mathcal{C}}_j$. By means of this normalization, it is ensured that the global cost distribution sums to unity. Consequently, the relative fraction of the total systemic effort attributable to node i is directly described by this metric, thereby providing an exact quantification of its topological relevance as a function of its responsiveness to intervention policies.

A fundamental parameter of the proposed algorithm is established as the number of steps T increase.

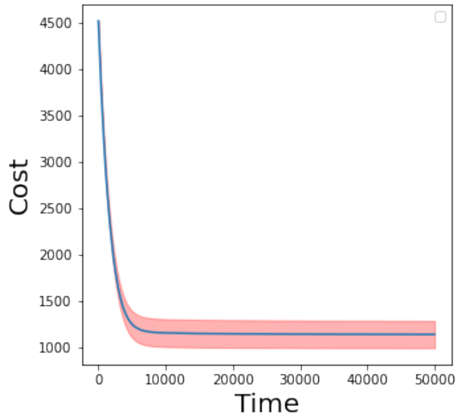

Figure 3.3: Cost on the Facebook ego network, with $n = 334$ and action set $\mathcal{A} = \{0, 1, ..., 9\}$, averaged over 1000 runs.

The evolution of the algorithm is reported in Figure 3.3. The input to the model constitutes an ego-network $\mathcal{G}$, that is a model of social network formed by an individual, the ego, and all the people with whom the ego has a social connection. The center node u of the ego-network (i.e., the "ego") is included in $\mathcal{G}$, so that $\mathcal{G}$ consists of node u's friends (the "alters"). The topological data representing the Facebook ego network[1], utilized as the input model for these computational simulations. The algorithm was simulated across 1000 independent iterations to evaluate the trajectory of the average intervention cost. Although the topological structure of the central ego-node remains invariant, the initial threshold parameters were stochastically resampled for each realization. This variation demonstrates that localized interventions within ego-networks are highly sensitive to the initial state parameters of the central node. As illustrated in Figure 3.3, the cost function exhibits a rapid initial decay, subsequently converging asymptotically to a minimal intervention bound. This asymptotic stabilization empirically substantiates the scalability and computational efficiency of the Markovian algorithm when applied to complex, highly clustered empirical topologies. Consequently, the randomized Markov chain is shown to systematically minimize policy costs over the execution horizon, achieving robust convergence toward the optimal sufficient intervention set.

[1]Retrieved from https://snap.stanford.edu/data/

Chapter 4

Optimal Seeding in Large-Scale Super-Modular Network Games

In this chapter, super-modular network games $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ with binary action sets $\mathcal{A}_i = \{0, 1\}$, for all agent $i \in \mathcal{V}$, are examined.

In this restricted binary setting, the unimodality assumption of Chapter 3 becomes a trivial consequence. A payoff function defined over a two-point domain inherently exhibits unimodality for any fixed configuration of the opponents' actions. Consequently, the convergence guarantees and the path-finding logic established inherently extend to the binary-action framework, preserving computational tractability without necessitating further structural assumptions.

For binary-action super-modular network games, the maximal best-response dynamic (2.16) is mathematically equivalent to the Linear Threshold Model (LTM) (Granovetter, 1978). This equivalence is established by Proposition (2.4.2) and Theorem (2.4.5), which demonstrate that the LTM corresponds to the synchronous dynamic evolution of the M-attracting $\overline{\mathcal{B}}$-path determined by the game's super-modular structure. Consequently, applying lower-bound restrictions

$$\mathcal{A}_i(\mathbf{h}) = \{x_i \in \mathcal{A}_i \mid x_i \geq h_i\},$$

via an intervention vector $\mathbf{h} \in \{0, 1\}^{\mathcal{V}}$ is equivalent to forcing a subset of agents

$$\mathcal{S} = \{i \in \mathcal{V} \mid h_i = 1\} \subseteq \mathcal{V}$$

to adopt action 1.
This theoretical formulation corresponds to the standard *Seeding Problem*.

The problem of identifying a minimal-cost seeding that guarantees at least a predefined fraction of players adopt a certain action in every Nash equilibrium is systematically studied. In the context of large-scale network games, two key analytical challenges are faced by a central planner during the design of interventions.

First, optimal intervention strategies do not scale well with the network size: e.g., many formulations of optimal seeding problems are inherently NP-hard, necessitating the development of approximation algorithms to achieve sub-optimal solutions (Como et al., 2021; Kempe et al., 2003; Messina et al., 2023).

Second, full information on the network structure is rarely available to the planner, as it is either computationally prohibitive to collect or severely constrained by privacy and proprietary issues. In such instances, it is required that the central planner rely on a stochastic network model that accurately matches the available statistical information to design the intervention (Parise and Ozdaglar, 2023). Furthermore, the necessity of exact topological data for designing effective interventions has been critically re-evaluated in the recent literature. It has been formally demonstrated that the theoretical value of complete network information for diffusion processes is frequently inflated; specifically, augmenting a randomly selected seed set by a marginal fraction can often yield larger cascades than optimally targeted seeding derived from exact network topologies (Akbarpour et al., 2023). In a similar vein, it has been established that contagion outcomes and optimal interventions in large populations can be rigorously predicted and designed by utilizing graphons as a stochastic network formation model, thereby entirely bypassing the requirement for precise network realizations (Erol et al., 2023). Motivated by these limitations and recent findings, the optimization problem is formulated within this chapter to rely exclusively on aggregate network statistics.

Building upon the local mean-field approximation of linear threshold dynamics on configuration model random graphs (Rossi et al., 2017), an optimization problem is formulated that relies exclusively on network statistics. It is demonstrated that this optimization problem assumes the structure of a linear program subject to an infinite set of constraints. Subsequently, a methodology is presented to approximate the solution using standard linear programs with finitely many constraints.

4.1 Binary-Action Semi-Anonymous Games

Networks are modeled as finite directed multi-graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, where $\mathcal{V}$ is the node set, $\mathcal{E}$ is the set of directed links, and $\tau : \mathcal{E} \to \mathcal{V}$ and $\sigma : \mathcal{E} \to \mathcal{V}$ are the maps associating to each link its tail and head node, respectively. Let $n = |\mathcal{V}|$ be the network order, and W in $\mathbb{R}^{\mathcal{V}\times\mathcal{V}}$ be the adjacency matrix, with entries $W_{i,j} = |\{e \in \mathcal{E} : \tau(e) = i, \sigma(e) = j\}|$.

Let $\mathbf{d} = W'\mathbf{1}$ and $\mathbf{k} = W\mathbf{1}$ denote the in- and out-degree vectors, respectively. It is assumed that $\mathcal{G}$ contains no self-loops, i.e., $\tau(e) \neq \sigma(e)$ for every e in $\mathcal{E}$. Given a network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, binary-action semi-anonymous (BASA) games are considered Arditti et al., 2023; Jackson et al., 2008; Ravazzi et al., 2023, i.e., strategic games $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ with player set $\mathcal{V}$, whereby every player i in $\mathcal{V}$ has action set $\mathcal{A}_i = \{0, 1\}$ and utility function

$$u_i\left(x_i, \mathbf{x}_{-i}\right) = \mathfrak{u}_i\left(x_i, \sum_{j\in\mathcal{V}} W_{i,j}x_j\right), \qquad \forall\, \mathbf{x} \in \mathcal{A} \tag{4.1}$$

where $\mathfrak{u}_i : \mathcal{A}_i \times \mathbb{Z}_+ \to \mathbb{R}$.

The payoff of the players depends on their action and on the weighted number of neighbors choosing action 1, $\sum_{j\in\mathcal{V}} W_{i,j}x_j$, rather than on the specific action profile of individual peers.

It is further assumed that the function

$$f_i\left(\sum_{j\in\mathcal{V}} W_{i,j}x_j\right) = \mathfrak{u}_i\left(1, \sum_{j\in\mathcal{V}} W_{i,j}x_j\right) - \mathfrak{u}_i\left(0, \sum_{j\in\mathcal{V}} W_{i,j}x_j\right) \tag{4.2}$$

returning the net utility gain of a player i, when unilaterally switching from action 0 to action 1, is a non-decreasing function in the number $\sum_{j\in\mathcal{V}} W_{i,j}x_j$ of the out-neighbors that are playing action 1.

For BASA games, this is equivalent to the increasing difference property (2.12), hence super-modularity. Consequently, BASA games exhibiting a non-decreasing function $f_i\left(\sum_{j\in\mathcal{V}} W_{i,j}x_j\right)$ defined in (4.2) are designated as BASASM games.

The following result states that the best response correspondence admits a threshold structure, and that a minimal Nash equilibrium of Γ exists.

Lemma 4. *For a BASASM game* Γ *on a network* $\mathcal{G}$*, let*

$$r_i = \inf\{\theta \in \mathbb{Z} \ : \ f_i(\theta) \geq 0\}, \tag{4.3}$$

for every player i *in* $\mathcal{V}$*. Then:*

(i) the utility of a player i *in* $\mathcal{V}$ *satisfies*

$$u_i(1, \mathbf{x}_{-i}) \geq u_i(0, \mathbf{x}_{-i}) \Longleftrightarrow \sum_{j \in \mathcal{V}} W_{i,j} x_j \geq r_i; \tag{4.4}$$

(ii) there exists a minimal Nash equilibrium *of* Γ*, i.e.,* $\underline{\mathbf{x}}^*$ *in* $\mathcal{A}^*$ *such that* $\underline{\mathbf{x}}^* \leq \mathbf{x}^*$ *for every other Nash equilibrium* $\mathbf{x}^*$ *in* $\mathcal{A}^*$*.*

Proof.

(i) For every player $i \in \mathcal{V}$, action 1 dominates action 0 if and only if

$$u_i(1, \mathbf{x}_{-i}) - u_i(0, \mathbf{x}_{-i}) \geq 0,$$

and by considering a utility of the form (4.1), the following is obtained

$$\mathrm{u}_i \left(1, \sum_{j \in \mathcal{V}} W_{i,j} x_j\right) - \mathrm{u}_i \left(0, \sum_{j \in \mathcal{V}} W_{i,j} x_j\right) \geq 0.$$

By (4.2) this is equivalent to $f_i\left(\sum_j W_{i,j} x_j\right) \geq 0$. Since f_i is non-decreasing, this is in turn equivalent to $\sum_{j \in \mathcal{V}} W_{i,j} x_j \geq r_i$.

(ii) That $f_i(\cdot)$ is non-decreasing for every i in $\mathcal{V}$ implies super-modularity of the game Γ. Existence of a minimal Nash equilibrium $\underline{\mathbf{x}}^*$ then follows from standard results for super-modular games 2.3.1. ■

Remark 5. *Lemma 4 implies that the minimal Nash equilibrium* $\underline{\mathbf{x}}^*$ *of a BASASM game* $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i \in \mathcal{V}}, \{u_i\}_{i \in \mathcal{V}})$ *depends solely on the couple* $(\mathcal{G}, \mathbf{r})$ *of the network* $\mathcal{G}$ *and the threshold vector* $\mathbf{r}$ *defined in* (4.3)*.*

4.2 The Optimal Seeding Problem

In this chapter, seeding interventions are considered, which consist of selecting a subset of players

$$\mathcal{S} \subseteq \mathcal{V},$$

termed the *seed set*, who are compelled to play action 1 independently of the actions taken by other players. Formally, let the intervention space be defined as $\mathcal{H} = 2^{\mathcal{V}}$, the power set of players. An intervention parameter is therefore:

$$\mathbf{h} = \mathcal{S} \subseteq \mathcal{V}.$$

Given $\mathcal{S} \subseteq \mathcal{V}$, let the seeded game be defined as $\Gamma(\mathcal{S}) = (\mathcal{V}, \{\mathcal{A}_i(\mathcal{S})\}_{i\in\mathcal{V}}, \{u_i^{\mathcal{S}}\}_{i\in\mathcal{V}})$, where the action space is

$$\mathcal{A}(\mathcal{S}) = \left\{\mathbf{x} \in \{0,1\}^{\mathcal{V}} \ : \ x_i = 1 \text{ for all } i \in \mathcal{S}\right\}.$$

Remark 6. *Seeding interventions, instead of restricting action sets, can be formulated by* ***incentive-based implementation****,*

$$u_i^{\mathcal{S}}(\mathbf{x}) = u_i(\mathbf{x}) + M\mathbf{1}_{[i\in\mathcal{S},\ x_i=1]}(x_i),$$

for sufficiently large $M > 0$ that exceeds any possible gain from deviation, action 0 becomes strictly dominated for seeded players. This corresponds to an intervention wherein players i in $\mathcal{S}$ are assigned the utility function $u_i^{\mathcal{S}}(x_i, \mathbf{x}_{-i}) = x_i$, while the remaining players i in $\mathcal{V} \setminus \mathcal{S}$ retain the utility function specified in (4.1).

Remark 7. *Lemma 4 implies that including a player i in the seeding $\mathcal{S}$ effectively amounts to replacing its threshold r_i with a value $r_i^{\mathcal{S}} = 0$.*

Let $\gamma_i \geq 0$ be the cost associated with including player i in the seeding $\mathcal{S}$. Without loss of generality (Remark 7), it is assumed that $\gamma_i = 0$ for every agent i such that $r_i = 0$. For a given tolerance value ϵ in $[0, 1]$, a seeding $\mathcal{S}$ of minimal aggregate cost is sought, which guarantees

$$\frac{1}{n}\sum_{i\in\mathcal{V}} \underline{x}_i^*(\mathcal{S}) \geq 1 - \epsilon, \tag{4.5}$$

i.e., that the minimal Nash equilibrium of the game $\Gamma(\mathcal{S})$ has all but at most a fraction ϵ of players playing action 1.

Hence, the subsequent optimization problem is formulated

$$\begin{array}{ll} \underset{\mathcal{S} \subseteq \mathcal{V}}{\text{minimize}} & \sum_{i \in \mathcal{S}} \gamma_i \\ \text{subject to} & \frac{1}{n} \sum_{i \in \mathcal{V}} \underline{x}_i^*(\mathcal{S}) \geq 1 - \epsilon \end{array} \tag{4.6}$$

Problem (4.6) is always feasible since $\underline{\mathbf{x}}^*(\mathcal{V}) = \mathbf{1}$ (c.f. Remark 5), so that (4.5) is always satisfied by $\mathcal{S} = \mathcal{V}$.

The minimal equilibrium is established as the primary quantity of interest because it mathematically represents the worst-case coordination outcome of the network game. Due to the super-modular structure of the game, multiple Nash equilibria may exist, ordered from a minimal to a maximal configuration. By enforcing the constraint that the targeted adoption fraction is satisfied even in the minimal Nash equilibrium, global systemic activation is mathematically guaranteed regardless of the equilibrium ultimately selected by the decentralized system. Consequently, designing interventions based on this strict lower bound provides a robust guarantee against coordination failures.

An instance of the optimal seeding problem (4.6) is thus fully parameterized by the tolerance value ϵ and the triple $(\mathcal{G}, \mathbf{r}, \gamma)$, comprising the network $\mathcal{G}$, the threshold vector $\mathbf{r}$ dictating the minimal Nash equilibrium (Remark (5)), and cost vector γ.

Within the binary-action framework explored in this chapter, the intervention mechanism is constrained to the selection of a seed set $\mathcal{S}$.

Accordingly, the target class of games Γ^* corresponds directly to the feasibility constraint formulated in Equation 4.5. It is formally established that Γ^* encompasses the set of all seeded games $\Gamma(\mathcal{S})$ wherein the minimal Nash equilibrium, denoted by $\underline{\mathbf{x}}^*(\mathcal{S})$, yields an aggregate adoption level of at least $1 - \epsilon$. Thus, forcing the network game into the target class Γ^* via seeding ensures that, irrespective of the initial configuration, the dynamics converge to a state where the fraction of non-adopting agents is strictly bounded above by the tolerance parameter ϵ.

Example 4.2.1 (Optimal seeding in network coordination games)**.** *Let the network coordination game* (2.3) *on a graph $\mathcal{G}$ be considered, where agents i choose their binary action $x_i \in \{0, 1\}$ so as to maximize their utilities*

$$u_i(x_i, \mathbf{x}_{-i}) = \sum_{j \in \mathcal{V}} W_{i,j} \left((1 - x_i)(1 - x_j) + x_i x_j \right) + b_i x_i$$

The utility function above can be rewritten in the form (4.2) *with*

$$f_i\left(\sum_{j\in\mathcal{V}} W_{i,j}x_j\right) = \mathfrak{u}_i\left(1, \sum_{j\in\mathcal{V}} W_{i,j}x_j\right) - \mathfrak{u}_i\left(0, \sum_{j\in\mathcal{V}} W_{i,j}x_j\right)$$
$$= 2\sum_{j\in\mathcal{V}} W_{i,j}x_j - k_i + b_i$$

that is increasing in $\sum_{j\in\mathcal{V}} W_{i,j}x_j$.

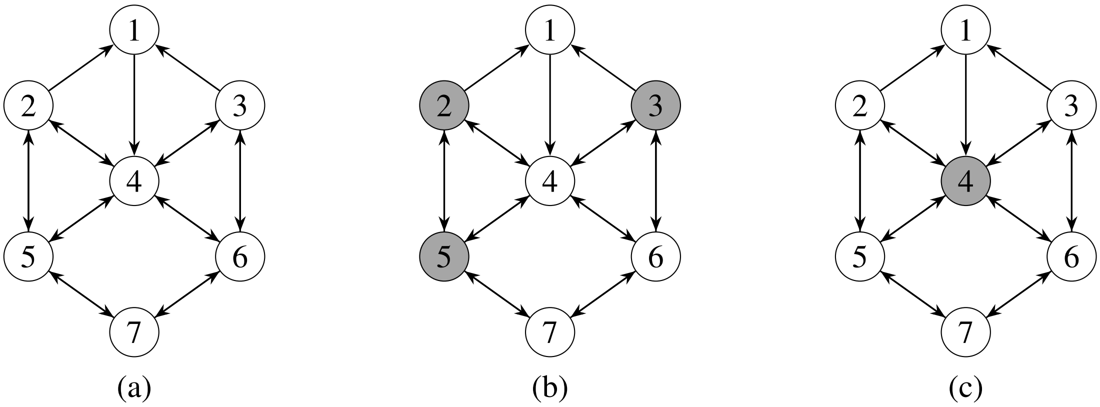

Figure 4.1: Network for optimal seeding in network coordination games.

Let the network coordination game Γ *on the network* $\mathcal{G}$ *displayed in Figure 4.1(a) be considered, with bias parameter* $\mathbf{b} = \mathbf{0}$. *The in- and out-degree vectors are respectively,*

$$\mathbf{d} = (2, 2, 2, 5, 3, 3, 2), \qquad \mathbf{k} = (1, 3, 3, 4, 3, 3, 2)$$

while the threshold vector is

$$\mathbf{r} = (0, 1, 1, 2, 1, 1, 1)$$

Let the cost vector be $\gamma = (2, 1, 1, 4, 1, 1, 2)$ *and the tolerance value* ϵ *in* $[0, 1/7)$. *Then,* $\mathcal{S}^* = \{2, 3, 5\}$ *(see Figure 4.1(b)) can be shown to be an optimal seeding with aggregate cost* 3 *(also* $\{2, 3, 6\}$, $\{2, 5, 6\}$, *and* $\{3, 5, 6\}$ *are optimal seedings). On the other hand,* $\mathcal{S} = \{4\}$ *(see Figure 4.1(c)) satisfies* (4.5)*, however its cost is* 4*, hence it is not an optimal seeding.*

4.3 Mean-Field Approximations

The exact resolution of the optimal seeding problem (4.6) is known to be NP-hard (Como et al., 2021; Kempe et al., 2003). The proposed methodology involves initially formulating an equivalent dynamical representation of constraint (4.5), and subsequently approximating these evolutionary dynamics, which represent the expected fraction of agents playing action 1, via a recursive equation in the spirit of Rossi et al., 2017. This equation, termed the Local Mean-Field dynamic, dictates the evolution of expected activation on infinite (locally tree-like) random networks possessing the same statistical properties as the original network. A concentration theorem is proven: for a generic instance of the network, the probability that the activation process remains bounded close to the Local Mean-Field dynamic converges to one exponentially fast with respect to the network size. Consequently, the analysis is reduced to examining the fixed points and corresponding trajectories of a scalar autonomous system. Finally, the optimization problem is reformulated within this framework. Beyond analytical simplification, the primary advantage of this approach lies in enabling the design of optimal seeding strategies for families of large-scale networks based solely on aggregate statistics, thereby bypassing the requirement for full topological knowledge.

Seeded Linear Threshold Dynamics

For a given binary-action semi-anonymous super-modular (BASASM) game $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$, parameterized by the network topology $\mathcal{G}$ and the intrinsic threshold vector $\mathbf{r}$, a modified threshold vector $\mathbf{r}^{\mathcal{S}}$ is defined with respect to a designated seed set $\mathcal{S} \subseteq \mathcal{V}$. The components of this seeded threshold vector are formally assigned as follows:

$$r_i^{\mathcal{S}} = \begin{cases} 0 & \text{if } i \in \mathcal{S} \\ r_i & \text{if } i \notin \mathcal{S}. \end{cases}$$

The *Seeded Linear Threshold Dynamics* over the configuration space $\mathcal{A}$ are dictated by the application of the LTM transition operator, evaluated utilizing the modified thresholds $\mathbf{r}^{\mathcal{S}}$. The discrete-time state evolution is thus governed by:

$$\mathbf{x}(t+1) = \Phi_{\mathbf{r}^{\mathcal{S}}}(\mathbf{x}(t)), \qquad \forall t \geq 0,$$

The operator $\Phi_{\mathbf{r}^{\mathcal{S}}}$ is defined strictly in accordance with the standard LTM mechanics, the theoretical properties of which are detailed in Appendix B.

Remark 8. *The* Seeded Linear Threshold Dynamics *on the strategy profile space $\mathcal{A}$ is equivalently described by the Linear Threshold Dynamics* (B.1) *defined on the seeded strategy space $\mathcal{A}(\mathcal{S})$. A detailed treatment of the Linear Threshold Dynamics can be found in Appendix B*

Then, the following result is achieved.

Proposition 4.3.1. *Given a BASASM game on a network $\mathcal{G}$. Then,*

$$\frac{1}{n}\sum_{i\in\mathcal{V}} \underline{x}_i^*(\mathcal{S}) \geq 1-\epsilon,$$

holds true if and only if the linear threshold dynamics (B.1) *is such that*

$$\frac{1}{n}\sum_{i\in\mathcal{V}} x_i(t) \geq 1-\epsilon\ , \qquad \forall\, t \geq n, \tag{4.7}$$

for every $\mathbf{x}(0)$ in $\mathcal{A}(\mathcal{S})$.

Proof. Let $\mathbf{x}_{\min} \in \mathcal{A}(\mathcal{S})$ denote the minimal configuration in the seeded action space, defined such that $x_i \in \mathcal{S}$ and $x_i = 0$ for all $i \notin \mathcal{S}$. As established in Example 2.4.2, the Linear Threshold Model (LTM) acts as the discrete-time synchronous maximal best response dynamic for this class of coordination games. According to Lemma (7) in Appendix B, the operator $\Phi_{\mathbf{r}}$ is monotone and non-decreasing component-wise. This is theoretically grounded in Proposition (2.4.1), which confirms that best responses in super-modular games are non-decreasing. Because the operator $\Phi_{\mathbf{r}}^{\mathcal{S}}$ is monotone non-decreasing, initializing the dynamics at $\mathbf{x}(0) = \mathbf{x}_{\min}$ produces a monotonically non-decreasing sequence of states, $\mathbf{x}(t) \leq \mathbf{x}(t+1)$ for all $t \geq 0$. As discussed Example 2.4.2 this sequence must converge to the minimal fixed point (i.e., the minimal Nash equilibrium of the seeded game), $\underline{\mathbf{x}}^*(\mathcal{S})$, in at most n steps. Therefore, when $\mathbf{x}(0) = \mathbf{x}_{\min}$, it follows that $\mathbf{x}(t) = \underline{\mathbf{x}}^*(\mathcal{S})$ for all $t \geq n$.

$\Rightarrow$ Assume that $\frac{1}{n}\sum_{i\in\mathcal{V}} \underline{x}_i^*(\mathcal{S}) \geq 1-\epsilon$.

For any arbitrary initial configuration $\mathbf{x}(0)$ in $\mathcal{A}(\mathcal{S})$, it holds by definition that $\mathbf{x}(0) \geq \mathbf{x}_{\min}$. By the monotonicity of the dynamics, the trajectory starting from $\mathbf{x}(0)$ will always dominate the trajectory starting from $\mathbf{x}_{\min}$. Consequently, for all $t \geq n$, it holds $\mathbf{x}(t) \geq \underline{\mathbf{x}}^*(\mathcal{S})$.

Summing over the components yields:

$$\frac{1}{n}\sum_{i\in\mathcal{V}} x_i(t) \geq \frac{1}{n}\sum_{i\in\mathcal{V}} \underline{x}_i^*(\mathcal{S}) \geq 1-\epsilon.$$

$\Leftarrow$ Assume that $\frac{1}{n}\sum_{i\in\mathcal{V}} x_i(t) \geq 1-\epsilon$ for all $t \geq n$, for every $\mathbf{x}(0)$ in $\mathcal{A}(\mathcal{S})$.

This condition must natively hold for the specific minimal initial configuration $\mathbf{x}(0) = \mathbf{x}_{\min}$. As established, the trajectory from $\mathbf{x}_{\min}$ converges exactly to $\underline{\mathbf{x}}^*(\mathcal{S})$ by time step n. It directly follows that:

$$\frac{1}{n}\sum_{i\in\mathcal{V}} \underline{x}_i^*(\mathcal{S}) = \frac{1}{n}\sum_{i\in\mathcal{V}} x_i(n) \geq 1-\epsilon.$$

This concludes the proof. ■

By Proposition 4.3.1, the minimal Nash equilibrium is established to coincide precisely with the steady-state configuration of the LTM when initiated from the absolute minimal state. Due to the monotonicity of the LTM transition operator, all feasible activation sequences are strictly lower-bounded by this specific trajectory. Consequently, a deterministic guarantee is provided: if the targeted adoption fraction is satisfied by the minimal equilibrium, the required systemic activation is achieved within finite iterations, strictly independent of the initial configuration.

Mean-field Approximation of the Linear Threshold Dynamics

Agent Types. The statistical representation of the network via agent types is theoretically adopted from the framework established by Rossi et al., 2017. In large-scale systems, typically only statistical information regarding the network's topology (e.g., degree distribution) is available; furthermore, acquiring more detailed topological knowledge is generally computationally unmanageable.

For the analysis of such diffusion processes, a large, random network model, specifically, a directed version of the so-called configuration model, is considered to effectively incorporate the degree distribution of the agents.

The objective is to predict the behavior of networked systems strictly in a statistical sense. Consequently, the distribution of agents with specific features is modeled by assigning to every agent i in $\mathcal{V}$ a type ω_i that uniquely determines its in-degree d_i, out-degree k_i, threshold r_i and associated cost γ_i.
Let Ω be a countable set of types and let d, k, r, γ in $\mathbb{R}^\Omega$ be such that

$$d_{\omega_i} = d_i, \quad k_{\omega_i} = k_i, \quad r_{\omega_i} = r_i, \quad \gamma_{\omega_i} = \gamma_i, \quad \forall i \in \mathcal{V}.$$

The vector $p^{(0)}$ in $[0,1]^{\Omega}$ is then defined by

$$p^{(0)}_{\omega} = \frac{1}{n}|\{i \in \mathcal{V} \,:\, \omega_i = \omega\}|, \qquad \forall \omega \in \Omega, \tag{4.8}$$

which denotes the empirical distribution of types and is referred to as the initial *statistic* of the problem prior to any intervention.

Remark 9. *The statistic $p^{(0)}_{\omega}$ is computed after degree and threshold sequences have been established. A prior distribution may be utilized to sample the degrees and thresholds, under the assumption that the empirical statistics converge to the prior as $n \to \infty$.*

Analogously, a vector-valued function $\eta : \Omega \to [0,1]$ can be associated to every seeding $\mathcal{S} \subseteq \mathcal{V}$, with η_{ω} denoting the fraction of agents of type ω that are included in the seeding. Consistently with such interpretation, η must satisfy

$$0 \leq \eta_{\omega} \leq p^{(0)}_{\omega}, \qquad \forall \omega \in \Omega. \tag{4.9}$$

Moreover, since the number of nodes of type ω included in the seeding has to be an integer, η needs to satisfy

$$n\eta_{\omega} \in \mathbb{Z}_{+}, \qquad \forall \omega \in \Omega. \tag{4.10}$$

A pair (n, η) is considered feasible if conditions (4.9) - (4.10) are met. Now, let

$$\Omega(\omega) = \{\omega' \in \Omega : d_{\omega} = d_{\omega'},\ k_{\omega} = k_{\omega'},\ \gamma_{\omega} = \gamma_{\omega'},\ r_{\omega'} > r_{\omega}\}$$

denote the set of types ω' that differ from type ω only in the threshold and such that $r_{\omega'} > r_{\omega}$. By decreasing the agent thresholds, an intervention generates a new network statistics $p(\eta)$ defined by

$$p_{\omega}(\eta) = \begin{cases} p^{(0)}_{\omega} - \eta_{\omega}, & \text{if } 0 < r_{\omega} \leq k_{\omega}, \\ p^{(0)}_{\omega} + \sum\limits_{\omega' \in \Omega(\omega)} \eta_{\omega'}, & \text{if } r_{\omega} = 0. \end{cases} \tag{4.11}$$

for every ω in Ω.

Remark 10. *Note that the null statistical seeding, denoted $\eta^{(0)}$, has entries*

$$\eta^{(0)}_{\omega} = 0, \qquad \forall \omega \in \Omega.$$

With such intervention, $p(\eta^{(0)}) = p^{(0)}$.

Well Posedness. Given a statistic $p(\eta)$ and a finite set of nodes $\mathcal{V}$ of cardinality $n = |\mathcal{V}|$, a vector of types ω in $\Omega^{\mathcal{V}}$ with statistics $p(\eta)$ exists if and only if

$$np_\omega(\eta) \in \mathbb{Z}_+, \qquad \forall \omega \in \Omega. \tag{4.12}$$

Remark 11. *Note that, since the intervention modifies only the agent thresholds and not the degrees, it follows that*

$$\langle p(\eta), d\rangle = \langle p^{(0)}, d\rangle.$$

The Configuration Model. The construction of the directed configuration model follows the methodologies derived from Rossi et al., 2017.

The directed configuration model is generated in two phases.
First, a finite set of n independent nodes $\mathcal{V}$ is defined, with each node $i \in \mathcal{V}$ assigned an in-degree d_i and an out-degree k_i. For a well-defined topological structure to emerge, these sequences must be strictly compatible; that is, the total sum of in-degrees must perfectly equate to the total sum of out-degrees, yielding a global edge cardinality of $|\mathcal{E}| = \sum_{i\in\mathcal{V}} d_i = \sum_{i\in\mathcal{V}} k_i$.

Within this framework, each agent can be conceptualized as possessing d_i distinct inbound half-edges "plugs" and k_i distinct outbound half-edges "sockets". All such half-edges across the network are subsequently indexed by the integer set $\{1, 2, \ldots, |\mathcal{E}|\}$. Furthermore, each node i can be assigned its activation threshold sequence r_i (assumed to be known a priori) and its initial state $x_i(0)$.

For a network $\mathcal{G}$ to exist with node set $\mathcal{V}$ and in- and out-degrees $d_i = d_{\omega_i}$ and $k_i = k_{\omega_i}$ for all i in $\mathcal{V}$, it is necessary and sufficient that

$$\langle p(\eta), d\rangle = \langle p(\eta), k\rangle\,, \tag{4.13}$$

$$d_\omega + k_\omega \leq n\langle p(\eta), d\rangle, \qquad \forall \omega \in \Omega \ : \ p_\omega(\eta) > 0. \tag{4.14}$$

Eq. (4.13) guarantees that the total out-degree equals the total in-degree, while (4.14) ensures the existence of a set of links compatible with the degree distribution and that does not contain any self-loops.

A pair $(n, p(\eta))$ is said to be compatible if conditions (4.12) - (4.14) are met.
In particular, this is guaranteed if both (n, η) and $(n, p^{(0)})$ are compatible pairs. At this stage, only the wiring, represented by the set $\mathcal{E}$, which dictates the nodal connections, remains to be determined.

Second, a specific realization of the network graph is generated by selecting a permutation π uniformly at random from the symmetric group defined over the index set $\{1, 2, \dots, |\mathcal{E}|\}$. The outbound and inbound half-edges are then pairwise connected in strict accordance with the bijective mapping induced by π: $(i; \pi(i))$. The resultant collection of paired indices (i, j) constitutes the directed edge set $\mathcal{E}$, thereby completely defining the multi-graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ upon which the diffusion dynamics are subsequently analyzed. This random graph model is referred to as the *directed configuration model*

For a consistent pair $(n, p(\eta))$, the following definition formalizes the idea of sampling uniformly from the set of problems with statistics $p(\eta)$.

Definition 4.3.1. *Let $\mathcal{V} = \{1, \cdots, n\}$ and let the type profile ω in $\Omega^{\mathcal{V}}$ have empirical distribution $p(\eta)$. Let $l = n\langle p(\eta), d\rangle$ and $\mathcal{E} = \{1, \cdots, l\}$.*

Define $\tau : \mathcal{E} \to \mathcal{V}$ by letting $\tau(e)$ in $\mathcal{V}$ be the unique value such that

$$\sum_{i=1}^{\tau(e)-1} k_i < e \leq \sum_{i=1}^{\tau(e)} k_i,$$

for all e in $\mathcal{E}$. Let $\sigma : \mathcal{E} \to \mathcal{V}$, $\sigma(e) = \tau(\pi(e))$ for e in $\mathcal{E}$, where π is sampled uniformly from the set of permutations of $\mathcal{E}$ such that

$$\tau(e) \neq \tau(\pi(e)), \qquad \forall e \in \mathcal{E}. \tag{4.15}$$

Then, the (resampled) configuration model $\mathcal{C}_{n,p(\eta)}$ *is the triple $(\mathcal{G}, r, \gamma)$ of the random network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, the threshold vector $\mathbf{r}$, and the vector of costs γ with entries $r_i = r_{\omega_i}$ and $\gamma_i(\cdot) = \gamma_{\omega_i}(\cdot)$, respectively, for all i in $\mathcal{V}$.*

In the remainder of this section, the seeding problem (4.6) is reformulated utilizing the mean-field approximation. The derivation of the subsequent local mean-field recursive equations is directly adopted from Rossi et al., 2017. To better formalize the result, let

$$\varphi_{k,r}(z) = \sum_{u=r}^{k} \binom{k}{u} z^u (1-z)^{k-u}, \qquad \forall\, 0 \leq r \leq k,$$

and two maps $\psi_{p(\eta)}, \phi_{p(\eta)} : [0,1] \to [0,1]$ be defined by

$$\psi_{p(\eta)}(z) = \sum_{\omega \in \Omega} p_\omega(\eta) \varphi_{k_\omega, r_\omega}(z),$$

$$\phi_{p(\eta)}(z) = \frac{1}{\langle p(\eta), d\rangle} \sum_{\omega \in \Omega} p_\omega(\eta) d_\omega \varphi_{k_\omega, r_\omega}(z). \tag{4.16}$$

These two maps are polynomials with non-negative coefficients that depend on the network statistics $p(\eta)$, hence, in particular they are monotonically non-decreasing.

The polynomial maps defined in Equation (4.16) describe the macroscopic dynamics of the system. Specifically, $\psi_{p(\eta)}(z)$ maps the expected fraction of active edges to the expected fraction of active nodes at the subsequent time step. The function $\phi_{p(\eta)}(z)$ represents the degree-weighted probability that an outbound link points to an active node. These operators recursively determine whether the threshold condition for global adoption is satisfied within the configuration model.

The next result is instrumental for its analysis.

Proposition 4.3.2. *For every intervention* $\eta : \Omega \to [0,1]^{\mathcal{V}}$ *satisfying* (4.9)*,*

$$\phi_{p(\eta)}(z) = \phi_{p^{(0)}}(z) + \sum_{\omega \in \Omega} a_\omega(z)\eta_\omega, \tag{4.17}$$

where

$$a_\omega(z) = \frac{d_\omega \left(1 - \varphi_{k_\omega, r_\omega}(z)\right)}{\langle p^{(0)}, d \rangle} \geq 0. \tag{4.18}$$

Proof.

$$\phi_{p(\eta)}(z) = \frac{1}{\langle p^{(0)}, d \rangle} \sum_{\omega \in \Omega} p_\omega(\eta) d_\omega \varphi_{k_\omega, r_\omega}(z)$$

By substituting (4.11) into this expression, the following is obtained

$$\frac{1}{\langle p^{(0)}, d \rangle} \left[\sum_{\omega \in \Omega, r_\omega \geq 1} d_\omega \left(p_\omega^{(0)} - \eta_\omega\right) \varphi_{k_\omega, r_\omega}(z) + \sum_{\omega \in \Omega, r_\omega = 0} d_\omega \left(p_\omega^{(0)} + \sum_{\omega' \in \Omega(\omega)} \eta_{\omega'} \right) \varphi_{k_\omega, r_\omega}(z) \right]$$

Notice that $\eta_\omega = 0$ for each type $\omega \in \Omega$ such that $r_\omega = 0$, since it is the type whose agents are already in the seeding. Also $\varphi_{k_\omega, 0}(z) = 1$ by property of binomial coefficient. Then,

$$\begin{aligned} &= \frac{1}{\langle p^{(0)}, d \rangle} \sum_{\omega \in \Omega} d_\omega p_\omega^{(0)} \varphi_{k_\omega, r_\omega}(z) + d_\omega \left(1 - \varphi_{k_\omega, r_\omega}(z)\right) \eta_\omega = \\ &= \phi_{p^{(0)}}(z) + \frac{1}{\langle p^{(0)}, d \rangle} \sum_{\omega \in \Omega} d_\omega \left(1 - \varphi_{k_\omega, r_\omega}(z)\right) \eta_\omega. \end{aligned}$$ ■

The main result of this section, which serves as the theoretical foundation for the subsequent problem reformulation, can now be formalized

Proposition 4.3.3. *Let $p(\eta)$ be a probability distribution on Ω with finite moments $\langle p(\eta), d\rangle$, $\langle p(\eta), k\rangle$, $\langle p(\eta), d^2\rangle$, and $\langle p(\eta), k^2\rangle$. Let $(n, p^{(n)}(\eta))$ be a sequence of compatible pairs such that*

$$\begin{aligned} p^{(n)}(\eta) &\overset{n\to+\infty}{\longrightarrow} p(\eta), \\ \langle p^{(n)}(\eta), d\rangle &\overset{n\to+\infty}{\longrightarrow} \langle p(\eta), d\rangle \\ \langle p^{(n)}(\eta), d^2\rangle &\overset{n\to+\infty}{\longrightarrow} \langle p^{(n)}(\eta), d^2\rangle, \\ \langle p^{(n)}(\eta), k^2\rangle &\overset{n\to+\infty}{\longrightarrow} \left\langle p(\eta), k^2\right\rangle . \end{aligned}$$

Let ϵ in $(0,1]$. If

$$\phi_{p(\eta)}(z) > z, \quad \forall z \in [0, \psi^{-1}_{p(\eta)}(1-\epsilon)], \tag{4.19}$$

then (4.5) *holds true on a fraction of problems drawn from the (resampled) configuration model ensemble $\mathcal{C}_{n,p^{(n)}(\eta)}$ that converges to 1 as n grows large.*

Proof. Let

$$Y(t) = \frac{1}{n}\sum_{i\in\mathcal{V}} x_i(t), \quad Z(t) = \frac{1}{|\mathcal{E}|}\sum_{i\in\mathcal{V}} d_i x_i(t)$$

denote the fraction of state-1 agents and the fraction of links pointing to state-1 agents in the linear threshold model (B.2).
Define $z(t)$ and $y(t)$ recursively by putting $y(0) = z(0) = 0$ and, for $t \geq 0$,

$$y(t+1) = \psi_{p(\eta)}(z(t)) \tag{4.20}$$

$$z(t+1) = \phi_{p(\eta)}(z(t)). \tag{4.21}$$

Then, Theorem B.2.1 implies that $Z(t)$ and $Y(t)$ are arbitrarily close to $z(t)$ and $y(t)$, respectively, on all but an asymptotically vanishing fraction of networks $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, where $\sigma(e) = \tau(\pi(e))$ for a uniform permutation π of $\mathcal{E}$.
Since π satisfies (4.15) with probability bounded away from 0 asymptotically, this implies that $Z(t)$ and $Y(t)$ are arbitrarily close to $z(t)$ and $y(t)$, respectively, on the configuration model $\mathcal{C}_{n,p^{(n)}(\eta)}$ with probability approaching 1 as n grows large. This asymptotic concentration bound, which guarantees that the stochastic diffusion process remains arbitrarily close to the deterministic recursion on configuration model graphs, is an established theorem derived from Rossi et al., 2017.
By recursive equation (4.21), $z(t+1) = \phi_{p(\eta)}(z(t))$, since condition (4.19) holds

$$\phi_{p(\eta)}(z) > z, \quad \forall z \in [0, \psi^{-1}_{p(\eta)}(1-\epsilon)],$$

imply that there exists a finite time τ such that $z(t) > \psi^{-1}_{p(\eta)}(1-\epsilon)$ for all $t \geq \tau$.

Equation (4.20) then implies that

$$y(t+1) = \psi_{p(\eta)}(z(t)) > \psi_{p(\eta)}\left(\psi_{p(\eta)}^{-1}(1-\epsilon)\right) = 1-\epsilon, \qquad \forall t > \tau.$$

since $\psi_{p(\eta)}$ is monotonically non-decreasing.

This implies that, with probability converging to 1 as n grows large $\frac{1}{n}\sum_{i\in\mathcal{V}} x_i(t) = Y(t) > 1-\epsilon$ for $t > \tau$, i.e., (4.7) holds true, thus completing the proof. ■

Proposition 4.3.3 states that, given an intervention η with associated statistics $p(\eta)$ satisfying (4.19), ensures that the fraction of links pointing towards state-1 agents will be at least $\psi_{p(\eta)}^{-1}(1-\epsilon)$ after a certain number of iterations, By (4.20), the fraction of nodes in state 1 will be at least $1-\epsilon$, thus, as far as the network is sufficiently large, the Seeded Linear Threshold Dynamics will converge almost surely to a configuration with at least a fraction $1-\epsilon$ of agents in state 1. So interventions η will drive the system towards a desirable configuration, achieving the goal of the planner.

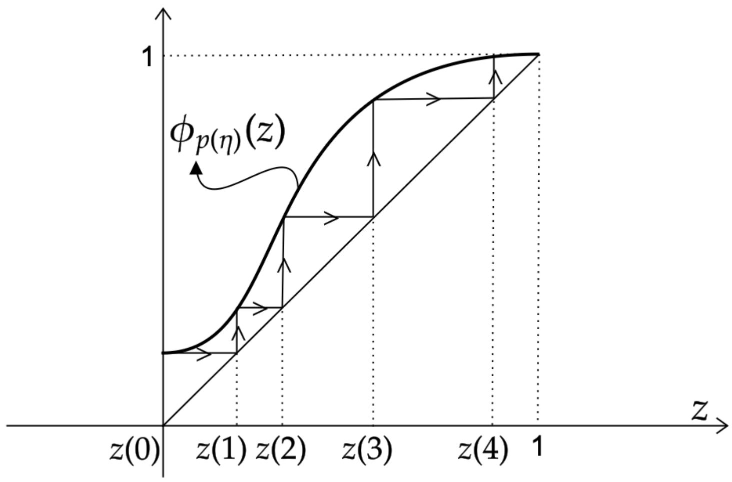

Figure 4.2: The evolution of fraction of links pointing to state-1 agents in the mean-field local approximation according to (4.21) for a statistics $p(\eta)$.

An example of dynamics of $z(t)$ with initial condition $z(0) = 0$ is illustrated in Figure 4.2. For this example, the fraction of links pointing to state-1 agents (and therefore also the fraction of state-1 agents) converges close to 1 in four time steps.

Mean-Field Optimization

In this subsection, the optimization problem (4.6) is reformulated within the statistical framework. While the underlying mean-field approximations build heavily upon the foundational theorems of Rossi et al., 2017, the specific integration of targeted cost-minimization via dimension-independent linear programming constitutes the novel, independent contribution of this chapter. In this framework, the goal of the planner is to find the minimal cost intervention η such that (4.19) holds true, so that Proposition 4.3.3 guarantees that on almost all large-scale networks with statistics $p(\eta)$ at least a fraction $1-\epsilon$ nodes will converge to state 1.

To formalize the problem, the cost of an intervention η must be defined.

Since γ_ω is the cost associated with including a player of type ω in the seeding, it follows by definition of η that the number of nodes of type ω in the seeding is $n\eta_\omega$.Therefore, a cost can be associated with intervention η as follows:

$$C(\eta) = \sum_{\omega\in\Omega} \eta_\omega \gamma_\omega \geq 0,$$

with $C(\eta^{(0)}) = 0$. The optimization problem (4.6) is subsequently reformulated. The *Mean-field Optimal Seeding Problem* is defined by

$$\begin{aligned}
\underset{\eta\,:\,\mathbb{Z}_+ \to [0,1]^\Omega}{\text{minimize}} \quad & C(\eta) \\
\text{subject to} \quad & p_\omega(\eta) = \begin{cases} p_\omega^{(0)} - \eta_\omega, \text{ if } 0 < r_\omega \leq k_\omega, \\ p_\omega^{(0)} + \sum\limits_{\omega'\in\Omega(\omega)} \eta_{\omega'}, \text{ if } r_\omega = 0. \end{cases} \quad \forall \omega \in \Omega, \\
& \phi_{p(\eta)}(z) > z, \quad \forall z \in \left[0, \psi^{-1}_{p(\eta)}(1-\epsilon)\right]
\end{aligned} \tag{4.22}$$

Proposition 4.3.2 states that $\phi_{p(\eta)}$ is linear in η, hence constraint (4.19),

$$\phi_{p(\eta)}(z) = \phi_{p^{(0)}}(z) + \sum_{\omega\in\Omega} a_\omega(z)\eta_\omega > z, \quad \forall z \in \left[0, \psi^{-1}_{p(\eta)}(1-\epsilon)\right]$$

is linear in η. Therefore, in contrast with the original problem (4.6), which is a non-linear program whose size scales with the number of agents, problem (4.22) is a linear program whose size scales with the number of types. However, (4.22) is still highly challenging, because the domain on which constraint (4.19) needs to be satisfied depends on the intervention variable η through the function $\psi_{p(\eta)}$. Moreover, the constraint (4.19) must be satisfied for a continuum of z, leading to an infinite-dimensional optimization problem.

4.4 Approximation Methods

In this section, a numerical method is proposed for small tolerance values $\epsilon > 0$ to identify near-optimal solutions and solve (4.22),

$$\begin{aligned} \underset{\eta\,:\,\mathbb{Z}_+ \to [0,1]^{\Omega}}{\text{minimize}} \quad & C(\eta) \\ \text{subject to} \quad & p_\omega(\eta) = \begin{cases} p_\omega^{(0)} - \eta_\omega, \text{ if } 0 < r_\omega \le k_\omega, \\ p_\omega^{(0)} + \sum\limits_{\omega' \in \Omega(\omega)} \eta_{\omega'}, \text{ if } r_\omega = 0 \end{cases}, \\ & \phi_{p(\eta)}(z) > z, \quad \forall z \in \left[0, \psi_{p(\eta)}^{-1}(1-\epsilon)\right] \end{aligned}$$

A primary technical challenge arises from the fact that the domain requiring the constraint $\phi_{p(\eta)}(z) > z$ to be satisfied depends upon η.

To circumvent the computational complexity introduced by this dependence, an alternative optimization problem is formulated, which is subsequently proven to be a relaxation of the original formulation

$$\phi_{p(\eta)}(z) > z, \quad \forall z \in [0, 1-\alpha_\epsilon], \tag{4.23}$$

where $d_{\min} = \min\limits_{i \in \mathcal{V}} d_i$ and

$$\alpha_\epsilon = \frac{\epsilon d_{\min}}{\langle p^{(0)}, d \rangle}. \tag{4.24}$$

The *relaxed optimization problem* is consequently formulated as follows:

$$\begin{aligned} \underset{\eta\,:\,\mathbb{Z}_+ \to [0,1]^{\Omega}}{\text{minimize}} \quad & C(\eta) \\ \text{subject to} \quad & p_\omega(\eta) = \begin{cases} p_\omega^{(0)} - \eta_\omega, \text{ if } 0 < r_\omega \le k_\omega, \\ p_\omega^{(0)} + \sum\limits_{\omega' \in \Omega(\omega)} \eta_{\omega'}, \text{ if } r_\omega = 0 \end{cases}, \\ & \phi_{p(\eta)}(z) > z, \qquad \forall z \in [0, 1-\alpha_\epsilon] \end{aligned} \tag{4.25}$$

This problem is computationally more tractable than (4.22), given that α_ϵ is independent of the intervention variable η.

The subsequent result formalizes the relationship between the two problems.

Proposition 4.4.1. *Let $\mathcal{P}_\epsilon$ denote the set of feasible intervention η for problem* (4.22)*and $\mathcal{P}_{\alpha_\epsilon}$ denote the set of feasible η for problem* (4.25). *Then:*

i) $\mathcal{P}_{\alpha_\epsilon} \subseteq \mathcal{P}_\epsilon$.

ii) $\lim_{\epsilon \to 0} \mathcal{P}_\epsilon \setminus \mathcal{P}_{\alpha_\epsilon} = \emptyset$.

Proof.

i) By imposing that

$$\phi_{p(\eta)}(z) > z, \quad \forall z \in [0, 1 - \alpha_\epsilon],$$

the recursion $z(t+1) = \phi_{p(\eta)}(z(t))$ implies that there exists a finite time τ such that $z(t) > 1 - \alpha_\epsilon$ for all $t \geq \tau$, i.e., with high probability, the dynamics converge to a configuration where at most a fraction α_ϵ of links point to agents in state 0. This implies that the fraction of agents in state 0 is upper bounded by $\alpha_\epsilon \langle p, d \rangle / d_{\min} = \epsilon$. This, in turn, proves that $\psi_{p(\eta)}(z) > 1 - \epsilon$, so that

$$z > \psi_{p(\eta)}^{(-1)}(1 - \epsilon)$$

and constraint in (4.19) is satisfied.

ii) It is noted that $\psi_{p(\eta)}^{-1}(1) = 1$ for all η. By continuity of $\psi_{p(\eta)}$ and by definition of α_ϵ,

$$\psi_{p(\eta)}^{-1}(1 - \epsilon) \overset{\epsilon \to 0^+}{\longrightarrow} 1 \,, \quad 1 - \alpha_\epsilon \overset{\epsilon \to 0^+}{\longrightarrow} 1 \,,$$

so that the set of feasible interventions for (4.22)and (4.25) get arbitrarily close in the limit of vanishing ϵ. ■

While the two optimization problems (4.22) and (4.25) have the same objective function, Proposition 4.4.1 establishes that the set of feasible interventions for the former problem is a superset of the feasible interventions for the latter one, hence the solution of (4.25) is feasible for (4.22). Moreover, for small values of ϵ, the two sets of feasible interventions converge to each other. This suggests that the solutions of problem (4.25), beyond being feasible for (4.22), should also be close to optimal when ϵ vanishes. This is a key for the main result Theorem 4.4.2.

Focus is now shifted to problem (4.25), introducing a discretization to obtain a finite-dimensional approximation. Specifically, the interval $[0, 1-\alpha_\epsilon]$ is discretized into $N+1$ equally spaced points

$$z_i = \frac{(1-\alpha_\epsilon)i}{N}, \quad i = 0, 1, \cdots, N,$$

requiring that

$$\phi_{p(\eta)}(z_i) - z_i \geq \Delta_N, \quad \forall i = 0, 1, \cdots, N, \tag{4.26}$$

for a properly chosen $\Delta_N > 0$. The subsequent result relies on the regularity of $\phi_{p(\eta)}$ for every intervention η, proving instrumental in establishing the appropriate value for Δ_N.

Lemma 5. *For every feasible η,*

$$\left|\frac{d}{dz}(\phi_{p(\eta)}(z) - z)\right| \leq \frac{d_{\max} 2^{k_{\max}+1} k_{\max}}{\langle p^{(0)}, d\rangle} + 1, \quad \forall z \in [0, 1].$$

where $k_{\max} = \max_{i\in\mathcal{V}} k_i$ and $d_{\max} = \max_{i\in\mathcal{V}} d_i$.

Proof. By Proposition (4.3.2), it follows that

$$\left|\frac{d}{dz}(\phi_{p(\eta)}(z))\right| = \left|\sum_\omega \frac{d_\omega}{\langle p^{(0)}, d\rangle} p_\omega(\eta) \frac{d}{dz}\varphi_{k_\omega, r_\omega}(z)\right|. \tag{4.27}$$

Moreover,

$$\begin{aligned}
\left|\frac{d}{dz}\varphi_{k_\omega, r_\omega}(z)\right| &= \sum_{u=r_\omega}^{k_\omega} \binom{k_\omega}{u} z^{u-1}(1-z)^{k_\omega-u-1}\,|-k_\omega z + u| \\
&\leq \sum_{u=r_\omega}^{k_\omega} \binom{k_\omega}{u} |-k_\omega z + u| \leq 2\sum_{u=r_\omega}^{k_\omega} \binom{k_\omega}{u} k_{\max} \\
&\leq 2k_{\max} \sum_{u=0}^{k_\omega} \binom{k_\omega}{u} = 2k_{\max} 2^{k_{\max}},
\end{aligned}$$

where the two inequalities follow from the fact that z belongs to $[0, 1]$. By substituting this expression into (4.27), it follows that

$$\begin{aligned}
\left|\frac{d}{dz}(\phi_{p(\eta)}(z) - z)\right| &\leq \frac{2^{k_{\max}+1} k_{\max}}{\langle p^{(0)}, d\rangle} \sum_\omega d_\omega p_\omega(\eta) + 1 \\
&\leq \frac{2^{k_{\max}+1} k_{\max} d_{\max}}{\langle p^{(0)}, d\rangle} \sum_\omega p_\omega(\eta) + 1 \\
&= \frac{2^{k_{\max}+1} k_{\max} d_{\max}}{\langle p^{(0)}, d\rangle} + 1,
\end{aligned}$$

■

For a given N in $\mathbb{N}$, the *discretized optimization problem* is defined as follows

$$
\begin{aligned}
&\underset{\eta:\Omega\to[0,1]^{\Omega}}{\text{minimize}} \quad C(\eta)\\
&\text{subject to} \qquad p_\omega(\eta)=\begin{cases} p_\omega^{(0)}-\eta_\omega, \text{ if } 0<r_\omega\le k_\omega,\\ p_\omega^{(0)}+\sum\limits_{\omega'\in\Omega(\omega)}\eta_{\omega'}, \text{ if } r_\omega=0\end{cases}, \\
&\qquad\qquad\qquad \phi_{p(\eta)}(z)\ge z_i+\Delta_N, \qquad \forall i=0,1,\cdots,N,
\end{aligned} \tag{4.28}
$$

where

$$
\Delta_N=\frac{1-\alpha_\epsilon}{2N}\left(\frac{d_{\max}2^{k_{\max}+1}k_{\max}}{\langle p^{(0)},d\rangle}+1\right). \tag{4.29}
$$

The next result establishes the relation between the discretized problem (4.28) and the original problem (4.22). It states that (4.28) is a relaxation of (4.22), that is, the interventions that are feasible for (4.28) are also feasible for (4.22).

Theorem 4.4.2. *Let $\epsilon>0$ and N in $\mathbb{N}$, and let η^* and η^N be the solutions of* (4.22)*and* (4.28)*, respectively. Then:*

i) *η^N is feasible for* (4.22)*;*

ii) $\lim\limits_{\epsilon\to 0}\lim\limits_{N\to+\infty} C(\eta^N)=C(\eta^*)$.

Proof.

i) Since $\phi_{p(\eta^N)}(z)$ is differentiable, by the Mean Value Theorem, for any $z\in[z_i,z_{i+1}]$, there exists some ζ between z and z_i such that

$$
\underbrace{\phi_{p(\eta^N)}(z)-z}_{F(z)}=\underbrace{\phi_{p(\eta^N)}(z_i)-z_i}_{F(z_i)}+\underbrace{\frac{d}{dz}(\phi_{p(\eta)}(\zeta)-\zeta)}_{F'(\zeta)}(z-z_i)
$$

By taking absolute values, it follows that

$$
\phi_{p(\eta^N)}(z)-z\ge\phi_{p(\eta^N)}(z_i)-z_i-\left|\frac{d}{dz}(\phi_{p(\eta)}(\zeta)-\zeta\right||z-z_i|
$$

Because the discretization is equally spaced with $z_{i+1}-z_i=\frac{(1-\alpha_\epsilon)}{N}$, any $z\in(z_i,z_{i+1})$ satisfies the inequality $|z-z_i|<\frac{(1-\alpha_\epsilon)}{2N}$, thus,

$$
\begin{aligned}
\phi_{p(\eta^N)}(z)-z&\ge\phi_{p(\eta^N)}(z_i)-z_i-\left|\frac{d}{dz}(\phi_{p(\eta)}(\zeta)-\zeta\right||z-z_i|\\
&>\phi_{p(\eta^N)}(z_i)-z_i-\left|\frac{d}{dz}(\phi_{p(\eta)}(\zeta)-\zeta)\right|\frac{(1-\alpha_\epsilon)}{2N}.
\end{aligned}
$$

It is observed that η^N constitutes a solution to the discretized optimization problem (4.28); thus, η^N satisfies the constraints defined in (4.26),

$$\phi_{p(\eta^N)}(z_i) - z_i \geq \Delta_N, \quad \forall i = 0, \cdots, N.$$

This implies that, for every z in $[0, 1-\alpha_\epsilon]$,

$$\phi_{p(\eta^N)}(z) - z > \Delta_N - \frac{(1-\alpha_\epsilon)}{2N} \left| \frac{d}{dz} \left(\phi_{p(\eta^N)}(\zeta) - \zeta \right) \right| . \tag{4.30}$$

Plugging the explicit expression of Δ_N given in (4.29) into this equation and using Lemma 5, it is obtained that

$$\phi_{p(\eta^N)}(z) - z > 0, \quad \forall z \in [0, 1-\alpha_\epsilon].$$

Hence, η^N satisfies the relaxed constraint (4.23) and is therefore feasible for relaxed optimization problem (4.25).
Proposition 4.4.1(i) implies that η^N is also feasible for problem (4.22).

ii) Let η^{α_ϵ} be the solution of the relaxed optimization problem (4.25). Consider the case where there exists a type q in Ω with positive empirical frequency $p_q(\eta^{\alpha_\epsilon}) > 0$ and positive threshold $r_q > 0$. Now, let

$$\xi_N = \frac{\Delta_N}{\min\limits_{\zeta \in [0, \alpha_\epsilon]} a_q(\zeta)}, \tag{4.31}$$

well defined because $a_\omega(z) > 0$ for every z in $[0, \alpha_\epsilon]$ and type ω in Ω (cf (4.18)). Note now from (4.29) and (4.31) that

$$\Delta_N \overset{N \to +\infty}{\longrightarrow} 0, \qquad \xi_N \overset{N \to +\infty}{\longrightarrow} 0. \tag{4.32}$$

Let q' in Ω be the type differing from q only in the threshold value $r_{q'} = 0$. Hence, there always exists a sufficiently large N such that

$$p_\omega(\hat{\eta}) = \begin{cases} p_\omega(\eta^{\alpha_\epsilon}) - \xi_N, & \text{if } \omega = q, \\ p_\omega(\eta^{\alpha_\epsilon}) + \xi_N, & \text{if } \omega = q', \\ p_\omega(\eta^{\alpha_\epsilon}) & \text{otherwise}, \end{cases} \tag{4.33}$$

satisfies $p(\hat{\eta}) \in [0, 1]^\Omega$ and is therefore a statistics. Moreover, Proposition 4.3.2 implies that, for every z in $[0, 1-\alpha_\epsilon]$,

$$\begin{aligned} \phi_{p(\hat{\eta})}(z) &= \phi_{p(\eta^{\alpha_\epsilon})}(z) + a_q(z) \cdot \xi_N & (4.34) \\ &\geq \phi_{p(\eta^{\alpha_\epsilon})}(z) + \Delta_N \geq z + \Delta_N, & (4.35) \end{aligned}$$

where the last inequality follows from the fact that η^{α_ϵ} is feasible for problem (4.25) (in fact, it is its solution) and therefore satisfies (4.23). Eq. (4.35) implies that $\hat{\eta}$ satisfies (4.26) and is therefore feasible for problem (4.28), whose solution is η^N. Moreover, η^N is feasible for problem (4.25), whose solution is η^{α_ϵ}. Therefore,

$$C(\hat{\eta}) \geq C(\eta^N) \geq C(\eta^{\alpha_\epsilon}). \tag{4.36}$$

Now, note from (4.33) that

$$C(\hat{\eta}) = C(\eta^{\alpha_\epsilon}) + \gamma_q \xi_N \overset{N \to +\infty}{\longrightarrow} C(\eta^{\alpha}).$$

This, together with (4.37), proves that

$$C(\eta^N) \overset{N \to +\infty}{\longrightarrow} C(\eta^{\alpha_\epsilon}). \tag{4.37}$$

Finally, Proposition 4.4.1(ii) and continuity of $C(\eta)$ together prove that

$$C(\eta^{\alpha_\epsilon}) \overset{\epsilon \to 0}{\longrightarrow} C(\eta^*).$$

Hence, for this case the proof is concluded by combining this equation with (4.37).

The case is now considered wherein $p(\eta^{\alpha_\epsilon})$ is such that $p_\omega(\eta^{\alpha_\epsilon}) = 0$ for every type ω in Ω with positive threshold $r_\omega > 0$. Observe from (4.24) that $\alpha_\epsilon > 0$ because $\epsilon > 0$. Since η^{α_ϵ} is optimal for problem (4.25) and $\alpha_\epsilon > 0$, it follows that $\eta^{\alpha_\epsilon} = \eta^{(0)}$. Therefore, it follows from (4.16) that $\phi_{p^{(0)}}(z) = 1$ for every z in $[0, \alpha_\epsilon]$. This, together with (4.32) and with the fact that $C(\eta^{(0)}) = 0$, implies that, as N grows large, $\eta^N = \eta^* = \eta^{(0)}$. The proof is now concluded. ■

Theorem 4.4.2 states that the solutions of (4.28) are feasible and close to optimal for problem (4.22) as long as N grows large and ϵ vanishes. Remarkably, (4.28) is a linear program with finitely many constraints whose size scales with the number of agent types, much easier to solve than the original statistical problem (4.22), and, even more importantly, of problem (4.6).

4.5 Computational Results

In this section, numerical simulations are presented to validate the solutions proposed in Section 4.4. To this end, the topology of the social network Epinions.com[1] is utilized. The network contains $n = 26,588$ nodes and $100,120$ undirected links, so that $d_i = k_i$ for every i in $\mathcal{V}$. It is assumed that every node i has unitary cost $\gamma_i = 1$ and threshold $r_i = \lfloor \Theta_i k_i \rfloor$, with Θ_i i.i.d. random variables drawn with uniform probability from $\{0.25, 0.5, 0.75\}$. This defines a triple $(\mathcal{G}, \mathbf{r}, \gamma)$. The solution η^N of the discretized problem (4.28) is then found with $N = 10000$ and $\epsilon = 0.1$ for the statistics $p^{(0)}$ extracted from $(\mathcal{G}, \mathbf{r}, \gamma)$.

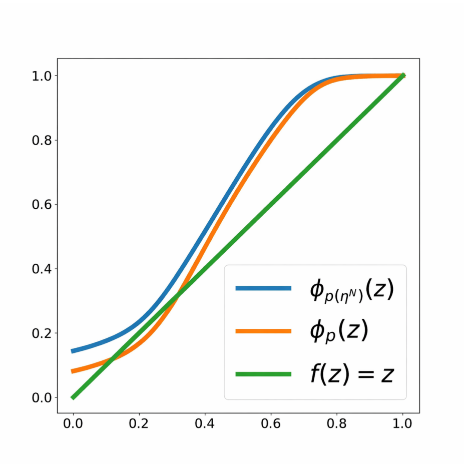

Figure 4.3: The recursion functions for the problem presented in Section .

Figure 4.3 illustrates the recursion functions $\phi_{p^{(0)}}$ and $\phi_{p(\eta^N)}$ for the original statistics $p^{(0)}$ and for the seeded statistics $p(\eta^N)$. Note that $\phi_{p(\eta^N)}(z) > z$ for all z in $[0, 1]$, so that the dynamics for a game sampled from $\mathcal{C}_{n,p(\eta^N)}$ is expected to converge to $\mathbf{x} = \mathbf{1}$ for every initial condition.

To validate this, two instances drawn from $\mathcal{C}_{n,p^{(0)}}$ and $\mathcal{C}_{n,p(\eta^N)}$ are generated respectively, and the linear threshold dynamics are simulated with initial condition $y(0) = 0$. Figure 4.4 (*left*) illustrates the dynamical behaviour of the average state $Y(t)$ and the one predicted by the recursion (4.20), denoted $y(t)$, with initial conditions $Y(0) = y(0) = 0$. As expected, since the order of the network is large, the two are very close to each other. Moreover, the average state converges to 1 under statistics $p(\eta^N)$, as predicted by Figure 4.3.

[1]Retrieved from https://networkrepository.com

A seeding is then applied, consistent with η^N to the Epinions network. Despite the lack of theoretical guarantees for this network, the right panel of Figure 4.4 shows that the dynamics converges to a configuration with almost all ones. This suggests that designing the optimal seeding based on the problem statistics is a valid approach even for real networks that are not generated from a configuration model ensemble.

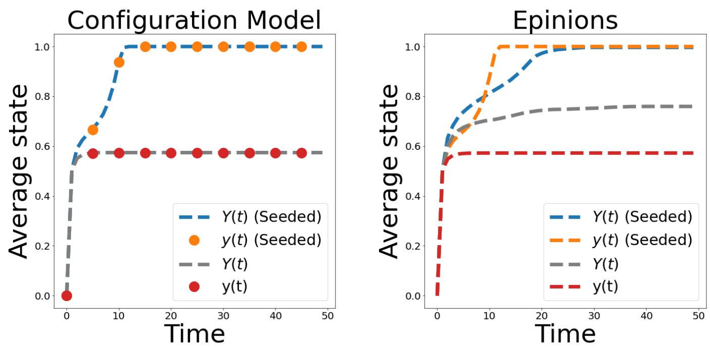

Figure 4.4: The TM dynamics on the Epinions.com network (*right*) and on the Configuration model with same statistics (*left*). The plot illustrates the fraction of state-1 adopters $Y(t)$ for the original and the seeded statistics p, η^N, compared with the outputs $y(t)$ of the recursion (4.20).

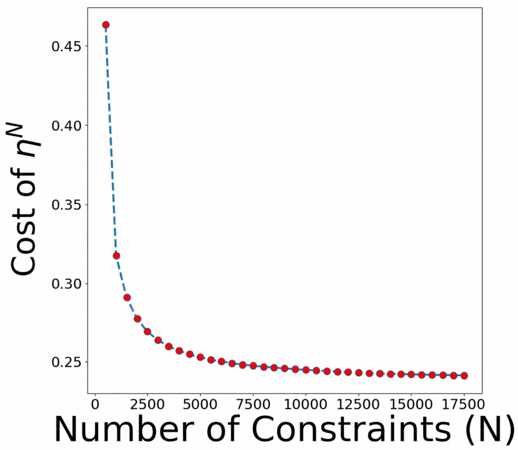

Figure 4.5: $C(\eta^N)$ as a function of N for the statistics p' in Section 4.5.

Finally, to investigate the sensitivity of η^N with respect to N, Problem (4.28) is solved with $\epsilon = 0.1$ for different values of N for a given statistics $p'_{2,3,0,1} = p'_{2,3,1,1} = p'_{4,3,0,1} = p'_{4,3,1,1} = 1/12$, $p'_{2,3,2,1} = p'_{4,3,2,1} = 1/3$. Figure 4.5 illustrates the total cost of η^N as a function of N. However, notice that N, and therefore the complexity of (4.28), does not scale with n.

Chapter 5

Optimal Intervention in Large-Scale Super-Modular Network Games

In this chapter, the problem of optimal intervention in diffusion processes governed by the linear threshold model is investigated, with a specific focus on the Least-Cost Influence Problem (LCIP). Rather than restricting interventions to binary seeding decisions, a more flexible framework is considered wherein fractional incentives can be assigned to individuals by a central planner. These incentives effectively reduce agents' adoption thresholds, thereby capturing more realistic policy instruments such as targeted subsidies, discounts, or information campaigns. The underlying objective is to minimize the total intervention cost while ensuring a prescribed level of diffusion (i.e., maximizing the spread of adoption). The scalability of solutions to these problems remains a fundamental challenge in social and economic networks, where the number of agents can be extremely large and only aggregate information may be available. This makes classical combinatorial or network-level optimization approaches computationally infeasible or practically inapplicable. To overcome inherent scalability limitations, and in continuity with the approach developed in Chapter 4, the LCIP is studied on general directed multigraphs utilizing a formulation specifically tailored for large-scale random networks. This methodology builds upon analytical tools derived from the study of diffusion processes on random graphs, specifically configuration models. By leveraging mean-field approximations, the diffusion dynamics are described at a statistical level, thereby characterizing aggregate adoption behavior in terms of degree distributions and threshold heterogeneity. This perspective enables a significant simplification of the problem. In particular, the LCIP is recast as a linear program whose dimension is independent of the size of the underlying network.

5.1 Intervention Dynamics

Given a finite directed multi-graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$ nd a threshold vector $\mathbf{r}$, the *linear threshold model* introduced in Appendix B is recalled:

$$\mathbf{x}(t+1) = \Phi_{\mathbf{r}}(\mathbf{x}(t)), \qquad \forall t \geq 0,$$

To influence the dynamics, a central planner can assign an intervention $\mathbf{h} \in \mathbb{Z}_+^{\mathcal{V}}$ that effectively decreases the agents' adoption thresholds. The *linear threshold model with intervention* is thus naturally defined as:

$$\mathbf{x}(t+1) = \Phi_{\mathbf{r}-\mathbf{h}}(\mathbf{x}(t)), \qquad \forall t \geq 0. \tag{5.1}$$

Implementing an intervention h_i is equivalent to decreasing the threshold of player i by h_i units. Because thresholds must remain non-negative, an intervention is feasible if:

$$0 \leq h_i \leq r_i, \qquad \forall i \in \mathcal{V}.$$

The implementation of an intervention vector $\mathbf{h}$ is associated with a separable cost functional, defined as $C(\mathbf{h}) = \sum_{i \in \mathcal{V}} \gamma_i(h_i)$, where $\gamma_i : \mathbb{Z}_+ \to \mathbb{R}_+$ denotes the individual cost function for node i. It is assumed that γ_i is monotonically non-decreasing and $\gamma_i(0) = 0$ for all $i \in \mathcal{V}$. The separable structure of the cost function is strictly required to preserve the linearity of the objective function in the subsequent macroscopic formulation (Problem 5.13); without it, the mean-field approximation would yield a non-linear optimization problem.

The objective of the planner is to design a minimal-cost intervention such that, for a predefined tolerance value ϵ in $(0, 1]$, the fraction of active agents is at least $1 - \epsilon$ after n iterations, regardless of the initial condition. This is formally expressed as:

$$\frac{1}{n} \sum_{i \in \mathcal{V}} x_i^{\mathbf{h}}(t) \geq 1 - \epsilon, \qquad \forall\, t \geq n, \tag{5.2}$$

for every $\mathbf{x}(0) \in \mathcal{A}$.

Here, $x_i^{\mathbf{h}}(t) \in \{0, 1\}$ denotes the binary state of agent i at time step t under the intervention vector $\mathbf{h}$. The network configuration vector $\mathbf{x}^{\mathbf{h}}(t) = (x_1^{\mathbf{h}}(t), \ldots, x_n^{\mathbf{h}}(t))$ is generated by the iterative application of the modified Linear Threshold Model operator $\Phi_{\mathbf{r}-\mathbf{h}}$, such that $\mathbf{x}^{\mathbf{h}}(t) = (\Phi_{\mathbf{r}-\mathbf{h}})^t(\mathbf{x}(0))$. Therefore, $x_i^{\mathbf{h}}(t)$ represents the deterministic trajectory of the i-th node's activation status when the network dynamics are governed by the effectively reduced thresholds $r_i - h_i$.

Remark 12. *By* (5.1) *and the monotonicity property established in Lemma 7, requiring* (5.2) *to hold for every initial condition is equivalent to:*

$$\frac{1}{n}\sum_{i\in\mathcal{V}}((\Phi_{\mathbf{r}-\mathbf{h}})^n(\mathbf{0}))_i \geq 1-\epsilon. \tag{5.3}$$

The optimization problem addressed by the planner is thus formulated as:

$$\begin{array}{ll} \underset{\mathbf{h}\in\mathbb{Z}_+^{\mathcal{V}}}{\text{minimize}} & \sum_{i\in\mathcal{V}}\gamma_i(h_i) \\ \text{subject to} & \frac{1}{n}\sum_{i\in\mathcal{V}} x_i^{\mathbf{h}}(t) \geq 1-\epsilon \end{array} \tag{5.4}$$

In the literature, this objective is known as the Least-Cost Influence Problem (LCIP). By allowing integer incentives h_i that lower individual adoption thresholds r_i, the LCIP generalizes binary seed selection problems. The target class of games Γ^* is defined as the family of modified games $\Gamma(\mathbf{h})$ that satisfy the macroscopic diffusion bound in Equation 5.2. Constraining the induced game to Γ^* mathematically guarantees that the system achieves an aggregate activation fraction of at least $1-\epsilon$ universally across all initial configurations.

Remark 13. *Problem* (5.4) *constitutes a variant of the LCIP evaluated on graphs with integer weights. The general LCIP is known to be NP-complete even in simplified settings (see Günneç et al., 2020b). If the cost functions are defined as* $\gamma_i(h_i) = c_i$ *for a positive constant* c_i *for all* $h_i > 0$*, the optimal strategy restricts interventions to either* $h_i = 0$ *or* $h_i = r_i$ *(so that necessarily* $x_i(t) = 1$ *for every positive time* t*). In this scenario, the framework reduces to the optimal seeding problem analyzed in Chapter 4. Furthermore, if* $c_i = 1$ *for all* i *in* $\mathcal{V}$*, the objective simplifies to identifying a minimal-cardinality subset of agents whose forced activation ensures the fractional adoption bound* $1-\epsilon$ *is met for all* $t \geq n$*.*

Optimal Intervention in BASASM Games

The Linear Threshold Model can be interpreted as the best response dynamics in so-called semi anonymous network games with strategic complements, whereby the agents $i \in \mathcal{V}$ have utilities that are increasing functions of the number of their out-neighbors taking the same action. Given a network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, BASASM games are considered $\Gamma = (\mathcal{V}, \{\mathcal{A}_i\}_{i\in\mathcal{V}}, \{u_i\}_{i\in\mathcal{V}})$ with utility function

$$u_i\left(x_i, \mathbf{x}_{-i}\right) = \mathfrak{u}_i\left(x_i, \sum_{j\in\mathcal{V}} W_{i,j} x_j\right), \qquad \forall\, \mathbf{x} \in \mathcal{A} \tag{5.5}$$

Consider intervention $\mathbf{h} \in \mathcal{H} = \mathbb{R}^V$ that modifies utility

$$u_i^{\mathbf{h}}(x_i, \mathbf{x}_{-i}) = \mathfrak{u}_i\left(x_i, \sum_{j\in\mathcal{V}} W_{i,j} x_j + h_i\right)$$

Remark 14. *Intervening on player i with intervention h_i effectively amounts to requiring $\sum_j W_{i,j} x_j \geq r_i - h_i$; i.e., decreasing the threshold of node i by h_i.*
Since the action of the agents $\mathbf{x} \in \{0,1\}^{\mathcal{V}}$ and the threshold r_i are piecewise constant, any vector whose components take on values $h_i = 0, \ldots, r_i$, $\forall i \in \mathcal{V}$ is considered, without loss of generality, a feasible intervention.

Example 5.1.1. *Consider* network coordination game (2.3) *on a weighted graph $\mathcal{G}$ where agents i choose their binary action $x_i \in \{0,1\}$ and utilities*

$$u_i(x_i, \mathbf{x}_{-i}) = \sum_{j\in\mathcal{V}} W_{i,j}\left(\left(1 - x_i\right)\left(1 - x_j\right) + x_i x_j\right) + b_i x_i$$

A parameter $\mathbf{h} \in \mathbb{R}_+^{\mathcal{V}}$, corresponding to the game intervention, is selected by the central regulator. Let the game

$$\Gamma(\mathbf{h}) = \left(\mathcal{V}, \left\{\mathcal{A}_i(\mathbf{h})\right\}_{i\in\mathcal{V}}, \left\{u_i^{\mathbf{h}}\right\}_{i\in\mathcal{V}}\right)$$

be considered, with no restriction on the action sets, i.e., $\mathcal{A}_i(\mathbf{h}) = \mathcal{A}_i$, for all $i \in \mathcal{V}$ and utility function

$$\begin{aligned} u_i^{\mathbf{h}}(x_i, \mathbf{x}_{-i}) &= u_i(\mathbf{x}) + h_i x_i = \\ &= \sum_{j\in\mathcal{V}} W_{i,j}\left(\left(1 - x_i\right)\left(1 - x_j\right) + x_i x_j\right) + (b_i + h_i) x_i \end{aligned}$$

Lemma 6. *Consider an initial configuration $x(0) \in \mathcal{A} = \{0,1\}^{\mathcal{V}}$, the collection of threshold sequences $\mathbf{r} \in \mathbb{Z}_+^{\mathcal{V}}$ and the collection of payoff parameter sequences* $\mathbf{h}$. *The linear threshold model and the deterministic maximal best response dynamic are equivalent (describe the same dynamic $x(t) \in \mathcal{A}$) if for every i the sequences* $\mathbf{r}$ *and* $\mathbf{h}$ *are such that*

$$r_i = \left\lceil \frac{k_i - b_i - h_i}{2} \right\rceil .$$

Proof. In the deterministic maximal best response dynamic $x_i(t+1) = 1$ if and only if the argmax in gives either $\{1\}$ or $\{0,1\}$, which happens if

$$\begin{aligned} \mathfrak{u}_i\left(1, \sum_{j\in\mathcal{V}} W_{i,j}x_j + \mathbf{h}\right) \geq \mathfrak{u}_i\left(0, \sum_{j\in\mathcal{V}} W_{i,j}x_j + \mathbf{h}\right) &\iff \\ \sum_{j\in\mathcal{V}} W_{i,j}x_j + b_i + h_i \geq \sum_{j\in\mathcal{V}} W_{i,j}(1-x_j) &\iff \\ \sum_{j\in\mathcal{V}} W_{i,j}x_j + b_i + h_i \geq k_i - \sum_{j\in\mathcal{V}} W_{i,j}x_j &\iff \\ \sum_{j\in\mathcal{V}} W_{i,j}x_j \geq \frac{k_i - b_i - h_i}{2}&. \end{aligned}$$

Since $\sum_{j\in\mathcal{V}} W_{i,j}x_j$ is integer, then $\sum_{j\in\mathcal{V}} W_{i,j}x_j \geq \left\lceil \frac{k_i-b_i-h_i}{2} \right\rceil$. ■

The linear threshold model is slightly more flexible than the deterministic best response with parameters $h_i \in [-k_i, k_i]$. Each threshold r_i in the LTM belongs to $\{0, 1, \cdots, k_i+1\}$: thresholds are allowed to be explicitly 0 or k_i+1, corresponding to stubborn behaviors. The "ties resolution" of the deterministic best response, maps $b_i = k_i$ to $r_i = 0$, but there is no counterpart for $k_i + 1$.

5.2 Mean-Field Approximations

To solve problem (5.4), the same methodological approach previously adopted in Chapter 4 to address the optimal seeding problem in super-modular games is utilized. Specifically, the linear threshold model is analyzed on a large random network comprising heterogeneous agents. A vector-valued function $\eta : \mathbb{Z}_+ \to [0,1]^\Omega$ is associated with every intervention $\mathbf{h}$, where $\eta_\omega(\nu)$ denotes the fraction of agents whose initial type is ω and whose threshold is decreased by ν units due to the intervention.

Consistent with this interpretation, η must satisfy the following condition for every type $\omega \in \Omega$;

$$\begin{cases} \eta_\omega(\nu) = 0, & \forall \nu > r_\omega, \\ \sum_{\nu=0}^{r_\omega} \eta_\omega(\nu) = p_\omega^{(0)}. \end{cases} \tag{5.6}$$

Furthermore, given that the number of nodes of type ω experiencing a threshold decrease of ν units must strictly be an integer, η is required to satisfy:

$$n\eta_\omega(\nu) \in \mathbb{Z}_+, \qquad \forall \nu = 0, 1, \cdots, r_\omega, \ \forall \omega \in \Omega. \tag{5.7}$$

A pair (n, η) is deemed feasible if conditions (5.6) - (5.7) are satisfied.

Given an instance of the Optimal Intervention Problem (5.4), specified by the triple $(\mathcal{G}, r, \gamma)$ of the random network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, the threshold vector $\mathbf{r}$, and of the cost functions γ the initial statistics vector $p^{(0)}$ in $[0,1]^\Omega$ is extracted as follows:

$$p_\omega^{(0)} = \frac{1}{n} |\{i \in \mathcal{V} \, : \, \omega_i = \omega\}|, \qquad \forall \omega \in \Omega.$$

Through the reduction of agent thresholds, a new network statistic $p(\eta)$ is generated by the intervention, defined as:

$$p_\omega(\eta) = p_\omega^{(0)} + \sum_{\omega' \in \Omega(\omega)} \eta_{\omega'}(r_{\omega'} - r_\omega) - \sum_{\nu=1}^{r_\omega} \eta_\omega(\nu), \tag{5.8}$$

for every ω in Ω. Specifically, the second term accounts for agents of types ω' with $r_{\omega'} > r_\omega$ whose threshold is reduced by $r_{\omega'} - r_\omega$ units, whereas the final term represents the fraction of agents of type ω whose threshold is reduced, resulting in a change of their respective types.

Given a statistic $p(\eta)$ and a finite set of nodes $\mathcal{V}$ of cardinality $n = |\mathcal{V}|$, a vector of types ω in $\Omega^\mathcal{V}$ with statistic $p(\eta)$ exists if and only if

$$np_\omega(\eta) \in \mathbb{Z}_+, \qquad \forall \omega \in \Omega. \tag{5.9}$$

Remark 15. *The intervention η, is implemented via a probabilistic approach wherein every node $i \in \mathcal{V}$ of type ω_i, transitions with probability $\eta_{\omega_i}(\nu)/p_{\omega_i}$ to a new type maintaining the identical in-degree d_i and out-degree k_i but featuring a modified threshold $r_i - \nu$, with $\nu = 0, 1, \cdots, r_\omega$. This effectively maps the node to a set of types sharing all parameters with ω_i, excluding the threshold.*

Proposition 5.2.1. *For every intervention $\eta : \mathbb{Z}_+ \to [0,1]^{\mathcal{V}}$ satisfying the feasibility conditions* (5.6)*, the perturbed recursion function evaluates to:*

$$\phi_{p(\eta)}(z) = \phi_{p^{(0)}}(z) + \sum_{\omega \in \Omega} \sum_{\nu=1}^{r_\omega} a_\omega(\nu, z) \eta_\omega(\nu), \tag{5.10}$$

where the linear coefficients are given by:

$$a_\omega(\nu, z) = \frac{d_\omega \left(\varphi_{k_\omega, (r_\omega - \nu)}(z) - \varphi_{k_\omega, r_\omega}(z) \right)}{\langle p^{(0)}, d \rangle} \geq 0. \tag{5.11}$$

Proof. The Local Mean-Field recursive function governing the fraction of active edges under an intervention η is initially established as:

$$\phi_{p(\eta)}(z) = \frac{1}{\langle p(\eta), d \rangle} \sum_{\omega \in \Omega} p_\omega(\eta) d_\omega \varphi_{k_\omega, r_\omega}(z) \tag{5.12}$$

It is first observed that the intervention operates strictly by modifying the agents' activation thresholds without altering the underlying network topology. Consequently, the expected in-degree of the network is invariant under the intervention, implying that $\langle p(\eta), d \rangle = \langle p^{(0)}, d \rangle$, (Remark 11).

Substituting the explicitly defined post-intervention statistic $p_\omega(\eta)$ from Equation (5.8) into the recursive function, the expression is expanded into three distinct summation terms:

$$\phi_{p(\eta)}(z) = \frac{1}{\langle p(\eta), d \rangle} \sum_{\omega \in \Omega} \left[p_\omega^{(0)} + \sum_{\omega' \in \Omega(\omega)} \eta_{\omega'}(r_{\omega'} - r_\omega) - \sum_{\nu=1}^{r_\omega} \eta_\omega(\nu) \right] d_\omega \varphi_{k_\omega, r_\omega}(z)$$

The first term corresponds to the pre-intervention network statistic and is trivially identified as the baseline recursive function:

$$\frac{1}{\langle p^{(0)}, d \rangle} \sum_{\omega \in \Omega} p_\omega^{(0)} d_\omega \varphi_{k_\omega, r_\omega}(z) = \phi_{p^{(0)}}(z)$$

The second term represents the probability mass transferred into type ω from types ω' with strictly higher thresholds, and requires careful re-indexing to expose its structural form:

$$\frac{1}{\langle p^{(0)}, d\rangle} \sum_{\omega\in\Omega} \sum_{\omega'\in\Omega(\omega)} \eta_{\omega'}(r_{\omega'} - r_\omega) d_\omega \varphi_{k_\omega, r_\omega}(z)$$

By the definition of the set $\Omega(\omega)$, the condition $\omega' \in \Omega(\omega)$ implies that the type ω' is identical to ω in all attributes $(d_\omega = d_{\omega'},\ k_\omega = k_{\omega'},\ \gamma_\omega = \gamma_{\omega'})$ except that its threshold is strictly greater $(r_{\omega'} > r_\omega)$.

Let the magnitude of the threshold reduction be defined as $\nu = r_{\omega'} - r_\omega$. Since both thresholds are non-negative integers, ν takes values in the range $\{1, \ldots, r_{\omega'}\}$. The order of summation can thus be exchanged by fixing the source type ω' as the outer index and summing over all possible reduction magnitudes ν. Under this transformation, the properties of $\Omega(\omega)$ allow the substitution $d_\omega = d_{\omega'},\ k_\omega = k_{\omega'}$, while the threshold index becomes $r_\omega = r_{\omega'} - \nu$:

$$\frac{1}{\langle p^{(0)}, d\rangle} \sum_{\omega\in\Omega} \sum_{\nu=1}^{r_\omega} \eta_\omega(\nu) d_\omega \varphi_{k_\omega, r_\omega - \nu}(z)$$

The third term represents the probability mass exiting type ω due to threshold reductions:

$$-\frac{1}{\langle p^{(0)}, d\rangle} \sum_{\omega\in\Omega} \sum_{\nu=1}^{r_\omega} \eta_\omega(\nu) d_\omega \varphi_{k_\omega, r_\omega}(z)$$

By aggregating all contributions yields (5.12). This confirms the specified linear structure, where the sensitivity coefficient is identified as:

$$a_\omega(\nu, z) = \frac{d_\omega \left(\varphi_{k_\omega,(r_\omega-\nu)}(z) - \varphi_{k_\omega, r_\omega}(z)\right)}{\langle p^{(0)}, d\rangle}$$

To conclude the proof, the non-negativity of $a_\omega(\nu, z)$ must be established. The function $\varphi_{k,r}(z)$ defines the cumulative probability of obtaining at least r successes in a binomial distribution with k trials and success probability z. Because $\nu \geq 1$, the requirement of achieving at least $r_\omega - \nu$ successes is a strictly weaker condition than achieving at least r_ω successes. Consequently, the cumulative probability must satisfy $\varphi_{k_\omega,(r_\omega-\nu)}(z) - \varphi_{k_\omega, r_\omega}(z)$ for all $z \in [0, 1]$. Given that expected degrees are strictly non-negative, it necessarily follows that $a_\omega(\nu, z) \geq 0$. ∎

Remark 16. *Notice that this is a generalization of Proposition 4.3.2, indeed for the special case of seeding problem optimal intervention on each agent is either not to intervene, or to set the threshold to* 0.

Interventions, thus, satisfy $\eta_\omega(\nu) = 0$ *for each* $\nu = 1, \dots, r_\omega - 1$. *So the linear coefficient* (5.11)*, restrict to be*

$$
\begin{aligned}
a_\omega(r_\omega, z) &= \frac{d_\omega \left(\varphi_{k_\omega,(r_\omega - r_\omega)}(z) - \varphi_{k_\omega, r_\omega}(z)\right)}{\langle p^{(0)}, d\rangle} \\
&= \frac{d_\omega \left(\varphi_{k_\omega,0}(z) - \varphi_{k_\omega, r_\omega}(z)\right)}{\langle p^{(0)}, d\rangle} \\
&= \frac{d_\omega \left(1 - \varphi_{k_\omega, r_\omega}(z)\right)}{\langle p^{(0)}, d\rangle}
\end{aligned}
$$

that is $a_\omega(z)$ *in the expression* (4.18).

Mean-Field Optimization

In this subsection, the optimization problem (5.4) is reformulated within the statistical framework. Within this paradigm, the planner's objective is to identify the minimal-cost intervention η satisfying (4.19), such that Proposition 4.3.3 guarantees that, in almost all large-scale networks with statistic $p(\eta)$ at least a fraction $1 - \epsilon$ nodes will converge to state 1. To formalize the problem, the cost of an intervention must be explicitly defined. Given that $\gamma_\omega(\nu)$ represents the cost of decreasing the threshold of an agent of type ω by ν, it follows from the definition of η that the number of nodes of type ω undergoing a threshold reduction of ν units is $n\eta_\omega(\nu)$. Therefore, a cost is associated with the intervention η as follows:

$$
C(\eta) = \sum_{\omega \in \Omega} \sum_{\nu=0}^{r_\omega} \eta_\omega(\nu)\gamma_\omega(\nu) \geq 0,
$$

with $C(\eta^{(0)}) = 0$. The optimization problem is subsequently reformulated as:

$$
\begin{aligned}
\underset{\eta : \mathbb{Z}_+ \to [0,1]^\Omega}{\text{minimize}} \quad & C(\eta) \\
\text{subject to} \quad & \sum_{\nu=0}^{r_\omega} \eta_\omega(\nu) = p_\omega^{(0)}, \qquad \forall \omega \in \Omega, \\
& \phi_{p(\eta)}(z) > z, \qquad \forall z \in \left[0, \psi_{p(\eta)}^{-1}(1-\epsilon)\right].
\end{aligned}
\tag{5.13}
$$

5.3 Approximation Methods

While the original LCIP (5.4) is a non-linear program whose size scales with the number of agents, Proposition 5.2.1 states that $\phi_{p(\eta)}$ is linear in η. Consequently, the constraint

$$\phi_{p(\eta)}(z) > z, \qquad \forall z \in \left[0, \psi^{-1}_{p(\eta)}(1-\epsilon)\right]$$

is linear in η, rendering the local mean-field reformulation (5.13) a linear program whose dimensionality scales purely with the number of types.

Remark 17. *The local mean-field reformulation of the WTSS proposed in Chapter 4 is also a linear program, but the two problems have different sizes, since in* (5.13) *the planner chooses for every type $\omega \in \Omega$ the fraction of agents whose threshold is decreased by ν units, with $\nu = 0, \ldots, r_\omega$, while in the WTSS $\nu \in \{0, r_\omega\}$. The advantage of the LCIP is that the objective of steering the network towards a desirable configuration can be achieved with considerably smaller budget, as the example in Section 5.4 illustrates.*

Therefore, in contrast with the original problem (5.4), which is a non-linear program whose size scales with the number of agents, problem (5.13) is a linear program whose size scales with the number of types.
An alternative optimization problem is defined, which is subsequently proven to constitute a relaxation of the original formulation. Consider the constraint

$$\phi_{p(\eta)}(z) > z, \quad \forall z \in [0, 1-\alpha_\epsilon], \tag{5.14}$$

where $d_{\min} = \min_{i \in \mathcal{V}} d_i$ and

$$\alpha_\epsilon = \frac{\epsilon d_{\min}}{\langle p^{(0)}, d \rangle}.$$

Now, consider the optimization problem

$$\begin{aligned} \underset{\eta : \mathbb{Z}_+ \to [0,1]^\Omega}{\text{minimize}} \quad & C(\eta) \\ \text{subject to} \quad & \sum_{\nu=0}^{r_\omega} \eta_\omega(\nu) = p^{(0)}_\omega, \qquad \forall \omega \in \Omega, \\ & \phi_{p(\eta)}(z) > z, \quad \forall z \in [0, 1-\alpha_\epsilon] \end{aligned} \tag{5.15}$$

This problem is easier to solve than (5.13) since α_ϵ does not depend on η. The next result illustrates the relation between the two problems.

Proposition 5.3.1. *Let $\mathcal{P}_\epsilon$ denote the set of feasible η for problem* (5.13) *and $\mathcal{P}_{\alpha_\epsilon}$ denote the set of feasible η for problem* (5.15)*. Then:*

i) $\mathcal{P}_{\alpha_\epsilon} \subseteq \mathcal{P}_\epsilon$.

ii) $\lim_{\epsilon \to 0} \mathcal{P}_\epsilon \setminus \mathcal{P}_{\alpha_\epsilon} = \emptyset$.

Proof. The proof is analogous to that of Proposition (4.4.1), since the higher dimensionality of the intervention variables does not affect the properties and dynamic convergence of the recursion functions

$$y(t+1) = \psi_{p(\eta)}(z(t)), \quad z(t+1) = \phi_{p(\eta)}(z(t))$$

but impacts only the speed of convergence and associated cost (Remark 17). ■

By discretizing the interval $[0, 1-\alpha_\epsilon]$ into $N+1$ equally spaced points

$$z_i = \frac{(1-\alpha_\epsilon)i}{N}, \quad i = 0, 1, \cdots, N,$$

and requiring that

$$\phi_{p(\eta)}(z_i) - z_i \geq \Delta_N, \quad \forall i = 0, 1, \cdots, N, \tag{5.16}$$

the discretized version of Problem (5.15) is defined,

$$\begin{aligned} \underset{\eta : \mathbb{Z}_+ \to [0,1]^\Omega}{\text{minimize}} \quad & C(\eta) \\ \text{subject to} \quad & \sum_{\nu=0}^{r_\omega} \eta_\omega(\nu) = p_\omega^{(0)}, & \forall \omega \in \Omega, \\ & \phi_{p(\eta)}(z) \geq z_i + \Delta_N, & \forall i = 0, 1, \cdots, N, \end{aligned} \tag{5.17}$$

retrieving that it is a relaxation of (5.13), alongside feasibility guarantees of its solutions and closeness to optimality as $\epsilon \to 0$.

Theorem 5.3.2. *Let $\epsilon > 0$ and N in $\mathbb{N}$, and let η^* and η^N be the solutions of* (5.13) *and* (5.17)*, respectively. Then:*

i) η^N *is feasible for* (5.13)*;*

ii) $\lim_{\epsilon \to 0} \lim_{N \to +\infty} C(\eta^N) = C(\eta^*)$.

Proof.

i) The proof is identical to that of Theorem 4.4.2(i), leveraging the structural properties of $\phi_{p(\eta)}(z)$, and is therefore omitted.

ii) Let η^{α_ϵ} be the solution of problem (5.15). Consider the scenario in which there exists a type q in Ω characterized by a positive empirical frequency $p_q(\eta^{\alpha_\epsilon}) > 0$ and a strictly positive threshold $r_q > 0$. Now, let

$$\beta_N = \frac{\Delta_N}{\min\limits_{\zeta \in [0,\alpha_\epsilon]} a_q(r_q, \zeta)}, \tag{5.18}$$

well defined because $a_\omega(\nu, z) > 0$ for every z in $[0, \alpha_\epsilon]$, $\nu > 0$ and type ω in Ω (cf (5.11)). Note now from (5.18) that

$$\Delta_N \overset{N\to+\infty}{\longrightarrow} 0, \qquad \beta_N \overset{N\to+\infty}{\longrightarrow} 0. \tag{5.19}$$

Let q' in Ω be the type differing from q only in the threshold value $r_{q'} = 0$. Hence, there always exists a sufficiently large N such that

$$p_\omega(\hat{\eta}) = \begin{cases} p_\omega(\eta^{\alpha_\epsilon}) - \beta_N, & \text{if } \omega = q, \\ p_\omega(\eta^{\alpha_\epsilon}) + \beta_N, & \text{if } \omega = q', \\ p_\omega(\eta^{\alpha_\epsilon}) & \text{otherwise}, \end{cases} \tag{5.20}$$

satisfies $p(\hat{\eta}) \in [0,1]^\Omega$ and is therefore a statistics.

Moreover, Proposition 5.2.1 implies that, for every z in $[0, 1-\alpha_\epsilon]$,

$$\begin{aligned} \phi_{p(\hat{\eta})}(z) &= \phi_{p(\eta^{\alpha_\epsilon})}(z) + a_q(r_q, z) \cdot \beta_N \\ &\geq \phi_{p(\eta^{\alpha_\epsilon})}(z) + \Delta_N \geq z + \Delta_N, \end{aligned} \tag{5.21}$$

where the last inequality follows from the fact that η^{α_ϵ} is feasible for problem (5.15) (in fact, it is its solution) and therefore satisfies (5.14). Eq. (5.21) implies that $\hat{\eta}$ satisfies (5.16) and is therefore feasible for problem (5.17), whose solution is η^N. Moreover, η^N is feasible for problem (5.15), whose solution is η^{α_ϵ}. Therefore,

$$C(\hat{\eta}) \geq C(\eta^N) \geq C(\eta^{\alpha_\epsilon}). \tag{5.22}$$

Now, note from (5.20) that

$$C(\hat{\eta}) = C(\eta^{\alpha_\epsilon}) + c_q(\omega_q)\beta_N \overset{N\to+\infty}{\longrightarrow} C(\eta^\alpha).$$

This proves that

$$C(\eta^N) \overset{N\to+\infty}{\longrightarrow} C(\eta^{\alpha_\epsilon}). \tag{5.23}$$

Finally, Proposition 5.3.1(ii) and continuity of $C(\eta)$ together prove that

$$C(\eta^{\alpha_\epsilon}) \overset{\epsilon\to 0}{\longrightarrow} C(\eta^*).$$

Hence, for this case the proof is concluded.

The case is now considered wherein $p(\eta^{\alpha_\epsilon})$ is such that $p_\omega(\eta^{\alpha_\epsilon}) = 0$ for every type ω in Ω with positive threshold $r_\omega > 0$.
Observe that $\alpha_\epsilon > 0$ because $\epsilon > 0$.

Since η^{α_ϵ} is optimal for problem (5.15) and $\alpha_\epsilon > 0$, it follows that $\eta^{\alpha_\epsilon} = \eta^{(0)}$. Therefore, it follows from (4.16) that $\phi_{p^{(0)}}(z) = 1$ for every z in $[0, \alpha_\epsilon]$. This, together with (5.19) and with the fact that $C(\eta^{(0)}) = 0$, implies that, as N grows large, $\eta^N = \eta^* = \eta^{(0)}$. The proof is now concluded. ■

Theorem 5.3.2 states that the solutions of (5.17) are feasible and close to optimal for problem (5.13) as long as N grows large and ϵ vanishes. Remarkably, (5.17) is a linear program with finitely many constraints whose size scales with the number of agent types, much easier to solve than the original statistical problem (5.13), and, even more importantly, of problem (5.4).

5.4 Computational Results

In this section, numerical experiments on the real-world networks Epinions[1] and Power Grid Network[2] are presented.

The network Epinions contains 26588 nodes and 100120 undirected links. For every type ω in Ω, the threshold is set to $r_\omega = \lfloor k_\omega/2 \rfloor$ and cost $\gamma_\omega(\eta) = \eta$. This defines a triple $(\mathcal{G}, \rho, \gamma)$. The statistics $p^{(0)}$ is extracted, and the optimal statistical intervention η^N is derived by solving problem (5.17) with $\epsilon = 0.1$, $N = 100$ and $\Delta = 0.05$ and the optimal statistical intervention $\eta^{N,s}$ for the related WTSS problem studied in Chapter 4. Such problem is obtained by setting $\gamma_\omega(0) = 0$ and $\gamma_\omega(\eta) = r_\omega$ for $\nu > 0$, so that for the planner it is optimal either not to intervene on an agent or to set the threshold to 0. A random intervention complying with the optimal statistical one is then implemented, and a comparison is made between the fraction of state-1 agents observed on the network under such interventions and the fraction of state-1 agents predicted by recursion (4.20).

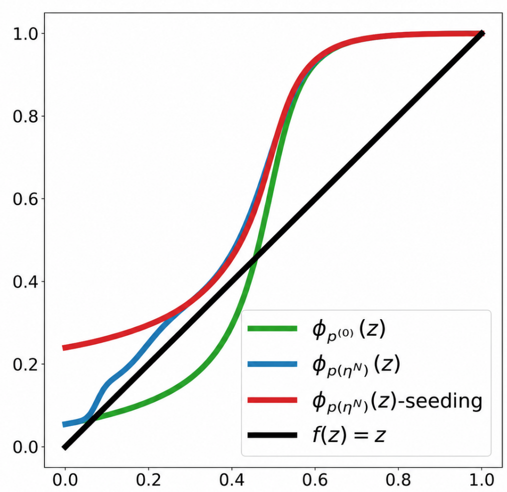

Figure 5.1: Evolution function for the fraction of links pointing to state-1 agents with initial statistics $p^{(0)}$, with optimal intervention statistics $p(\eta^N)$, and with optimal seeding $p(\eta^N)$-seeding, for the Epinions network.

[1]Retrieved from https://networkrepository.com
[2]Retrieved from https://websites.umich.edu/∼ mejn/netdata/

Fig. 5.1 illustrates the recursion function driving the evolution of the fraction of links pointing to state-1 agents in the local mean-field approximation with the initial statistics $p^{(0)}$, with statistics $p(\eta^N)$ corresponding to the optimal intervention, and with the statistics $p(\eta_s^N)$ corresponding to the optimal seeding.

Intuitively, the seeding is suboptimal whenever the curve $\phi_{p^{(0)}}(z)$ is above z for small values of z and below z for large values of z. While in such a case the intervention η^N allows to target the curve in precise regions of the domain, seeding interventions are forced to increase the entire curve, especially for small values of z, producing suboptimal intervention costs.

To confirm the validity of the statistical approach, random interventions complying with the statistical interventions η^N and η_s^N are implemented on the network, and the TM is simulated with initial condition $x(0) = \mathbf{0}$ on the Epinions network.

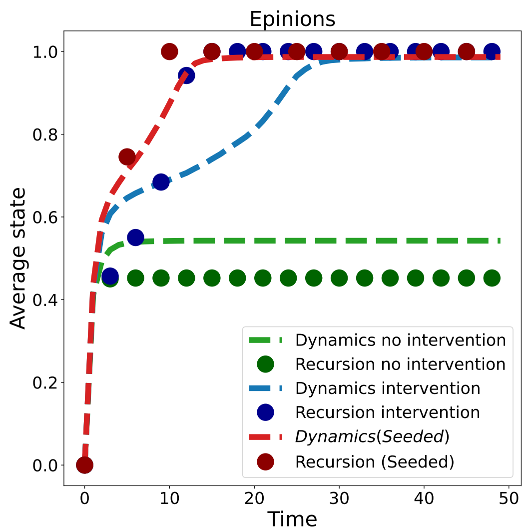

Figure 5.2: The evolution of the fraction of agents in state 1 on the Epinions network versus the one predicted by the recursion (4.20).

Figure 5.2 illustrates the evolution of the fraction of state-1 agents on the networks versus the one predicted by recursion (4.20) based only on the statistics. Note that the seeding intervention, due to its higher cost, is able to drive the agents towards state 1 faster than the optimal intervention. However, also the optimal intervention drives the system towards the desired configurations in a finite number of steps, as predicted by (4.20).

It is remarked that the proposed interventions are designed without relying on the exact network knowledge, suggesting that designing interventions based on the problem statistics is a valid approach even for real networks that are not generated from a configuration model ensemble. Moreover, these interventions apply to an arbitrary threshold distribution, in contrast with those provided by Kempe et al., 2003, which assumes random and uniformly distributed thresholds.

To empirically validate the proposed theoretical framework, numerical experiments were conducted on the Power Grid Network. This infrastructure is mathematically formalized as an undirected, unweighted graph representing the topology of the Western States Power Grid of the United States, comprising an order of $n = 4941$ nodes and a size of $|\mathcal{E}| = 6594$ edges. The Least-Cost Influence Problem (LCIP) is analyzed on this specific topology to establish a comparative baseline against the algorithmic approach previously developed by Cordasco et al., 2015. Consistent with their methodology, the initial activation thresholds r_i were sampled uniformly at random from the discrete set $\{1, \ldots, k_i\}$. Furthermore, the individual intervention cost functions were defined as linear, such that $\gamma_i(\nu) = \nu$ and the macroscopic tolerance parameter was strictly fixed at $\epsilon = 0.3$. Utilizing these structural assumptions, the initial aggregate statistics $p^{(0)}$ were extracted. Subsequently, the optimal statistical intervention η^N was derived by solving the linear programming relaxation formalized in Problem (5.17), with discretization parameters configured at $N = 100$ and $\Delta = 0.05$. The experimental parameters and primary metrics are systematically summarized in the following table:

Parameter / Metric	**Value / Mathematical Description**
Network Topology	Western States Power Grid (Undirected, Unweighted)
Network Order	$n = 4941$
Network Size	$\|\mathcal{E}\| = 6594$
Threshold Distribution	$r_i \sim \mathcal{U}(\{1, \ldots, k_i\})$
Cost Function	$\gamma_i(\nu) = \nu$
Tolerance	$\epsilon = 0.3$
Discretization Parameters	$N = 100$, $\Delta = 0.05$
Evaluation Framework	Expected cost averaged over 10 independent instances
Performance Comparison	$\approx 25\%$ cost reduction relative to Cordasco et al., 2015

Table 5.1: Experimental Configuration and Results for the Least-Cost Influence Problem (LCIP) on the Western States Power Grid Network

By evaluating the expected cost averaged over 10 independent realizations of the random thresholds, it is demonstrated that the proposed statistical interventions successfully satisfy the targeted constraint, driving the predefined fraction of agents into the active state (state 1). Notably, this approach yields a substantial improvement over the network-aware strategies proposed by Cordasco et al., 2015, achieving an approximate 25% reduction in the total intervention cost. It must be critically emphasized that this optimization is attained despite the relatively coarse discretization parameter ($N = 100$) utilized in the linear program. A primary property of this approach is that the optimal policy is derived independently of the exact network topology. Unlike benchmark methods requiring complete adjacency matrices, the local mean-field formulation depends exclusively on aggregate statistics. This formulation is therefore scalable to large-scale networks where only the statistical distributions of the underlying topology are known.

Chapter 6

Conclusions

Concluding Remarks This thesis establishes a comprehensive and computationally tractable framework for the formulation and resolution of optimal intervention problems within super-modular network games. The core premise demonstrated throughout this work is that strategic complementarities, which typically generate systemic unpredictability via multiple equilibria, can be systematically exploited by a central planner. By modifying individual incentives or applying action constraints, a regulator can steer non-cooperative agents toward a maximal, efficient configuration. A primary theoretical contribution is the linkage of intervention strategies to lattice-theoretic methods and maximal best response dynamics. It is formally proven that targeted parameter changes can reliably eliminate suboptimal equilibria and induce finite-time convergence to a designated target region. The characterization of M-attractivity via weak improvement paths provides a robust foundation for identifying optimal constraints using Markov chain algorithms.

Crucially, this work bridges the critical gap between theoretical mechanism design and practical algorithmic scalability. By translating the Linear Threshold Model into a Local Mean-Field dynamical system on random configuration graphs, the inherent NP-hardness of the Optimal Seeding Problem and the fractional Least-Cost Influence Problem (LCIP) is successfully bypassed. The reduction of these complex interventions to linear programs governed strictly by macroscopic network statistics provides a highly scalable tool for policy implementation.
Ultimately, this framework ensures that effective coordination and systemic stability can be achieved in large-scale social, economic, and infrastructure networks, even under conditions of strict topological uncertainty.

Future Research While this dissertation provides a robust framework for optimal intervention in super-modular network games, several critical avenues remain open for future theoretical exploration and practical application.

- **Mixed Externalities and Strategic Substitutability:** The analytical results presented herein rely heavily on the increasing differences property inherent to super-modular games. A natural and challenging extension involves adapting this framework to environments characterized by mixed externalities, where strategic complementarities coexist with strategic substitutes (e.g., peer adoption effects occurring alongside resource congestion or market competition). Future research could investigate the conditions under which targeted interventions can still prevent coordination failures when the strict monotonicity of best-response dynamics is relaxed.
- **Dynamic and Endogenous Network Formation:** The models developed in this work assume either a static deterministic topology or a fixed stochastic generation process, such as the directed configuration model. In many complex systems, however, the network structure evolves endogenously; agents continuously form, sever, or re-weight links in response to changing economic incentives or risk exposures. Extending the Local Mean-Field approximation to account for time-varying, state-dependent network formation would broaden the applicability of these intervention strategies.
- **Robust Intervention under Incomplete Information:** While the statistical approach to large-scale networks successfully circumvents the need for exact topological mapping, it remains reliant on accurate aggregate statistics, such as joint degree and threshold distributions. Future research should explore robust mechanism design wherein the central regulator operates under severe uncertainty or localized data perturbations. Integrating distributionally robust optimization techniques could yield intervention policies that guarantee a baseline level of systemic stability even when prior statistical estimations are flawed.
- **Continuous Action Spaces in Systemic Risk Modeling:** The transition from binary optimal seeding to the fractional incentives of the Least-Cost Influence Problem (LCIP) provides a highly flexible policy instrument. However, extending these methodologies to fully continuous action spaces would allow for highly granular modeling of continuous effort or capital allocation. This extension is particularly vital for modeling cascading defaults and systemic risk in large-scale financial and insurance networks, where continuous policy interventions, such as targeted liquidity injections or dynamic capital requirements, are required to stabilize interconnected institutions.

Appendix A

Elements of Graph Theory

In this section, basic notions of graph theory are presented, and the notation and terminology utilized throughout the remainder of this dissertation are introduced.

A.1 Network Topology and Centrality Measures

The interconnection structure of a networked system is naturally modeled by means of a graph. A network is described by a set of nodes, representing individual units, and a set of links encoding their pairwise interactions Nodes represent the individual units of the system, while links encode the interactions among them. Depending on the application, a link may represent information exchange, influence among agents, observation of neighboring states, or the possibility for a physical or virtual quantity to flow between nodes.

Definition A.1.1. *A (*directed*)* graph *is a pair*

$$\mathcal{G} = (\mathcal{V}, \mathcal{E})$$

where

- $\mathcal{V}$ *is the (finite or countable) set of* nodes *(also called* vertices*);*
- $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ *is the set of* links *(also called* edges*).*

Throughout the thesis, it is assumed that the set of nodes $\mathcal{V}$ is finite and

$$n = |\mathcal{V}|$$

denotes the *order* of the graph. Links (i, i) whose head node coincides with its tail node are referred to as *self-loops*.

The presence of the edge (i, j) is interpreted as a connection between node i and node j. Connections between nodes can be represented in a more compact way using the adjacency matrix $A \in \{0, 1\}^{\mathcal{V}\times\mathcal{V}}$, which is defined as follows:

$$A_{i,j} = \begin{cases} 1, & \text{if } (i, j) \in \mathcal{E}, \\ 0, & \text{if } (i, j) \notin \mathcal{E}. \end{cases}$$

A graph can be defined either through the pair $\mathcal{V}$ and $\mathcal{E}$, or through $\mathcal{V}$ and A.

Weighted graph

Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a positive weight $W_{i,j} > 0$ can be associated with each edge $(i, j) \in \mathcal{E}$. The weights can be gathered in a matrix W in $\mathbb{R}_+^{\mathcal{V}\times\mathcal{V}}$ named the *weight matrix* with the property that $W_{i,j} > 0 \iff (i, j) \in \mathcal{E}$, i.e., if (i, j) is a link or, in other words $W_{i,j} > 0 \iff A_{i,j} = 1$.

A graph equipped with a weight matrix $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ is called *weighted graph*. A weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ is referred to as:

- *undirected* when its weight matrix $W = W^T$ is symmetric, i.e., when a link (i, j) exists if and only if the link with reversed direction (j, i) exists.
- *simple* when it is undirected, unweighted, and the weight matrix W has zero diagonal, equivalently if $\mathcal{G}$ contains no self-loops.

Given a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, the following notions are formalized.

- The in-neighborhood and the out-neighborhood of a node $i \in \mathcal{V}$ are, respectively, the sets

  $$\mathcal{N}_i^- = \{j \in \mathcal{V} \mid (j, i) \in \mathcal{E}\}, \quad \mathcal{N}_i = \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\}$$

  Nodes in $\mathcal{N}_i^-$ and $\mathcal{N}_i$ are referred to, respectively, as in-neighbors and out-neighbors of node i in $\mathcal{G}$.

- The in-degree and out-degree of a node i are defined, respectively, as

$$d_i = \sum_{j \in \mathcal{V}} W_{j,i}, \quad \text{and} \quad k_i = \sum_{j \in \mathcal{V}} W_{i,j}.$$

  The shorter term "degree" is frequently utilized alongside the compact notation:

$$\mathbf{d} = W^T \mathbb{1}, \quad \mathbf{k} = W \mathbb{1}.$$

- $\mathcal{G}$ is called balanced if $\mathbf{d} = \mathbf{k}$.

It is noted that in undirected graphs, no distinction exists between in- and out-neighbors, in- and out-neighborhoods, or in- and out-degrees.

Multigraphs

It often proves convenient to consider notions of network structure that generalize the directed weighted graphs considered thus far.

A primary step is allowing multiple links to connect the same pair of nodes. Consequently, the notion of a link must be modified so it is no longer strictly identified with an ordered pair of nodes.

A (*weighted, directed*) *multigraph* is a tupla

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$$

where $\mathcal{V}$ is the node set, $\mathcal{E}$ is the set of directed links, and $\tau : \mathcal{E} \to \mathcal{V}$ and $\sigma : \mathcal{E} \to \mathcal{V}$ are the maps associating to each link its tail and head node, respectively. It is assumed that $\mathcal{G}$ contains no self-loops, i.e., $\tau(e) \neq \sigma(e)$ for every e in $\mathcal{E}$.

To every multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \tau, \sigma)$, an *adjacency matrix* W in $\mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ is associated, whose entries are defined as:

$$W_{i,j} = |\{e \in \mathcal{E} : \ \tau(e) = i, \sigma(e) = j\}|.$$

Reachability and Connected Components

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a weighted graph. The fundamental concepts regarding connectivity are formally defined as follows:

- A *walk* from node i to node j is a finite sequence of nodes $\gamma = (i_0, i_1, \ldots, i_l)$ such that $i_0 = i, i_l = j$, and $(i_{h-1}, i_h) \in \mathcal{E}$ for all $h = 1, \ldots, l$, i.e., there is a link between every two consecutive nodes. The parameter l denotes the length of the walk. Walks of length 0 represent a transition from a node to itself.
- A node j is said to be reachable from i if there exists a walk from i to j.
- A walk $\gamma = (i_0, i_1, \ldots, i_l)$ such that $(i_{h-1}, i_h) \neq (i_{k-1}, i_k)$ for all $1 \leq h < k \leq l$, except for possibly $i_0 = i_l$, is called a *path*.

A weighted graph $\mathcal{G}$ is called *strongly connected* if given any two nodes i and j, node i is reachable from node j. In numerous applications, it is advantageous to formalize notions of connectedness with respect to specific subsets of nodes. Given a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ and node subset $\mathcal{U} \subseteq \mathcal{V}$, $\mathcal{U}$ is said to be *globally reachable* if for every $i \in \mathcal{V} \setminus \mathcal{U}$ there exists a walk from i to some node j in $\mathcal{U}$. Notice that $\mathcal{G}$ is strongly connected if and only if every subset of nodes $\mathcal{U} \subseteq \mathcal{V}$ is globally reachable.

Finally, another important concept related to connectedness of a graph is that of *period*. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the period $per_{\mathcal{G}}(i)$ of a node i in $\mathcal{V}$ is defined as the greatest common divisor of the lengths of all paths starting and ending in i (with $per_{\mathcal{G}}(i) = 1$ if no such paths exist). The common period of all nodes in a strongly connected graph $\mathcal{G}$ is denoted by $per_{\mathcal{G}}$. A strongly connected graph $\mathcal{G}$ is aperiodic if $per_{\mathcal{G}} = 1$.

Proposition A.1.1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a simple graph. Then*

- *If $\mathcal{G}$ is connected, then necessarily*

$$|\mathcal{E}| \geq |\mathcal{V}| - 1$$

- *If $\mathcal{G}$ has no cycle, then necessarily*

$$|\mathcal{E}| \leq |\mathcal{V}| - 1$$

Algebraic Graph Theory

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a weighted graph. Matrix power of W encode interesting information on the walks on the graph $\mathcal{G}$, as shown below. First, the weight of a length-l walk $\gamma = (i_0, i_1, ..., i_l)$ is defined as the product of its l link weights:

$$W_\gamma = \prod_{1 \leq h \leq l} W_{i_{h-1}, i_h}$$

with the convention that length-0 walks have unitary weight. In the special case of unweighted graphs all walks have unitary weight. Observe that a walk of length $l = 1$ from node i to node j is nothing but a link and its weight is simply the link weight $W_{i,j}$. The following simple but important property states that the (i, j)-th entry of the matrix power W^l coincides with the sum of the weights of length-l walks from i to j.

Proposition A.1.2. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a weighted graph. Then, for every length $l \geq 0$ and every two nodes $i, j \in \mathcal{V}$,*

$$\left(W^l\right)_{i,j} = \sum_{\gamma \in \Lambda_{i,j}^{(l)}} W_\gamma$$

where $\Lambda_{i,j}^{(l)} = \{length - l\ walks\ from\ i\ to\ j\}$. Moreover, if $\mathcal{G}$ is unweighted, the right-hand side coincides the number of length-l walks from i to j.

Proof. By induction on $l \geq 0$. For $l = 0$, the statement is trivial as length-0 walks start and end in the same node and have unitary weight by convention. Then, assume the claim to be true for some l. Observe that the set of all length-$(l+1)$ walks from node i to node j can be partitioned according to the l-th visited node: if $\gamma = (i, \ldots, k, j)$ is a walk of length $(l + 1)$, we have that $W_\gamma = W_{\tilde{\gamma}} W_{k,j}$ where $\tilde{\gamma}$ is the subwalk of γ from i to k. The following computation can be performed

$$\begin{aligned} \sum_{\gamma \in \Lambda_{i,j}^{(l+1)}} W_\gamma &= \sum_{k \in \mathcal{V}} W_{k,j} \sum_{\tilde{\gamma} \in \Lambda_{i,k}^{(l)}} W_{\tilde{\gamma}} \\ &= \sum_{k \in \mathcal{V}} \left(W^l\right)_{i,k} W_{k,j} \\ &= \left(W^{l+1}\right)_{i,j} \end{aligned}$$

where the second equality follows from the inductive assumption. ∎

Corollary A.1.3. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a weighted graph. Then, for every length $l \geq 0$ and every two nodes $i, j \in \mathcal{V}$,*

(i) $\left(W^l\right)_{i,j} > 0$ *if and only if there exists a walk of length l from i to j;*

(ii) $\mathcal{G}$ *is connected if and only if for every two nodes i and j in $\mathcal{V}$, there exists $l > 0$ such that* $\left(W^l\right)_{i,j} > 0$*.*

Proof. Proof are immediate consequences of Proposition A.1.2 ■

Network Centrality

Measures capturing the importance of a node's position in a weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ are analyzed. These are referred to as centrality measures.

The simplest notion is degree centrality, whereby the importance of a node i is determined solely by its degree. In non-balanced networks, a choice must be made as to whether the in-degree d_i or the out-degree k_i should be used. Both options may be appropriate depending on the context. For instance, in the case of a Twitter account, the number of followers (i.e., its in-degree) can be considered a relevant measure of influence. Similarly, the number of citations received may be interpreted as a measure of the impact of a research article. Both the number of followers in Twitter and the number of citations in a citation network are examples of in-degree centrality measures.

A natural extension of (in-)degree centrality is given by *eigenvector centrality*. The key idea is that connections from nodes with high centrality contribute more to the centrality of a node than connections from nodes with low centrality. Formally, the centrality ζ_i of node i is defined to be proportional to the sum of the centralities of its in-neighbors j, that is,

$$\zeta_i \propto \sum_{j \in \mathcal{V}} W_{j,i} \zeta_j .$$

If the proportionality constant is denoted by $\frac{1}{\lambda} > 0$, and $\zeta \in \mathbb{R}^n$ represents the vector of node centralities, the relation can be rewritten as

$$\lambda \zeta = W^T \zeta .$$

Thus, ζ is an eigenvector of W^T associated with the eigenvalue $\lambda > 0$.
By selecting $\lambda = \lambda_W$, the dominant eigenvalue of W, it follows that W^T admits a corresponding non-negative eigenvector satisfying

$$\zeta = \frac{1}{\lambda_W} W^T \zeta .$$

Under the assumption that $\mathcal{G}$ is connected and by imposing the normalization $\zeta^T \mathbf{1} = 1$, the vector ζ is uniquely determined and is referred to as the *eigenvector centrality* of $\mathcal{G}$.

The notion of centrality can be further extended to account for the global structure of the network by incorporating the infinite summation of discounted walks.

This concept, formalized as the Katz-Bonacich centrality (Bonacich, 1987), generalizes degree centrality by measuring the influence or prestige of a node based on all paths originating from or pointing toward it. Depending on the theoretical application, this measure can be formulated using the adjacency matrix W to capture influence via outgoing walks, or its transpose W^T to capture prestige via incoming walks.

Given a decay factor δ satisfying $\delta\rho(W) < 1$ and a nonnegative vector $\mathbf{b}$ representing the intrinsic baseline weight of the nodes, the standard Katz-Bonacich centrality vector $\zeta(W, \delta, \mathbf{b})$ is defined as the solution to $\zeta = \mathbf{b} + \delta W\zeta$, which is explicitly given by:

$$\zeta(W, \delta, \mathbf{b}) = (\mathbb{1} - \delta W)^{-1}\mathbf{b}$$

This formulation is particularly relevant in the study of network games, where the matrix W encodes the outgoing observation channels of the agents.

Alternatively, a normalized, prestige-based variant of the Katz centrality can be established by allowing nodes to possess intrinsic centrality independently of their in-neighbors. Let a parameter $\beta \in (0, 1]$ and a nonnegative vector μ be introduced to represent this intrinsic centrality. A standard choice is $\mu = \mathbf{1}$, implying that all nodes are assigned identical intrinsic centrality. This specific Katz centrality vector is then defined as the solution $\zeta(\beta)$ of:

$$\zeta(\beta) = \left(\frac{1-\beta}{\lambda_W}\right) W^T\zeta(\beta) + \beta\mu.$$

It is observed that, for every $0 < \beta \leq 1$, the dominant eigenvalue of $\frac{1}{\lambda_W}(1 - \beta)W^T$ is strictly smaller than 1, implying that the matrix $\left(\mathbb{1} - \frac{1}{\lambda_W}(1 - \beta)W^T\right)$ is invertible. Consequently, the vector is well defined, unique, and can be expressed as

$$\zeta(\beta) = \left(\mathbb{1} - \frac{1}{\lambda_W}(1 - \beta)W^T\right)^{-1} \beta\mu. \tag{A.1}$$

A.2 Markov Chains

In this section, Markov chains are introduced as stochastic network dynamics. These processes evolve over a state space that can be associated with the nodes of a graph, where the edges represent the possible transitions between states. To formalize these dynamics, the following definitions are established:

Definition A.2.1. *A non-negative square matrix P satisfying*

$$P\mathbf{1} = \mathbf{1} \tag{A.2}$$

(i.e., every row sums to 1*) is defined as a stochastic matrix.*

In the context of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, the *normalized weight matrix*, whose entries are defined

$$P_{i,j} = \frac{W_{i,j}}{k_i}, \qquad \forall i, j \in \mathcal{V}$$

i.e., the (i, j)-the entry of the normalized weight matrix P is the weight of the link (i, j) normalized by the out-weightdegree $k_i = \sum_{j \in \mathcal{V}} W_{i,j}$ of its tail node i, is a stochastic matrix.

Notice that any stochastic matrix P can be interpreted as the normalized weight matrix of a normal graph $\mathcal{G}_P = (\mathcal{V}, \mathcal{E}, P)$, where the node set $\mathcal{V} = \mathcal{X}$ and the link set consists of pairs (i, j) for which $P_{i,j} > 0$. The following terminology is standardized: a stochastic matrix P is referred to as irreducible if $\mathcal{G}_P$ is strongly connected, and it is referred to as aperiodic if $\mathcal{G}_P$ is so.

Definition A.2.2. *A discrete-time Markov chain is a stochastic process $X(t)_{t=0,1,\ldots}$ defined on a state space $\mathcal{X}$ (coinciding with a node set $\mathcal{V}$), where the future state depends strictly on the present state (the Markov property). For any states $i, j \in \mathcal{X}$:*

$$\mathbb{P}(X(t+1) = j | X(0) = i_0, \ldots, X(t) = i_t) = P_{i_t,j} \tag{A.3}$$

where P is a stochastic matrix.

The process $X(t)$ is usually referred to as a discrete-time Markov chain with transition probability matrix P. In order to determine the probability distribution of the trajectories of a Markov chain $X(t)$, one needs to specify both a transition probability matrix P and an *initial probability distribution* to be denoted by $\pi(0)$, i.e., a probability vector with entries corresponding to the elements of $\mathcal{X}$ such that

$$\mathbb{P}(X(0) = i) = \pi_i(0), \quad i \in \mathcal{X}$$

Then, the Markov chain's trajectory satisfies

$$\mathbb{P}(X(0) = i_0, X(1) = i_1, \dots, X(t) = i_t) = \pi_{i_0}(0) \prod_{1 \le s \le t} P_{i_{s-1}, i_s}, \quad t \ge 0.$$

From the equation above, recursive formulas for the marginal probability distribution $\pi(t)$ of $X(t)$, for every time $t \ge 0$, whose entries

$$\pi_i(t) = \mathbb{P}(X(t) = i), \quad i \in \mathcal{X},$$

coincide with the probability that the Markov chain is in a specific node i at time t, as well as for the t-step transition probability matrix $P(t)$, whose entries

$$P_{i,j}(t) = \mathbb{P}(X(t) = j | X(0) = i), \quad i, j \in \mathcal{X},$$

specify the conditional probability that the chain is in node j at time t given that it started in node i at time 0. These recursive equations read as:

$$\pi(t+1) = P^T \pi(t)$$

and, respectively,

$$P(t+1) = P(t)P, \quad P(t+1) = PP(t),$$

Stationarity, convergence, and ergodicity

Let P be a stochastic matrix and let

$$\pi = P^T \pi$$

be an invariant probability distribution for P. If the initial state of the Markov chain $X(0)$ has probability distribution $\pi(0) = \pi$, then

$$\pi(t) = \pi, \quad t \ge 0,$$

i.e., each $X(t)$ has the same marginal probability distribution π as $X(0)$. That justifies the term *stationary distribution* when referring to π, if the initial distribution of the Markov chain is π, it does not matter how long one waits; the marginal probability distribution at any time t will invariably be π.

Proposition A.2.1. *Let P be an irreducible and aperiodic stochastic matrix and $\pi = P^T \pi$ be its unique stationary probability vector. Let $\pi(t)$, for t $= 0, 1, \dots,$ be the probability distribution vector of a Markov chain with transition probability matrix P. Then,*

$$\lim_{t \to \infty} \pi(t) = \pi,$$

Proof. The convergence of an irreducible and aperiodic Markov chain to its unique stationary distribution π is a foundational result in the theory of stochastic processes. Irreducibility guarantees the existence of a unique stationary distribution, while aperiodicity ensures that the state probabilities do not oscillate, meaning the limit $\lim_{t\to\infty} \pi(t)$ exists and equals π for any arbitrary initial distribution $\pi(0)$. The reader is referred to standard texts, such as Levin and Peres, 2017, for a formal proof utilizing coupling methods or the Perron-Frobenius theorem. ■

Proposition A.2.1 states that, if a Markov chain $X(t)$ has irreducible and aperiodic transition probability matrix P, then the marginal probability distribution of $X(t)$ converges to the stationary probability vector π as t grows large.

This means that, the probability of observing the Markov chain $X(t)$ in a given subset $\mathcal{S} \subseteq \mathcal{X}$ approaches the aggregate stationary probability of $\mathcal{S}$

$$\pi_{\mathcal{S}} = \sum_{i\in\mathcal{S}} \pi_i,$$

in the limit as t grows large.

Time-reversible chains

In this section, time-reversible Markov chains are considered (i.e., Markov chains with a reversible transition probability matrix P).

A stochastic matrix P is said to be reversible with respect to a probability distribution π if it satifies the detailed balance equation

$$\pi_i P_{i,j} = \pi_j P_{j,i}, \quad \forall i, j \in \mathcal{X}.$$

Summing up both sides gives:

$$\pi_i = \sum_{j\in\mathcal{X}} \pi_i P_{i,j} = \sum_{j\in\mathcal{X}} \pi_j P_{j,i}, \quad \forall i \in \mathcal{X}$$

so that, if P is reversible with respect to a probability distribution π, then necessarily $\pi = P^T\pi$ is a stationary probability distribution.

Consider a Markov chain with transition probability matrix P reversible with respect to a stationary distribution π, and assume that the initial state $X(0)$ has distribution π.

Then, for all $i, j \in \mathcal{X}$, it holds that

$$\mathbb{P}(X(0) = i, X(1) = j) = \pi_i P_{i,j} = \pi_j P_{j,i} = \mathbb{P}(X(1) = i, X(0) = j),$$

i.e., the pair $(X(0), X(1))$ has the same joint probability distribution as the reversed pair $(X(1), X(0))$. The argument above can be generalized for all $t \geq 0$:

$$(X(0), X(1), \ldots, X(t-1), X(t)) = (X(t), X(t-1), \ldots, X(1), X(0))$$

where $=$ stands for equality in distribution. I.e., reversing the arrow of time does not change the probability distribution of a stationary reversible Markov chain.

This property is known as time-reversibility of the chain $X(t)$. Discrete-time simple random walks over undirected graphs are time-reversible Markov chains.

Appendix B

The Linear Threshold Model

The spread of new ideas and behaviors, the adoption of innovative technologies, and the default contagion of financial institutions are dynamic processes which may exhibit cascading effects in social, economic, and technological networks.

These phenomena generally depend on the topology of the network as well as on the nature of the local agents' dynamics. A popular model proposed to describe such diffusion phenomena is the *Linear Threshold Model*. The number of acquaintances utilizing a new technology is assessed by each agent, and adoption occurs if a specific, personally defined threshold is met. The adoption process can continue in rounds until either the entire network has adopted the new technology or no further adoption is possible.

Cascades are one core problem addressed in the literature, alongside the identification of influential agents and the dynamic progression of adoption. From a control perspective, the possibility of changing the decision of some agents or modifying their adoption threshold is of significant theoretical and practical interest. In this section, the Linear Threshold Model is formally introduced.

B.1 Network Dynamics

Given a network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ with node set $\mathcal{G}$ and edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, each agent on the network is endowed with a time-varying binary state $x_i(t) \in \{0, 1\}$, dependent on a discrete time variable t.

The states 0 and 1, denoting inactivity and activity respectively, may assume specific meanings depending on the application. The evolution of each nodal state is determined by a local deterministic rule, which is dependent upon a personal, endogenous variable $r_i \in \{0, 1, \dots, k_i\}$, called activation thresholds.
In the Linear Threshold Model (LTM), an agent activates if the number of active out-neighbors is greater than or equal to its current activation threshold; otherwise, the agent deactivates.

The *Linear Threshold Model* is the discrete-time dynamical system on the configuration space $\mathcal{X} = \{0, 1\}^{\mathcal{V}}$ defined by

$$\mathbf{x}(t+1) = \Phi_{\mathbf{r}}(\mathbf{x}(t)), \qquad \forall t \geq 0, \tag{B.1}$$

where $\Phi_{\mathbf{r}} : \mathcal{X} \to \mathcal{X}$ is the map defined by

$$(\Phi_{\mathbf{r}}(\mathbf{x}))_i = \begin{cases} 1 & \text{if } \sum\limits_{j \in \mathcal{V}} W_{i,j} x_j \geq r_i \\ 0 & \text{if } \sum\limits_{j \in \mathcal{V}} W_{i,j} x_j < r_i \end{cases} \tag{B.2}$$

Given an initial configuration $\mathbf{x}(0) \in \{0, 1\}^{\mathcal{V}}$ and the endogenous activation thresholds, the network configuration evolves deterministically through the synchronous application of local rules.

In the linear threshold model, agents possess no memory of previous states and compute new states at each iteration: depending on the local condition, they can switch state from 0 to 1 and viceversa.

Remark 18. *A node with $r_i = 0$ activates irrespective of the local condition ($x_i(t+1) = 1$), while a node with $r_i = k_i + 1$ invariably deactivatese ($x_i(t+1) = 0$). These two special activation threshold values can be chosen to enforce stubbornness. Thresholds smaller than 0 or larger than $k_i + 1$ are not necessary and have been excluded on purpose.*

The following lemma captures some basic monotonicity properties of the TM dynamics that prove particularly useful in its analysis.

Lemma 7. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ the network topology and $\mathbf{r}$ a threshold vector. Let $\mathbf{x}(t)$ and $\mathbf{x}'(t)$ be the state vectors of the Threshold Model* (B.2) *on $\mathcal{G}$. Then,*

(i) if $\mathbf{x}'(0) \geq \mathbf{x}(0)$, then $\mathbf{x}'(t) \geq \mathbf{x}(t)$ for all $t \geq 0$.

(ii) if $r_i \leq (1 - x_i(0))k_i$ for all i, then $\mathbf{x}(t)$ is non-decreasing in t; hence, in particular, it is eventually constant.

The map $\Phi_{\mathbf{r}}(\cdot)$ is non-decreasing component-wise

Proof. (i) It is assumed $\mathbf{x}'(0) \geq \mathbf{x}(0)$. The proof proceeds by induction.

Suppose $\mathbf{x}'(t) \geq \mathbf{x}(t)$ for some $t > 0$. The LTM update rule states

$$x_i(t+1) = 1 \iff \sum_{j \in \mathcal{V}} W_{i,j} x_j(t) \geq r_i.$$

Because edge weights are non-negative ($W_{i,j} \geq 0$) and $x'_j(t) \geq x_j(t)$ by the inductive hypothesis, it follows that

$$\sum_{j \in \mathcal{V}} W_{i,j} x'_j(t) \geq \sum_{j \in \mathcal{V}} W_{i,j} x_j(t).$$

Therefore, if the threshold is met for the configuration $\mathbf{x}(t)$, it is strictly also met for $\mathbf{x}'(t)$. This implies $\mathbf{x}'(t+1) \geq \mathbf{x}(t+1)$.

(ii) The objective is to show that $\mathbf{x}(t)$ is non-decreasing in time, meaning $\mathbf{x}(t+1) \geq \mathbf{x}(t)$, for all $t \geq 0$. Let $\mathbf{x}(0)$ the initial configuration, then for any node i:

- if $x_i(0) = 0$, then trivially $x_i(1) \geq x_i(0)$ since states are binary.
- If $x_i(0) = 1$, the hypothesis $r_i \leq (1 - x_i(0))k_i$ yields $r_i \leq 0$. Since thresholds are non-negative, $r_i = 0$. By the update rule, a threshold of 0 guarantees activation, so $x_i(1) = 1 \geq x_i(0)$.

Since $\mathbf{x}(1) \geq \mathbf{x}(0)$, part (i) implies $\mathbf{x}(2) \geq \mathbf{x}(1)$, and by induction $\mathbf{x}(t+1) \geq \mathbf{x}(t)$ for all t. ■

Permanent Activation Linear Threshold Model

Dynamics (B.2) allows agents to switch back from 1 to 0.

The *Permanent Activation Linear Threshold Model* (PALTM) is a variant of the linear threshold model characterized by permanent active states; once a node is activated, it remains active indefinitely.

In such one-switch-only dynamics, at every iterative step $t+1$ an agent selects action 1 if the number of active out-neighbors is at least equal to its activation threshold or if action 1 was played at step t.

$$x_i(t+1) = \begin{cases} 1 & \text{if } \sum\limits_{j \in \mathcal{V}} W_{i,j} x_j(t) \geq r_i \text{ or } x_i(t) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{B.3}$$

The PALTM (B.3) is *progressive*, meaning each agent retains a memory of its previous state; this ensures convergence to a final set of adopters within at most $|\mathcal{V}|$ iterations since at every iteration either the set of adopters monotonically increases by at least one agent or it stops growing from that point onwards. Since there are $|\mathcal{V}|$ agents, after at most $|\mathcal{V}|$ steps the final set of adopters is established.

B.2 Local Mean-Field Approximations

In this section, the activation processes obtained via the LTM with variable thresholds are analyzed. The analysis is conducted on a random graph ensemble, specifically the directed configuration model. Within the ensemble, the average activation can be approximated by a scalar recursive equation, referred to as the Local Mean-Field dynamical system. Activation dynamics on a large network are analyzed, with a primary focus on aggregate behaviors.

The fraction of active agents is:

$$Y(t) = \frac{1}{n} \sum_{i \in \mathcal{V}} x_i(t)$$

whereas the corresponding fraction of "active edges" is:

$$Z(t) = \frac{1}{|\mathcal{E}|} \sum_{i \in \mathcal{V}} d_i x_i(t)$$

These two quantities will be approximated by the mean-field recursive equation. To sample uniformly from the set of problems possessing identical statistics, the following notions are introduced. For a probability distribution p on Ω and f in $\mathbb{R}_+^\Omega$, let

$$\langle p, f \rangle = \sum_{\omega \in \Omega} p_\omega f_\omega \in [0, +\infty]$$

be the expected value of f with respect to p. It is demonstrated that, with high probability, a random problem sampled from the configuration model $\mathcal{C}_{n,p}$ can be accurately approximated by a deterministic linear program. Let

$$\varphi_{k,r}(z) = \sum_{u=r}^{k} \binom{k}{u} z^u (1-z)^{k-u}, \qquad \forall 0 \leq r \leq k.$$

Then, define the maps $\psi_p, \phi_p : [0,1] \to [0,1]$ as

$$\psi_p(z) = \sum_{\omega \in \Omega} p_\omega \varphi_{k_\omega, r_\omega}(z) \tag{B.4}$$

$$\phi_p(z) = \frac{1}{\langle p, d \rangle} \sum_{\omega \in \Omega} p_\omega d_\omega \varphi_{k_\omega, r_\omega}(z). \tag{B.5}$$

The *Local Mean Field* dynamical system is

$$\begin{cases} y(t+1) & = \psi_p(z(t)) \\ z(t+1) & = \phi_p(z(t)) \end{cases}$$

with a given initial condition $z(0) = z_0$ and $y(0) = y_0$.

Theorem B.2.1. *Let $\mathcal{G}$ be a network sampled from configuration model ensemble $\mathcal{C}_{p,n}$ of size n and statistics p. Let $x(t)$, for $t \geq 0$ be the state vector of the TM dynamics on $\mathcal{G}$, let $Y(t) = \frac{1}{n}\sum\limits_{i\in\mathcal{V}} x_i(t)$ and let $y(t)$ be the output of the recursion. Then, for $\epsilon > 0$ and $n \geq \frac{d_{max}k_{max}^{2t+3}}{\langle p,d\rangle\epsilon}$ it holds true*

$$|Y(t) - y(t)| \leq \epsilon$$

for all but at most a fraction $2e^{-\epsilon^2\beta n}$ of networks $\mathcal{G}$ from the configuration model ensemble $\mathcal{C}_{p,n}$, where $\beta = (32\langle p,d\rangle d_{max}^{2t})^{-1}$.

Proof. This concentration bound relies on approximating the discrete diffusion process on the random configuration model via a deterministic differential equation framework. By concentration inequalities adapted to the configuration model, it can be demonstrated that the empirical average state $Y(t)$ deviates from the theoretical local mean-field limit $y(t)$ by more than ϵ with a probability that vanishes exponentially with the network size n. For the complete technical derivation of this specific concentration bound, the reader is referred to the methodology established in Rossi et al., 2017 ■

Analysis of the Local Mean-Field dynamics

In the previous subsection, the LTM was analyzed on a large random network, the directed version of the configuration model. The goal was to estimate the fraction of active agents after each update round, yielding an accurate Local Mean Field (LMF) approximation. The Local Mean Field approach consists in a scalar dynamical system. It is recalled that in the previous chapter, the joint in-, out-degree distribution was denoted by p. Here, the main properties of the LMF systemare studied. The following lemmas, collect these properties.

Lemma 8. *The function $\phi(z)$ is smooth and monotonically non-decreasing. The minimum is $\phi(0) = \frac{1}{\langle p,d\rangle}\sum\limits_{\omega\in\Omega:r_\omega=0} p_\omega d_\omega$, the maximum is $\phi(1) = 1-\frac{1}{\langle p,d\rangle}\sum\limits_{\omega\in\Omega:r_\omega=k_\omega+1} p_\omega d_\omega$ and there exist at least one fixed point $x^+ \in [0,1]$. If there exists a triple $d, k > 0$ and $r \neq \{0, k+1\}$ such that $p_\omega > 0$, then $\phi(x)$ is strictly increasing.*

Let the set P_ϕ contain the fixed points of $\phi(z)$

$$P_\phi = \{z \in [0,1] \mid z = \phi(z)\}$$

and assume it contains only isolated points. This non-empty set defines a partition on the interval $[0,1]$ and, since $\phi(z)$ is monotone, in each sub-interval of the partition the state sequence $z(t)$ is monotone.

Corollary B.2.2. *Assume the set P_ϕ contains isolated points. Consider the augmented set $P_\phi \cup \{0,1\}$. The elements define a partition on the interval $[0,1]$ and, under the function $\phi(z)$, any sub-interval maps into itself. Let $[z_i, z_{i+1}]$ be a sub-interval of the partition, then $\forall z \in (z_i, z_{i+1})$, either $\phi(z) > z$ or $\phi(z) < z$.*

Note that previous Corollary excludes limit cycles in the state dynamic. Since the sequence $z(t)$ admits a limit because is monotone and bounded, and the limit must be a fixed points. Each fixed point z^* in P_ϕ can be stable, unstable or half stable. Suppose $z^* \in (0,1)$ and consider the two sub-intervals of the partition adjacent to z^*. The fixed point z^* is said to be stable if trajectories are converging to z^* from both sub-intervals, it is unstable if trajectories move away from both sides. Otherwise it is a hybrid called half-stable: sequences started on one side of it converge to it whereas on the other side move away. Fixed points possibly in 0 or 1 are either stable or unstable. A simple sufficient condition for the stability of fixed points involves the first derivative of $\phi(z)$.

Lemma 9. *Let z^* be a fixed point of the state function ϕ. If $\phi'(z^*) < 1$, then z^* is stable. If $\phi'(z^*) > 1$, then z^* is unstable.*

The monotonicity of the state function $\phi(z)$ of the autonomous system implies that the fraction of active edges evolves as a monotone sequence, from the initial condition z_0 until the first reachable fixed point. The output function $\psi(z)$ has the same properties described in Lemma 8 for $\phi(z)$. In particular $\psi(z)$ is non-decreasing, thus $y(t)$ is also a monotone sequence.

Bibliography

Acemoglu, D., & Ozdaglar, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, *1*(1), 3–49.

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, *105*(2), 564–608.

Acemoglu, D., Ozdaglar, A., & Yildiz, E. (2011). Diffusion of innovations in social networks. *2011 50th IEEE conference on decision and control and European control conference*, 2329–2334.

Aghaeeyan, A., & Ramazi, P. (2025). From discrete to continuous binary best-response dynamics: Discrete fluctuations almost surely vanish with population size. *IEEE Transactions on Automatic Control*.

Akbarpour, M., Malladi, S., & Saberi, A. (2023). Just a few seeds more: The inflated value of network data for diffusion. *Unpublished manuscript, Stanford University*.

Allen, F., & Gale, D. (2000). Financial contagion. *Journal of political economy*, *108*(1), 1–33.

Arditti, L., Como, G., & Fagnani, F. (2023). On the separability of functions and games. *IEEE Transactions on Control of Network Systems*, *11*(2), 831–841.

Ballester, C., Calvó-Armengol, A., & Zenou, Y. (2006). Who's who in networks. wanted: The key player. *Econometrica*, *74*(5), 1403–1417.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The diffusion of microfinance. *Science*, *341*(6144), 1236498.

Beaman, L., BenYishay, A., Magruder, J., & Mobarak, A. M. (2021). Can network theory-based targeting increase technology adoption? *American Economic Review*, *111*(6), 1918–1943.

Bloch, F. (2016). Targeting and pricing in social networks.

Bonacich, P. (1987). Power and centrality: A family of measures. *American journal of sociology*, *92*(5), 1170–1182.

Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E., & Havlin, S. (2010). Catastrophic cascade of failures in interdependent networks. *Nature*, *464*(7291), 1025–1028.

Bullmore, E., & Sporns, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, *10*(3), 186–198.

Chen, N. (2009). On the approximability of influence in social networks. *SIAM Journal on Discrete Mathematics*, *23*(3), 1400–1415.

Chen, W., Castillo, C., & Lakshmanan, L. V. (2013). *Information and influence propagation in social networks*. Morgan & Claypool Publishers.

Como, G., Durand, S., & Fagnani, F. (2021). Optimal targeting in super-modular games. *IEEE Transactions on Automatic Control*, *67*(12), 6366–6380.

Cordasco, G., Gargano, L., Rescigno, A. A., & Vaccaro, U. (2015). Optimizing spread of influence in social networks via partial incentives. *International colloquium on structural information and communication complexity*, 119–134.

Demaine, E. D., Hajiaghayi, M., Mahini, H., Malec, D. L., Raghavan, S., Sawant, A., & Zadimoghadam, M. (2014). How to influence people with partial incentives. *Proceedings of the 23rd international conference on World wide web*, 937–948.

Elliott, M., Golub, B., & Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, *104*(10), 3115–3153.

Erol, S., Parise, F., & Teytelboym, A. (2023). Contagion in graphons. *Journal of Economic Theory*, *211*, 105673.

Gai, P., & Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the royal society A: Mathematical, physical and engineering sciences*, *466*(2120), 2401–2423.

Galeotti, A., Golub, B., & Goyal, S. (2020). Targeting interventions in networks. *Econometrica*, *88*(6), 2445–2471.

Goyal, A., Lu, W., & Lakshmanan, L. V. (2011). Simpath: An efficient algorithm for influence maximization under the linear threshold model. *2011 IEEE 11th international conference on data mining*, 211–220.

Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, *83*(6), 1420–1443.

Gunnec, D. (2012). Integrating social network effects in product design and diffusion.

Günneç, D., Raghavan, S., & Zhang, R. (2020a). A branch-and-cut approach for the least cost influence problem on social networks. *Networks*, *76*(1), 84–105.

Günneç, D., Raghavan, S., & Zhang, R. (2020b). Least-cost influence maximization on social networks. *INFORMS Journal on Computing*, *32*(2), 289–302.

Jackson, M. O., et al. (2008). *Social and economic networks* (Vol. 3). Princeton university press Princeton.

Jackson, M. O., & Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of economic theory*, *71*(1), 44–74.

Jackson, M. O., & Zenou, Y. (2015). Games on networks. In *Handbook of game theory with economic applications* (pp. 95–163, Vol. 4). Elsevier.

Kempe, D., Kleinberg, J., & Tardos, É. (2003). Maximizing the spread of influence through a social network. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146.

Kuhn, H. W., & Tucker, A. W. (2013). Nonlinear programming. In *Traces and emergence of nonlinear programming* (pp. 247–258). Springer.

Levin, D. A., & Peres, Y. (2017). *Markov chains and mixing times* (Vol. 107). American Mathematical Soc.

Messina, S., Como, G., Durand, S., & Fagnani, F. (2023). Optimal intervention in non-binary super-modular games. *IEEE Control Systems Letters*, *7*, 2353–2358.

Milgrom, P., & Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, 1255–1277.

Morris, S. (2000). Contagion. *The Review of Economic Studies*, *67*(1), 57–78.

Motter, A. E., & Lai, Y.-C. (2002). Cascade-based attacks on complex networks. *Physical review E*, *66*(6), 065102.

Parise, F., & Ozdaglar, A. (2023). Graphon games: A statistical framework for network games and interventions. *Econometrica*, *91*(1), 191–225.

Raghavan, S., & Zhang, R. (2019). A branch-and-cut approach for the weighted target set selection problem on social networks. *INFORMS Journal on Optimization*, *1*(4), 304–322.

Ramazi, P., & Cao, M. (2017). Asynchronous decision-making dynamics under best-response update rule in finite heterogeneous populations. *IEEE Transactions on Automatic Control*, *63*(3), 742–751.

Ramazi, P., & Cao, M. (2020). Convergence of linear threshold decision-making dynamics in finite heterogeneous populations. *Automatica*, *119*, 109063.

Ramazi, P., Riehl, J., & Cao, M. (2016). Networks of conforming or nonconforming individuals tend to reach satisfactory decisions. *Proceedings of the National Academy of Sciences*, *113*(46), 12985–12990.

Ravazzi, C., Como, G., Garetto, M., Leonardi, E., & Tarable, A. (2023). Asynchronous semianonymous dynamics over large-scale networks. *SIAM Journal on Applied Dynamical Systems*, *22*(2), 1300–1343.

Richardson, M., & Domingos, P. (2002). Mining knowledge-sharing sites for viral marketing. *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 61–70.

Riehl, J., Ramazi, P., & Cao, M. (2018a). Incentive-based control of asynchronous best-response dynamics on binary decision networks. *IEEE Transactions on Control of Network Systems*, *6*(2), 727–736.

Riehl, J., Ramazi, P., & Cao, M. (2018b). A survey on the analysis and control of evolutionary matrix games. *Annual Reviews in Control*, *45*, 87–106.

Rossi, W. S., Como, G., & Fagnani, F. (2017). Threshold models of cascades in large-scale networks. *IEEE Transactions on Network Science and Engineering*, *6*(2), 158–172.

Sakhaei, N., Maleki, Z., & Ramazi, P. (2023). Equilibration analysis and control of coordinating decision-making populations. *IEEE Transactions on Automatic Control*, *69*(8), 5065–5080.

Tarski, A. (1955). A lattice-theoretical fixpoint theorem and its applications.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations research*, *26*(2), 305–321.

Topkis, D. M. (1998). *Supermodularity and complementarity*. Princeton university press.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, *19*(3), 305–321.