

Player Capability and Locally Suboptimal Behavior in Strategic Games

By

Yichen Yang

B.A., University of Cambridge (2018)

M.Eng., University of Cambridge (2018)

Submitted to the Department of Electrical Engineering and Computer Science in Partial
Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2023

©2023 Yichen Yang. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Yichen Yang
Department of Electrical Engineering and Computer Science
May 16, 2023

Certified by: Martin C. Rinard
Department of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by: Leslie A. Kolodziejcki
Professor of Department of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Player Capability and Locally Suboptimal Behavior in Strategic Games

by

Yichen Yang

Submitted to the Department of Electrical Engineering and Computer Science
on May 16, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Game theory has a profound influence across many different disciplines, including economics, social science, logic, and computer science. Research in game theory has surfaced many interesting phenomena on how strategic players interact in various game settings. In this thesis, I consider two topics in game theory. The research in both topics surfaces and characterizes interesting phenomena of how strategic players interact in game theoretic settings.

The first topic is the emergence of locally suboptimal behavior in finitely repeated games. Locally suboptimal behavior refers to players playing suboptimally in some rounds of the repeated game (i.e., not maximizing their payoffs in those rounds) while maximizing their total payoffs in the whole repeated game. The emergence of locally suboptimal behavior reflects some fundamental psychological and social phenomena, such as delayed gratification, threats, and incentivized cooperation. The central research question in this part is when can locally suboptimal behavior arise from rational play in finitely repeated games. To this end, we prove the first sufficient and necessary condition that provides a complete mathematical characterization of when locally suboptimal behavior can arise for 2-player finitely repeated games. We also present an algorithm for the computational problem of, given an arbitrary game, deciding if locally suboptimal behavior can arise in the corresponding finitely repeated games. This addresses the practical side of the research question.

The second topic is the impact of player capability on game outcome. Varying player capabilities can significantly affect the outcomes of strategic games. Developing a comprehensive understanding of how different player capabilities affect the dynamics and overall outcomes of strategic games is therefore an important long-term research goal in the field. We propose a general framework for quantifying varying player capability and studying how different player capabilities affect game outcomes. We introduce a new game model based on network congestion games and study how player capabilities affect social welfare at Nash equilibria in this context. The results in this part surface an interesting phenomenon that in some situations, increasing player capabilities may deliver a worse overall outcome of the game. We characterize

when such phenomena happen for the games we study.

We further extend the new game model introduced above with incomplete information on player capability and multi-round play. We establish (algorithmic) game theoretic properties in these extensions, regarding the existence of different types of equilibrium solutions and the complexity of finding equilibrium solutions. These extensions model aspects of interactions between strategic agents that lead to phenomena such as concealment and deception.

Thesis Supervisor: Martin C. Rinard

Title: Professor of Electrical Engineering and Computer Science

Acknowledgments

Throughout the course of my PhD, I am fortunate to have interacted with many fantastic people, the majority of which are here at MIT. They provided me enormous amount of help, inspirations, guidance, friendship, mentorship, and companionship. I want to express my deepest gratitude to these people here.

First of all, I want to thank my advisor Martin Rinard. Martin is an extraordinary person. Since I first arrived at MIT, I have been constantly impressed and inspired by his creative, critical, and sometimes crazy thinking, by his seemingly endless energy, by the multitude of rare knowledge he has, and by his approaches to different kinds of things in research and outside of research. While being inspired, I have also learned and channeled parts of these to myself, or at least I hope I had. Being his student has definitely changed me in broader aspects beyond academically. And I believe this experience will be beneficial to me for my whole life.

I want to thank my two thesis committee members, Asu Ozdaglar and Pablo Parrilo. Despite their busy schedule, they have provided me valuable advice, feedbacks, and discussions from the proposal to the defense of this thesis.

I want to thank all the people in our group, especially Kai Jia, Jiasi Shen, Jose Cambronero, Shivam Handa, Charles Jin, Sara Achour, Thurston Dang, Jürgen Cito, Jeff Perkins, Nikos Vasilakis, Deokhwan Kim, Malavika Samak, Austin Gadiant, Farid Arthaud, and Idan Orzech. I always enjoyed the group lunches and office chats with this diverse and fantastic group of people. It is this group that makes my research journey not so lonely. I was especially fortunate to collaborate with Kai Jia on the work in Chapter 3 of this thesis. Kai also helped check my proofs of the other parts of the thesis, which I am very grateful of. Kai has a super quick and brilliant mind. I really enjoyed working with him.

I want to thank all the people I have collaborated with during earlier years of my PhD, in works that are not included in this thesis: Jeevana Inala, Osbert Bastani, Mangpo Phothilimthana, Remy Wang, Evan Pu, Max Willsey, James Paulos, Sudip Roy, Jacques Pienaar, and Armando Solar-Lezama. These wonderful collaborators

introduced me to research in a wide range of fields in computer science, spanning over AI and systems. Together with the theoretical work in this thesis, I got to do research and publish papers in a wide range of fields through these wonderful collaborations.

I am fortunate to meet many other friends at MIT, especially Yiqiu Wang, Sijie Chen, Yunming Zhang, Lijie Fan, Lingxiao Li, Alex Lew, Yi Tian, and Tiancheng Yu. I also want to particularly thank my 5-year roommate Aniruddh Raghu, and my best friends throughout the years, Qiujia Li, Daren Chen, and Yibo Gao, for all their friendship and supports.

Finally and most importantly, I want to thank my parents for their unconditional support throughout my whole life. And I want to thank my partner, Zhiyu Li, for all her emotional support and companionship throughout the journey. I cannot imagine what it would be like without her.

Contents

1	Introduction	15
1.1	Emergence of Locally Suboptimal Behavior in Finitely Repeated Games	16
1.1.1	Sufficient and Necessary Conditions for 2-Player Games	19
1.1.2	Effect of Changing from Pure Strategies to Mixed Strategies on the Emergence of Local Suboptimality	20
1.1.3	Computational Aspects	22
1.1.4	Generalization to n -Player Games	22
1.1.5	Modeled Phenomena	23
1.2	Impact of Player Capability on Game Outcome in Congestion Games	23
1.2.1	A Framework for Quantifying Player Capabilities and Studying the Impact of Player Capability on Game Outcome	25
1.2.2	Complexity Results for DNC and DNCDA	26
1.2.3	Impact of Player Capability on Game Outcome in DNCDA and GMG	27
1.2.4	Modeled Phenomena	27
1.3	Network Congestion Games with Incomplete Information on Player Capability and Multi-Round Play	28
1.3.1	Introduction of Extensions of DNC	30
1.3.2	Existence of Equilibrium Solutions and Complexity Results for the New Game Models	30
1.3.3	Emergence of Locally Suboptimal Play in Repeated DNC with Private Capability	31

1.3.4	Modeled Phenomena	32
1.4	Roadmap	32
2	Sufficient and Necessary Condition for the Emergence of Locally Suboptimal Behavior in Finitely Repeated Games	33
2.1	The Model	34
2.2	Motivating Examples	36
2.2.1	Example Game where Local Suboptimality Occurs	36
2.2.2	Example Game where SPE with Local Suboptimality Strictly Dominates SPEs without Local Suboptimality	37
2.2.3	Example Games Demonstrating Difference Between Local Suboptimality and the Property in Folk Theorems	38
2.2.4	Example Game where Local Suboptimality Only Occurs with Large T	40
2.3	The Pure Strategy Case	42
2.4	The General Case	46
2.5	The Pure Strategy Against Mixed Strategy Case	54
2.6	Effect of Changing from Pure Strategies to Mixed Strategies on the Emergence of Local Suboptimality	60
2.6.1	Changing from Pure Strategies to Mixed Strategies when the Other Player Can Only Use Pure Strategies	61
2.6.2	Changing from Pure Strategies to Mixed Strategies when the Other Player Can Use Mixed Strategies	67
2.6.3	Changing Both Players from Pure Strategies to Mixed Strategies	72
2.6.4	Discussion	74
2.7	Computational Aspects	75
2.7.1	Compute All Nash Equilibria	77
2.7.2	Determine the Uniqueness of Payoffs at Equilibrium	77
2.7.3	Check the Existence of Required Off-Nash Strategy Profiles	78
2.8	Generalization to n -Player Games	81

2.9	Related Work	86
3	On the Impact of Player Capability on Congestion Games	89
3.1	Background	89
3.2	Models	91
3.2.1	Distance-Bounded Network Congestion Game	91
3.2.2	Distance-Bounded Network Congestion Game with Default Action	93
3.2.3	Gold and Mines Game	95
3.3	Complexity Results	97
3.3.1	Complexity Results for DNC	97
3.3.2	Complexity Results for DNCDA	103
3.4	Impact of Player Capability on Social Welfare	107
3.4.1	Capability Preference Properties	107
3.4.2	Sufficient and Necessary Conditions for Capability Preference Properties in DNCDA	109
3.4.3	Sufficient and Necessary Conditions for Capability Preference Properties in GMG	112
3.4.4	Complete Characterization for the Alternating Ordering Game	117
3.4.5	Full Proofs for the Alternating Ordering Game	121
3.5	Related Work	136
4	Network Congestion Games with Incomplete Information on Player Capability and Multi-Round Play	139
4.1	Models	139
4.1.1	DNC with Mixed Capability	140
4.1.2	DNC with Private Capability	140
4.1.3	Sequential DNC	141
4.1.4	Sequential DNC with Private Capability	142
4.1.5	Repeated DNC	142
4.1.6	Repeated DNC with Private Capability	143

4.2	Existence of Equilibrium Solutions	143
4.2.1	DNC-private	145
4.2.2	seq-DNC	145
4.2.3	seq-DNC-private	147
4.2.4	rep-DNC	149
4.2.5	rep-DNC-private	149
4.3	Complexity Results	151
4.3.1	DNC-mixed	151
4.3.2	DNC-private	151
4.3.3	rep-DNC	152
4.3.4	seq-DNC	152
4.4	Emergence of Locally Suboptimal Play in Repeated DNC with Private Capability	154
4.5	Related Work	158
5	Future Work	159
5.1	Minimum Number of Rounds for Local Suboptimality to Emerge . . .	159
5.2	Sufficient and Necessary Condition for Local Suboptimality for n -Player Games	160
5.3	Emergence of Local Suboptimality in Other Contexts	161
5.4	Impact of Player Capability on Learning Dynamics	162
5.5	Impact of Player Capability in Other Game Models	164
6	Conclusion	165
6.1	Emergence of Locally Suboptimal Behavior in Finitely Repeated Games	165
6.2	Impact of Player Capability on Game Outcome	167

List of Figures

3-1	An example of the DNCDA game and a program to represent a strategy.	94
3-2	Resource layout for the alternating ordering game, a specific version of GMG. Each dot (resp. cross) is a gold (resp. mine).	95
3-3	The DNC instance corresponding to a four-player quadratic threshold game. The distance bound $b = 19$. Non-unit-length edges have labels to indicate their lengths. Dashed gray edges correspond to the S_i^{out} strategies.	99
3-4	Illustration of the DNC instance corresponding to a 3-partition problem. Double-line edges are slow edges, dashed edges are fast edges, and other edges have no delay. Non-unit-length edges have labels to indicate their lengths. Deciding whether the total delay can be bounded by $6m - 3$ is NP-complete.	100
3-5	Illustration of the DNC instance corresponding to a 3-partition problem, for proving NP-hardness of computing the worst social welfare among PNEs. Double-line edges are slow edges, dashed edges are fast edges, and other edges except the ones that connect to the source s have no delay. Non-unit-length edges have labels to indicate their lengths.	102
3-6	Illustration of the DNCDA instance corresponding to a four-player quadratic threshold game. The distance bound is $b = 13$. Non-zero-length edges have labels to indicate their lengths.	104

3-7	Illustration of the DNCDA instance corresponding to a 3-partition problem. Double-line edges are slow edges, dashed edges are fast edges, and other edges have no delay. Non-zero-length edges have labels to indicate their lengths. Deciding whether the total delay can be bounded by $6m - 3$ is NP-complete.	105
3-8	Illustration of the DNCDA instance corresponding to a 3-partition problem, for proving NP-hardness of computing the worst social welfare among PNEs. Double-line edges are slow edges, dashed edges are fast edges, and other edges except the ones that connect to the source s have no delay. Non-zero-length edges have labels to indicate their lengths.	106
3-9	Counterexamples when $d(\cdot)$ does not meet the conditions. Dashed arrows denote unit-length edges and solid arrows denote zero-length edges (default action). Every edge shares the same delay function $d(\cdot)$	111
3-10	Counterexample for the case $\rho < 1$ for r_g and $\rho > 1$ for r_m . Each dot is a resource. The upper (lower) line is $y = 1$ ($y = 0$). Block B and B' are centrosymmetric. The direction of increasing x is from left to right.	114
3-11	Counterexample for the case $\rho > 1$ for r_g and $\rho < 1$ for r_m	114
3-12	Resource layout for the alternating ordering game. Each dot (resp. cross) is a gold (resp. mine). The dashed lines represent a PNE when $b = 2$ (with $-2 + \rho < \mu < -\rho$).	118
3-13	W_{Equil} , W_{best} , POA varying with b . $M = 10, \rho = 0.2$	119
3-14	Characteristics of PNEs over the ρ - μ landscape.	119
4-4	The seq-DNC instance corresponding to a four-player quadratic threshold game. All players share the same distance bound $b = 19$. Non-unit-length edges have labels to indicate their lengths. Dashed gray edges correspond to the S_i^{out} strategies.	153

List of Tables

2.1	Example stage game G where local suboptimality occurs in an SPE of $G(2)$	36
2.2	Example stage game where in the repeated game, some SPE in which local suboptimality occurs strictly dominates all SPEs where local suboptimality does not occur.	37
2.3	Example stage game where local suboptimality can occur, but not all feasible and individually rational payoffs can be attained in the repeated game.	39
2.4	Example stage game where every possible strategy profile is an NE. .	40
2.5	Example stage game where local suboptimality only occurs with large T .	40
2.6	Example stage game in matrix form, row player is player 1, column player is player 2. For all $a \in A_1$, for all $\sigma_2 \in \Delta A_2$ that is a best response to a , a is also a best response to σ_2	46
2.7	Example stage game in matrix form, row player is player 1, column player is player 2. When both players can only use pure strategies, there is no Nash equilibrium. When player 1 can use mixed strategies and player 2 can only use pure strategies, there are multiple Nash equilibria, and local suboptimality can occur in repeated games. . .	64
2.8	Example stage game in matrix form, row player is player 1, column player is player 2. $ V_1^{p,p} = 1$, $ V_2^{p,p} = 1$, $ V_1^{m,p} > 1$, and $ V_2^{m,p} > 1$. .	64

2.9	Example stage game in matrix form, row player is player 1, column player is player 2. When player 1 can play mixed strategies and player 2 can only play pure strategies, there is no Nash equilibrium. When both players can use mixed strategies, there are multiple Nash equilibria, and local suboptimality can occur in repeated games.	69
2.10	Example stage game in matrix form, row player is player 1, column player is player 2. $ V_1^{m,p} = 1$, $ V_2^{m,p} = 1$, $ V_1^{m,m} > 1$, and $ V_2^{m,m} > 1$.	69
2.11	Example stage game in matrix form, row player is player 1, column player is player 2. $ V_1^{m,p} > 1$, $ V_2^{m,p} = 1$. Local suboptimality can never occur in the mixed-pure case, but can occur in the mixed-mixed case.	70
2.12	Example stage game in matrix form, row player is player 1, column player is player 2. For all $a \in A_2$, for all $\sigma_1 \in \Delta A_1$ that is a best response to a , a is also a best response to σ_1	70
3.1	Necessary and sufficient conditions on the delay or payoff functions such that the capability preference properties hold universally. . . .	109
4.1	Existence results for different equilibrium classes. NE: Nash equilibrium, PNE: pure Nash equilibrium, SPE: subgame-perfect equilibrium, PSPE: pure strategy subgame-perfect equilibrium, SE: sequential equilibrium, PSE: pure strategy sequential equilibrium.	144

Chapter 1

Introduction

Since the pioneering work by John von Neumann [87] and the following foundational book by von Neumann and Morgenstern [89], game theory has had a profound influence across many different disciplines, including economics, social science, logic, and computer science. Research in game theory has surfaced many interesting phenomena on how strategic players interact in various game settings. In this thesis, I consider two topics in game theory. The research in both topics surfaces and characterizes interesting phenomena of how strategic players interact in game theoretic settings.

The first topic is the emergence of locally suboptimal behavior in finitely repeated games (Section 1.1). Locally suboptimal behavior refers to players playing suboptimally in some rounds of the repeated game (i.e., not maximizing their payoffs in those rounds) while maximizing their total payoffs in the whole repeated game. The emergence of locally suboptimal behavior reflects some fundamental psychological and social phenomena, such as delayed gratification, threats, and incentivized cooperation. The central research question in this part is when can locally suboptimal behavior arise from rational play in finitely repeated games. To this end, we prove the first sufficient and necessary condition that provides a complete mathematical characterization of when locally suboptimal behavior can arise for 2-player finitely repeated games.

The second topic is the impact of player capability on game outcome (Sections 1.2 and 1.3). Varying player capabilities can significantly affect the outcomes of strategic

games. Developing a comprehensive understanding of how different player capabilities affect the dynamics and overall outcomes of strategic games is therefore an important long-term research goal in the field. We propose a general framework for quantifying varying player capability and studying how different player capabilities affect game outcomes (Section 1.2). We introduce a new game model based on network congestion games and study how player capabilities affect social welfare at Nash equilibria in this context. One of the interesting results here is that in some situations, increasing player capabilities can deliver a worse outcome. This interesting phenomenon is somewhat counter-intuitive, since one would expect that increasing player capabilities is beneficial. We characterize when such phenomena happen for the games we study.

We further extend the new game model introduced above with incomplete information on player capability and multi-round play and study (algorithmic) game theoretic properties in these extensions (Section 1.3). These extensions model aspects of interactions between strategic agents that lead to phenomena such as concealment and deception.

1.1 Emergence of Locally Suboptimal Behavior in Finitely Repeated Games

Repeated games are widely studied in the literature of game theory [8, 61]. A *repeated game* is a game in which a set of players repeatedly play the same *stage game* for a number of rounds. A stage game $G = (n, A_1, \dots, A_n, u_1, \dots, u_n)$ in normal-form consists of n players, each player i 's action space A_i , and each player i 's payoff function $u_i : A \rightarrow \mathbb{R}$, where $A = A_1 \times A_2 \times \dots \times A_n$. A finitely repeated game $G(T)$ refers to the game where G is played repeatedly for T rounds, where T is a positive integer. Player i 's total payoff in the repeated game $G(T)$ is $U_i = \sum_{t=1}^T u_i(\mathbf{a}^t)$, where \mathbf{a}^t are the actions played by each player in round t . Each player can observe the *history* of play, i.e. the actions played by each player in all previous rounds, and decide their strategy in the next round according to the play history. Therefore, a strategy for player

i for the repeated game $G(T)$ specifies which actions to take in each round for all possible histories of play in the previous rounds. Subgame-perfect equilibrium (SPE) is a widely adopted refinement of Nash equilibrium (NE) for extensive-form games. SPE was originally introduced by [82, 83] to eliminate NEs that involve non-credible threats off the equilibrium path. A strategy profile is an SPE if, for any possible history of play at any point of the game, the strategy profile given this history forms an NE for the subgame starting from this point.

A widely known result that appears in many textbooks and lecture notes is that if the stage game G has a unique Nash equilibrium payoff for every player, then in any SPE of any finitely repeated game $G(T)$ with any T rounds, the strategy profile at each round forms an NE of the stage game G [40, 44, 71]. This is proved using backward induction. It is also known that there are stage games G where in some SPEs of the repeated game $G(T)$ for some T , the strategy profile at some round does not form a stage-game NE (Section 2.2.1 presents an example). Such off-(stage-game)-Nash play occurs due to ‘*threats*’ between players that are stated implicitly through players’ strategies. For example, player A does not play their stage-game best response in some round because the other players’ strategies state that if A does so, then the other players will play according to a stage-game NE that gives a lower payoff to A in the later rounds.

We define such off-(stage-game)-Nash plays in repeated games as *local suboptimality*. The emergence of local suboptimality reflects some fundamental psychological and social phenomena, such as delayed gratification, threats, and incentivized cooperation. As we have seen, for some stage games, local suboptimality can occur in some SPE of some repeated games; for other stage games, local suboptimality can never occur in any SPE of any repeated games. Therefore, we can partition the set of all stage games \mathcal{G} into two disjoint subsets \mathcal{G}_{LS} and \mathcal{G}_{LO} . \mathcal{G}_{LS} is the set of stage games G where local suboptimality occurs in some SPE of $G(T)$ for some T ; \mathcal{G}_{LO} is the set of stage games G where local suboptimality never occurs in any SPE of $G(T)$ for any T (LO stands for locally optimal). Our goal in this research is to completely characterize which stage games belong to \mathcal{G}_{LS} and \mathcal{G}_{LO} . The central research question

we aim to tackle is:

Question 1.1.1. *What is a sufficient and necessary condition on the stage game G that ensures that, for all T and all subgame-perfect equilibria of the repeated game $G(T)$, the strategy profile at every round of $G(T)$ forms a Nash equilibrium of the stage game G ?*

The answer to Question 1.1.1 completely characterizes \mathcal{G}_{LS} and \mathcal{G}_{LO} . As we have discussed, a sufficient condition for Question 1.1.1 is widely known (uniqueness of Nash equilibrium payoff for each player). However, this condition is not necessary; in fact, no previous work establishes a sufficient and necessary condition. A large body of work focuses on Folk Theorems [14, 13, 48, 84, 47], where the property of interest is: all feasible (i.e., the payoff profile lies in the convex hull of the set of all possible payoff profiles of the stage game) and individually rational (i.e., the payoff of each player is at least their minmax payoff in the stage game) payoff profiles can be attained in the equilibrium of the repeated game. As we show in Section 2.2.3, the property considered in Folk Theorems and the local suboptimality property we consider in this work are different, and the two properties do not have direct implications in either direction. Therefore, the conditions established for Folk Theorems in the literature do not solve the problem we consider.

In addition to the complete mathematical characterization of the partitioning between \mathcal{G}_{LS} and \mathcal{G}_{LO} , we also consider the computational aspect of the problem:

Question 1.1.2. *Given an arbitrary stage game G , how to (algorithmically) decide if there exists some T and some SPE of $G(T)$ where local suboptimality occurs? Is this problem decidable?*

A naive approach is to enumerate over T , solve for all SPEs for each $G(T)$, and check if off-Nash behavior occurs. Such an approach is not only computationally inefficient, but also not guaranteed to terminate due to the unboundedness of T . In fact, we show that there are stage games where local suboptimality only occurs in repeated games with very large T , and we can construct games where this minimum

T for local suboptimality to occur can be arbitrarily large (Section 2.2.4). These facts motivate the study of Question 1.1.2.

We summarize the results of this part in the following sections.

1.1.1 Sufficient and Necessary Conditions for 2-Player Games

A main theoretical contribution of this part is that we prove sufficient and necessary conditions for Question 1.1.1 for 2-player games. We prove the conditions for three cases: 1) only pure strategies are allowed (Theorem 2.3.1), 2) the general case where mixed strategies are allowed (Theorem 2.4.1), and 3) one player can only use pure strategies and the other player can use mixed strategies (Theorem 2.5.1). From the perspective of partitioning the set of stage games \mathcal{G} , denote $\mathcal{G}_{LS}^{p,p}$ as the set of stage games G where local suboptimality occurs in some SPE of $G(T)$ for some T when both players can only use pure strategies, $\mathcal{G}_{LO}^{p,p}$ as the set of stage games G where local suboptimality never occurs in any SPE of $G(T)$ for any T when both players can only use pure strategies, $\mathcal{G}_{LS}^{m,m}$ and $\mathcal{G}_{LO}^{m,m}$ as the corresponding partitioning when both players can use mixed strategies, and $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LO}^{m,p}$ as the corresponding partitioning when player 1 can use mixed strategies and player 2 can only use pure strategies. Essentially, we obtain complete mathematical characterizations of the partitioning of \mathcal{G} for cases (1), (2), and (3) above: 1) $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LO}^{p,p}$, 2) $\mathcal{G}_{LS}^{m,m}$ and $\mathcal{G}_{LO}^{m,m}$, and 3) $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LO}^{m,p}$.

As an example, the following theorem establishes a sufficient and necessary condition for the general case where mixed strategies are allowed. This theorem uses the following notations: A_i is the strategy space of player i in the stage game G , ΔS is the set of probability distributions over set S , $\text{Nash}(G)$ is the set of all Nash equilibria of the stage game G , and V_i is the set of payoff values attainable at $\text{Nash}(G)$ for player i .

Theorem (restating Theorem 2.4.1). *For general 2-player games (mixed strategies allowed), a sufficient and necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

1. $|V_1| > 1, |V_2| > 1$, and there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2$ where $(\hat{\sigma}_1, \hat{\sigma}_2) \notin \text{Nash}(G)$, OR
2. $|V_1| > 1, |V_2| = 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{\sigma}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$, OR
3. same as (2) but exchange player 1 and 2.

This theorem shows that: when there are multiple payoffs attainable at stage-game NEs for both players ($|V_1| > 1, |V_2| > 1$), the requirement for local suboptimality to occur in some SPE of $G(T)$ for some T is that there exists a strategy profile $(\hat{\sigma}_1, \hat{\sigma}_2)$ that is not a stage-game NE; when there are multiple payoffs attainable at stage-game NEs for one player but only a unique payoff attainable for the other player, the requirement for local suboptimality to occur is that there exists a strategy profile where the player who has a unique payoff attainable plays a best response and the player who has multiple payoffs attainable does not play a best response; in all other cases, local suboptimality never occurs in any SPE of $G(T)$ for any T .

This is the first sufficient and necessary condition for off-(stage-game)-Nash plays to occur in SPEs of 2-player finitely repeated games. And we prove the sufficient and necessary condition for each case regarding whether players have access to mixed strategies or not. All results in this part apply to general normal-form stage games G , not restricted to any specific types of games.

1.1.2 Effect of Changing from Pure Strategies to Mixed Strategies on the Emergence of Local Suboptimality

Our results on the sufficient and necessary conditions discussed above provide complete mathematical characterizations of $\mathcal{G}_{LS}^{p,p}$, $\mathcal{G}_{LS}^{m,p}$, and $\mathcal{G}_{LS}^{m,m}$ (and therefore $\mathcal{G}_{LO}^{p,p}$, $\mathcal{G}_{LO}^{m,p}$, and $\mathcal{G}_{LO}^{m,m}$). Based on these results, we further study the effect of changing from pure strategies to mixed strategies on the emergence of local suboptimality. We aim to answer the following question: under what conditions on the stage game G will allowing players to play mixed strategies change whether local suboptimality can ever

occur in some repeated game $G(T)$? Essentially, we aim to study the relationships between $\mathcal{G}_{LS}^{p,p}$, $\mathcal{G}_{LS}^{m,p}$, and $\mathcal{G}_{LS}^{m,m}$.

We prove that $\mathcal{G}_{LS}^{p,p} \subseteq \mathcal{G}_{LS}^{m,p} \subseteq \mathcal{G}_{LS}^{m,m}$ (Theorems 2.6.2, 2.6.10 and 2.6.17), i.e., if local suboptimality can occur before the change, then after changing any player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can still occur. This is because we prove that any SPE of the repeated game before the change is still an SPE after the change, and any strategy profile that is not a stage-game NE before the change is still not a stage-game NE after the change. So allowing players to play mixed strategies can never prohibit the emergence of local suboptimality.

On the other hand, we show that $\mathcal{G}_{LS}^{p,p} \neq \mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,p} \neq \mathcal{G}_{LS}^{m,m}$ (so $\mathcal{G}_{LS}^{p,p}$ is a proper subset of $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,p}$ is a proper subset of $\mathcal{G}_{LS}^{m,m}$), i.e., there are games where local suboptimality can never occur before the change, but after changing one player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can occur. We present complete characterizations of the sets $\mathcal{G}_{LS}^{m,p} \setminus \mathcal{G}_{LS}^{p,p}$, $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{m,p}$, and $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{p,p}$, by proving sufficient and necessary conditions on the stage game G such that local suboptimality can never occur before the change but can occur after the change (Theorems 2.6.8, 2.6.16 and 2.6.23). Our characterizations are fine-grained based on $|V_1|$ and $|V_2|$, the number of payoff values attainable at stage-game NEs for each player. For example, we show that under certain preconditions on $|V_1|$ and $|V_2|$, $\mathcal{G}_{LS}^{p,p} = \mathcal{G}_{LS}^{m,p}$; under other preconditions on $|V_1|$ and $|V_2|$, $\mathcal{G}_{LS}^{p,p} \neq \mathcal{G}_{LS}^{m,p}$, and for each of such cases, we present an example stage game G where $G \notin \mathcal{G}_{LS}^{p,p}$ and $G \in \mathcal{G}_{LS}^{m,p}$. These examples demonstrate different mechanisms of how changing a player from pure-strategies-only to mixed-strategies-allowed can lead to the emergence of local suboptimality. We perform the same fine-grained analyses on $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{p,p}$ as well.

The results in this part have a conceptual connection with the study of player capability in Chapter 3: whether players have access to mixed strategies can be viewed as a form of player capability. From this perspective, the results in this part provide insights on how player capabilities affect the emergence of local suboptimality

in 2-player finitely repeated games.

1.1.3 Computational Aspects

We propose an algorithm for deciding Question 1.1.2 for 2-player games for the general case where mixed strategies are allowed and analyze the computational complexity of this algorithm. This shows that Question 1.1.2 is decidable for 2-player games where mixed strategies are allowed. This algorithm provides a method for computationally deciding if local suboptimality can ever happen for a given stage game. The algorithm is based on the sufficient and necessary condition established in Theorem 2.4.1. We design several efficient methods for checking different parts of the condition by utilizing properties we prove for general games. Naive methods for checking these parts of the condition take exponential time in the worst case, whereas our methods for checking these parts of the condition take polynomial time in the worst case.

1.1.4 Generalization to n -Player Games

For n -player games, we prove a sufficient and necessary condition for Question 1.1.1 for the pure strategy case (i.e., only pure strategies are allowed for all players); for the general case where mixed strategies are allowed, we prove a separate sufficient condition and a separate necessary condition for Question 1.1.1. These conditions are both tighter than what is previously known in the literature (again, only a sufficient condition is known previously, i.e., there is a unique Nash equilibrium payoff for each player [40, 44, 71]).

The proof of a sufficient and necessary condition for the 2-player case when mixed strategies are allowed relies on some properties that we prove to hold for 2-player games (Lemma 2.4.3 and the subsequent arguments in the proof of Theorem 2.4.1 that uses Lemma 2.4.3 to show there exists a connected component of off-Nash strategy profiles). It is not clear whether similar properties hold for n -player games. Therefore, the questions of 1) what is a sufficient and necessary condition for n -player games when mixed strategies are allowed, and 2) is Question 1.1.2 decidable for n -player

games when mixed strategies are allowed, remain open.

1.1.5 Modeled Phenomena

The games we study in this part model aspects of interactions between strategic agents that lead to phenomena such as delayed gratification, threats, and incentivized cooperation. As we have discussed, the locally suboptimal behaviors considered in this part occur due to ‘threats’ between players. Our results completely characterize when such threats can and cannot happen in subgame-perfect equilibria of 2-player finitely repeated games. For such threats to happen, at least one of the players needs to have more than one payoff attainable at stage-game Nash equilibria, and there needs to exist an off-(stage-game)-Nash strategy profile that satisfies certain requirements based on the number of payoffs attainable at Nash equilibria for each player. Such threats lead to delayed gratification, since the player who is threatened sacrifices some payoff in earlier rounds to obtain higher payoff in later rounds. In some situations, such threats can incentivize cooperations, where both players obtain a better total payoff than the best total payoffs they can obtain if there are no threats (Section 2.2.2 presents an example).

1.2 Impact of Player Capability on Game Outcome in Congestion Games

Varying player capabilities can significantly affect the outcomes of strategic games. Developing a comprehensive understanding of how different player capabilities affect the dynamics and overall outcomes of strategic games is therefore an important long-term research goal in the field. Central questions include characterizing, and ideally precisely quantifying, player capabilities, then characterizing, and ideally precisely quantifying, how these different player capabilities interact with different game characteristics to influence or even fully determine individual and/or group dynamics and outcomes.

The type of player capability we consider in this research is the size of the strategy space: higher capability means players have access to a larger strategy space. We anticipate a range of mechanisms for characterizing player capabilities, from simple numerical parameters through to complex specifications of available player behavior. In this work, we propose a framework of using programs in a domain-specific language (DSL) to compactly represent player strategies. Bounding the sizes of the programs available to the players creates a natural capability hierarchy, with more capable players able to deploy more diverse strategies defined by larger programs. Building on this foundation, we study the effect of increasing or decreasing player capabilities on game outcomes such as social welfare at equilibrium. To the best of our knowledge, this research presents the first systematic analysis of the effect of different player capabilities on the outcomes of strategic games.

We introduce a new game, the *Distance-bounded Network Congestion game (DNC)*, as the basis of our study. DNC is a variant of the widely studied network congestion games [29, 78]. A network congestion game consists of a set of players and a directed graph where each edge is associated with a delay function. The goal of each player is to plan a path that minimizes the delay from a source vertex to a sink vertex. The delay of a path is the sum of the delays on the edges in the path, with the delay on each edge depending (only) on the number of players choosing the edge. The game is symmetric when all players share the same source and sink. DNC is a symmetric network congestion game in which each player is subject to a *distance bound* — i.e., a bound on the number of edges that a player can use.

We instantiate our framework on two variants of DNC where we define simple DSLs that compactly represent the strategy spaces. The first game is the *Distance-bounded Network Congestion game with Default Action (DNCDA)*. In this game, each node has a default outgoing edge that does not count towards the distance bound. Hence a strategy can be compactly represented by specifying only the non-default choices. The second game is the *Gold and Mines Game (GMG)*, where there are gold and mine sites placed on parallel horizontal lines, and a player uses a program to compactly describe the line that they choose at each horizontal location.

Covering a gold site gives a player a positive payoff, whereas covering a mine site gives a player a negative payoff. Both these payoffs depend on how many players cover that resource together. We show that GMG is a special form of DNCDA. Our analysis is then centered around the following research question:

Question 1.2.1. *How does varying player capability affect social welfare at equilibrium for the games we consider? What are the conditions for each type of relationship between player capability and social welfare at equilibrium to hold for these games?*

A line of work in algorithmic game theory establishes complexity results in the context of network congestion games and their variants, including the complexity of finding a pure Nash equilibrium [29, 3], the complexity of finding the best/worst social welfare at pure Nash equilibria [33], and the complexity of welfare maximization (not restricted to Nash equilibria) [62]. Since DNC and DNCDA are new variants of network congestion games, we aim to establish the complexity results for these new game models:

Question 1.2.2. *What are the computational complexities of 1) finding a pure Nash equilibrium, 2) finding the best/worst social welfare at pure Nash equilibrium, and 3) finding the best social welfare across all pure strategy profiles (not only at equilibrium), in the context of DNC and DNCDA?*

We summarize the results of this part in the following sections.

1.2.1 A Framework for Quantifying Player Capabilities and Studying the Impact of Player Capability on Game Outcome

We present a general framework for quantifying varying player capabilities and studying how different player capabilities affect the outcome of strategic games. In this framework, we use programs in a Domain-Specific Language (DSL) to compactly represent player strategies. We define a player’s capability as the maximum size of the

programs they can use, with larger maximum program sizes corresponding to larger strategy spaces.

We propose four *capability preference* properties that characterize the impact of player capability on social welfare at equilibrium. We call a game *capability-positive* (resp. *capability-negative*) if social welfare at equilibrium does not decrease (resp. increase) when players become more capable. A game is *max-capability-preferred* (resp. *min-capability-preferred*) if the worst social welfare at equilibrium when players have maximal (resp. minimal) capability is at least as good as any social welfare at equilibrium when players have lower (resp. higher) capability. These are general properties applicable to any types of games.

1.2.2 Complexity Results for DNC and DNCDA

Fabrikant et al. [29] shows that finding a pure Nash equilibrium (PNE) is in P for symmetric network congestion games but PLS-complete for asymmetric network congestion games. DNC is a symmetric network congestion game with a bound on the number of edges each player can use. We prove the following complexity results for DNC:

- Finding a PNE in DNC is PLS-complete. This is interesting because finding a PNE in symmetric network congestion games is in P [29]. So with the addition of a distance bound, the problem becomes harder.
- Computing the best/worst social welfare among PNEs of a DNC is NP-hard.
- Computing the best social welfare among all pure strategy profiles of a DNC is NP-hard.

We further prove that DNCDA has the same complexity results as DNC: finding a PNE is PLS-complete; computing the best/worst social welfare among PNEs is NP-hard; computing the best social welfare among all pure strategy profiles is NP-hard.

1.2.3 Impact of Player Capability on Game Outcome in DNCDA and GMG

We instantiate our framework and study the impact of player capability on game outcome in the context of DNCDA and GMG.

For DNCDA, we focus on a restricted version where all edges share the same delay function $d(\cdot)$, which we call *distance-bounded network congestion game with default action and shared delay (DNCIDAS)*. We prove sufficient and necessary conditions on $d(\cdot)$ under which each of the four capability preference properties holds *universally* (i.e., for all network configurations of DNCIDAS). This means that if $d(\cdot)$ satisfies the proven condition, then the target property (e.g., capability-positive) holds for all possible network configurations; if $d(\cdot)$ does not satisfy the proven condition, then for any such delay function, we can always find a network configuration where the target property (e.g., capability-positive) does not hold for the game. Similarly for GMG, we prove sufficient and necessary conditions on $r_g(\cdot)$ and $r_m(\cdot)$, the payoff functions for gold and mines, under which each of the four capability preference properties hold *universally* (i.e., for all game layout). The results on these sufficient and necessary conditions are summarized in Table 3.1.

Finally, for a specific version of GMG called the alternating ordering game, we fully characterize how social welfare at equilibrium varies with player capability by proving the functional form of $W_{\text{equil}}(b)$, where b is the player capability and W_{equil} is the social welfare at equilibrium, in terms of the game parameters. We identify situations where social welfare at equilibrium increases, stays the same, or decreases as players become more capable. This result provides insights and intuitions on the factors that affect whether increasing player capability is beneficial or not.

1.2.4 Modeled Phenomena

The research in this part surfaces an interesting phenomenon that in some situations, increasing player capabilities may deliver a worse overall outcome of the game. This phenomenon occurs since players engage in harmful/wasteful competitions due to

their selfish nature. And the level of competition increases as players become more capable, which leads to a decrease in the overall social welfare. For situations where such harmful competitions prevail, regulators may consider imposing restrictions on the power/capabilities of the players. Game/rule designers may also improve the designs by reducing the opportunities for such harmful competitions.

1.3 Network Congestion Games with Incomplete Information on Player Capability and Multi-Round Play

In the previous part where we study the impact of player capability on game outcomes, we introduced a new game model, the Distance-bounded Network Congestion game (DNC). DNC has a natural hierarchy of player capabilities as measured by the distance bound. This hierarchy of player capabilities allows us to study the impact of player capability on social welfare at equilibrium in the context of DNC. In this part, we extend the original DNC model in a variety of directions to obtain a richer set of game models with hierarchies of player capabilities. These new game models open up the space of research on new phenomena involving player capabilities.

The first extensions incorporate incomplete information on player capabilities. In almost all interactions between intelligent agents (e.g., humans, robots, AI), it is often the case that each agent has only partial information on the capabilities of the other agents. Concealing one's capability from opponents or even misleading opponents into believing false capability information is a common strategy in non-cooperative or competitive environments. Incorporating incomplete information on player capabilities into DNC can therefore generate new game models that enable the study of important social phenomena such as concealment and deception. To model incomplete information on player capabilities, we adopt the Bayesian game theory framework as introduced by the pioneering work by Harsanyi [49]. In this framework, each player has a set of possible *types*. At the start of the game, nature samples and

assigns a type for each player from a prior (joint) distribution. The prior distribution is known to all players, while the type assignment of each player is only known by themselves. The goal of each player is to maximize their *expected* payoff given their belief on the other players' types and the strategies of the other players. In this research, we let different player types correspond to different player capabilities.

The second type of extensions incorporates multi-round play. In many contexts, agents engage in multiple interactions, with these interactions potentially continuous over a period of time and/or repeated over longer time horizons. Such continuous/repeated interactions allow agents to 1) use past interactions to infer future behaviors of other agents, and 2) make representations of their own future behavior to influence the present behavior of other agents. These dynamics give rise to many social phenomena/concepts, such as credibility, cooperation, discovering and exploiting weaknesses, deception, threats, and treachery. Games with multi-round play can often model the dynamics that give rise to these phenomena, enabling us to study the above phenomena in a formal setup.

In this research, we extend the DNC model with two different types of multi-round play. The first type is sequential play within a single network congestion game, where instead of choosing a complete path through the network at the start of the game, each player only chooses the next edge in their path in each round of the game. The second type is repeated play, where players repeatedly play the same game for a number of rounds.

The main research problem we tackle in this part is to study the (algorithmic) game theoretic properties of these new game models, regarding the existence of different types of equilibrium solutions and the complexity of finding equilibrium solutions:

Question 1.3.1. *Do different types of equilibrium solutions exist for each game model? What is the complexity of finding a pure Nash equilibrium for these game models?*

We summarize the results of this part in the following sections.

1.3.1 Introduction of Extensions of DNC

We introduce the following new game models that extend the original DNC model with incomplete information on player capability and multi-round play:

- DNC with mixed capability (DNC-mixed): in the original DNC, all players in the same game have the same capability (distance bound), and we study what happens when player capabilities vary across games. In this extension, different players can have different distance bounds within the same game.
- DNC with private capability (DNC-private): players are uncertain about the distance bounds of the other players. We adopt the settings of Bayesian games [49]: players know a prior distribution from which the capabilities of each player is drawn, but they do not know the actual capabilities of the other players.
- Sequential DNC (seq-DNC): sequential version of DNC, where in each round of the game, every player simultaneously chooses the next edge in their paths.
- Sequential DNC with private capability (seq-DNC-private): sequential version of DNC where players are uncertain about the distance bounds of the other players.
- Repeated DNC (rep-DNC): players repeatedly play the same DNC for a finite number of rounds.
- Repeated DNC with private capability (rep-DNC-private): repeated DNC where players are uncertain about the distance bounds of the other players.

This set of new game models provides a broader context for studying the impact of player capability on the dynamics and outcomes of strategic games.

1.3.2 Existence of Equilibrium Solutions and Complexity Results for the New Game Models

For the types of equilibrium solutions, we consider 1) Nash equilibrium, 2) subgame-perfect equilibrium [82], and 3) sequential equilibrium [57]. The existence of mixed

strategy versions of each type of equilibrium follows from prior work [66, 57, 83]. We prove the existence (or non-existence) of 1) pure Nash equilibrium, 2) pure strategy subgame-perfect equilibrium, and 3) pure strategy sequential equilibrium, for each of the 6 new game models. Table 4.1 summarizes the results.

We prove the complexity of finding a PNE for DNC-mixed, DNC-private, rep-DNC, and seq-DNC: all four are PLS-complete. PNE does not in general exist for seq-DNC-private and rep-DNC-private.

These results establish the game theoretic properties of the new game models, regarding the existence of different types of equilibrium solutions and the complexity of finding a pure Nash equilibrium. This provides a basis for future research on this richer sets of game models involving varying player capabilities.

1.3.3 Emergence of Locally Suboptimal Play in Repeated DNC with Private Capability

In the first part of this thesis, we studied the emergence of locally suboptimal play in finitely repeated games with complete information. In such complete information games, local suboptimality occurs due to ‘threats’ between players. In this part, we introduced games with incomplete information on player capabilities and multi-round play. For such games, there can be another type of motivation for locally suboptimal play: players may sacrifice some payoff in earlier rounds to hide their capability from other players, in order to get better payoff in the future and maximize their total payoff.

We present an example rep-DNC-private game where local suboptimality emerges from rational play and provide a complete characterization of how it occurs. We prove that in *any* sequential equilibrium of this game, one player plays some strictly dominated strategy in all but the last round and switches to the dominant strategy only in the last round. Therefore, local suboptimality occurs in *every* sequential equilibrium of this game. In contrast, for finitely repeated games with complete information (the game model we consider in the first part of this thesis), players repeatedly playing

the same stage-game Nash equilibrium is always a subgame-perfect equilibrium of the repeated game, so there always exists a subgame-perfect equilibrium where local suboptimality does not occur. Therefore, incomplete information on player capabilities makes such universal occurrence of local suboptimality possible; with complete information, such universal occurrence of local suboptimality can never happen in finitely repeated games.

1.3.4 Modeled Phenomena

The games we study in this part model aspects of interactions between strategic agents that lead to phenomena such as concealment and deception. From the analysis of an example repeated DNC with private capability game, we show that players sacrifice some payoff in earlier rounds to hide their capability from other players, in order to get better payoff in the future and maximize their total payoff. In fact, we show that such concealment/deception occurs in *every* sequential equilibria of this game. Therefore, in some situations concealment/deception always emerges from rational play.

1.4 Roadmap

This thesis is organized as follows. Chapter 2 presents our research on the emergence of locally suboptimal behavior in finitely repeated games. Chapter 3 presents our research on the impact of player capability on game outcome in the context of DNC and its variants. Chapter 4 presents the extensions of the DNC model with incomplete information on player capability and multi-round play and establishes their (algorithmic) game theoretic properties. Chapter 5 discusses future work directions. Chapter 6 concludes the thesis.

Acknowledgement of Collaborations The research presented in Chapter 3 is a joint work with Kai Jia and Martin Rinard, with Kai and I sharing equal contributions. The research in Chapters 2 and 4 is joint work with Martin Rinard.

Chapter 2

Sufficient and Necessary Condition for the Emergence of Locally Suboptimal Behavior in Finitely Repeated Games

This chapter presents our research on the emergence of locally suboptimal behavior in finitely repeated games. Section 2.1 presents the notations and formal definitions of the game model we consider and local suboptimality. Section 2.2 presents several example games that demonstrate different aspects of local suboptimality to motivate our study. Sections 2.3 to 2.5 present the results on the sufficient and necessary conditions for local suboptimality to occur for 2-player games, for cases where 1) only pure strategies are allowed (Section 2.3), 2) mixed strategies are allowed (Section 2.4), and 3) one player can only use pure strategies and the other player can use mixed strategies (Section 2.5). Section 2.6 presents the results on the effect of changing from pure strategies to mixed strategies on the emergence of local suboptimality. Section 2.7 considers the computational aspects of the problem. Section 2.8 considers the generalization to n -player games. Section 2.9 discusses the related work.

2.1 The Model

A stage game $G = (n, A_1, \dots, A_n, u_1, \dots, u_n)$ in normal form consists of n players, each player i 's strategy space A_i , and each player i 's payoff function $u_i : A \rightarrow \mathbb{R}$, where $A = A_1 \times A_2 \times \dots \times A_n$. We assume n and A are finite. Throughout this chapter, we use a to denote pure strategies (or actions) in the stage game and σ to denote mixed strategies in the stage game, e.g. $a_i \in A_i$ denotes a pure strategy for player i and $\sigma_i \in \Delta A_i$ denotes a mixed strategy for player i , both for the stage game, where ΔS denotes the set of probability distributions over set S . We use $S_{\sigma_i} = \{a \in A_i, \sigma_i(a) > 0\}$ to denote the support for mixed strategy σ_i . A strategy profile $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ is a set of strategies for all players. In general, we use bold symbols to represent collections over players. For convenience, we use $u_i(\boldsymbol{\sigma})$ to denote the expected payoff of player i under the (mixed) strategy profile $\boldsymbol{\sigma}$. A strategy σ_i is a *best response* to the strategy profile of the other players $\boldsymbol{\sigma}_{-i}$ if $u_i(\sigma_i, \boldsymbol{\sigma}_{-i}) = \max_{\sigma'_i \in \Delta A_i} u_i(\sigma'_i, \boldsymbol{\sigma}_{-i})$. A strategy profile $\boldsymbol{\sigma}$ is a *Nash equilibrium* (NE) if for all player $i \in [n]$ ($[n]$ denotes the set $\{1, \dots, n\}$), σ_i is a best response to $\boldsymbol{\sigma}_{-i}$. We use $\text{Nash}(G)$ to denote the set of all Nash equilibria of the stage game G . We use $V_i = \{u_i(\boldsymbol{\sigma}) \mid \boldsymbol{\sigma} \in \text{Nash}(G)\}$ to denote the set of payoff values attainable at Nash equilibria for player i .

We use $G(T)$ to denote the game where G is played repeatedly for T rounds, where T is a positive integer. Denote the *outcome* in round $t \in [T]$ as $\mathbf{a}^t \in A$. Player i 's total payoff in the repeated game $G(T)$ is $U_i = \sum_{t=1}^T u_i(\mathbf{a}^t)$. A strategy of player i in $G(T)$ specifies which actions to take in each round given any history of play in the previous rounds. Formally, denote a *history* of play in the first k rounds as $h(k) = (\mathbf{a}^1, \dots, \mathbf{a}^k)$, and the set of all possible k -round histories as $H(k) = A^k$ ($H(0)$ denotes the singleton set containing the empty history). A (mixed) strategy of player i in $G(T)$ can be represented as $\mu_i : H \rightarrow \Delta A_i$, where $H = \cup_{k=0}^{T-1} H(k)$ is the set of all histories. This form of representation is also commonly known as *behavior strategies*. We use $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)$ to denote strategy profiles of $G(T)$, and the concept of best response and Nash equilibrium are defined in the same way as for the stage game.

In this research, we focus on *subgame-perfect equilibria* (SPE) of $G(T)$. SPE

was originally introduced by [82, 83] to eliminate NEs that involve non-credible threats off the equilibrium path. Given a strategy μ_i of player i for $G(T)$, denote $\mu_{i|h(k)}$ as the resulting strategy for subgame $G(T - k)$ obtained by conditioning μ_i on some history $h(k)$. Formally, given $h(k) = (\mathbf{a}^1, \dots, \mathbf{a}^k)$, $\mu_{i|h(k)}$ is given by: 1) $\mu_{i|h(k)}(h(0)) = \mu_i(\mathbf{a}^1, \dots, \mathbf{a}^k)$; 2) for any $t < T - k$ and any $(\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^t) \in H(t)$, $\mu_{i|h(k)}(\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^t) = \mu_i(\mathbf{a}^1, \dots, \mathbf{a}^k, \mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^t)$. And denote $\boldsymbol{\mu}_{|h(k)} = (\mu_{1|h(k)}, \dots, \mu_{n|h(k)})$. A strategy profile $\boldsymbol{\mu}$ is an SPE if for all $0 \leq k < T$ and all $h(k) \in H(k)$, $\boldsymbol{\mu}_{|h(k)}$ is an NE of $G(T - k)$. We use $\text{SPE}(G, T)$ to denote the set of all SPEs of the repeated game $G(T)$.

The phenomenon we are interested in is when in some SPE of the repeated game $G(T)$, the behavior strategy profile in some round does not form an NE of the stage game G . In other words, some player uses a locally suboptimal strategy in some round, in the sense that the strategy is not a best response for that round, as part of an SPE in the repeated game. We formally define this phenomenon of local suboptimality as follows.

Definition 2.1.1 (Local suboptimality). *Local suboptimality occurs in some SPE $\boldsymbol{\mu}$ of some repeated game $G(T)$ if there exists some $0 \leq k < T$ and play history $h(k) \in H(k)$ where $\boldsymbol{\mu}(h(k)) = (\mu_1(h(k)), \dots, \mu_n(h(k))) \notin \text{Nash}(G)$, i.e. the behavior strategy profile at some round does not form an NE of the stage game.*

We refer to such behavior strategy profiles that do not form an NE of the stage game as *off-(stage-game)-Nash* plays, or *off-Nash* plays in short. The emergence of locally suboptimal behavior reflects some fundamental psychological and social phenomena, such as delayed gratification, threats, and incentivized cooperation.

Denote the set of all stage games G as \mathcal{G} (\mathcal{G} is an infinite set). \mathcal{G} can be partitioned into two disjoint subsets \mathcal{G}_{LS} and \mathcal{G}_{LO} . \mathcal{G}_{LS} is the set of stage games G where local suboptimality occurs in some SPE of $G(T)$ for some T ; \mathcal{G}_{LO} is the set of stage games G where local suboptimality never occurs in any SPE of $G(T)$ for any T (LO stands for locally optimal). Our central research questions stated in Question 1.1.1 and Question 1.1.2 are essentially about solving the following problems: 1) completely

characterize \mathcal{G}_{LS} (thus \mathcal{G}_{LO}) using mathematical conditions, and 2) given any stage game G , algorithmically determine if G is in \mathcal{G}_{LS} or \mathcal{G}_{LO} .

2.2 Motivating Examples

In this section, we present several example games that motivate our study on the emergence of local suboptimality in finitely repeated games.

2.2.1 Example Game where Local Suboptimality Occurs

We first present a simple example game where local suboptimality occurs to give a flavor of how such phenomena arise.

	a_2	b_2	
a_1	(3,1)	(0,1)	
b_1	(2,1)	(1,1)	

Table 2.1: Example stage game G where local suboptimality occurs in an SPE of $G(2)$.

Example 2.2.1. *Table 2.1 presents an example stage game G where local suboptimality occurs in an SPE of the repeated game $G(2)$. The game is represented in matrix form. Row player chooses from actions a_1 and b_1 , column player chooses from actions a_2 and b_2 . In each entry of the matrix, the first value is the payoff of the row player, and the second value is the payoff of the column player.*

The strategy profile (b_1, a_2) is not an NE of the stage game G . However, the following is an SPE of the 2-round repeated game $G(2)$, in which the strategy profile in the first round is (b_1, a_2) :

- *In the second round, if the row player plays a_1 in the first round, play (b_1, b_2) ; else, play (a_1, a_2) .*
- *In the first round, play (b_1, a_2) .*

Notice that although the row player can obtain an additional payoff of 1 in the first round by switching to play a_1 in the first round, they will lose a payoff of 2 in the second round. This is why the above strategy profile is an SPE of the repeated game. Intuitively, the column player ‘threatens’ the row player by stating (implicitly through the column player’s strategy) that if the row player deviates in the first round, the column player will play according to the stage-game NE that gives a lower payoff to the row player in the second round.

2.2.2 Example Game where SPE with Local Suboptimality Strictly Dominates SPEs without Local Suboptimality

Here we present an example game where in the repeated game, some SPE in which local suboptimality occurs strictly dominates all SPEs where local suboptimality does not occur.

	a_2	b_2	c_2
a_1	(3,3)	(0,4)	(0,0)
b_1	(4,0)	(2,2)	(0,1)
c_1	(0,0)	(1,0)	(1,1)

Table 2.2: Example stage game where in the repeated game, some SPE in which local suboptimality occurs strictly dominates all SPEs where local suboptimality does not occur.

Example 2.2.2. Table 2.2 presents the example stage game in matrix form. This stage game G has three Nash equilibria: (b_1, b_2) , (c_1, c_2) , and a mixed NE (σ_1, σ_2) where $\sigma_1(b_1) = \sigma_1(c_1) = \sigma_2(b_2) = \sigma_2(c_2) = 0.5$. The payoffs of each of the above NEs are: $(2, 2)$, $(1, 1)$, and $(1, 1)$ respectively. Therefore, for a T -round repeated game $G(T)$, in any SPE where local suboptimality does not occur, the total payoff of each player is at most $2T$. We argue that for any $T > 2$, the following strategy profile is an SPE of $G(T)$:

- In the first $T - 2$ rounds, the row player plays a_1 and the column player plays a_2 (note that this strategy profile is not an NE of the stage game G). If any player

deviates to other actions in any round, the two players immediately switch to play (c_1, c_2) for the rest of the game.

- In the last 2 rounds, players play (b_1, b_2) .

The total payoff of each player under the above SPE is $3T - 2$. For all $T > 2$, $3T - 2 > 2T$. Therefore, the above SPE in which local suboptimality occurs strictly dominates any SPE in which local suboptimality does not occur.

2.2.3 Example Games Demonstrating Difference Between Local Suboptimality and the Property in Folk Theorems

Under the theme of analyzing equilibrium solutions in repeated games, a large body of work focuses on Folk Theorems [14, 13, 48, 84, 47], where the property of interest is: all *feasible* (i.e., the payoff profile lies in the convex hull of the set of all possible payoff profiles of the stage game) and *individually rational* (i.e., the payoff of each player is at least their minmax payoff in the stage game) payoff profiles can be attained in equilibria of the repeated game. Here we show that the property considered in Folk Theorems and the local suboptimality property we consider in this research are different, and the two properties do not have direct implications in either direction. We present 1) an example game where local suboptimality can occur, but not all feasible and individually rational payoffs can be attained in the repeated game, and 2) an example game where all feasible and individually rational payoffs can be attained in the repeated game, but local suboptimality cannot occur.

Example 2.2.3. Table 2.3 presents an example stage game G where local suboptimality can occur in the repeated game, but not all feasible and individually rational payoffs can be attained in the repeated game. This example is taken from [14]. This game contains 3 players. Player 1 selects rows (a_1, b_1, c_1) , player 2 selects columns (a_2, b_2) , and player 3 selects matrices (a_3, b_3) . While [14] analyzes this example with only pure strategies allowed, we consider the general case where mixed strategies are allowed. There are three Nash equilibria: (i) (a_1, a_2, a_3) ; (ii) (a_1, b_2, b_3) ; (iii) $(a_1, \sigma_2, \sigma_3)$ where

$\sigma_2(a_2) = \sigma_2(b_2) = 0.5$, $\sigma_3(a_3) = 0.25$, $\sigma_3(b_3) = 0.75$. These equilibria achieve payoffs of $(3, 3, 3)$, $(2, 2, 2)$, $(1.5, 1.5, 1.5)$ respectively. Following a similar idea in Example 2.2.1, it is easy to construct an SPE in a repeated game where local suboptimality occurs. For example, the following strategy profile is an SPE of $G(4)$:

- In the first round, play (a_1, a_2, b_3) .
- In the last three rounds, if players play in the first round is (a_1, a_2, b_3) , play (a_1, a_2, a_3) ; otherwise, play (a_1, b_2, b_3) .

In this SPE, the first round play does not form a stage-game NE. Therefore, local suboptimality occurs.

	a_2	b_2		a_2	b_2
a_1	$(3,3,3)$	$(0,0,0)$		a_1	$(1,1,1)$ $(2,2,2)$
b_1	$(0,0,0)$	$(0,0,0)$		b_1	$(0,1,1)$ $(0,1,1)$
c_1	$(0,1,1)$	$(0,0,0)$		c_1	$(0,1,1)$ $(0,0,0)$
	a_3			b_3	

Table 2.3: Example stage game where local suboptimality can occur, but not all feasible and individually rational payoffs can be attained in the repeated game.

For this stage game G , each player's minmax payoff is 0. We follow a similar argument as [14]. Denote $w_i(T)$ as the worst payoff that player i can get in any SPE of $G(T)$, the T -round repeated game. We claim that for $i = 2, 3$, $w_i(T)/T \geq 0.5$, therefore not all feasible and individually rational payoffs can be approximated. We use induction. The claim is true for $T = 1$. Suppose $w_i(T-1) \geq 0.5(T-1)$. Consider the strategy profile μ in $G(T)$ that attains $w_2(T)$ and $w_3(T)$. Notice that player 2 and 3 always get the same payoff in this game, so $w_2(T)$ and $w_3(T)$ will be attained at the same time. Consider the behavior strategy profile in the first round $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ as specified in μ . If $\sigma_1(c_1) \cdot \sigma_2(b_2) \leq 0.5$, then player 3 playing b_3 in the first round gives them at least a total payoff of $0.5 + w_3(T-1)$. This implies $w_3(T) \geq 0.5 + w_3(T-1)$ and we are done by the induction hypothesis. If $\sigma_1(c_1) \cdot \sigma_2(b_2) > 0.5$, then player 2 playing a_2 in the first round gives them at least a total payoff of $0.5 + w_2(T-1)$. This implies $w_2(T) \geq 0.5 + w_2(T-1)$ and we are done by the induction hypothesis.

Example 2.2.4. Table 2.4 presents an example stage game G where all feasible and individually rational payoffs can be attained in the repeated game, but local suboptimality cannot occur. For this stage game G , the set of all feasible and individually rational payoff profiles is $\{(u_1, u_2) \mid 0 \leq u_1 \leq 1, 0 \leq u_2 \leq 1\}$. All these feasible and individually rational payoffs can be attained in the stage game G itself, which is also a one-round repeated game $G(1)$. But every possible strategy profile in G is an NE of G , so local suboptimality cannot occur.

	a_2	b_2	
a_1	(0,0)	(1,0)	
b_1	(0,1)	(1,1)	

Table 2.4: Example stage game where every possible strategy profile is an NE.

2.2.4 Example Game where Local Suboptimality Only Occurs with Large T

One of the research questions we consider is in the computational aspect: given an arbitrary stage game G , how to (algorithmically) decide if there exists some T and some SPE of $G(T)$ where local suboptimality occurs? A naive approach is to enumerate over T , solve for all SPEs for each $G(T)$, and check if off-(stage-game)-Nash behavior occurs. Here, we present a construction of games where local suboptimality only occurs with arbitrarily large T . This means that the naive approach above might need to check an arbitrarily large number of T 's before returning a result.

	a_2	b_2	
a_1	(3,2)	(α ,1)	
b_1	(3,2)	(2,2)	
c_1	(α ,1)	(2,2)	

Table 2.5: Example stage game where local suboptimality only occurs with large T .

Example 2.2.5. Table 2.5 presents a construction of stage games where local suboptimality only occurs with arbitrarily large T . We claim that for any $\alpha < 2$, local suboptimality cannot occur with any $T < \frac{1}{2}(2 - \alpha)$, and local suboptimality can occur with any $T > 3 - \alpha$. Therefore, as α becomes smaller, we have games where local suboptimality only occurs with arbitrarily large T . We present the proof as follows.

It is easy to see that the set of Nash equilibria of this stage game G is

- (σ_1, a_2) where $\sigma_1(a_1) = \lambda, \sigma_1(b_1) = 1 - \lambda$ for all $0 \leq \lambda \leq 1$,
- (b_1, σ_2) where $\sigma_2(a_2) = \lambda, \sigma_2(b_2) = 1 - \lambda$ for all $0 \leq \lambda \leq 1$,
- (σ_1, b_2) where $\sigma_1(b_1) = \lambda, \sigma_1(c_1) = 1 - \lambda$ for all $0 \leq \lambda \leq 1$.

It follows that $V_1 = [2, 3]$, the continuous range from 2 to 3, and $V_2 = \{2\}$.

First, we show that for any $T > 3 - \alpha$, local suboptimality can occur. Here is an SPE where local suboptimality occurs:

- The first round strategy profile is $(\hat{\sigma}_1, b_2)$ where $\hat{\sigma}_1(a_1) = \frac{1}{4}$ and $\hat{\sigma}_1(c_1) = \frac{3}{4}$. $(\hat{\sigma}_1, b_2) \notin \text{Nash}(G)$ since player 1 is not playing a best response (but player 2 is playing a best response).
- If player 1's first round play is b_1 or c_1 , we let the players play a stage game Nash equilibrium that achieves $u_1 = 2$ (minimum payoff for player 1) in all the remaining $T - 1$ rounds.
- If player 1's first round play is a_1 , we let the players play a sequence of stage game Nash equilibria that achieves a total payoff $U_1 = 2(T - 1) + 2 - \alpha$ in the remaining $T - 1$ rounds. This is possible since $T - 1 > 2 - \alpha$ and V_1 contains the continuous interval between 2 and 3.

Now let T^* be the smallest T such that local suboptimality can occur in $G(T)$. Let μ^* to be any SPE of $G(T^*)$ where local suboptimality occurs. Denote $\sigma^* = (\sigma_1^*, \sigma_2^*)$ to be the first round strategy profile in μ^* . It follows that $\sigma^* \notin \text{Nash}(G)$, and all strategy profiles in all later rounds in μ^* belongs to $\text{Nash}(G)$. Since $|V_2| = 1$, player 2 must play a best response in σ^* . Therefore, σ_1^* must assign positive probabilities

in both a_1 and c_1 , since otherwise $\sigma^* \in \text{Nash}(G)$. For σ_2^* , either $\sigma_2^*(a_2) \geq 0.5$ or $\sigma_2^*(b_2) \geq 0.5$. If $\sigma_2^*(b_2) \geq 0.5$, we have $u_1(b_1, \sigma_2^*) - u_1(a_1, \sigma_2^*) \geq 0.5(2 - \alpha)$. Denote $U_1(\mu_{|b_1}^*)$ as the expected total payoff for player 1 in the last $T^* - 1$ rounds given player 1 plays b_1 in the first round, and similarly for $U_1(\mu_{|a_1}^*)$. For μ^* to be an SPE, we must have $U_1(\mu_{|a_1}^*) - U_1(\mu_{|b_1}^*) \geq u_1(b_1, \sigma_2^*) - u_1(a_1, \sigma_2^*) \geq 0.5(2 - \alpha)$. But we also have $U_1(\mu_{|a_1}^*) - U_1(\mu_{|b_1}^*) \leq 3(T^* - 1) - 2(T^* - 1) = T^* - 1$, so $T^* > 0.5(2 - \alpha)$. The same argument can be applied to the case where $\sigma_2^*(a_2) \geq 0.5$. Therefore, local suboptimality cannot occur with any $T < \frac{1}{2}(2 - \alpha)$.

2.3 The Pure Strategy Case

We start by considering 2-player games where players can only use pure strategies. In this case, a strategy of player i in $G(T)$ is $\mu_i : H \rightarrow A_i$. We use $\text{Nash}^{p,p}(G)$, $\text{SPE}^{p,p}(G, T)$, and $V_i^{p,p}$ to denote the corresponding concepts of $\text{Nash}(G)$, $\text{SPE}(G, T)$, and V_i when both players can only use pure strategies. We use $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LO}^{p,p}$ to denote the partition of the set of all stage games \mathcal{G} when both players can only use pure strategies. For any stage game G where $\text{Nash}^{p,p}(G) = \emptyset$, there is no SPE in the repeated game $G(T)$ for any T , so $G \in \mathcal{G}_{LO}^{p,p}$ since local suboptimality can never occur. The following theorem presents a complete mathematical characterization of $\mathcal{G}_{LS}^{p,p}$. As we will see, the key requirement for local suboptimality to occur is the ability of some player to ‘threaten’ the other player to play off stage-game NEs in some rounds.

Theorem 2.3.1 (2-player, pure strategy). *For 2-player pure-strategy-only games, a sufficient and necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

1. $|V_1^{p,p}| > 1$, $|V_2^{p,p}| > 1$, and there exists some $\hat{a}_1 \in A_1, \hat{a}_2 \in A_2$ where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$, OR
2. $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, and there exists $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 , OR

3. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$, and there exists $\hat{a}_1 \in A_1, \hat{a}_2, a'_2 \in A_2$ where $u_2(\hat{a}_1, \hat{a}_2) < u_2(\hat{a}_1, a'_2)$ and \hat{a}_1 is a best response to \hat{a}_2 .

Proof. First we show the condition is sufficient, by showing if the condition is satisfied, then we can construct some T and some SPE where local suboptimality occurs.

If the condition is satisfied, then at least one of (1),(2),(3) must be satisfied. If (1) is satisfied, there exists $\mathbf{a}_1, \mathbf{a}'_1 \in \text{Nash}^{p,p}(G)$ where $u_1(\mathbf{a}_1) > u_1(\mathbf{a}'_1)$ and $\mathbf{a}_2, \mathbf{a}'_2 \in \text{Nash}^{p,p}(G)$ where $u_2(\mathbf{a}_2) > u_2(\mathbf{a}'_2)$ (note that $\mathbf{a}_1, \mathbf{a}'_1$ do not need to be different from $\mathbf{a}_2, \mathbf{a}'_2$). From the given $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$, let $\delta_1 = \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)$ and $\delta_2 = \max_{a_2 \in A_2} u_2(\hat{a}_1, a_2) - u_2(\hat{a}_1, \hat{a}_2)$. We set $T = 2k + 1$ and $k \geq \max\left(\frac{\delta_1}{u_1(\mathbf{a}_1) - u_1(\mathbf{a}'_1)}, \frac{\delta_2}{u_2(\mathbf{a}_2) - u_2(\mathbf{a}'_2)}\right)$. We argue that the following strategy profile is an SPE of $G(T)$:

- In the first round, play (\hat{a}_1, \hat{a}_2) ,
- For later $2k$ rounds, if the first round play is (\hat{a}_1, \hat{a}_2) , players play their corresponding strategy according to $(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_1, \mathbf{a}_2, \dots)$; if the first round play is (a'_1, \hat{a}_2) where $a'_1 \neq \hat{a}_1$, players play their corresponding strategy according to $(\mathbf{a}'_1, \mathbf{a}_2, \mathbf{a}'_1, \mathbf{a}_2, \dots)$; if the first round play is (\hat{a}_1, a'_2) where $a'_2 \neq \hat{a}_2$, players play their corresponding strategy according to $(\mathbf{a}_1, \mathbf{a}'_2, \mathbf{a}_1, \mathbf{a}'_2, \dots)$; otherwise, players play their corresponding strategy according to $(\mathbf{a}_1, \mathbf{a}_1, \dots)$ (or any sequence of NEs).

For every subgame in the game tree starting from the second round or later, the above strategy profile forms an NE, since a stage game NE is played in every round. So we only need to show that the above strategy profile forms an NE for the root game $G(T)$. The total payoff for player 1 under the above strategy profile is $U_1 = u_1(\hat{a}_1, \hat{a}_2) + k \cdot (u_1(\mathbf{a}_1) + u_1(\mathbf{a}_2))$. If player 1 unilaterally deviates on the first round play (and possibly later rounds as well), the new total payoff $U'_1 \leq \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) + k \cdot (u_1(\mathbf{a}'_1) + u_1(\mathbf{a}_2)) \leq U_1$ due to our choice of k . The same argument applies for player 2, so the above strategy profile forms an NE for the root game. Local suboptimality occurs here since the first round behavior strategy profile $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$.

If (2) is satisfied, there exists $\mathbf{a}_1, \mathbf{a}'_1 \in \text{Nash}^{p,p}(G)$ where $u_1(\mathbf{a}_1) > u_1(\mathbf{a}'_1)$. Let $\delta = \max_{a_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)$. We set $T = k + 1$ and $k \geq \frac{\delta}{u_1(\mathbf{a}_1) - u_1(\mathbf{a}'_1)}$. We argue that the following strategy profile is an SPE of $G(T)$:

- In the first round, play (\hat{a}_1, \hat{a}_2) ,
- For the later k rounds, if player 1's first round play is \hat{a}_1 , players play their corresponding strategy according to $(\mathbf{a}_1, \mathbf{a}_1, \dots)$; if player 1's first round play is $\mathbf{a}'_1 \neq \hat{a}_1$, players play their corresponding strategy according to $(\mathbf{a}'_1, \mathbf{a}'_1, \dots)$.

Again, for every subgame in the game tree starting from the second round or later, the above strategy profile forms an NE. So we only need to show that the above strategy profile forms an NE for the root game. Player 2 cannot deviate to get a higher total payoff since \hat{a}_2 is the best response to \hat{a}_1 and u_2 is the same under all $\text{Nash}^{p,p}(G)$. For player 1, the total payoff under the above strategy profile is $U_1 = u_1(\hat{a}_1, \hat{a}_2) + k \cdot u_1(\mathbf{a}_1)$. If player 1 unilaterally deviates, the new total payoff $U'_1 \leq \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) + k \cdot u_1(\mathbf{a}'_1) \leq U_1$ due to our choice of k . So the above strategy profile forms an NE for the root game. Local suboptimality occurs here since $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$.

The same construction exchanging player 1 and 2 works for the case when (3) is satisfied. This finishes the proof that the condition is sufficient.

To prove this condition is necessary, we prove that if the condition is not satisfied, then for any T and any SPE $\boldsymbol{\mu}$ of $G(T)$, local suboptimality does not occur, i.e., the strategy profile at each round must form an NE of the stage game. The condition is not satisfied means all of (1),(2),(3) are false. This can be divided into the following disjoint cases.

1. $|V_1^{p,p}| = 0, |V_2^{p,p}| = 0$. Here, $\text{Nash}^{p,p}(G) = \emptyset$, so there is no SPE in the repeated game $G(T)$ for any T . Therefore, local suboptimality can never occur.
2. $|V_1^{p,p}| = 1, |V_2^{p,p}| = 1$. Using backward induction, we know that in any SPE, the strategy profile at each round must form an NE of the stage game.
3. $|V_1^{p,p}| > 1, |V_2^{p,p}| > 1$. Since (1) is false, there does not exist \hat{a}_1, \hat{a}_2 where

$(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$, i.e., $A = \text{Nash}^{p,p}(G)$ (Example 2.3.2 shows an example of such games). Therefore, it trivially follows that in any SPE of $G(T)$, the strategy profile at each round must form an NE of the stage game.

4. $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$. Since (2) is false, there does not exist $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 (Example 2.3.3 shows an example of such games). This means that for any \hat{a}_1, \hat{a}_2 where \hat{a}_2 is a best response to \hat{a}_1 , \hat{a}_1 is a best response to \hat{a}_2 , and thus $(\hat{a}_1, \hat{a}_2) \in \text{Nash}^{p,p}(G)$. Now we can use backward induction to prove that in any SPE, the strategy profile at each round must form an NE of the stage game. The strategy profiles in the last round must form NEs. Given that the strategy profiles in the last k rounds must all form NEs, consider the $(k + 1)$ -to-last round. Player 2 must play a best response in this round, since their play in this round does not affect the total payoff they get in the final k rounds. And since all strategy profiles where player 2 plays best response is an NE of the stage game, the induction step is complete.

5. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$. The same proof for case 4 applies here.

This finishes the proof that the above condition is necessary. \square

A reader may wonder whether there exists stage games G that belong to cases 3 and 4 in the above proof for the necessity of the condition. We present here example games that belong to each case.

Example 2.3.2. Table 2.4 presents an example stage game G where $|V_1^{p,p}| > 1$, $|V_2^{p,p}| > 1$, and there does not exist \hat{a}_1, \hat{a}_2 where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{p,p}(G)$, i.e., $A = \text{Nash}^{p,p}(G)$.

Example 2.3.3. Table 2.6 presents an example stage game G where $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, and there does not exist $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 . The set of pure Nash equilibria of G is: (a_1, a_2) , (b_1, a_2) , (b_1, b_2) , and (c_1, b_2) . Therefore, $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| = 1$. We can see that for all strategy profiles where player 2 plays a best response, player 1 also plays a best response.

	a_2	b_2
a_1	(3,2)	(1,1)
b_1	(3,2)	(2,2)
c_1	(1,1)	(2,2)

Table 2.6: Example stage game in matrix form, row player is player 1, column player is player 2. For all $a \in A_1$, for all $\sigma_2 \in \Delta A_2$ that is a best response to a , a is also a best response to σ_2 .

From the constructions of SPEs where local suboptimality occurs used in the above proof, we can obtain the following corollary regarding the value of T above which local suboptimality can occur if the condition in Theorem 2.3.1 is satisfied:

Corollary 2.3.4. *For 2-player pure-strategy-only games, given a stage game G :*

1. *If $|V_1^{p,p}| > 1$, $|V_2^{p,p}| > 1$, and there exists some $\hat{a}_1 \in A_1, \hat{a}_2 \in A_2$ where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}(G)$, then for all $T \geq 2 \cdot \max\left(\frac{\delta_1}{\max(V_1^{p,p}) - \min(V_1^{p,p})}, \frac{\delta_2}{\max(V_2^{p,p}) - \min(V_2^{p,p})}\right) + 1$ where $\delta_1 = \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)$ and $\delta_2 = \max_{a_2 \in A_2} u_2(\hat{a}_1, a_2) - u_2(\hat{a}_1, \hat{a}_2)$, there exists some SPE of $G(T)$ where local suboptimality occurs.*
2. *If $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, and there exists $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 , then for all $T \geq \frac{\max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)}{\max(V_1^{p,p}) - \min(V_1^{p,p})} + 1$, there exists some SPE of $G(T)$ where local suboptimality occurs.*
3. *If $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$, and there exists $\hat{a}_1 \in A_1, \hat{a}_2, a'_2 \in A_2$ where $u_2(\hat{a}_1, \hat{a}_2) < u_2(\hat{a}_1, a'_2)$ and \hat{a}_1 is a best response to \hat{a}_2 , then for all $T \geq \frac{\max_{a_2 \in A_2} u_2(\hat{a}_1, a_2) - u_2(\hat{a}_1, \hat{a}_2)}{\max(V_2^{p,p}) - \min(V_2^{p,p})} + 1$, there exists some SPE of $G(T)$ where local suboptimality occurs.*

2.4 The General Case

Now we consider the general case for 2-player games where mixed strategies are allowed. We use $\text{Nash}^{m,m}(G)$, $\text{SPE}^{m,m}(G, T)$, and $V_i^{m,m}$ to denote the corresponding concepts of $\text{Nash}(G)$, $\text{SPE}(G, T)$, and V_i when both players can use mixed strategies. Since mixed Nash equilibrium always exists for G [65], $|V_1^{m,m}| \geq 1$, $|V_2^{m,m}| \geq 1$. We

use $\mathcal{G}_{LS}^{m,m}$ and $\mathcal{G}_{LO}^{m,m}$ to denote the partition of the set of all stage games \mathcal{G} when both players can use mixed strategies. The following theorem presents a complete mathematical characterization of $\mathcal{G}_{LS}^{m,m}$.

Theorem 2.4.1 (2-player, general case). *For general 2-player games (mixed strategies allowed), a sufficient and necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

1. $|V_1^{m,m}| > 1$, $|V_2^{m,m}| > 1$, and there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2$ where $(\hat{\sigma}_1, \hat{\sigma}_2) \notin \text{Nash}^{m,m}(G)$, OR
2. $|V_1^{m,m}| > 1$, $|V_2^{m,m}| = 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{\sigma}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$, OR
3. same as (2) but exchange player 1 and 2.

We first establish some useful lemmas.

Lemma 2.4.2. *For any two-player game G , if there exists $\sigma_1 \in \Delta A_1$ and $\sigma_2 \in \Delta A_2$ where $(\sigma_1, \sigma_2) \notin \text{Nash}^{m,m}(G)$, then there exists $a_1 \in A_1$ and $a_2 \in A_2$ where $(a_1, a_2) \notin \text{Nash}^{m,m}(G)$.*

Proof. $(\sigma_1, \sigma_2) \notin \text{Nash}^{m,m}(G)$ implies that there exists some $a'_1 \in A_1$ where $u_1(\sigma_1, \sigma_2) < u_1(a'_1, \sigma_2)$, or there exists some $a'_2 \in A_2$ where $u_2(\sigma_1, \sigma_2) < u_2(\sigma_1, a'_2)$. We consider the case of there exists some a'_1 where $u_1(\sigma_1, \sigma_2) < u_1(a'_1, \sigma_2)$, and the same argument applies to the other case. Since $u_1(\sigma_1, \sigma_2) \geq \min_{a_1 \in S_{\sigma_1}} u_1(a_1, \sigma_2)$, there exists some a_1, a'_1, σ_2 where $u_1(a_1, \sigma_2) < u_1(a'_1, \sigma_2)$. So $u_1(a'_1, \sigma_2) - u_1(a_1, \sigma_2) = \sum_{a_2 \in S_{\sigma_2}} \sigma_2(a_2) \cdot (u_1(a'_1, a_2) - u_1(a_1, a_2)) > 0$, which means there exists some a_2 where $u_1(a'_1, a_2) - u_1(a_1, a_2) > 0$. This $(a_1, a_2) \notin \text{Nash}^{m,m}(G)$, which finishes the proof. \square

Lemma 2.4.3. *For any two-player game G , define*

$$I = \left\{ (i, j) \mid i \in A_1, j \in A_2, u_2(i, j) = \max_{j' \in A_2} u_2(i, j') \right\}$$

as the set of pure strategy profiles where player 2 plays a best response. If

1. $|V_1^{m,m}| > 1$, $|V_2^{m,m}| = 1$, and
2. there does not exist $\hat{a}_1 \in A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{a}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to \hat{a}_1 , and
3. there does not exist $a_1 \in A_1$ and $a_2, a'_2 \in A_2$ where $a_2 \neq a'_2$ and both a_2 and a'_2 are best responses to a_1 ,

then,

- (a). $I \subseteq \text{Nash}^{m,m}(G)$,
- (b). for each $i \in A_1$, there is a unique $j \in A_2$ such that $(i, j) \in I$,
- (c). there exists $b \in \mathbb{R}$ such that for all $(i, j) \in I$, $u_2(i, j) = b$, and for all $(i', j') \notin I$, $u_2(i', j') < b$,
- (d). there exists $(i, j), (i', j') \in I$ such that $u_1(i, j) \neq u_1(i', j')$, and $i \neq i', j \neq j'$.

Proof. (2) directly implies (a). From the definition of I , for each $i \in A_1$, there is at least one $j \in A_2$ such that $(i, j) \in I$. This combines with (3) implies (b).

Since $|V_2^{m,m}| = 1$ and $I \subseteq \text{Nash}^{m,m}(G)$, for all $(i, j) \in I$, $u_2(i, j) = b$ where b is the only element in $V_2^{m,m}$. It then follows from the definition of I that for all $(i', j') \notin I$, $u_2(i', j') < b$. So (c) follows.

For (d), assume in contradiction that all $u_1(i, j)$ for $(i, j) \in I$ are the same. Since $|V_1^{m,m}| > 1$, there exists $\sigma, \sigma' \in \text{Nash}^{m,m}(G)$ such that $u_1(\sigma) \neq u_1(\sigma')$. Denote \mathcal{S}_σ as the set of pure strategy profiles (i, j) that occur with non-zero probability under the strategy profile σ . Then at least one of \mathcal{S}_σ and $\mathcal{S}_{\sigma'}$ needs to contain elements not in I , since otherwise $u_1(\sigma) = u_1(\sigma')$. WLOG, let \mathcal{S}_σ contain elements not in I . By (c), $u_2(\sigma) < b$, which contradicts with $|V_2^{m,m}| = 1$. So there exists $(i, j), (i', j') \in I$ such that $u_1(i, j) \neq u_1(i', j')$. For such $(i, j), (i', j')$, if $i = i'$, then (b) implies that $j = j'$, which contradicts with $u_1(i, j) \neq u_1(i', j')$. So $i \neq i'$. And since $I \subseteq \text{Nash}^{m,m}(G)$ (due to (a)), $(i, j), (i', j')$ are both NEs with different payoffs for player 1, so $j \neq j'$. Therefore, (d) follows. □

Now we are ready to prove Theorem 2.4.1.

Proof of Theorem 2.4.1. Following the same argument as the proof for the pure strategy case (Theorem 2.3.1), we can show the condition is necessary.

We prove the condition is sufficient by showing if the condition is satisfied, we can always construct some T and some SPE where local suboptimality occurs. If the condition is satisfied, then at least one of (1),(2),(3) must be satisfied. We consider each case here.

If (1) is satisfied, there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2$ where $(\hat{\sigma}_1, \hat{\sigma}_2) \notin \text{Nash}^{m,m}(G)$. By Lemma 2.4.2, there exists $\hat{a}_1 \in A_1$ and $\hat{a}_2 \in A_2$ where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{m,m}(G)$. Then we can use the same construction that is used in the proof of the pure strategy case here (see the proof of sufficiency in Theorem 2.3.1, the part that handles the case where (1) is satisfied).

The rest of the proof focus on the case when (2) is satisfied. The same argument applies for the case where (3) is satisfied. We first notice that all games that satisfy (2) can be categorized into the following 3 disjoint cases:

- (a). There exists $\hat{a}_1 \in A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{a}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to \hat{a}_1 .
- (b). (a) is false, and there exists $a_1 \in A_1$ and $a_2, a'_2 \in A_2$ where $a_2 \neq a'_2$ and both a_2 and a'_2 are best responses to a_1 .
- (c). Both (a) and (b) are false.

We consider each case here.

Case (a). We can use the same construction that is used in the proof of sufficiency for the pure strategy case (Theorem 2.3.1), the part that handles the case where (2) is satisfied.

Case (b). (a) is false implies that for all $a_1 \in A_1$, for all $\sigma_2 \in \Delta A_2$ that is a best response to a_1 , a_1 is also a best response to σ_2 . Table 2.6 is an example of such

games. We know that there exists some $a^* \in A_1$, $b_1^*, b_2^* \in A_2$ where $b_1^* \neq b_2^*$ and both b_1^* and b_2^* are best responses to a^* . Therefore, a^* is a best response to both b_1^* and b_2^* . WLOG, let $u_1(a^*, b_1^*) \geq u_1(a^*, b_2^*)$. Denote $\sigma_\lambda \in \Delta A_2$ as the mixed strategy for player 2 which assigns $\sigma_\lambda(b_1^*) = \lambda$ and $\sigma_\lambda(b_2^*) = 1 - \lambda$. Then for all $0 \leq \lambda \leq 1$, $(a^*, \sigma_\lambda) \in \text{Nash}^{m,m}(G)$.

Since (2) is satisfied, there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{\sigma}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$. We construct $T = 1 + T_1 + T_2$ and a strategy profile μ^* for $G(T)$ with the following structure:

- In the first round, play $(\hat{\sigma}_1, \hat{\sigma}_2)$.
- For the later rounds, if player 1's first round play is $i \in S_{\hat{\sigma}_1}$, players play their corresponding strategies according to SPE μ^i of $G(T - 1)$; otherwise, players play their corresponding strategies according to SPE μ^\perp of $G(T - 1)$.

Since $|V_1^{m,m}| > 1$, let $\sigma^{\min}, \sigma^{\max} \in \text{Nash}^{m,m}(G)$ such that $u_1(\sigma^{\min}) = \min(V_1^{m,m})$ and $u_1(\sigma^{\max}) = \max(V_1^{m,m})$, so $u_1(\sigma^{\max}) > u_1(\sigma^{\min})$. We construct the SPEs $\mu^\perp, \{\mu^i\}_{i \in S_{\hat{\sigma}_1}}$ as follows:

- μ^\perp is players playing σ^{\min} repeatedly for $T_1 + T_2$ rounds.
- For all μ^i , the last T_2 rounds consist of players repeated playing σ^{\max} . T_2 is chosen to be large enough such that $U_1(\mu^i) - U_1(\mu^\perp) > \max_{a, a' \in A_1} u_1(a, \hat{\sigma}_2) - u_1(a', \hat{\sigma}_2)$ for every $i \in S_{\hat{\sigma}_1}$. This makes sure that in μ^* , player 1 deviating to any $i \notin S_{\hat{\sigma}_1}$ in the first round will reduce their total payoff in $G(T)$.
- The first T_1 rounds strategies for each μ^i adopt the following structure:
 - In the first round, play $(a^*, \sigma_{\lambda_i})$, where λ_i is a parameter to be set for each i .
 - In the latter $T_1 - 1$ rounds, if player 2 plays b_1^* in the first round, players repeatedly play σ^{\max} ; otherwise, players repeatedly play σ^{\min} .

Pick $i^m \in S_{\hat{\sigma}_1}$ such that $u_1(i^m, \hat{\sigma}_2) = \max_{i \in S_{\hat{\sigma}_1}} u_1(i, \hat{\sigma}_2)$. We set $\lambda_{i^m} = 0$. For each $i \in S_{\hat{\sigma}_1} \setminus \{i^m\}$, we set λ_i such that $U_1(\mu^i) + u_1(i, \hat{\sigma}_2) = U_1(\mu^{i^m}) + u_1(i^m, \hat{\sigma}_2)$.

This makes sure that in $\boldsymbol{\mu}^*$, player 1 choosing any $i \in S_{\hat{\sigma}_1}$ in the first round will obtain the same total payoff in $G(T)$. We argue that with large enough T_1 , such choice of λ_i 's is always possible. Consider the difference between two sides of the equation as a function of λ_i , $f(\lambda_i) = U_1(\boldsymbol{\mu}^i) - U_1(\boldsymbol{\mu}^{i^m}) + u_1(i, \hat{\sigma}_2) - u_1(i^m, \hat{\sigma}_2)$. By choosing $T_1 \geq \frac{\max_{i \in S_{\hat{\sigma}_1}} u_1(i^m, \hat{\sigma}_2) - u_1(i, \hat{\sigma}_2)}{\max(V_1^{m,m}) - \min(V_1^{m,m})} + 1$, we have $f(0) \leq 0$ and $f(1) \geq 0$. Since $f(\lambda_i)$ is a continuous function, there must exist some $\lambda_i \in [0, 1]$ such that $f(\lambda_i) = 0$ as desired.

One can easily verify that $\boldsymbol{\mu}^\perp$ and $\{\boldsymbol{\mu}^i\}_{i \in S_{\hat{\sigma}_1}}$ are SPEs of $G(T - 1)$. In addition, their construction ensures that $\boldsymbol{\mu}^*$ is an NE of the root game $G(T)$. Therefore, $\boldsymbol{\mu}^*$ is an SPE of $G(T)$ where the first round strategy profile does not form an NE of the stage game G .

Case (c). Since both (a) and (b) are false, and $|V_1^{m,m}| > 1$, $|V_2^{m,m}| = 1$, applying Lemma 2.4.3, we know that there exists $(i_1, j_1), (i_2, j_2) \in I$ such that $i_1 \neq i_2, j_1 \neq j_2$, where I is the set of pure strategy profiles where player 2 plays a best response, as defined in Lemma 2.4.3. Take such $(i_1, j_1), (i_2, j_2)$, Lemma 2.4.3 further implies that for all $j \neq j_1, u_2(i_1, j) < b$, and for all $j \neq j_2, u_2(i_2, j) < b$, where b is the only element in $V_2^{m,m}$. Denote $\hat{\sigma}_\lambda \in \Delta A_1$ as the mixed strategy for player 1 which assigns $\hat{\sigma}_\lambda(i_1) = \lambda$ and $\hat{\sigma}_\lambda(i_2) = 1 - \lambda$. Denote $J(\lambda) = \{a_2 \mid a_2 \in A_2, a_2 \text{ is a best response to } \hat{\sigma}_\lambda\}$ as the set of best response pure strategies for player 2 against $\hat{\sigma}_\lambda$. It is helpful to consider a geometric interpretation of $J(\lambda)$. For each $j \in A_2$, $u_2(\hat{\sigma}_\lambda, j) = \lambda \cdot u_2(i_1, j) + (1 - \lambda) \cdot u_2(i_2, j)$ is a linear function in λ . We can plot the function $f_j(\lambda) = u_2(\hat{\sigma}_\lambda, j)$ for each $j \in A_2$, which gives $|A_2|$ straight lines within domain $[0, 1]$. $J(\lambda)$ is then the set of lines that attains the maximum value at λ . We know that $J(0) = \{j_2\}$ and $J(1) = \{j_1\}$, so there must exist some $\lambda_1 \in (0, 1)$ where $|J(\lambda_1)| > 1$, which corresponds to some intersection point. Take such λ_1 and $\hat{j}_1, \hat{j}_2 \in J(\lambda_1)$ where $\hat{j}_1 \neq \hat{j}_2$. Denote $\hat{\sigma}_\rho \in \Delta A_2$ as the mixed strategy for player 2 which assigns $\hat{\sigma}_\rho(\hat{j}_2) = \rho$ and $\hat{\sigma}_\rho(\hat{j}_1) = 1 - \rho$. Then for all $\rho \in [0, 1]$, $\hat{\sigma}_\rho$ is a best response to $\hat{\sigma}_{\lambda_1}$, which implies that $\hat{\sigma}_{\lambda_1}$ is not a best response to $\hat{\sigma}_\rho$. This is because if $\hat{\sigma}_{\lambda_1}$ is a best response to $\hat{\sigma}_\rho$, then $(\hat{\sigma}_{\lambda_1}, \hat{\sigma}_\rho) \in \text{Nash}^{m,m}(G)$, but $u_2(\hat{\sigma}_{\lambda_1}, \hat{\sigma}_\rho) < b$, which contradicts with $|V_2^{m,m}| = 1$.

Now we show that we can always construct some T and some SPE $\boldsymbol{\mu}^*$ of $G(T)$ where the first round strategy profile is $(\hat{\sigma}_{\lambda_1}, \hat{\sigma}_{\rho})$ for some ρ . Since the above argument shows that $(\hat{\sigma}_{\lambda_1}, \hat{\sigma}_{\rho}) \notin \text{Nash}^{m,m}(G)$, local suboptimality occurs in $\boldsymbol{\mu}^*$. We treat the case where $u_1(i_1, \hat{j}_1) = u_1(i_2, \hat{j}_1)$ and $u_1(i_1, \hat{j}_1) \neq u_1(i_2, \hat{j}_1)$ separately.

If $u_1(i_1, \hat{j}_1) = u_1(i_2, \hat{j}_1)$, we construct $\boldsymbol{\mu}^*$ as:

- In the first round, play $(\hat{\sigma}_{\lambda}, \hat{j}_1)$. (\hat{j}_1 is $\hat{\sigma}_{\rho}$ with $\rho = 0$)
- For the later rounds, if player 1's first round play is i_1 or i_2 , players repeatedly play $\boldsymbol{\sigma}^{\max}$; otherwise, players repeatedly play $\boldsymbol{\sigma}^{\min}$.

Again, $\boldsymbol{\sigma}^{\min}, \boldsymbol{\sigma}^{\max} \in \text{Nash}^{m,m}(G)$ such that $u_1(\boldsymbol{\sigma}^{\min}) = \min(V_1^{m,m})$ and $u_1(\boldsymbol{\sigma}^{\max}) = \max(V_1^{m,m})$. T is chosen to be large enough such that player 1 deviating to any $i \notin \{i_1, i_2\}$ in the first round will reduce their total payoff in $G(T)$. It can be easily checked that $\boldsymbol{\mu}^*$ is an SPE of $G(T)$.

If $u_1(i_1, \hat{j}_1) \neq u_1(i_2, \hat{j}_1)$, WLOG, assume $u_1(i_1, \hat{j}_1) > u_1(i_2, \hat{j}_1)$. We construct $T = 1 + T_1 + T_2$ and $\boldsymbol{\mu}^*$ with the following structure:

- In the first round, play $(\hat{\sigma}_{\lambda}, \hat{\sigma}_{\rho})$.
- For the later rounds, if the first round play is (i_1, \hat{j}_2) , players play their corresponding strategy according to SPE $\boldsymbol{\mu}^1$ of $G(T - 1)$; if the first round play is (i_2, \hat{j}_2) , (i_1, \hat{j}_1) or (i_2, \hat{j}_1) , players play their corresponding strategy according to SPE $\boldsymbol{\mu}^2$ of $G(T - 1)$; otherwise, players play their corresponding strategy according to SPE $\boldsymbol{\mu}^{\perp}$ of $G(T - 1)$.
 - $\boldsymbol{\mu}^{\perp}$ is players play $\boldsymbol{\sigma}^{\min}$ repeatedly for $T_1 + T_2$ rounds.
 - $\boldsymbol{\mu}^2$ is players play $\boldsymbol{\sigma}^{\max}$ repeatedly for $T_1 + T_2$ rounds.
 - $\boldsymbol{\mu}^1$ is players play $\boldsymbol{\sigma}^{\min}$ repeatedly for T_1 rounds, and then $\boldsymbol{\sigma}^{\max}$ for T_2 rounds.

T_2 is chosen to be large enough such that player 1 deviating to any $i \notin \{i_1, i_2\}$ in the first round will reduce their total payoff in $G(T)$. In order for $\boldsymbol{\mu}^*$ to be an NE of the root game, player 1 choosing i_1 and i_2 in the first round need to yield the same

total payoff in $G(T)$. The total payoff of player 1 achieved by choosing i_1 in the first round under $\boldsymbol{\mu}^*$ is $\rho \cdot \left(u_1(i_1, \hat{j}_2) + U_1(\boldsymbol{\mu}^1) \right) + (1 - \rho) \cdot \left(u_1(i_1, \hat{j}_1) + U_1(\boldsymbol{\mu}^2) \right)$, and the total payoff by choosing i_2 is $\rho \cdot \left(u_1(i_2, \hat{j}_2) + U_1(\boldsymbol{\mu}^2) \right) + (1 - \rho) \cdot \left(u_1(i_2, \hat{j}_1) + U_1(\boldsymbol{\mu}^2) \right)$. Consider the difference between these two quantities as a function of ρ , $g(\rho) = \rho \cdot \left(U_1(\boldsymbol{\mu}^1) - U_1(\boldsymbol{\mu}^2) \right) + \rho \cdot \left(u_1(i_1, \hat{j}_2) - u_1(i_2, \hat{j}_2) \right) + (1 - \rho) \cdot \left(u_1(i_1, \hat{j}_1) - u_1(i_2, \hat{j}_1) \right)$. We have $g(0) > 0$. By choosing a large enough T_1 , $g(1) < 0$. Since $g(\rho)$ is a continuous function, there must exist some $\rho \in (0, 1)$ such that $g(\rho) = 0$. With this value of ρ , $\boldsymbol{\mu}^*$ is an SPE of $G(T)$ as desired. □

Again, from the constructions of SPEs where local suboptimality occurs used in the above proof, we can obtain the following corollary regarding the value of T above which local suboptimality can occur if the condition in Theorem 2.4.1 is satisfied:

Corollary 2.4.4. *For general 2-player games (mixed strategies allowed), given a stage game G :*

1. *If $|V_1^{m,m}| > 1$, $|V_2^{m,m}| > 1$, and there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2$ where $(\hat{\sigma}_1, \hat{\sigma}_2) \notin \text{Nash}^{m,m}(G)$, then by Lemma 2.4.2, there exists $\hat{a}_1 \in A_1$ and $\hat{a}_2 \in A_2$ where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{m,m}(G)$, for all $T \geq 2 \cdot \max \left(\frac{\delta_1}{\max(V_1^{m,m}) - \min(V_1^{m,m})}, \frac{\delta_2}{\max(V_2^{m,m}) - \min(V_2^{m,m})} \right) + 1$ where $\delta_1 = \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)$ and $\delta_2 = \max_{a_2 \in A_2} u_2(\hat{a}_1, a_2) - u_2(\hat{a}_1, \hat{a}_2)$, there exists some SPE of $G(T)$ where local suboptimality occurs.*
2. *If $|V_1^{m,m}| > 1$, $|V_2^{m,m}| = 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{\sigma}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to $\hat{\sigma}_1$, then*
 - (a) *If there exists $\hat{a}_1 \in A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{a}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to \hat{a}_1 , then for all $T \geq \frac{\max_{a_1 \in A_1} u_1(a_1, \hat{\sigma}_2) - u_1(\hat{a}_1, \hat{\sigma}_2)}{\max(V_1^{m,m}) - \min(V_1^{m,m})} + 1$, there exists some SPE of $G(T)$ where local suboptimality occurs.*
 - (b) *If there does not exist $\hat{a}_1 \in A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{a}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to \hat{a}_1 , and there exists $a_1 \in A_1$ and $a_2, a'_2 \in A_2$ where $a_2 \neq a'_2$ and both a_2 and a'_2 are best responses to a_1 , then*

for all $T \geq 3 + \frac{\max_{a,a' \in S_{\hat{\sigma}_1}} u_1(a, \hat{\sigma}_2) - u_1(a', \hat{\sigma}_2)}{\max(V_1^{m,m}) - \min(V_1^{m,m})} + \frac{\max_{a,a' \in A_1} u_1(a, \hat{\sigma}_2) - u_1(a', \hat{\sigma}_2)}{\max(V_1^{m,m}) - \min(V_1^{m,m})}$, there exists some SPE of $G(T)$ where local suboptimality occurs.

(c) If there does not exist $\hat{a}_1 \in A_1, \hat{\sigma}_2 \in \Delta A_2, a'_1 \in A_1$ where $u_1(\hat{a}_1, \hat{\sigma}_2) < u_1(a'_1, \hat{\sigma}_2)$ and $\hat{\sigma}_2$ is a best response to \hat{a}_1 , and there does not exist $a_1 \in A_1$ and $a_2, a'_2 \in A_2$ where $a_2 \neq a'_2$ and both a_2 and a'_2 are best responses to a_1 . By Lemma 2.4.3, there exists $(i_1, j_1), (i_2, j_2) \in I$ such that $i_1 \neq i_2, j_1 \neq j_2$, where I is the set of pure strategy profiles where player 2 plays a best response. Denote $\hat{\sigma}_\lambda \in \Delta A_1$ as the mixed strategy for player 1 which assigns $\hat{\sigma}_\lambda(i_1) = \lambda$ and $\hat{\sigma}_\lambda(i_2) = 1 - \lambda$. By the proof of Theorem 2.4.1, there exists $\lambda_1 \in (0, 1), \hat{j}_1, \hat{j}_2 \in A_2$ where \hat{j}_1, \hat{j}_2 are both best responses to $\hat{\sigma}_{\lambda_1}$ and $\hat{j}_1 \neq \hat{j}_2$. If $u_1(i_1, \hat{j}_1) = u_1(i_2, \hat{j}_1)$, then for all $T \geq \frac{\max_{a_1 \in A_1} u_1(a_1, \hat{j}_1) - u_1(i_1, \hat{j}_1)}{\max(V_1^{m,m}) - \min(V_1^{m,m})} + 1$, there exists some SPE of $G(T)$ where local suboptimality occurs. If $u_1(i_1, \hat{j}_1) \neq u_1(i_2, \hat{j}_1)$, then for all $T \geq 3 + \frac{\max_{a_1 \in A_1, i \in \{i_1, i_2\}, j \in \{\hat{j}_1, \hat{j}_2\}} u_1(a_1, j) - u_1(i, j)}{\max(V_1^{m,m}) - \min(V_1^{m,m})} + \frac{|u_1(i_1, \hat{j}_2) - u_1(i_2, \hat{j}_2)|}{\max(V_1^{m,m}) - \min(V_1^{m,m})}$, there exists some SPE of $G(T)$ where local suboptimality occurs.

3. same as (2) but exchange player 1 and 2.

2.5 The Pure Strategy Against Mixed Strategy Case

To complete the picture, we also analyze the case where one player can only use pure strategies while the other player can use mixed strategies. Without loss of generality, we consider the case where player 1 can use mixed strategies and player 2 can only use pure strategies in both the stage game and the repeated games. We use $\text{Nash}^{m,p}(G)$, $\text{SPE}^{m,p}(G, T)$, and $V_i^{m,p}$ to denote the corresponding concepts of $\text{Nash}(G)$, $\text{SPE}(G, T)$, and V_i when player 1 can use mixed strategies and player 2 can only use pure strategies. We use $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LO}^{m,p}$ to denote the partition of the set of all stage games \mathcal{G} when player 1 can use mixed strategies and player 2 can only use pure strategies. For stage games G where $\text{Nash}^{m,p}(G) = \emptyset$, there is no SPE

in the repeated game $G(T)$ for any T , so we categorize such stage games G to $\mathcal{G}_{LO}^{m,p}$ since local suboptimality can never occur. The following theorem presents a complete mathematical characterization of $\mathcal{G}_{LS}^{m,p}$. As we will see, the sufficient and necessary condition for local suboptimality to occur in this case is different from both the pure strategy case and the general case.

Theorem 2.5.1 (2-player, pure strategy against mixed strategy). *For 2-player games where player 1 can use mixed strategies and player 2 can only use pure strategies, a sufficient and necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

1. $|V_1^{m,p}| > 1$, $|V_2^{m,p}| > 1$, and there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where $(\hat{\sigma}_1, \hat{a}_2) \notin \text{Nash}^{m,p}(G)$, OR

2. $|V_1^{m,p}| > 1$, $|V_2^{m,p}| = 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where

(a) $u_1(\hat{\sigma}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$, and

(b) \hat{a}_2 is a best response to $\hat{\sigma}_1$, and

(c) if $\hat{\sigma}_1$ has more than one support (not a pure strategy), denote the set of possible differences in u_1 between pairs of NEs in the stage game as $D = \{u_1(\sigma) - u_1(\sigma') \mid \sigma, \sigma' \in \text{Nash}^{m,p}(G)\}$, there exists an action a from the support of $\hat{\sigma}_1$, i.e., $a \in S_{\hat{\sigma}_1}$, such that, for every $a' \in S_{\hat{\sigma}_1} \setminus a$, there exists some integer $n_{a'} \geq 0$ and $d_k^{a'} \in D$, $k = 1, \dots, n_{a'}$ such that $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = \sum_{k=1}^{n_{a'}} d_k^{a'}$,

OR

3. $|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2, a'_2 \in A_2$ where $u_2(\hat{\sigma}_1, \hat{a}_2) < u_2(\hat{\sigma}_1, a'_2)$ and $\hat{\sigma}_1$ is a best response to \hat{a}_2 .

We first establish a useful lemma.

Lemma 2.5.2. *For 2-player stage game G and T -round repeated game $G(T)$ where player 1 can use mixed strategies and player 2 can only use pure strategies, for any*

SPE $\boldsymbol{\mu}$ of $G(T)$ where the strategy profile at each round forms a stage-game NE, player 1's total payoff in the repeated game $U_1(\boldsymbol{\mu}) = \sum_{k=1}^T c_k$ for some $c_k \in V_1^{m,p}$, $k = 1, \dots, T$, i.e., player 1's total payoff in the repeated game equals the sum of some sequence of stage-game NE payoffs.

Proof. We prove by induction on T . The proposition trivially holds for $T = 1$. Given that the proposition holds for $T = K - 1$, consider $T = K$. For any SPE $\boldsymbol{\mu}$ of $G(K)$ where the strategy profile at each round forms a stage-game NE, denote the first round strategy profile as $(\hat{\sigma}_1, \hat{a}_2)$ where $\hat{\sigma}_1 \in \Delta A_1$, $\hat{a}_2 \in A_2$. Since $(\hat{\sigma}_1, \hat{a}_2) \in \text{Nash}^{m,p}(G)$, for all $a, a' \in S_{\hat{\sigma}_1}$, $u_1(a, \hat{a}_2) = u_1(a', \hat{a}_2) = u_1(\hat{\sigma}_1, \hat{a}_2)$. Denote $\boldsymbol{\mu}_{|(a_1, a_2)}$ as the strategy profile starting from round 2 given that players play (a_1, a_2) in the first round, for $\boldsymbol{\mu}$ to be a Nash equilibrium of the root game, we have for all $a, a' \in S_{\hat{\sigma}_1}$, $u_1(a, \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}) = u_1(a', \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a', \hat{a}_2)}) = U_1(\boldsymbol{\mu})$. By the induction hypothesis, for any $a \in S_{\hat{\sigma}_1}$, $U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}) = \sum_{k=1}^{K-1} c_k$ for some $c_k \in V_1^{m,p}$, $k = 1, \dots, K - 1$. Therefore, $U_1(\boldsymbol{\mu}) = u_1(a, \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}) = u_1(\hat{\sigma}_1, \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}) = \sum_{k=1}^K c_k$ for some $c_k \in V_1^{m,p}$, $k = 1, \dots, K$. This completes the induction step. \square

Proof of Theorem 2.5.1. First we show the condition is sufficient, by showing if the condition is satisfied, we can construct some T and some SPE where local suboptimality occurs. If the condition is satisfied, at least one of (1),(2),(3) must be satisfied. We consider each case here.

If (1) is satisfied, we can use the same construction that is used for the proofs of the pure strategy case and the general case (see the proofs of Theorems 2.3.1 and 2.4.1, the parts that handle the case where (1) is satisfied). If (3) is satisfied, we can use the same construction that is used for the proof of the pure strategy case (see the proofs of Theorem 2.3.1, the parts that handle the case where (2) is satisfied).

If (2) is satisfied, then there exists $\hat{\sigma}_1 \in \Delta A_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where (a), (b), and (c) are satisfied. If $|S_{\hat{\sigma}_1}| = 1$, i.e., $\hat{\sigma}_1$ is a pure strategy, we can use the same construction that is used for the proof of the pure strategy case (see the proofs of Theorem 2.3.1, the parts that handle the case where (2) is satisfied). If $|S_{\hat{\sigma}_1}| > 1$, i.e., $\hat{\sigma}_1$ is not a pure strategy, we know from (c) that there exists an action $a \in$

$S_{\hat{\sigma}_1}$ such that, for every $a' \in S_{\hat{\sigma}_1} \setminus a$, there exists some integer $n_{a'} \geq 0$ and $d_k^{a'} \in D$, $k = 1, \dots, n_{a'}$ such that $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = \sum_{k=1}^{n_{a'}} d_k^{a'}$. Let $a, n_{a'}$ for every $a' \in S_{\hat{\sigma}_1} \setminus a$, $d_k^{a'} \in D$ for every $a' \in S_{\hat{\sigma}_1} \setminus a$ and $k = 1, \dots, n_{a'}$ be such a set of assignments. We consider a game with $T = 1 + \sum_{a' \in S_{\hat{\sigma}_1} \setminus a} n_{a'} + n_{\perp}$ where n_{\perp} is a large enough integer. We construct an SPE μ^* consisting of segments. Denote μ^t as all the t -th round behavior strategy profiles in μ^* and $\mu^{t_1:t_2}$ as all the behavior strategy profiles between the t_1 -th round and the t_2 -th round in μ^* . μ^* is divided into segments: $\mu^1, \mu^{2:T_1}, \mu^{T_1+1:T_2}, \dots, \mu^{T_{|S_{\hat{\sigma}_1}|-1}+1:T_{|S_{\hat{\sigma}_1}|}}$. Each segment ending at T_i for $i = 1, \dots, |S_{\hat{\sigma}_1}| - 1$ corresponds to one of $a' \in S_{\hat{\sigma}_1} \setminus a$. We denote the segment corresponding to a' as $\mu^{a'}$ for every $a' \in S_{\hat{\sigma}_1} \setminus a$; $\mu^{a'}$ has $n_{a'}$ rounds. We denote the final segment as μ^{\perp} , which has n_{\perp} rounds. Each $\mu^{a'}$ for every $a' \in S_{\hat{\sigma}_1} \setminus a$ and μ^{\perp} only depend on the play in the first round, not depending on any later rounds. we construct μ^* as follows:

- In the first round, play $\mu^1 = (\hat{\sigma}_1, \hat{a}_2)$.
- For segment $\mu^{a'}$ for each $a' \in S_{\hat{\sigma}_1} \setminus a$, if player 1's first round play is a' , play the sequence of stage-game NEs $\Sigma_1 = (\sigma_1^1, \dots, \sigma_1^{n_{a'}})$; otherwise, play the sequence of stage-game NEs $\Sigma_0 = (\sigma_0^1, \dots, \sigma_0^{n_{a'}})$. The sequences Σ_1 and Σ_0 are chosen according to $u_1(\sigma_1^k) - u_1(\sigma_0^k) = d_k^{a'}$ for all $k = 1, \dots, n_{a'}$.
- Let $\sigma^{\min}, \sigma^{\max} \in \text{Nash}^{m,p}(G)$ such that $u_1(\sigma^{\min}) = \min(V_1^{m,p})$ and $u_1(\sigma^{\max}) = \max(V_1^{m,p})$, so $u_1(\sigma^{\max}) > u_1(\sigma^{\min})$. For segment μ^{\perp} , if player 1's first round play is some $a \in S_{\hat{\sigma}_1}$, play $(\sigma^{\max}, \sigma^{\max}, \dots)$; if player 1's first round play is some $a \notin S_{\hat{\sigma}_1}$, play $(\sigma^{\min}, \sigma^{\min}, \dots)$.

This construction ensures that: 1) player 1 choosing any action $a \in S_{\hat{\sigma}_1}$ in the first round results in the same total payoff in $G(T)$, and 2) player 1 choosing any action $a \notin S_{\hat{\sigma}_1}$ in the first round results in a lower total payoff in $G(T)$ as long as n_{\perp} is large enough. This ensures that such μ^* is an SPE of $G(T)$ and the first round play does not form a stage-game NE.

To prove this condition is necessary, we prove that if the condition is not satisfied, then for any T and any SPE μ of $G(T)$, local suboptimality does not occur, i.e., the strategy profile at each round must form an NE of the stage game. The condition is not satisfied means all of (1),(2),(3) are false. This can be divided into the following disjoint cases.

1. $|V_1^{m,p}| = 0, |V_2^{m,p}| = 0$. Here, $\text{Nash}^{m,p}(G) = \emptyset$, so there is no SPE in the repeated game $G(T)$ for any T . Therefore, local suboptimality can never occur.

2. $|V_1^{m,p}| = 1, |V_2^{m,p}| = 1$. Using backward induction, we know that in any SPE, the strategy profile at each round must form an NE of the stage game.

3. $|V_1^{m,p}| > 1, |V_2^{m,p}| > 1$. Since (1) is false, there does not exist $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where $(\hat{\sigma}_1, \hat{a}_2) \notin \text{Nash}^{m,p}(G)$. Therefore, it trivially follows that in any SPE of $G(T)$, the strategy profile at each round must form an NE of the stage game.

4. $|V_1^{m,p}| = 1, |V_2^{m,p}| > 1$. Since (3) is false, there does not exist $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2, a'_2 \in A_2$ where $u_2(\hat{\sigma}_1, \hat{a}_2) < u_2(\hat{\sigma}_1, a'_2)$ and $\hat{\sigma}_1$ is a best response to \hat{a}_2 . This means that for any $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where $\hat{\sigma}_1$ is a best response to \hat{a}_2 , \hat{a}_2 is a best response to $\hat{\sigma}_1$, and thus $(\hat{\sigma}_1, \hat{a}_2) \in \text{Nash}^{m,p}(G)$. Now we can use the same backward induction argument as in the proof of Theorem 2.3.1 (the part that proves the condition is necessary, case 4) to prove that in any SPE, the strategy profile at each round must form an NE of the stage game.

5. $|V_1^{m,p}| > 1, |V_2^{m,p}| = 1$. Since (2) is false, we know that:

(*) For all pure strategy $a_1 \in A_1$ and $a_2 \in A_2$, if a_2 is a best response to a_1 , then a_1 is also a best response to a_2 , i.e., $(a_1, a_2) \in \text{Nash}^{m,p}(G)$.

(**) For all $\sigma_1 \in \Delta A_1$ and $a_2 \in A_2$ where a_2 is a best response to σ_1 and σ_1 is not a best response to a_2 and σ_1 is not a pure strategy, denote the set of possible differences in u_1 between pairs of NEs in the stage game as $D = \{u_1(\sigma) - u_1(\sigma') \mid \sigma, \sigma' \in \text{Nash}^{m,p}(G)\}$, then there does not exist an action from the support of σ_1 $a \in S_{\sigma_1}$ such that, for every $a' \in S_{\sigma_1} \setminus a$, there exists some integer $n_{a'} \geq 0$ and $d_k^{a'} \in D, k = 1, \dots, n_{a'}$ such that $u_1(a, a_2) - u_1(a', a_2) = \sum_{k=1}^{n_{a'}} d_k^{a'}$.

Now we can use backward induction to prove that in any SPE, the strategy profile at each round must form an NE of the stage game. The strategy profiles in the last round must form stage-game NEs. Given that the strategy profiles in the last k rounds must all form stage-game NEs, consider the $(k+1)$ -to-last round. Denote the strategies played in this round as $(\hat{\sigma}_1, \hat{a}_2)$. \hat{a}_2 must be a best response to $\hat{\sigma}_1$, since player 2's play in this round does not affect the total payoff they get in the final k rounds. If $\hat{\sigma}_1$ is a pure strategy, then according to (*), $(\hat{\sigma}_1, \hat{a}_2)$ forms a stage-game NE, which completes the induction step.

If $\hat{\sigma}_1$ is a mixed strategy (more than one support), assume on the contrary that the strategy profile at this round does not form a stage-game NE, then $\hat{\sigma}_1$ is not a best response to \hat{a}_2 . By (**), there does not exist an action from the support of $\hat{\sigma}_1$ $a \in S_{\hat{\sigma}_1}$ such that, for every $a' \in S_{\hat{\sigma}_1} \setminus a$, there exists some integer $n_{a'} \geq 0$ and $d_k^{a'} \in D$, $k = 1, \dots, n_{a'}$ such that $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = \sum_{k=1}^{n_{a'}} d_k^{a'}$. Denote $\boldsymbol{\mu}_{|(a_1, a_2)}^{-k}$ as the strategy profile in the last k rounds given players played (a_1, a_2) in the first round. For $(\hat{\sigma}_1, \hat{a}_2)$ to be part of a Nash equilibrium of the $(k+1)$ -round repeated game, we have for all $a, a' \in S_{\hat{\sigma}_1}$, $u_1(a, \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}^{-k}) = u_1(a', \hat{a}_2) + U_1(\boldsymbol{\mu}_{|(a', \hat{a}_2)}^{-k})$. By Lemma 2.5.2, for each $a \in S_{\hat{\sigma}_1}$, $U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}^{-k}) = \sum_{t=1}^k c_t^a$ for some $c_t^a \in V_1^{m,p}$, $t = 1, \dots, k$. Therefore, taking an arbitrary $a \in S_{\hat{\sigma}_1}$, for every $a' \in S_{\hat{\sigma}_1} \setminus a$, $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = U_1(\boldsymbol{\mu}_{|(a', \hat{a}_2)}^{-k}) - U_1(\boldsymbol{\mu}_{|(a, \hat{a}_2)}^{-k}) = \sum_{t=1}^k c_t^{a'} - \sum_{t=1}^k c_t^a = \sum_{t=1}^k d_t^{a'}$, where $d_t^{a'} = c_t^{a'} - c_t^a \in D$. This produces a contradiction. Therefore, the strategy profile at the $(k+1)$ -to-last round also forms a stage-game NE. This completes the induction step. \square

Again, we can obtain the following corollary regarding the value of T above which local suboptimality can occur if the condition in Theorem 2.5.1 is satisfied:

Corollary 2.5.3. *For 2-player games where player 1 can use mixed strategies and player 2 can only use pure strategies, given a stage game G :*

1. *If $|V_1^{m,p}| > 1$, $|V_2^{m,p}| > 1$, and there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where $(\hat{\sigma}_1, \hat{a}_2) \notin \text{Nash}(G)$, then by Lemma 2.4.2, there exists $\hat{a}_1 \in A_1$ and $\hat{a}_2 \in A_2$ where $(\hat{a}_1, \hat{a}_2) \notin \text{Nash}^{m,p}(G)$, for all $T \geq 2 \cdot \max \left(\frac{\delta_1}{\max(V_1^{m,p}) - \min(V_1^{m,p})}, \frac{\delta_2}{\max(V_2^{m,p}) - \min(V_2^{m,p})} \right) +$*

1 where $\delta_1 = \max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{a}_1, \hat{a}_2)$ and $\delta_2 = \max_{a_2 \in A_2} u_2(\hat{a}_1, a_2) - u_2(\hat{a}_1, \hat{a}_2)$, there exists some SPE of $G(T)$ where local suboptimality occurs.

2. If $|V_1^{m,p}| > 1$, $|V_2^{m,p}| = 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where

(a) $u_1(\hat{\sigma}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$, and

(b) \hat{a}_2 is a best response to $\hat{\sigma}_1$, and

(c) if $\hat{\sigma}_1$ has more than one support (not a pure strategy), denote the set of possible differences in u_1 between pairs of NEs in the stage game as $D = \{u_1(\sigma) - u_1(\sigma') \mid \sigma, \sigma' \in \text{Nash}^{m,p}(G)\}$, there exists an action from the support of $\hat{\sigma}_1$ $a \in S_{\hat{\sigma}_1}$ such that, for every $a' \in S_{\hat{\sigma}_1} \setminus a$, there exists some integer $n_{a'} \geq 0$ and $d_k^{a'} \in D, k = 1, \dots, n_{a'}$ such that $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = \sum_{k=1}^{n_{a'}} d_k^{a'}$,

then if $\hat{\sigma}_1$ is a pure strategy, i.e., $|S_{\hat{\sigma}_1}| = 1$, for all $T \geq \frac{\max_{a_1 \in A_1} u_1(a_1, \hat{a}_2) - u_1(\hat{\sigma}_1, \hat{a}_2)}{\max(V_1^{m,p}) - \min(V_1^{m,p})} + 1$, there exists some SPE of $G(T)$ where local suboptimality occurs; if $|S_{\hat{\sigma}_1}| > 1$, for all $T \geq 2 + \sum_{a' \in S_{\hat{\sigma}_1} \setminus a} n_{a'} + \frac{\max_{i, i' \in A_1} u_1(i, \hat{a}_2) - u_1(i', \hat{a}_2)}{\max(V_1^{m,p}) - \min(V_1^{m,p})}$, where a and $n_{a'}$ for every $a' \in S_{\hat{\sigma}_1} \setminus a$ are given in (c), there exists some SPE of $G(T)$ where local suboptimality occurs.

3. If $|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$, and there exists $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2, a'_2 \in A_2$ where

$u_2(\hat{\sigma}_1, \hat{a}_2) < u_2(\hat{\sigma}_1, a'_2)$ and $\hat{\sigma}_1$ is a best response to \hat{a}_2 , then for all $T \geq \frac{\max_{a_2 \in A_2} u_2(\hat{\sigma}_1, a_2) - u_2(\hat{\sigma}_1, \hat{a}_2)}{\max(V_2^{m,p}) - \min(V_2^{m,p})} +$

1, there exists some SPE of $G(T)$ where local suboptimality occurs.

2.6 Effect of Changing from Pure Strategies to Mixed Strategies on the Emergence of Local Suboptimality

In Sections 2.3 to 2.5, we established sufficient and necessary conditions on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality

occurs for 2-player games, for cases where: 1) both players can only use pure strategies, 2) one player can only use pure strategies and the other player can use mixed strategies, and 3) both players can use mixed strategies. Essentially, we established a complete characterization of $\mathcal{G}_{LS}^{p,p}$, $\mathcal{G}_{LS}^{m,p}$, and $\mathcal{G}_{LS}^{m,m}$ (and therefore $\mathcal{G}_{LO}^{p,p}$, $\mathcal{G}_{LO}^{m,p}$, and $\mathcal{G}_{LO}^{m,m}$). Based on these results, in this section we study the effect of changing from pure strategies to mixed strategies on the emergence of local suboptimality. We aim to answer the following question: under what conditions on the stage game G will allowing players to play mixed strategies change whether local suboptimality can ever occur in some repeated game $G(T)$?

Essentially, we aim to study the relationships between $\mathcal{G}_{LS}^{p,p}$, $\mathcal{G}_{LS}^{m,p}$, and $\mathcal{G}_{LS}^{m,m}$. For example, are there stage games G where $G \notin \mathcal{G}_{LS}^{p,p}$ and $G \in \mathcal{G}_{LS}^{m,p}$, i.e., local suboptimality can never occur when both players can only use pure strategies but can occur when player 1 obtains access to mixed strategies? What is a complete characterization of such stage games? And in the other direction, are there stage games G where $G \in \mathcal{G}_{LS}^{p,p}$ and $G \notin \mathcal{G}_{LS}^{m,p}$? We also study the corresponding questions for the relationships between $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,m}$ and between $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LS}^{m,m}$.

For the simplicity of descriptions, we refer to the case where both players can only use pure strategies as the *pure-pure* case, the case where player 1 can use mixed strategies and player 2 can only use pure strategies as the *mixed-pure* case, and the case where both players can use mixed strategies as the *mixed-mixed* case.

2.6.1 Changing from Pure Strategies to Mixed Strategies when the Other Player Can Only Use Pure Strategies

We first analyze the situation for changing from the pure-pure case to the mixed-pure case, i.e., study the relationship between $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LS}^{m,p}$.

We first establish a useful theorem:

Theorem 2.6.1. *For all 2-player stage games G , for all $T \in \mathbb{Z}^+$, for all $\mu \in SPE^{p,p}(G, T)$, $\mu \in SPE^{m,p}(G, T)$.*

Proof. Assume in contradiction that there exists some G^* , T^* , and $\mu^* \in SPE^{p,p}(G^*, T^*)$

where $\boldsymbol{\mu}^* \notin \text{SPE}^{m,p}(G^*, T^*)$. Let k^* to be the largest k where $\boldsymbol{\mu}_{|h(k)}^*$ is not an NE of $G^*(T - k)$ for some history $h(k)$ in the mixed-pure case. Then one of the players must be able to unilaterally change their strategy in the first round of $\boldsymbol{\mu}_{|h(k)}^*$ to obtain a higher total payoff in $G^*(T - k)$. Player 2 cannot do so since they can still only play pure strategies. For player 1, they cannot do so when they can only play pure strategies, which means any alternative actions in the first round of $\mu_{1|h(k)}^*$ cannot lead to a higher total payoff in $G^*(T - k)$. But this means that any alternative mixed strategies in the first round of $\mu_{1|h(k)}^*$ also cannot lead to a higher total payoff in $G^*(T - k)$. This produces a contradiction. Therefore, for all 2-player stage games G , for all $T \in \mathbb{Z}^+$, for all $\boldsymbol{\mu} \in \text{SPE}^{p,p}(G, T)$, $\boldsymbol{\mu} \in \text{SPE}^{m,p}(G, T)$. \square

Now we present Theorems 2.6.2 and 2.6.8 and Lemmas 2.6.3 and 2.6.4, which together completely characterize the relationship between $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LS}^{m,p}$.

Theorem 2.6.2. $\mathcal{G}_{LS}^{p,p} \subseteq \mathcal{G}_{LS}^{m,p}$, i.e., for 2-player stage games G , if there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the pure-pure case, then there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case.

Proof. If there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the pure-pure case, denote T^* to be such a T and $\boldsymbol{\mu}^*$ to be such an SPE of $G(T^*)$. By Theorem 2.6.1, $\boldsymbol{\mu}^*$ is also an SPE of $G(T^*)$ in the mixed-pure case. Since any pure strategy profiles not in $\text{Nash}^{p,p}(G)$ are also not in $\text{Nash}^{m,p}(G)$, local suboptimality occurs in $\boldsymbol{\mu}^*$ in the mixed-pure case. Therefore, there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case. \square

Lemma 2.6.3. For all stage games G where $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| > 1$, if $G \notin \mathcal{G}_{LS}^{p,p}$, then $G \notin \mathcal{G}_{LS}^{m,p}$.

Proof. If $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$, and $G \notin \mathcal{G}_{LS}^{p,p}$, by Theorem 2.3.1, for all $a_2 \in A_2$, for all $a_1 \in A_1$ that is a best response to a_2 , a_2 is also a best response to a_1 . Then for any $(\sigma_1, a_2) \in \text{Nash}^{m,p}(G)$, for any $a_1 \in S_{\sigma_1}$, a_1 is a best response to a_2 , so a_2 is a best response to a_1 , which means $(a_1, a_2) \in \text{Nash}^{p,p}(G)$ and $u_1(a_1, a_2) \in$

$V_1^{p,p}$. And since $u_1(\sigma_1, a_2) = u_1(a_1, a_2)$ for any $a_1 \in S_{\sigma_1}$, $u_1(\sigma_1, a_2) \in V_1^{p,p}$. Thus, $V_1^{m,p} \subseteq V_1^{p,p}$. Furthermore, since $\text{Nash}^{p,p}(G) \subseteq \text{Nash}^{m,p}(G)$, $V_1^{p,p} \subseteq V_1^{m,p}$. Therefore, $|V_1^{p,p}| = |V_1^{m,p}|$. So $|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$. We argue that there cannot exist $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2, a'_2 \in A_2$ where $u_2(\hat{\sigma}_1, \hat{a}_2) < u_2(\hat{\sigma}_1, a'_2)$ and $\hat{\sigma}_1$ is a best response to \hat{a}_2 . Assuming there exists such $\hat{\sigma}_1$ and \hat{a}_2 . Then for all $a_1 \in S_{\hat{\sigma}_1}$, a_1 is a best response to \hat{a}_2 , which means \hat{a}_2 is also a best response to a_1 . Then \hat{a}_2 is also a best response to $\hat{\sigma}_1$, which produces a contradiction. Therefore, by Theorem 2.5.1, there does not exist some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case, so $G \notin \mathcal{G}_{LS}^{m,p}$. \square

Lemma 2.6.4. *For all stage games G where $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| > 1$, if $G \in \mathcal{G}_{LS}^{m,p}$, then $G \in \mathcal{G}_{LS}^{p,p}$.*

Proof. If $|V_1^{p,p}| > 1$, $|V_2^{p,p}| > 1$, and $G \in \mathcal{G}_{LS}^{m,p}$, since $\text{Nash}^{p,p}(G) \subseteq \text{Nash}^{m,p}(G)$, $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| > 1$. Then since there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case, by Theorem 2.5.1, there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where $(\hat{\sigma}_1, \hat{a}_2) \notin \text{Nash}^{m,p}(G)$. Applying Lemma 2.4.2, there exists some $a_1 \in A_1$ and $a_2 \in A_2$ where $(a_1, a_2) \notin \text{Nash}^{p,p}(G)$. By Theorem 2.3.1, there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the pure-pure case, so $G \in \mathcal{G}_{LS}^{p,p}$. \square

The above two lemmas show that under certain preconditions on $|V_1^{p,p}|$ and $|V_2^{p,p}|$, $\mathcal{G}_{LS}^{p,p} = \mathcal{G}_{LS}^{m,p}$. Now the question is whether $\mathcal{G}_{LS}^{p,p}$ always equals $\mathcal{G}_{LS}^{m,p}$. Here we show that this is not the case. We present example stage games G where $G \in \mathcal{G}_{LS}^{m,p}$ but $G \notin \mathcal{G}_{LS}^{p,p}$ for each of the remaining cases regarding the values of $|V_1^{p,p}|$ and $|V_2^{p,p}|$. This shows that $\mathcal{G}_{LS}^{p,p} \neq \mathcal{G}_{LS}^{m,p}$. Combined with Theorem 2.6.2, this shows that $\mathcal{G}_{LS}^{p,p}$ is a proper subset of $\mathcal{G}_{LS}^{m,p}$.

Example 2.6.5. *Table 2.7 presents an example stage game G where $|V_1^{p,p}| = 0$, $|V_2^{p,p}| = 0$, $G \in \mathcal{G}_{LS}^{m,p}$, and $G \notin \mathcal{G}_{LS}^{p,p}$. When both players can only use pure strategies, there is no Nash equilibrium, so $|V_1^{p,p}| = 0$ and $|V_2^{p,p}| = 0$. By Theorem 2.3.1, $G \notin \mathcal{G}_{LS}^{p,p}$. When player 1 can use mixed strategies and player 2 can only use pure strategies, the set of Nash equilibria are: (σ_1, b_2) and (σ_1, c_2) for all σ_1 where $0.25 \leq$*

$\sigma_1(a_1) \leq 0.75$. So $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| = 1$. And (b_1, a_2) is a strategy profile where player 1 does not play a best response and player 2 plays a best response. Therefore, the condition in Theorem 2.5.1 is satisfied, so $G \in \mathcal{G}_{LS}^{m,p}$.

	a_2	b_2	c_2	d_2
a_1	(4,0)	(1,3)	(2,3)	(0,4)
b_1	(0,4)	(1,3)	(2,3)	(4,0)

Table 2.7: Example stage game in matrix form, row player is player 1, column player is player 2. When both players can only use pure strategies, there is no Nash equilibrium. When player 1 can use mixed strategies and player 2 can only use pure strategies, there are multiple Nash equilibria, and local suboptimality can occur in repeated games.

Example 2.6.6. Table 2.8 presents an example stage game G where $|V_1^{p,p}| = 1$, $|V_2^{p,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,p}$, and $G \notin \mathcal{G}_{LS}^{p,p}$. When both players can only use pure strategies, the only Nash equilibria are (a_1, a_2) and (b_1, c_2) , so $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| = 1$. By Theorem 2.3.1, $G \notin \mathcal{G}_{LS}^{p,p}$. When player 1 can use mixed strategies and player 2 can only use pure strategies, (σ_1, b_2) for all σ_1 where $0.25 \leq \sigma_1(a_1) \leq 0.75$ are Nash equilibria. So $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| > 1$. And $(b_1, a_2) \notin \text{Nash}^{m,p}(G)$. Therefore, the condition in Theorem 2.5.1 is satisfied, so $G \in \mathcal{G}_{LS}^{m,p}$.

	a_2	b_2	c_2
a_1	(4,4)	(1,3)	(0,0)
b_1	(0,0)	(1,3)	(4,4)

Table 2.8: Example stage game in matrix form, row player is player 1, column player is player 2. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| = 1$, $|V_1^{m,p}| > 1$, and $|V_2^{m,p}| > 1$.

Example 2.6.7. Table 2.6 presents an example stage game G where $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,p}$, and $G \notin \mathcal{G}_{LS}^{p,p}$. For this game G , the set of pure Nash equilibria is: (a_1, a_2) , (b_1, a_2) , (b_1, b_2) , and (c_1, b_2) . Therefore, $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| = 1$. We can see that for all pure strategy profiles where player 2 plays a best response, player 1 also plays a best response. So the condition in Theorem 2.3.1 is not satisfied, and $G \notin \mathcal{G}_{LS}^{p,p}$. The strategy profile (σ_1, a_2) where $\sigma_1(a_1) = 0.9$ and $\sigma_1(c_1) = 0.1$ is an

example strategy profile where player 2 plays a best response and player 1 does not play a best response. And $u_1(a_1, a_2) - u_1(c_1, a_2)$ can be expressed as the sum of some sequence of values in $D = \{u_1(\sigma) - u_1(\sigma') \mid \sigma, \sigma' \in \text{Nash}^{m,p}(G)\}$. So the condition in Theorem 2.5.1 is satisfied, and $G \in \mathcal{G}_{LS}^{m,p}$. Intuitively speaking, in the pure-pure case, although $|V_1^{p,p}| > 1$ makes player 1 potentially vulnerable to threats, there is no way to construct such a threat since there is no strategy profiles where player 1 does not play a best response but player 2 plays a best response. But in the mixed-pure case, such off-(stage-game)-Nash strategy profiles become available, which makes the threat possible.

Furthermore, the following theorem completely characterizes $\mathcal{G}_{LS}^{m,p} \setminus \mathcal{G}_{LS}^{p,p}$, i.e., the set of 2-player stage games G where 1) in the pure-pure case, local suboptimality can never occur, and 2) in the mixed-pure case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs.

Theorem 2.6.8. *For 2-player stage games G , $G \notin \mathcal{G}_{LS}^{p,p}$ and $G \in \mathcal{G}_{LS}^{m,p}$ if and only if*

1. $|V_1^{p,p}| = 0$, $|V_2^{p,p}| = 0$, and the condition in Theorem 2.5.1 is satisfied, OR
2. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| = 1$, and the condition in Theorem 2.5.1 is satisfied, OR
3. $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, there does not exist $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 , and the condition in Theorem 2.5.1 is satisfied.

Proof. To show the above condition is sufficient, we show that each of 1, 2, and 3 is sufficient.

1. $|V_1^{p,p}| = 0$ and $|V_2^{p,p}| = 0$ mean that G has no pure Nash equilibrium. So in the pure-pure case, there is no SPE for any repeated game $G(T)$, therefore local suboptimality can never occur. Furthermore, the condition in Theorem 2.5.1 is satisfied implies that in the mixed-pure case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs. Example 2.6.5 presents a example stage game G that satisfies this condition.

2. Since $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| = 1$, by Theorem 2.3.1, in the pure-pure case, local suboptimality can never occur. And the condition in Theorem 2.5.1 is satisfied implies that in the mixed-pure case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs. Example 2.6.6 presents an example stage game G that satisfies this condition.

3. Since $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, and there does not exist $\hat{a}_1, a'_1 \in A_1, \hat{a}_2 \in A_2$ where $u_1(\hat{a}_1, \hat{a}_2) < u_1(a'_1, \hat{a}_2)$ and \hat{a}_2 is a best response to \hat{a}_1 , by Theorem 2.3.1, in the pure-pure case, local suboptimality can never occur. And the condition in Theorem 2.5.1 is satisfied implies that in the mixed-pure case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs. Example 2.6.7 presents an example stage game G that satisfies this condition.

To show the condition is necessary, we split all stage games G into five disjoint cases based on $|V_1^{p,p}|$ and $|V_2^{p,p}|$:

(a) $|V_1^{p,p}| = 0$ and $|V_2^{p,p}| = 0$,

(b) $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| = 1$,

(c) $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| = 1$,

(d) $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| > 1$,

(e) $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| > 1$.

For cases (a), (b), and (c), a direct application of Theorems 2.3.1 and 2.5.1 implies that our target condition is necessary. For cases (d) and (e), Lemmas 2.6.3 and 2.6.4 show that $\mathcal{G}_{LS}^{m,p} \setminus \mathcal{G}_{LS}^{p,p}$ is empty under these two cases. Therefore, our target condition is necessary. □

2.6.2 Changing from Pure Strategies to Mixed Strategies when the Other Player Can Use Mixed Strategies

Next, we analyze the situation for changing from the mixed-pure case to the mixed-mixed case, i.e., study the relationship between $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,m}$.

Following the same argument as the proof of Theorem 2.6.1, we can prove the following theorem, which is useful for the results in this part.

Theorem 2.6.9. *For all 2-player stage games G , for all $T \in \mathbb{Z}^+$, for all $\boldsymbol{\mu} \in \text{SPE}^{m,p}(G, T)$, $\boldsymbol{\mu} \in \text{SPE}^{m,m}(G, T)$.*

Proof. Assume in contradiction that there exists some G^* , T^* , and $\boldsymbol{\mu}^* \in \text{SPE}^{m,p}(G^*, T^*)$ where $\boldsymbol{\mu}^* \notin \text{SPE}^{m,m}(G^*, T^*)$. Let k^* to be the largest k where $\boldsymbol{\mu}_{|h(k)}^*$ is not an NE of $G^*(T - k)$ for some history $h(k)$ in the mixed-mixed case. Then one of the players must be able to unilaterally change their strategy in the first round of $\boldsymbol{\mu}_{|h(k)}^*$ to obtain a higher total payoff in $G^*(T - k)$. Player 1 cannot do so since their strategy space is the same in the mixed-pure case and the mixed-mixed case. For player 2, they cannot do so when they can only play pure strategies, which means any alternative actions in the first round of $\boldsymbol{\mu}_{|h(k)}^*$ cannot lead to a higher total payoff in $G^*(T - k)$. But this means that any alternative mixed strategies in the first round of $\boldsymbol{\mu}_{|h(k)}^*$ also cannot lead to a higher total payoff in $G^*(T - k)$. This produces a contradiction. Therefore, for all 2-player stage games G , for all $T \in \mathbb{Z}^+$, for all $\boldsymbol{\mu} \in \text{SPE}^{m,p}(G, T)$, $\boldsymbol{\mu} \in \text{SPE}^{m,m}(G, T)$. \square

Now we present Theorems 2.6.10 and 2.6.16 and Lemma 2.6.11, which together completely characterize the relationship between $\mathcal{G}_{LS}^{m,p}$ and $\mathcal{G}_{LS}^{m,m}$.

Theorem 2.6.10. *$\mathcal{G}_{LS}^{m,p} \subseteq \mathcal{G}_{LS}^{m,m}$, i.e., for 2-player stage games G , if there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case, then there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-mixed case.*

Proof. If there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case, denote T^* to be such a T and $\boldsymbol{\mu}^*$ to be such an SPE of $G(T^*)$.

By Theorem 2.6.9, μ^* is also an SPE of $G(T^*)$ in the mixed-mixed case. For any strategy profile (σ_1, a_2) where $\sigma_1 \in \Delta A_1$ and $a_2 \in A_2$, if $(\sigma_1, a_2) \notin \text{Nash}^{m,p}(G)$, $(\sigma_1, a_2) \notin \text{Nash}^{m,m}(G)$. So local suboptimality occurs in μ^* in the mixed-mixed case. Therefore, there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-mixed case. \square

Lemma 2.6.11. *For all stage games G where $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| > 1$, if $G \in \mathcal{G}_{LS}^{m,m}$, then $G \in \mathcal{G}_{LS}^{m,p}$.*

Proof. If $|V_1^{m,p}| > 1$, $|V_2^{m,p}| > 1$, and $G \in \mathcal{G}_{LS}^{m,m}$, since $\text{Nash}^{m,p}(G) \subseteq \text{Nash}^{m,m}(G)$, $|V_1^{m,m}| > 1$ and $|V_2^{m,m}| > 1$. Then since there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-mixed case, by Theorem 2.4.1, there exists some $\hat{\sigma}_1 \in \Delta A_1, \hat{\sigma}_2 \in \Delta A_2$ where $(\hat{\sigma}_1, \hat{\sigma}_2) \notin \text{Nash}^{m,m}(G)$. Applying Lemma 2.4.2, there exists some $a_1 \in A_1$ and $a_2 \in A_2$ where $(a_1, a_2) \notin \text{Nash}^{p,p}(G)$. By Theorem 2.5.1, there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-pure case, so $G \in \mathcal{G}_{LS}^{m,p}$. \square

The above lemma shows that under certain preconditions on $|V_1^{m,p}|$ and $|V_2^{m,p}|$, $\mathcal{G}_{LS}^{m,p} = \mathcal{G}_{LS}^{m,m}$. Now the question is whether $\mathcal{G}_{LS}^{m,p}$ always equals $\mathcal{G}_{LS}^{m,m}$. Here we show that this is not the case. We present example stage games G where $G \in \mathcal{G}_{LS}^{m,m}$ but $G \notin \mathcal{G}_{LS}^{m,p}$ for each of the remaining cases regarding the values of $|V_1^{m,p}|$ and $|V_2^{m,p}|$. This shows that $\mathcal{G}_{LS}^{m,p} \neq \mathcal{G}_{LS}^{m,m}$. Combined with Theorem 2.6.10, this shows that $\mathcal{G}_{LS}^{m,p}$ is a proper subset of $\mathcal{G}_{LS}^{m,m}$.

Example 2.6.12. *Table 2.9 presents an example stage game G where $|V_1^{m,p}| = 0$, $|V_2^{m,p}| = 0$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{m,p}$. This game is essentially the stage game presented in Table 2.7 with column player being player 1 and row player being player 2. When player 1 can play mixed strategies and player 2 can only play pure strategies, there is no Nash equilibrium, so $|V_1^{m,p}| = 0$ and $|V_2^{m,p}| = 0$. By Theorem 2.5.1, $G \notin \mathcal{G}_{LS}^{m,p}$. When both players can use mixed strategies, $|V_1^{m,m}| = 1$ and $|V_2^{m,m}| > 1$. And (a_1, b_2) is a strategy profile where player 1 plays a best response and player 2 does not play a best response. Therefore, the condition in Theorem 2.4.1 is satisfied, so $G \in \mathcal{G}_{LS}^{m,m}$.*

	a_2	b_2
a_1	(0,4)	(4,0)
b_1	(3,1)	(3,1)
c_1	(3,2)	(3,2)
d_1	(4,0)	(0,4)

Table 2.9: Example stage game in matrix form, row player is player 1, column player is player 2. When player 1 can play mixed strategies and player 2 can only play pure strategies, there is no Nash equilibrium. When both players can use mixed strategies, there are multiple Nash equilibria, and local suboptimality can occur in repeated games.

Example 2.6.13. Table 2.10 presents an example stage game G where $|V_1^{m,p}| = 1$, $|V_2^{m,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{m,p}$. This game is essentially the stage game presented in Table 2.8 with column player being player 1 and row player being player 2. Following the arguments in Example 2.6.6, for this game, $|V_1^{m,p}| = 1$, $|V_2^{m,p}| = 1$, $|V_1^{m,m}| > 1$, $|V_2^{m,m}| > 1$, and $(a_1, b_2) \notin \text{Nash}^{m,m}(G)$. Therefore, by Theorems 2.4.1 and 2.5.1, $G \in \mathcal{G}_{LS}^{m,m}$ and $G \notin \mathcal{G}_{LS}^{m,p}$.

	a_2	b_2
a_1	(4,4)	(0,0)
b_1	(3,1)	(3,1)
c_1	(0,0)	(4,4)

Table 2.10: Example stage game in matrix form, row player is player 1, column player is player 2. $|V_1^{m,p}| = 1$, $|V_2^{m,p}| = 1$, $|V_1^{m,m}| > 1$, and $|V_2^{m,m}| > 1$.

Example 2.6.14. Table 2.11 presents an example stage game G where $|V_1^{m,p}| > 1$, $|V_2^{m,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{m,p}$. When player 1 can play mixed strategies and player 2 can only play pure strategies, the set of Nash equilibria is: 1) (σ_λ, a_2) for all $0 \leq \lambda \leq 1$ where $\sigma_\lambda(a_1) = \lambda$ and $\sigma_\lambda(b_1) = 1 - \lambda$, and 2) (σ'_θ, b_2) for all $0 \leq \theta \leq 1$ where $\sigma'_\theta(b_1) = \theta$ and $\sigma'_\theta(c_1) = 1 - \lambda$. So $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| = 1$. For any $\hat{\sigma}_1 \in \Delta A_1, \hat{a}_2 \in A_2$ where \hat{a}_2 is a best response to $\hat{\sigma}_1$ and $\hat{\sigma}_1$ is not a best response to \hat{a}_2 , if $\hat{a}_2 = a_2$, $S_{\hat{\sigma}_1}$ must include c_1 and at least one of a_1 and b_1 ; if $\hat{a}_2 = b_2$, $S_{\hat{\sigma}_1}$ must include a_1 and at least one of b_1 and c_1 . In all these cases, there

is some $a, a' \in S_{\hat{\sigma}_1}$ where $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2) = 0.5$. But $V_1 = \{2, 3\}$ only contains integer elements. So such $u_1(a, \hat{a}_2) - u_1(a', \hat{a}_2)$ can never be expressed as the sum of a sequence of values from $D = \{u_1(\sigma) - u_1(\sigma') \mid \sigma, \sigma' \in \text{Nash}^{m,p}(G)\}$. Therefore, the condition in Theorem 2.5.1 is not satisfied, whereas the condition in Theorem 2.4.1 is satisfied. So $G \in \mathcal{G}_{LS}^{m,m}$ and $G \notin \mathcal{G}_{LS}^{m,p}$. Intuitively speaking, in the mixed-pure case, although $|V_1^{m,p}| > 1$ makes player 1 potentially vulnerable to threats, there is no way to construct such a threat with the available strategy profiles. But in the mixed-mixed case, as player 2 obtains access to mixed strategies, new mechanisms for constructing threats become available (as are used in the proof of Theorem 2.4.1), which enables local suboptimality to occur.

	a_2	b_2	
a_1	(3,2)	(1.5,1)	
b_1	(3,2)	(2,2)	
c_1	(2.5,1)	(2,2)	

Table 2.11: Example stage game in matrix form, row player is player 1, column player is player 2. $|V_1^{m,p}| > 1$, $|V_2^{m,p}| = 1$. Local suboptimality can never occur in the mixed-pure case, but can occur in the mixed-mixed case.

Example 2.6.15. Table 2.12 presents an example stage game G where $|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{m,p}$. This game is essentially the stage game presented in Table 2.6 with column player being player 1 and row player being player 2. The same discussion in Example 2.6.7 applies here.

	a_2	b_2	c_2	
a_1	(2,3)	(2,3)	(1,1)	
b_1	(1,1)	(2,2)	(2,2)	

Table 2.12: Example stage game in matrix form, row player is player 1, column player is player 2. For all $a \in A_2$, for all $\sigma_1 \in \Delta A_1$ that is a best response to a , a is also a best response to σ_1 .

Furthermore, the following theorem completely characterizes $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{m,p}$, i.e., the set of 2-player stage games G where 1) in the mixed-pure case, local suboptimality

can never occur, and 2) in the mixed-mixed case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs.

Theorem 2.6.16. *For 2-player stage games G , $G \notin \mathcal{G}_{LS}^{m,p}$ and $G \in \mathcal{G}_{LS}^{m,m}$ if and only if*

1. $|V_1^{m,p}| = 0$, $|V_2^{m,p}| = 0$, and the condition in Theorem 2.4.1 is satisfied, OR
2. $|V_1^{m,p}| = 1$, $|V_2^{m,p}| = 1$, and the condition in Theorem 2.4.1 is satisfied, OR
3. $|V_1^{m,p}| > 1$, $|V_2^{m,p}| = 1$, the condition in Theorem 2.5.1 is not satisfied, and the condition in Theorem 2.4.1 is satisfied, OR
4. $|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$, the condition in Theorem 2.5.1 is not satisfied, and the condition in Theorem 2.4.1 is satisfied.

Proof. Same as in the proof of Theorem 2.6.8, a direct application of Theorems 2.4.1 and 2.5.1 shows that the above condition is sufficient. Examples 2.6.12 to 2.6.15 present example stage games G that belong to each of the cases 1, 2, 3, and 4.

To show the condition is necessary, we split all stage games G into five disjoint cases based on $|V_1^{m,p}|$ and $|V_2^{m,p}|$:

- (a) $|V_1^{m,p}| = 0$ and $|V_2^{m,p}| = 0$,
- (b) $|V_1^{m,p}| = 1$ and $|V_2^{m,p}| = 1$,
- (c) $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| = 1$,
- (d) $|V_1^{m,p}| = 1$ and $|V_2^{m,p}| > 1$,
- (e) $|V_1^{m,p}| > 1$ and $|V_2^{m,p}| > 1$.

For cases (a), (b), (c), and (d), a direct application of Theorems 2.4.1 and 2.5.1 implies that our target condition is necessary. For case (e), Lemma 2.6.11 shows that $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{m,p}$ is empty under this case. Therefore, our target condition is necessary. \square

2.6.3 Changing Both Players from Pure Strategies to Mixed Strategies

Finally, we analyze the situation for changing from the pure-pure case to the mixed-mixed case, i.e., study the relationship between $\mathcal{G}_{LS}^{p,p}$ and $\mathcal{G}_{LS}^{m,m}$. The results and example games for this situation can mostly be derived directly from the results in Sections 2.6.1 and 2.6.2.

Theorem 2.6.17. $\mathcal{G}_{LS}^{p,p} \subseteq \mathcal{G}_{LS}^{m,m}$, i.e., for 2-player stage games G , if there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the pure-pure case, then there exists some T and some SPE of $G(T)$ where local suboptimality occurs in the mixed-mixed case.

Proof. Theorems 2.6.2 and 2.6.10 directly imply this result. \square

Lemma 2.6.18. For all stage games G where $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| > 1$, if $G \in \mathcal{G}_{LS}^{m,m}$, then $G \in \mathcal{G}_{LS}^{p,p}$.

Proof. If $|V_1^{p,p}| > 1$, $|V_2^{p,p}| > 1$, since $\text{Nash}^{p,p}(G) \subseteq \text{Nash}^{m,p}(G) \subseteq \text{Nash}^{m,m}(G)$, $|V_1^{m,p}| > 1$, $|V_2^{m,p}| > 1$, $|V_1^{m,m}| > 1$ and $|V_2^{m,m}| > 1$. Therefore, by applying Lemmas 2.6.4 and 2.6.11, we have if $G \in \mathcal{G}_{LS}^{m,m}$, then $G \in \mathcal{G}_{LS}^{p,p}$. \square

The above lemma shows that under certain preconditions on $|V_1^{p,p}|$ and $|V_2^{p,p}|$, $\mathcal{G}_{LS}^{p,p} = \mathcal{G}_{LS}^{m,m}$. Here we present example stage games G where $G \in \mathcal{G}_{LS}^{m,m}$ but $G \notin \mathcal{G}_{LS}^{p,p}$ for each of the remaining cases regarding the values of $|V_1^{p,p}|$ and $|V_2^{p,p}|$. These examples are reused from Sections 2.6.1 and 2.6.2. This shows that $\mathcal{G}_{LS}^{p,p} \neq \mathcal{G}_{LS}^{m,m}$. Combined with Theorem 2.6.17, this shows that $\mathcal{G}_{LS}^{p,p}$ is a proper subset of $\mathcal{G}_{LS}^{m,m}$.

Example 2.6.19. Table 2.7 presents an example stage game G where $|V_1^{p,p}| = 0$, $|V_2^{p,p}| = 0$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{p,p}$.

Example 2.6.20. Table 2.8 presents an example stage game G where $|V_1^{p,p}| = 1$, $|V_2^{p,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{p,p}$.

Example 2.6.21. Table 2.6 presents an example stage game G where $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{p,p}$.

Example 2.6.22. Table 2.12 presents an example stage game G where $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$, $G \in \mathcal{G}_{LS}^{m,m}$, and $G \notin \mathcal{G}_{LS}^{p,p}$.

The following theorem completely characterizes $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{p,p}$, i.e., the set of 2-player stage games G where 1) in the pure-pure case, local suboptimality can never occur, and 2) in the mixed-mixed case, there exists some T and some SPE of $G(T)$ where local suboptimality occurs.

Theorem 2.6.23. For 2-player stage games G , $G \notin \mathcal{G}_{LS}^{p,p}$ and $G \in \mathcal{G}_{LS}^{m,m}$ if and only if

1. $|V_1^{p,p}| = 0$, $|V_2^{p,p}| = 0$, and the condition in Theorem 2.4.1 is satisfied, OR
2. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| = 1$, and the condition in Theorem 2.4.1 is satisfied, OR
3. $|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, the condition in Theorem 2.3.1 is not satisfied, and the condition in Theorem 2.4.1 is satisfied, OR
4. $|V_1^{p,p}| = 1$, $|V_2^{p,p}| > 1$, the condition in Theorem 2.3.1 is not satisfied, and the condition in Theorem 2.4.1 is satisfied.

Proof. Same as in the proofs of Theorems 2.6.8 and 2.6.16, a direct application of Theorems 2.3.1 and 2.4.1 shows that the above condition is sufficient. Examples 2.6.19 to 2.6.22 present example stage games G that belong to each of the cases 1, 2, 3, and 4.

To show the condition is necessary, again, we split all stage games G into five disjoint cases based on $|V_1^{p,p}|$ and $|V_2^{p,p}|$:

- (a) $|V_1^{p,p}| = 0$ and $|V_2^{p,p}| = 0$,
- (b) $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| = 1$,
- (c) $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| = 1$,
- (d) $|V_1^{p,p}| = 1$ and $|V_2^{p,p}| > 1$,
- (e) $|V_1^{p,p}| > 1$ and $|V_2^{p,p}| > 1$.

For cases (a), (b), (c), and (d), a direct application of Theorems 2.3.1 and 2.4.1 implies that our target condition is necessary. For case (e), Lemma 2.6.18 shows that $\mathcal{G}_{LS}^{m,m} \setminus \mathcal{G}_{LS}^{m,p}$ is empty under this case. Therefore, our target condition is necessary. \square

2.6.4 Discussion

The central question we focus on in this section is how changing a player (or both players) from pure-strategies-only to mixed-strategies-allowed can affect the emergence of local suboptimality. We present here an intuitive interpretation of the results established in this section.

First, if local suboptimality can occur before the change, then after changing any player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can still occur (Theorems 2.6.2, 2.6.10 and 2.6.17). So allowing players to play mixed strategies can never prohibit the emergence of local suboptimality.

On the other hand, it is possible that local suboptimality can never occur before the change, but after changing one player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can occur. Such phenomena can happen through two different mechanisms. The first one is through the introduction of new stage-game Nash equilibria. Before the change, there might be no stage-game NE, or there is only one payoff value attainable at stage-game NEs for each player ($V_1 = V_2 = 1$), which makes neither of the players vulnerable to potential threats that force them to play locally suboptimally. After allowing one (or both) player(s) to play mixed strategies, a new set of stage-game NEs becomes available. This makes some $|V_i| > 1$, making that player vulnerable to potential threats. Cases 1 and 2 in Theorems 2.6.8, 2.6.16 and 2.6.23 and Examples 2.6.5, 2.6.6, 2.6.12, 2.6.13, 2.6.19 and 2.6.20 belong to this type.

For the second mechanism, before the change, some player already have $|V_i| > 1$, which means they are potentially vulnerable to threats. However, there is no way of constructing such a threat in any SPE given the available strategy profiles, so local suboptimality cannot occur. After allowing one (or both) player(s) to play mixed strategies, with the newly available strategy profiles, it becomes possible to construct

such a threat, which makes it possible for local suboptimality to occur. We show that this can happen in the following cases:

- The player that is potentially vulnerable to threats before the change (i.e., $|V_i| > 1$) changes from pure-strategies-only to mixed-strategies-allowed. This change can open up vulnerabilities for themselves, regardless of whether the opponent has access to mixed strategies or not. Case 3 in Theorem 2.6.8 ($|V_1^{p,p}| > 1$, $|V_2^{p,p}| = 1$, player 1 changes from pure to mixed, player 2 can only play pure), case 4 in Theorem 2.6.16 ($|V_1^{m,p}| = 1$, $|V_2^{m,p}| > 1$, player 2 changes from pure to mixed, player 1 can play mixed), and the corresponding example games Examples 2.6.7 and 2.6.15 demonstrate this type.
- A player changes from pure-strategies-only to mixed-strategies-allowed when their opponent is potentially vulnerable to threats before the change (i.e., $|V_i| > 1$). Here, only if their opponent has access to mixed strategies will such change be useful to create threats that were not possible before the change. Case 3 in Theorem 2.6.16 and the corresponding example game Example 2.6.14 demonstrate this situation. Importantly, if their opponent only has access to pure strategies, changing from pure-strategies-only to mixed-strategies-allowed will not enable a player to construct threats if it was impossible to create threats before the change. This is shown in Lemma 2.6.3.

2.7 Computational Aspects

In this section, we consider the computational aspect of the problem: given an arbitrary 2-player stage game G , how to (algorithmically) decide if there exists some T and some SPE of $G(T)$ where local suboptimality occurs. We focus on the general case where mixed strategies are allowed. A naive approach is to enumerate over T and solve for all subgame-perfect equilibria for each $G(T)$. Such an approach is not only computationally inefficient, but also not guaranteed to terminate due to the unboundedness of T . This leaves open the question of whether the above problem is

decidable or not.

Theorem 2.4.1 proves a necessary and sufficient condition that is solely described on the stage game G , independent of T . Based on this condition, we present here a more efficient algorithm for deciding the above problem (Algorithm 1). This algorithm also shows that the above problem is decidable.

Algorithm 1 Given stage game G , decide if there exists some T and some SPE of $G(T)$ where local suboptimality occurs

Input: $G = \{u_1(i, j), u_2(i, j)\}_{i \in [|A_1|], j \in [|A_2|]}$

Output: True/False

```

1: function DECIDELOCALSUBOPTIMALITY( $G$ )
2:    $\text{Nash}(G) \leftarrow \text{FINDALLNASH}(G)$ 
3:    $\text{uniqueV1}, \text{uniqueV2} \leftarrow \text{ISVALUEUNIQUE}(\text{Nash}(G), G)$ 
4:   if  $\text{uniqueV1}$  then
5:     if  $\text{uniqueV2}$  then
6:       return False
7:     else
8:       return EXISTOFF1BEST( $G$ )
9:     end if
10:  else
11:    if  $\text{uniqueV2}$  then
12:      return EXISTOFF2BEST( $G$ )
13:    else
14:      return EXISTOFF( $G$ )
15:    end if
16:  end if
17: end function

```

The input to the algorithm is the stage game G , represented as a matrix of payoffs for every action profile $\{u_1(i, j), u_2(i, j)\}_{i \in [|A_1|], j \in [|A_2|]}$. Overall, the algorithm consists of three steps. The first step is to compute the set of all Nash equilibria of G . The second step is to determine if the set of payoff values attainable at $\text{Nash}(G)$ is unique for each player, so as to know which case of the condition in Theorem 2.4.1 we need to further check. The third step is to check if there exists the respective off-Nash strategy profiles required for each case. We now present the algorithm for each of the three steps.

2.7.1 Compute All Nash Equilibria

There are many existing algorithms for computing the set of all Nash equilibria of two-player normal form games [27, 7, 91, 10]. In general, any existing algorithm can be used here as long as it can handle degenerate games where there are an infinite number of Nash equilibria. An example of such algorithm can be found in [10]. It handles the potentially infinite number of Nash equilibria in degenerate games by computing the finite set of *extreme equilibria*; the set of all Nash equilibria is then completely described by polytopes obtained from subsets of the extreme equilibria.

Complexity In general, the problem of finding all Nash equilibria of a two-player normal form game is NP-hard (as [45] shows that deciding if a game has a unique Nash equilibrium is NP-hard). Therefore, any algorithm for FINDALLNASH takes exponential time in the worst case (unless P=NP).

2.7.2 Determine the Uniqueness of Payoffs at Equilibrium

ISVALUEUNIQUE returns two Boolean values; the first (resp. second) return value is True if $|V_1| = 1$ (resp. $|V_2| = 1$). This function can be achieved by evaluating payoffs at each extreme equilibrium and compare to see if there are more than one values for each player.

Complexity In general, a non-degenerate 2-player bimatrix game can have an exponential number of Nash equilibria [90, 74]. So ISVALUEUNIQUE can take exponential time. But in practice, this step can be done in the first step (FINDALLNASH) with a constant factor overhead, by evaluating the payoff of each extreme equilibrium immediately after the extreme equilibrium is computed in FINDALLNASH and comparing with the payoffs of the previous equilibria.

2.7.3 Check the Existence of Required Off-Nash Strategy Profiles

Based on the uniqueness of payoffs attainable at equilibrium for each player (uniqueV1 and uniqueV2), we need to check the existence of off-Nash strategy profiles with the requirements corresponding to each case.

If $|V_1| > 1$ and $|V_2| > 1$ (both uniqueV1 and uniqueV2 are False), we simply need to check the existence of an off-Nash strategy profile without further requirements. By Lemma 2.4.2, it suffices to check the existence of an off-Nash *pure* strategy profile. Algorithm 2 achieves this functionality.

Algorithm 2 Check the existence of an off-Nash strategy profile

Input: $G = \{u_1(i, j), u_2(i, j)\}_{i \in [|A_1|], j \in [|A_2|]}$

Output: True/False

```

1: function EXISTOFF( $G$ )
2:   for  $j = 1, \dots, |A_2|$  do
3:     if not all  $u_1(i, j)$  for  $i = 1, \dots, |A_1|$  are the same then
4:       return True
5:     end if
6:   end for
7:   for  $i = 1, \dots, |A_1|$  do
8:     if not all  $u_2(i, j)$  for  $j = 1, \dots, |A_2|$  are the same then
9:       return True
10:    end if
11:  end for
12:  return False
13: end function

```

If $|V_1| > 1$ and $|V_2| = 1$ (uniqueV1 is False and uniqueV2 is True), we need to check the existence of an off-Nash strategy profile where player 2 plays a best response. The following lemma proves that it suffices to check the existence of an off-Nash strategy profile where player 2 plays a *pure strategy* best response.

Lemma 2.7.1. *For any two-player game G , if there exists $\sigma_1 \in \Delta A_1$, $\sigma_2 \in \Delta A_2$, $a'_1 \in A_1$ where $u_1(\sigma_1, \sigma_2) < u_1(a'_1, \sigma_2)$ and σ_2 is a best response to σ_1 , then there exists $\sigma_1 \in \Delta A_1$, $a_2 \in A_2$, $a'_1 \in A_1$ where $u_1(\sigma_1, a_2) < u_1(a'_1, a_2)$ and a_2 is a best response to σ_1 .*

Proof. Let S_{σ_2} be the support of σ_2 . σ_2 is a best response to σ_1 implies that any $a \in S_{\sigma_2}$ is a best response to σ_1 . We have $u_1(a'_1, \sigma_2) - u_1(\sigma_1, \sigma_2) = \sum_{a \in S_{\sigma_2}} \sigma_2(a) \cdot (u_1(a'_1, a) - u_1(\sigma_1, a))$. Since $u_1(a'_1, \sigma_2) - u_1(\sigma_1, \sigma_2) > 0$, there exists some $a_2 \in S_{\sigma_2}$ where $u_1(a'_1, a_2) - u_1(\sigma_1, a_2) > 0$. This $\sigma_1 \in \Delta A_1$, $a_2 \in A_2$, $a'_1 \in A_1$ satisfies $u_1(\sigma_1, a_2) < u_1(a'_1, a_2)$ and a_2 is a best response to σ_1 . \square

Algorithm 3 presents a method for checking the existence of an off-Nash strategy profile where player 2 plays a pure strategy best response. The idea is as follows. For each possible pure strategy j of player 2, we construct linear programs with constraints on player 1's mixed strategy (represented by probabilities $\{x_{i'}\}_{i'=1}^{|A_1|}$) such that j is a best response. We aim to find for every action i of player 1 that is not a best response to j , if j can be a best response to a mixed strategy of player 1 containing i in its support. The presented linear program achieves this purpose. If we can find such a mixed strategy σ_1 for player 1, then (σ_1, j) is an instance of an off-Nash strategy profile where player 2 plays a best response, as desired. Since we exhaustively enumerate over all possible cases, this method is complete.

EXISTOFF1BEST can be implemented using the same algorithm, exchanging player 1 and 2.

Complexity EXISTOFF has complexity $\mathcal{O}(|A_1| \cdot |A_2|)$. EXISTOFF2BEST (and similarly EXISTOFF1BEST) involves solving $\mathcal{O}(|A_1| \cdot |A_2|)$ instances of polynomial-sized linear programs. Since a linear program can be solved in polynomial time, EXISTOFF2BEST is a polynomial time algorithm. A vanilla support enumeration algorithm (such as [27]), which enumerates over all possible supports (subsets of the action sets) for the mixed strategies σ_1 and σ_2 in the required strategy profile and solve for each case, requires exponential time since there is an exponential number of possible supports. The algorithm we present here is more efficient.

Overall, the computational bottleneck is step 1, since finding all Nash equilibria is NP-hard. Step 2 can be computed within step 1 with a constant factor overhead. Step 3 can be computed in polynomial time.

Algorithm 3 Check the existence of an off-Nash strategy profile where player 2 plays a best response

Input: $G = \{u_1(i, j), u_2(i, j)\}_{i \in [|A_1|], j \in [|A_2|]}$

Output: True/False

```
1: function EXISTOFF2BEST( $G$ )
2:   for  $j = 1, \dots, |A_2|$  do
3:      $c \leftarrow \max_i u_1(i, j)$ 
4:     for  $i \in [|A_1|]$  where  $u_1(i, j) < c$  do
5:        $\max\_xi \leftarrow$  solve the following linear program
```

maximize: x_i

subject to: $x_{i'} \geq 0, i' = 1, \dots, |A_1|$

$$\sum_{i'=1}^{|A_1|} x_{i'} = 1$$

$$\sum_{i'=1}^{|A_1|} x_{i'} \cdot u_2(i', j) \geq \sum_{i'=1}^{|A_1|} x_{i'} \cdot u_2(i', j'), j' = 1, \dots, |A_2|$$

```
6:       if  $\max\_xi > 0$  then
7:         return True
8:       end if
9:     end for
10:  end for
11:  return False
12: end function
```

2.8 Generalization to n -Player Games

In this section, we aim to tackle Question 1.1.1 for n -player games. As we will present in the followings, we prove a sufficient and necessary condition for Question 1.1.1 for the pure strategy case; for the general case where mixed strategies are allowed, we prove a separate sufficient condition and a separate necessary condition for Question 1.1.1.

Denote $I(G) = \{i \mid |V_i| = 1\}$ as the set of players that have a unique payoff attainable at $\text{Nash}(G)$. Given any strategy profile σ , denote $B(\sigma) = \{i \mid \sigma_i \text{ is a best response to } \sigma_{-i}\}$ as the set of players that plays a best response strategy.

Theorem 2.8.1 (n -player, pure strategy case). *For n -player games where only pure strategies are allowed, a sufficient and necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

There exists a strategy profile $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_n) \in A$ where $I(G) \subseteq B(\hat{\mathbf{a}})$ and $B(\hat{\mathbf{a}}) \neq [n]$.

Proof. We prove the condition is sufficient by showing if the condition is satisfied, we can always construct some T and some SPE where local suboptimality occurs.

We construct an SPE μ^* consisting of segments. WLOG, let the first m players be the ones that do not play their best response in $\hat{\mathbf{a}}$, i.e. $[m] = [n] \setminus B(\hat{\mathbf{a}})$. Denote μ^t as all the t -th round behavior strategy profiles in μ^* and $\mu^{t_1:t_2}$ as all the behavior strategy profiles between the t_1 -th round and the t_2 -th round in μ^* . μ^* is divided into segments: $\mu^1, \mu^{2:T_1}, \mu^{T_1+1:T_2}, \dots, \mu^{T_{m-1}+1:T_m}$. We set $\mu^1 = \hat{\mathbf{a}}$. We let the strategies in each segment only depend on the play in the first round, not depending on any other segments. So given a play in the first round, each $\mu^{T_{i-1}+1:T_i}$ is an SPE of the $(T_i - T_{i-1})$ -round subgame. We construct $\mu^{T_{i-1}+1:T_i}$ as follows. Let $\sigma^{\min}, \sigma^{\max} \in \text{Nash}(G)$ such that $u_i(\sigma^{\min}) = \min(V_i)$ and $u_i(\sigma^{\max}) = \max(V_i)$. If player i plays \hat{a}_i in the first round, players play according to σ^{\max} in all rounds in $\mu^{T_{i-1}+1:T_i}$; otherwise, players play according to σ^{\min} in all rounds in $\mu^{T_{i-1}+1:T_i}$. The above construction ensures that μ^* is an SPE. And the first round strategy profile in μ^* does not form a stage game Nash equilibrium.

We prove the condition is necessary by showing that if there exists some T and some SPE of $G(T)$ where local suboptimality occurs, there exists a strategy profile $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_n) \in A$ where $I(G) \subseteq B(\hat{\mathbf{a}})$ and $B(\hat{\mathbf{a}}) \neq [n]$.

Let T^* and some SPE $\boldsymbol{\mu}^*$ of $G(T^*)$ to be an instance where local suboptimality occurs. Let $k^* = \max \{k \mid \exists h(k) \text{ s.t. } \boldsymbol{\mu}^*(h(k)) \notin \text{Nash}(G)\}$ be the last round where off-Nash play occurs and let $\boldsymbol{\mu}^*(h^*(k^*)) \notin \text{Nash}(G)$ be a behavior strategy profile that does not form a stage-game Nash equilibrium. Denote $G_{|h^*(k^*)}$ as the subgame starting from $h^*(k^*)$. Consider $\boldsymbol{\mu}_{|h^*(k^*)}^*$, the strategy profile in the subgame $G_{|h^*(k^*)}$. By the above construction, all the behavior strategy profiles in $\boldsymbol{\mu}_{|h^*(k^*)}^*$ after the first round belong to $\text{Nash}(G)$. Therefore, for every player $i \in I(G)$, their total payoff in $G_{|h^*(k^*)}$ after the first round does not depend on what is played in the first round. So they must play their best responses in the first round, i.e. in $\boldsymbol{\mu}^*(h^*(k^*))$. So $I(G) \subseteq B(\boldsymbol{\mu}^*(h^*(k^*)))$. And since $\boldsymbol{\mu}^*(h^*(k^*)) \notin \text{Nash}(G)$, $B(\boldsymbol{\mu}^*(h^*(k^*))) \neq [n]$. So $\boldsymbol{\mu}^*(h^*(k^*))$ is a strategy profile that satisfies our target condition.

□

Theorem 2.8.2 (*n*-player, general case, sufficient condition). *For general n-player games (mixed strategies allowed), a sufficient condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

There exists a strategy profile $\hat{\boldsymbol{\sigma}} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ where $I(G) \subseteq B(\hat{\boldsymbol{\sigma}})$ and $B(\hat{\boldsymbol{\sigma}}) \neq [n]$, and

1. *there exists $\boldsymbol{\sigma}, \boldsymbol{\sigma}' \in \text{Nash}(G)$ where*

(a) *for all $\lambda \in [0, 1]$, $\lambda\boldsymbol{\sigma} + (1 - \lambda)\boldsymbol{\sigma}' \in \text{Nash}(G)$, and*

(b) *for some $i \in I(G)$, $\sigma_i \neq \sigma'_i$,*

OR

2. *for all $i \in [n] \setminus B(\hat{\boldsymbol{\sigma}})$,*

(a) *$|S_{\hat{\sigma}_i}| = 1$, i.e. $\hat{\sigma}_i$ is a pure strategy, or*

(b) *V_i contains a non-zero length continuous interval, or*

(c) denote the set of possible differences in u_i between pairs of NEs in the stage game as $D_i = \{u_i(\boldsymbol{\sigma}) - u_i(\boldsymbol{\sigma}') \mid \boldsymbol{\sigma}, \boldsymbol{\sigma}' \in \text{Nash}(G)\}$, there exists an action from the support of $\hat{\sigma}_i$ $a \in S_{\hat{\sigma}_i}$ such that, for every $a' \in S_{\hat{\sigma}_i} \setminus a$, there exists some integers $n_{\mathbf{a}_{I(G)}} \geq 0$ and $d_k^{\mathbf{a}_{I(G)}} \in D_i, k = 1, \dots, n_{\mathbf{a}_{I(G)}}$ for each $\mathbf{a}_{I(G)} \in \times_{i \in I(G)} S_{\hat{\sigma}_i}$ such that $u_i(a', \hat{\boldsymbol{\sigma}}_{-i}) - u_i(a, \hat{\boldsymbol{\sigma}}_{-i}) = \sum_{\mathbf{a}_{I(G)} \in \times_{i \in I(G)} S_{\hat{\sigma}_i}} \hat{\boldsymbol{\sigma}}(\mathbf{a}_{I(G)}) \cdot \sum_{k=1}^{n_{\mathbf{a}_{I(G)}}} d_k^{\mathbf{a}_{I(G)}}$.

Proof. We prove the condition is sufficient by showing if the condition is satisfied, we can always construct some T and some SPE where local suboptimality occurs.

We construct an SPE $\boldsymbol{\mu}^*$ consisting of segments. WLOG, let the first m players be the ones that do not play their best response in $\hat{\boldsymbol{\sigma}}$, i.e. $[m] = [n] \setminus B(\hat{\boldsymbol{\sigma}})$. Denote $\boldsymbol{\mu}^t$ as all the t -th round behavior strategy profiles in $\boldsymbol{\mu}^*$ and $\boldsymbol{\mu}^{t_1:t_2}$ as all the behavior strategy profiles between the t_1 -th round and the t_2 -th round in $\boldsymbol{\mu}^*$. $\boldsymbol{\mu}^*$ is divided into segments: $\boldsymbol{\mu}^1, \boldsymbol{\mu}^{2:T_1}, \boldsymbol{\mu}^{T_1+1:T_2}, \dots, \boldsymbol{\mu}^{T_{m-1}+1:T_m}$. We set $\boldsymbol{\mu}^1 = \hat{\boldsymbol{\sigma}}$. We let the strategies in each segment only depend on the play in the first round, not depending on any other segments. So given a play in the first round, each $\boldsymbol{\mu}^{T_{i-1}+1:T_i}$ is an SPE of the $(T_i - T_{i-1})$ -round subgame.

(1) is satisfied. We let the strategies in the segment ending at T_i only depend on the action played by player i in the first round. In the following, we denote $\boldsymbol{\mu}_{|a_i}^{T_{i-1}+1:T_i}$ as the strategy profile in the segment ending at T_i given player i plays a_i in the first round.

$\boldsymbol{\mu}^{T_{i-1}+1:T_i}$ is constructed as follows. Pick $a_i^m \in S_{\hat{\sigma}_i}$ such that $u_i(a_i^m, \hat{\boldsymbol{\sigma}}_{-i}) = \max_{a_i \in S_{\hat{\sigma}_i}} u_i(a_i, \hat{\boldsymbol{\sigma}}_{-i})$. We construct $\boldsymbol{\mu}^{T_{i-1}+1:T_i}$ such that for all $a'_i \in S_{\hat{\sigma}_i}$, $U_i(\boldsymbol{\mu}_{|a'_i}^{T_{i-1}+1:T_i}) - U_i(\boldsymbol{\mu}_{|a_i^m}^{T_{i-1}+1:T_i}) = u_i(a_i^m, \hat{\boldsymbol{\sigma}}_{-i}) - u_i(a'_i, \hat{\boldsymbol{\sigma}}_{-i})$. Let $\boldsymbol{\sigma}^{\min}, \boldsymbol{\sigma}^{\max} \in \text{Nash}(G)$ such that $u_i(\boldsymbol{\sigma}^{\min}) = \min(V_i)$ and $u_i(\boldsymbol{\sigma}^{\max}) = \max(V_i)$, so $u_i(\boldsymbol{\sigma}^{\max}) > u_i(\boldsymbol{\sigma}^{\min})$. By using $\boldsymbol{\sigma}, \boldsymbol{\sigma}'$ given in (1), we can construct SPEs that achieve a continuous range of values of U_i for all $i \notin I(G)$. Denote $j \in I(G)$ such that $\sigma_j \neq \sigma'_j$, and $a_j \in A_j$ such that $\sigma_j(a_j) \neq \sigma'_j(a_j)$. WLOG, let $\sigma_j(a_j) > \sigma'_j(a_j)$. We can construct SPEs $\boldsymbol{\mu}(\lambda)$ parameterized by λ as:

- In the first round, play $\lambda\sigma + (1 - \lambda)\sigma'$.
- If the first round play by player j is a_j , players play σ^{\max} in all later rounds; otherwise, players play σ^{\min} in all later rounds.

By varying λ from 0 to 1, we can obtain a continuous range of values for $U_i(\mu(\lambda))$. And by setting the number of rounds larger, the value range can be arbitrarily large. Then we can use $\mu(\lambda)$ for $\mu_{|a_i}^{T_{i-1}+1:T_i}$ for each $a_i \in S_{\hat{\sigma}_i}$ with separately and appropriately assigned λ 's, such that $U_i(\mu_{|a_i}^{T_{i-1}+1:T_i}) + u_i(a_i, \hat{\sigma}_{-i})$ is a constant across all $a_i \in S_{\hat{\sigma}_i}$. Furthermore, we can include a segment in $\mu^{T_{i-1}+1:T_i}$ where if some $a_i \in S_{\hat{\sigma}_i}$ is played by player i in the first round, players play according to σ^{\max} ; otherwise, players play according to σ^{\min} .

The above construction ensures that μ^* is an SPE. And the first round strategy profile in μ^* does not form a stage game Nash equilibrium.

(2) is satisfied. We construct $\mu^{T_{i-1}+1:T_i}$ based on which case of (a), (b), and (c) is satisfied. Again, let $\sigma^{\min}, \sigma^{\max} \in \text{Nash}(G)$ such that $u_i(\sigma^{\min}) = \min(V_i)$ and $u_i(\sigma^{\max}) = \max(V_i)$.

If (a) is satisfied, i.e. $\hat{\sigma}_i$ is a pure strategy, denote a_i as the support for $\hat{\sigma}_i$. $\mu^{T_{i-1}+1:T_i}$ is then constructed as: if player i plays a_i in the first round of μ^* , players play according to σ^{\max} in all rounds in $\mu^{T_{i-1}+1:T_i}$; otherwise, players play according to σ^{\min} in all rounds in $\mu^{T_{i-1}+1:T_i}$.

If (b) is satisfied, i.e. V_i contains a non-zero length continuous interval, we can use a similar construction as the case when (1) is satisfied, replacing $\mu(\lambda)$ with repetitions of stage game NE that achieves appropriate values of V_i in the continuous interval, such that $U_i(\mu_{|a_i}^{T_{i-1}+1:T_i}) + u_i(a_i, \hat{\sigma}_{-i})$ is a constant across all $a_i \in S_{\hat{\sigma}_i}$.

If (c) is satisfied, we let $\mu^{T_{i-1}+1:T_i}$ depend on the play of the set of players $i \cup I(G)$ in the first round. Taking a as the chosen action from the support of $\hat{\sigma}_i$ in condition (c). We further divide $\mu^{T_{i-1}+1:T_i}$ into $|S_{\hat{\sigma}_i}| - 1$ segments and denote $\mu^{T_{i-1}+1:T_i}[a']$ as the segment corresponding to $a' \in S_{\hat{\sigma}_i} \setminus a$. For each $\mathbf{a}_{I(G)} \in \times_{i \in I(G)} S_{\hat{\sigma}_i}$ and $a' \in S_{\hat{\sigma}_i} \setminus a$, we construct the segment $\mu^{T_{i-1}+1:T_i}[a']$ by setting $\mu_{|(a, \mathbf{a}_{I(G)})}^{T_{i-1}+1:T_i}[a']$ and $\mu_{|(a', \mathbf{a}_{I(G)})}^{T_{i-1}+1:T_i}[a']$ using

corresponding sequences of stage game NEs as specified by $\{d_k^{\mathbf{a}_{I(G)}}\}_{k=1}^{n_{\mathbf{a}_{I(G)}}$ given in (c). $\boldsymbol{\mu}_{|(a'', \mathbf{a}_{I(G)})}^{T_{i-1}+1:T_i}[a'] = \boldsymbol{\mu}_{|(a, \mathbf{a}_{I(G)})}^{T_{i-1}+1:T_i}[a']$ for $a'' \in S_{\hat{\sigma}_i} \setminus \{a'\}$. This construction ensures that $U_i(\boldsymbol{\mu}_{|a_i}^{T_{i-1}+1:T_i}) + u_i(a_i, \hat{\sigma}_{-i})$ is a constant across all $a_i \in S_{\hat{\sigma}_i}$.

The above construction ensures that $\boldsymbol{\mu}^*$ is an SPE. And the first round strategy profile in $\boldsymbol{\mu}^*$ does not form a stage game Nash equilibrium. □

Theorem 2.8.3 (*n*-player, general case, necessary condition). *For general n-player games (mixed strategies allowed), a necessary condition on the stage game G for there exists some T and some SPE of $G(T)$ where local suboptimality occurs is:*

There exists a strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ where $I(G) \subseteq B(\hat{\sigma})$ and $B(\hat{\sigma}) \neq [n]$.

Proof. We prove the above condition is necessary by showing that if there exists some T and some SPE of $G(T)$ where local suboptimality occurs, there exists a strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ where $I(G) \subseteq B(\hat{\sigma})$ and $B(\hat{\sigma}) \neq [n]$.

Let T^* and some SPE $\boldsymbol{\mu}^*$ of $G(T^*)$ to be an instance where local suboptimality occurs. Let $k^* = \max \{k \mid \exists h(k) \text{ s.t. } \boldsymbol{\mu}^*(h(k)) \notin \text{Nash}(G)\}$ be the last round where off-Nash play occurs and let $\boldsymbol{\mu}^*(h^*(k^*)) \notin \text{Nash}(G)$ be a behavior strategy profile that does not form a stage-game Nash equilibrium. Denote $G_{|h^*(k^*)}$ as the subgame starting from $h^*(k^*)$. Consider $\boldsymbol{\mu}_{|h^*(k^*)}^*$, the strategy profile in the subgame $G_{|h^*(k^*)}$. By the above construction, all the behavior strategy profiles in $\boldsymbol{\mu}_{|h^*(k^*)}^*$ after the first round belong to $\text{Nash}(G)$. Therefore, for every player $i \in I(G)$, their total payoff in $G_{|h^*(k^*)}$ after the first round does not depend on what is played in the first round. So they must play their best responses in the first round, i.e. in $\boldsymbol{\mu}^*(h^*(k^*))$. So $I(G) \subseteq B(\boldsymbol{\mu}^*(h^*(k^*)))$. And since $\boldsymbol{\mu}^*(h^*(k^*)) \notin \text{Nash}(G)$, $B(\boldsymbol{\mu}^*(h^*(k^*))) \neq [n]$. So $\boldsymbol{\mu}^*(h^*(k^*))$ is a strategy profile that satisfies our target condition. □

Remark In the 2-player case, we are able to prove some properties that hold for 2-player games (Lemma 2.4.3 and the subsequent arguments in the proof of Theorem 2.4.1 that uses Lemma 2.4.3 to show there exists a connected component of

off-Nash strategy profiles), which allows the proof of the sufficient and necessary condition for the general case where mixed strategies are allowed. It is not clear whether similar properties hold for n -player games. Therefore, the questions of 1) what is a sufficient and necessary condition for n -player games when mixed strategies are allowed, and 2) is Question 1.1.2 decidable for n -player games when mixed strategies are allowed, remain open problems.

2.9 Related Work

Under the theme of analyzing equilibrium solutions in repeated games, a large body of work focuses on Folk Theorems, where the property of interest is: all feasible and individually rational payoff profiles can be attained in equilibria of the repeated game. In the context of infinitely repeated games, the original Folk Theorem asserts that all feasible and individually rational (see Section 2.2.3 for the definitions) payoff profiles can be attained in Nash equilibria of infinitely repeated games with sufficiently little discounting. This result is widely known in the field but not formally published, which is why it is called Folk Theorem. [9, 80] show that the same result holds when we consider subgame-perfect equilibria and assume no discounting. [39] proves a sufficient condition for Folk Theorem for subgame-perfect equilibria in infinitely repeated games with discounting. [36, 35] consider a variation of Folk Theorem where they show that any feasible payoff profile that Pareto dominates a Nash equilibrium of the stage game can be attained in a subgame-perfect equilibrium of the infinitely repeated game with discounting. In the context of finitely repeated games, [14] obtained sufficient conditions for Folk Theorem for subgame-perfect equilibria, and later [84] establishes necessary and sufficient conditions for subgame-perfect equilibria. Both results rely on mixed strategies are observable, meaning that players can directly observe the mixed strategies (i.e., probability distributions) used by other players in previous rounds of the game, not just the realized actions in the previous rounds; [48] establishes sufficient conditions for subgame-perfect equilibrium without this assumption. [13] obtained sufficient conditions for Nash equilibria, and [47] es-

establishes sufficient and necessary conditions for Nash equilibria. Folk Theorem has also been studied in a broader class of repeated game models. [39] considers Folk Theorem for finitely repeated game with incomplete information. [37] considers infinitely repeated game with imperfect monitoring. [28] considers infinite horizon stochastic games with perfect monitoring, and later [41] considers infinite horizon stochastic games with imperfect monitoring.

A major difference between the above line of work and this work is that Folk Theorems consider the set of payoffs attainable, whereas this work considers the occurrence of off-(stage-game)-Nash play. As we demonstrate in Section 2.2.3, the property considered in Folk Theorems and the local suboptimality property considered in this work do not have direct implications in either direction. Therefore, unlike this research, none of the above research establishes a sufficient and necessary condition for off-(stage-game)-Nash play to occur in finitely repeated games.

Several works in the literature establish additional characterizations on the equilibrium value set in repeated games. When the preconditions of Folk Theorems do not hold, these results provide some characterizations on the equilibrium value set. [25] provides a complete characterization of the set of pure strategy SPE payoff profiles in the limit as the time horizon increases for finitely repeated games with perfect monitoring. [75, 76] characterize limiting behavior of the equilibrium value set of infinitely repeated games with imperfect monitoring as the discount factor approaches 1. [1] further proves properties of the equilibrium value set in infinitely repeated games with discounting and imperfect monitoring. Again, this line of work considers the set of payoffs attainable, whereas our work considers the occurrence of off-(stage-game)-Nash play. Unlike our research, none of the above research establishes a sufficient and necessary condition for off-(stage-game)-Nash play to occur in finitely repeated games.

Chapter 3

On the Impact of Player Capability on Congestion Games

This chapter presents our research on the impact of player capability on congestion games. Section 3.1 presents some background on network congestion games and the complexity class PLS. Section 3.2 introduces the game models we consider in this research (DNC, DNCDA, and GMG). Section 3.3 establishes the complexity results of DNC and DNCDA. Section 3.4 presents the study of the impact of player capability on social welfare in the context of DNCDA and GMG. More specifically, Section 3.4.1 introduces the four capability preference properties; Sections 3.4.2 and 3.4.3 prove the sufficient and necessary conditions for the capability preference properties in DNCDA and GMG respectively; Sections 3.4.4 and 3.4.5 present the complete characterization of how social welfare at equilibrium varies with player capability for the alternating ordering game. Section 3.5 discusses the related work.

3.1 Background

The results in this research are obtained in the context of network congestion games. We present some background on network congestion games in this section.

Network congestion games Network congestion games have been widely studied in the literature [78, 29, 33]. In this work, we focus on atomic network congestion games, where there is a discrete (and finite) set of players, as opposed to non-atomic network congestion games where there is a continuous population of players. An (atomic) network congestion game consists of a set of players $\mathcal{N} = \{1, \dots, n\}$ and a directed graph $(\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the set of vertices in the network and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges in the network. Each edge $e \in \mathcal{E}$ has a delay function $d_e : \mathbb{N} \mapsto \mathbb{R}$ that maps the number of players choosing that edge to the delay of that edge. Each player $i \in \mathcal{N}$ has a source vertex s_i and a sink vertex t_i . The goal of player i is to plan a path from s_i to t_i that minimizes the total delay of the edges on the path. The game is *symmetric* when all players share the same source and sink vertices, i.e., $s_1 = s_2 = \dots = s_n = s$ and $t_1 = t_2 = \dots = t_n = t$. The game is *asymmetric* when different players have different source and sink vertices.

[77] proves that every congestion game has a pure Nash equilibrium. Since network congestion games are congestion games, pure Nash equilibria always exist for network congestion games.

Polynomial local search (PLS) *Polynomial local search (PLS)* [53] is a complexity class for local optimization problems. A local optimization problem L has a set of instances I_L . For each instance $x \in I_L$, there is a set of feasible solutions $F_L(x)$. For each feasible solution $s \in F_L(x)$, there is a cost $c_L(s, x)$ and a neighborhood $N_L(s, x) \subseteq F_L(x)$. An instance of the local optimization problem L is: given $x \in I_L$, find an $s \in F_L(x)$ such that $\forall s' \in N_L(s, x), c_L(s, x) \leq c_L(s', x)$. A local optimization problem L belongs to PLS if there is a polynomial time algorithm for each of the following:

- Given any instance $x \in I_L$, find some feasible solution $s \in F_L(x)$.
- Given any instance $x \in I_L$ and any feasible solution $s \in F_L(x)$, compute the cost $c_L(s, x)$.
- Given any instance $x \in I_L$ and any feasible solution $s \in F_L(x)$, return some

$s' \in N_L(s, x)$ with $c_L(s', x) < c_L(s, x)$, or, if no such s' exists, return ‘done’.

[53] proves that PLS lies between the functional versions of P and NP, i.e., $FP \subseteq PLS \subseteq FNP$. They further prove that if a PLS problem is NP-hard, then $NP = coNP$.

PLS-completeness results are proved using *PLS-reductions*. Given a PLS-complete problem L_1 and a target problem L_2 , a PLS-reduction from L_1 to L_2 requires two polynomial time computable functions f and g such that:

- Given any instance $x_1 \in I_{L_1}$, $f(x_1)$ returns an instance of L_2 , i.e., $f(x_1) \in I_{L_2}$.
- Given any local optimum $s_2 \in F_{L_2}(f(x_1))$ for the instance $f(x_1)$ that is mapped from any $x_1 \in I_{L_1}$, $g(s_2)$ returns a local optimum for x_1 .

[29] proves the following results regarding the complexity of finding a pure Nash equilibrium in network congestion games:

Theorem 3.1.1 ([29]). *There is a polynomial algorithm for finding a pure Nash equilibrium in symmetric network congestion games.*

Theorem 3.1.2 ([29]). *Finding a pure Nash equilibrium in asymmetric network congestion games is PLS-complete.*

3.2 Models

This section presents the game models we consider in this research. In Section 3.2.1, we present a new network congestion game, the *Distance-bounded Network Congestion game (DNC)*, as the basis of our study. Sections 3.2.2 and 3.2.3 present two variants of DNC, the *Distance-bounded Network Congestion game with Default Action (DNCDA)* and the *Gold and Mines Game (GMG)*, where we instantiate our framework by defining simple DSLs that compactly represent the strategy spaces.

3.2.1 Distance-Bounded Network Congestion Game

DNC is a variant of the widely studied network congestion games [29, 78] (Section 3.1). DNC is a symmetric network congestion game in which each player is subject to a

distance bound — i.e., a bound on the number of edges that a player can use.

Definition 3.2.1. *An instance of the Distance-bounded Network Congestion game (DNC) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, b, (d_e)_{e \in \mathcal{E}})$ where:*

- \mathcal{V} is the set of vertices in the network.
- $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges in the network.
- $\mathcal{N} = \{1, \dots, n\}$ is the set of players.
- $s \in \mathcal{V}$ is the source vertex shared by all players.
- $t \in \mathcal{V}$ is the sink vertex shared by all players.
- $b \in \mathbb{N}$ is the bound of the path length.
- $d_e : \mathbb{N} \mapsto \mathbb{R}$ is a non-decreasing delay function on edge e .

We also require that the network has no negative-delay cycles, i.e., for each cycle \mathcal{C} , we require $\sum_{e \in \mathcal{C}} \min_{i \in \mathcal{N}} d_e(i) = \sum_{e \in \mathcal{C}} d_e(1) \geq 0$.

We only consider pure strategies (i.e., deterministic strategies) in this research. The strategy space of a single player contains all $s - t$ simple paths whose length does not exceed b :

$$\mathcal{L}_b \stackrel{\text{def}}{=} \left\{ (p_0, \dots, p_k) \left| \begin{array}{l} p_0 = s, p_k = t, (p_i, p_{i+1}) \in \mathcal{E}, k \leq b, \\ p_i \neq p_j \text{ for } i \neq j \end{array} \right. \right\}$$

In a DNC, as in a general congestion game, a player's goal is to minimize their delay. Let $s_i \stackrel{\text{def}}{=} (p_{i0}, \dots, p_{ik_i}) \in \mathcal{L}_b$ denote the strategy of player i where $i \in \mathcal{N}$. A strategy profile $\mathbf{s} = (s_1, \dots, s_n) \in \mathcal{L}_b^n$ consists of strategies of all players. Let $E_i \stackrel{\text{def}}{=} \{(p_{ij}, p_{i,j+1}) \mid 0 \leq j < k_i\}$ denote the corresponding set of edges on the path chosen by player i . The load on an edge $e \in \mathcal{E}$ is defined as the number of players that occupy this edge: $x_e \stackrel{\text{def}}{=} |\{i \mid e \in E_i\}|$. The delay experienced by player i is $c_i(\mathbf{s}) \stackrel{\text{def}}{=} \sum_{e \in E_i} d_e(x_e)$. A strategy profile \mathbf{s} is a *pure Nash equilibrium (PNE)* if no player can improve their delay by unilaterally changing strategy, i.e., $\forall i \in \mathcal{N} : c_i(\mathbf{s}) = \min_{s' \in \mathcal{L}_b} c_i(\mathbf{s}_{-i}, s')$. All

players experience infinite delay if the distance bound permits no feasible solution (i.e., when $\mathcal{L}_b = \emptyset$). *Social welfare* is defined as the the negative total delay of all players where a larger welfare value means on average players experience less delay: $W(\mathbf{s}) \stackrel{\text{def}}{=} -\sum_{i \in \mathcal{N}} c_i(\mathbf{s})$.

3.2.2 Distance-Bounded Network Congestion Game with Default Action

As we have discussed, we formulate capability restriction as limiting the size of the programs accessible to a player. Here, we propose a variant of DNC where we define a DSL to compactly represent the strategies. We will also show that the size of a program equals the length of the path generated by the program, which can be much smaller than the number of edges in the path. The new game, called *distance-bounded network congestion game with default action (DNCDA)*, requires that each vertex except the source or sink has exactly one outgoing zero-length edge as its default action. All other edges have unit length. A strategy in this game can be compactly described by the actions taken at *divergent points* where a unit-length edge is followed.


Definition 3.2.2. *An instance of DNCDA is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, b, (d_e)_{e \in \mathcal{E}}, (w_e)_{e \in \mathcal{E}})$ where:*

- $w_e \in \{0, 1\}$ is the length of edge e .
- All other symbols have the same meaning as in Definition 3.2.1.

Moreover, we require the following properties:

- A default action, denoted as $\text{DA}(\cdot)$, can be defined for every non-source, non-sink vertex $v \in \mathcal{V}/\{s, t\}$ such that:

$$\begin{aligned} & \left(v, \text{DA}(v) \right) \in \mathcal{E}, \quad w_{(v, \text{DA}(v))} = 0, \\ & \forall u \in \mathcal{V}/\{\text{DA}(v)\} : (v, u) \in \mathcal{E} \implies w_{(v, u)} = 1 \end{aligned}$$



```

if (u == s) { return 2; } else {
  if (u == 3) { return 4; } else {
    return DA(u);
  }
}

```

(a) Example graph structure. Solid arrows are default edges and dashed arrows are unit-length edges.

(b) The shortest program to represent the strategy $(s, 2, 3, 4, 5, t)$.

Figure 3-1: An example of the DNCDA game and a program to represent a strategy.

- *Edges from the source have unit length:* $\forall v \in \mathcal{V} : (s, v) \in \mathcal{E} \implies w_{(s,v)} = 1$
- *The subgraph of zero-length edges is acyclic. Equivalently, starting from any non-source vertex, one can follow the default actions to reach the sink.*

The strategy space of a player contains all $s - t$ simple paths whose length does not exceed b :

$$\mathcal{L}_b \stackrel{\text{def}}{=} \left\{ (p_0, \dots, p_k) \left| \begin{array}{l} p_0 = s, p_k = t, (p_i, p_{i+1}) \in \mathcal{E}, \sum_{i=0}^{k-1} w_{(p_i, p_{i+1})} \leq b, \\ p_i \neq p_j \text{ for } i \neq j \end{array} \right. \right\}$$

Note that the strategy spaces are strictly monotonically increasing up to the longest simple $s - t$ path. This is because for any path p whose length is $b \geq 2$, we can remove the last non-zero-length edge on p and follow the default actions to arrive at t , which gives a new path with length $b - 1$. Formally, we have:

Property 3.2.3. *Let \bar{b} be the length of the longest simple $s - t$ path in a DNCDA instance. For $1 \leq b < \bar{b}$, $\mathcal{L}_b \subsetneq \mathcal{L}_{b+1}$*

We define a Domain Specific Language (DSL) with the following context-free grammar [50] to describe the strategy of a player:

```

Program  $\rightarrow$  return DA(u);
          | if (u == V) { return V; } else {Program}
V  $\rightarrow$  v  $\in$   $\mathcal{V}$ 

```

A program p in this DSL defines a computable function $f_p : \mathcal{V} \mapsto \mathcal{V}$ with semantics similar to the C language where the input vertex is stored in the variable u , as




Figure 3-2: Resource layout for the alternating ordering game, a specific version of GMG. Each dot (resp. cross) is a gold (resp. mine).

illustrated in Figure 3-1. The strategy corresponding to the program p is a path (c_0, \dots, c_k) from s to t where:

$$c_0 = s \quad c_{i+1} = f_p(c_i) \text{ for } i \geq 0 \text{ and } c_i \neq t \quad k = i \text{ if } c_i = t$$

We define the capability of a player as the maximum size of programs that they can use. The size of a program is the depth of its parse tree. Due to the properties of DNCDA, the shortest program that encodes a path from s to t specifies the edge chosen at all divergent points in this path. The size of this program equals the length of the path. Hence the distance bound in the game configuration specifies the capability of each player in the game. To study the game outcome under different player capability constraints, we study DNCDA instances with different values of b .

3.2.3 Gold and Mines Game

We further introduce a particular form of DNCDA called Gold and Mines Game (GMG). It provides a new perspective on how to define the strategy space hierarchy in congestion games. It also enables us to obtain additional characterizations of how social welfare at equilibrium varies with player capability. Intuitively, as shown in Figure 3-2, a GMG instance consists of a few parallel horizontal lines and two types of resources: gold and mine. Resources are placed at distinct horizontal locations on the lines. A player's strategy is a piecewise-constant function to cover a subset of resources. The function is specified by a program using if-statements.

Definition 3.2.4. *An instance of GMG is a tuple $G = (\mathcal{E}, K, \mathcal{N}, r_g, r_m, b)$ where:*

- \mathcal{E} is the set of resources. Each resource $e \in \mathcal{E}$ is described by a tuple (x_e, y_e, α_e) , where (x_e, y_e) denotes the position of the resource in the x - y plane, and $\alpha_e \in$

$\{\text{gold}, \text{mine}\}$ denotes the type of the resource. Each resource has a distinct value of x , i.e. $x_e \neq x_{e'}$ for all $e \neq e'$.

- $K \in \mathbb{N}$ is the number of lines the resources can reside on. All resources are located on lines $y = 0, y = 1, \dots, y = K - 1$, i.e. $\forall e, y_e \in \{0, 1, \dots, K - 1\}$.
- $\mathcal{N} = \{1, \dots, n\}$ denotes the set of players.
- $r_g : \mathbb{N} \mapsto \mathbb{R}^+$ is the payoff function for gold. r_g is a positive function.
- $r_m : \mathbb{N} \mapsto \mathbb{R}^-$ is the payoff function for mine. r_m is a negative function.
- $b \in \mathbb{N}$ is the level in the strategy space hierarchy defined by the domain-specific language \mathcal{L} (defined below). The strategy space is then \mathcal{L}_b .

The strategy s_i of player i is represented by a function $f_i(\cdot)$ that conforms to a domain-specific language \mathcal{L} with the following grammar:

```

Program → return C; | if ( $x < t$ ) { return C; } else { Program }
C       → 0 | 1 | ... |  $K - 1$ 
t       ∈ ℝ

```

This DSL defines a natural strategy space hierarchy by restricting the number of if-statements in the program. A program with $b - 1$ if-statements represents a piecewise-constant function with at most b segments. We denote \mathcal{L}_b as the level b strategy space which includes functions with at most $b - 1$ if-statements.

Player i covers the resources that their function f_i passes: $E_i = \{e \mid f_i(x_e) = y_e\}$. The load on each resource is the number of players that covers it: $x_e = |\{i \mid e \in E_i\}|$. Each player's payoff is $u_i = \sum_{e \in E_i} r_e(x_e)$, where r_e is either r_g or r_m depending on the resource type. The social welfare is $W(\mathbf{s}) = \sum_{i \in \mathcal{N}} u_i$.

Proposition 3.2.5. *Any instance of GMG can be represented as an instance of DNCDA.*

Proof. To represent an instance of GMG as an instance of DNCDA, we first order all resources in the GMG instance by increasing order in x_e , such that $x_1 < x_2 < \dots < x_{|\mathcal{E}|}$. Then the network in the corresponding DNCDA has the following vertices. We

assign positions to each vertex in terms of coordinates in a 2-D plane. This is only to help understand the correspondence between the GMG instance and the DNCDA instance:

- $v_{i,j}$ at $(\frac{x_i+x_{i+1}}{2}, j)$ for all $i = 1, \dots, |\mathcal{E}| - 1$ and $j = 0, \dots, K - 1$.
- $v_{0,j}$ at $(x_1 - 1, j)$ and $v_{|\mathcal{E}|,j}$ at $(x_{|\mathcal{E}|} + 1, j)$ for all $j = 0, \dots, K - 1$.
- A source s and a sink t .

The corresponding DNCDA has the following edges:

- $(v_{i,j}, v_{i+1,j})$ for all $i = 0, \dots, |\mathcal{E} - 1|$ and $j = 0, \dots, K - 1$, with length $w_e = 0$. If $y_{i+1} = j$, then the delay function $d_e = -r_{\alpha_{i+1}}$; otherwise, $d_e = 0$.
- $(v_{i,j}, v_{i,j'})$ for all $i = 1, \dots, |\mathcal{E} - 1|$ and $(j, j') \in \{0, \dots, K - 1\}^2$, with $w_e = 1$ and $d_e = 0$.
- $(s, v_{0,j})$ and $(v_{|\mathcal{E}|,j}, t)$ for all $j = 0, \dots, K - 1$, with $w_e = 1$ and $d_e = 0$.

The distance bound in the corresponding DNCDA is the same as the maximum number of segments in GMG.

□

3.3 Complexity Results

In this section, we establish results on the computational complexities of 1) finding a pure Nash equilibrium, 2) finding the best/worst social welfare at pure Nash equilibrium, and 3) finding the best social welfare across all pure strategy profiles (not only at equilibrium), in the context of DNC and DNCDA.

3.3.1 Complexity Results for DNC

Lemma 3.3.1. *DNC belongs to PLS.*

Proof. DNC is a potential game where local minima of its potential function correspond to PNEs [77]. Clearly there are polynomial algorithms for finding a feasible solution or evaluating the potential function. We only need to show that computing the best response of some player i given the strategies of others is in P. For each $v \in \mathcal{V}$, we define $f(v, d)$ to be the minimal delay experienced by player i over all paths from s to v with length bound d . It can be recursively computed via $f(v, d) = \min_{u \in \mathcal{V}: (u, v) \in \mathcal{E}} (f(u, d - 1) + d_{(u, v)}(x_{(u, v)} + 1))$ where $x_{(u, v)}$ is the load on edge (u, v) caused by other players. The best response of player i is then $f(t, b)$. If there are cycles in the solution, we can remove them without affecting the total delay because cycles must have zero delay in the best response. \square


Theorem 3.3.2. *DNC is PLS-complete.*

Proof. We have shown that DNC belongs to PLS. Now we present a PLS-reduction from a PLS-complete game to finish the proof.


The *quadratic threshold game* [3] is a PLS-complete game in which there are n players and $n(n + 1)/2$ resources. The resources are divided into two sets $\mathcal{R}^{\text{in}} = \{r_{ij} \mid 1 \leq i < j \leq n\}$ for all unordered pairs of players $\{i, j\}$ and $\mathcal{R}^{\text{out}} = \{r_i \mid i \in \mathcal{N}\}$. For ease of exposition, we use r_{ij} and r_{ji} to denote the same resource. Player i has two strategies: $S_i^{\text{in}} = \{r_{ij} \mid j \in \mathcal{N}/\{i\}\}$ and $S_i^{\text{out}} = \{r_i\}$.

Extending the idea in [3], we reduce from the quadratic threshold game to DNC. To simplify our presentation, we assign positive integer weights to edges. Each weighted edge can be replaced by a chain of unit-length edges to obtain an unweighted graph.

Figure 3-3 illustrates the game with four players. We create $n(n + 1)/2$ vertices arranged as a lower triangle. We use v_{ij} to denote the vertex at the i^{th} row (starting from top) and j^{th} column (starting from left) where $1 \leq j \leq i \leq n$. The vertex v_{ij} is connected to $v_{i, j+1}$ with an edge of length i when $j < i$ and to $v_{i+1, j}$ with a unit-length edge when $i < n$. This design ensures that the shortest path from v_{i1} to v_{ni} is the right-down path. The resource r_{ij} is placed at the off-diagonal vertex v_{ij} , which can be implemented by splitting the vertex into two vertices connected by a unit-length edge with the delay function of r_{ij} . Note that this implies visiting a vertex v_{ij} incurs



(a) The graph structure



(b) Splitting the vertex containing resource r_{ij}

Figure 3-3: The DNC instance corresponding to a four-player quadratic threshold game. The distance bound $b = 19$. Non-unit-length edges have labels to indicate their lengths. Dashed gray edges correspond to the S_i^{out} strategies.

a distance of 1 where $i \neq j$. We then create vertices s_i and t_i for $1 \leq i \leq n$ with unit-length edges (s_i, v_{i1}) and (v_{ni}, t_i) . We connect s_i to t_i with an edge of length w_i , which represents the resource r_i . Let b be the distance bound. We will determine the values of w_i and b later. The source s is connected to s_i with an edge of length $b - w_i - 1$. Vertices t_i are connected to the sink t via unit-length edges.

We define the following delay functions for edges associated with s or t :

$$d_{(s, s_i)}(x) = \mathbb{1}_{x \geq 2} \cdot (|\mathcal{N}| + 1)R \quad d_{(t_i, t)}(x) = (|\mathcal{N}| - i)R$$

$$\text{where } R = \left(\sum_{r \in \mathcal{R}^{\text{in}} \cup \mathcal{R}^{\text{out}}} \max_{i \in \mathcal{N}} d_r(i) \right) + 1$$

We argue that player i chooses edges (s, s_i) and (t_i, t) in their best responses. Since R is greater than the maximum possible sum of delays of resources in the threshold game, a player's best response must first optimize their choice of edges

linked to s or t . If two players choose the edge (s, s_i) , one of them can improve their latency by changing to an unoccupied edge $(s, s_{i'})$. Therefore, we can assume the i^{th} player chooses edge (s, s_i) WLOG. Player i can also decrease their latency by switching from (t_j, t) to (t_{j+1}, t) for any $j < i$ unless their strategy is limited by the distance bound when $j = i$.

Player i now has only two strategies from s_i to t_i due to the distance bound, corresponding to their strategies in the threshold game: (i) following the right-down path, namely $(s_i, v_{i1}, \dots, v_{ii}, v_{i+1,i}, \dots, v_{ni}, t_i)$, where they occupy resources corresponding to S_i^{in} ; and (ii) using the edge (s_i, t_i) , where they occupy the resource $S_i^{\text{out}} = \{r_i\}$. Clearly PNEs in this DNC correspond to PNEs in the original quadratic threshold game.

Now we determine the values of w_i and b . The shortest paths from s_i to t_i should be either the right-down path or the edge (s_i, t_i) . This implies that $w_i = a_i + b_i + c_i$ where $a_i = i(i - 1) + 1$ is the total length of horizontal edges, $b_i = n + 1 - i$ is the total length of vertical edges, and $c_i = n - 1$ is the total length of edges inside v_{ij} for resources r_{ij} . Hence $w_i = i(i - 2) + 2n + 1$. The bound b should accommodate player n who has the longest path and is set as $b = w_n + 2 = n^2 + 3$. \square

Theorem 3.3.3. *Computing the best social welfare (i.e., minimal total delay) among PNEs of a DNC is NP-hard.*




Figure 3-4: Illustration of the DNC instance corresponding to a 3-partition problem. Double-line edges are slow edges, dashed edges are fast edges, and other edges have no delay. Non-unit-length edges have labels to indicate their lengths. Deciding whether the total delay can be bounded by $6m - 3$ is NP-complete.

Proof. We reduce from the strongly NP-complete 3-partition problem [42].

In the *3-partition problem*, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$.

The question Q_1 is: Can S be partitioned into m sets S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$? Note that due to the strong NP-completeness of 3-partition, we assume the numbers use unary encoding so that the DNC graph size is polynomial.

As in the proof of Theorem 3.3.2, we assign a weight $w_e \in \mathbb{Z}^+$ to each edge e . The DNC instance has two types of edges with non-zero delay: fast edge and slow edge, with delay functions $d_{\text{fast}}(x) = \mathbb{1}_{x \geq 1} + 2\mathbb{1}_{x \geq 2}$ and $d_{\text{slow}}(x) = 2$.

As illustrated in Figure 3-4, for each integer a_i , we create a pair of vertices (s_i, t_i) connected by a fast edge with $w_{(s_i, t_i)} = a_i$. We create a new vertex t_0 as the source while using t_{3m} as the sink. For $0 \leq i < 3m$, we connect t_i to t_{i+1} by a unit-length slow edge and t_i to s_{i+1} by a unit-length edge without delay. There are m players who can choose paths with length bounded by $b = T + 3m$.

We ask the question Q_2 : Is there a PNE in the above game where the total delay is no more than $m(6m - 3)$? Each player prefers an unoccupied fast edge to a slow edge but also prefers a slow edge to an occupied fast edge due to the above delay functions. Since $T/4 < a_i < T/2$, the best response of a player contains either 2 or 3 fast edges, contributing $6m - 2$ or $6m - 3$ to the total delay in either case. Best social welfare of $m(6m - 3)$ is only achieved when every player chooses 3 fast edges, which also means that their choices together constitute a partition of the integer set S in Q_1 . Therefore, Q_2 and Q_1 have the same answer. \square

Theorem 3.3.4. *Computing the optimal global welfare of pure strategies in DNC is NP-hard.*

Proof. We use the same reduction from the 3-partition problem as presented in the proof of Theorem 3.3.3. Recall that in the 3-partition problem, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$. The question Q_1 is: Can S be partitioned into m sets S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$?

We use the same construction of the DNC game instance as in the proof of Theorem 3.3.3. We ask the question Q_2 : Is there a pure strategy profile in the constructed game where the total delay is no more than $m(6m - 3)$? Following the same argu-

ment as the proof of Theorem 3.3.3, we can see that the optimal global welfare of any “centralized” solution (where players cooperate to minimize total delay instead of selfishly minimizing their own delay) achieves $m(6m - 3)$ if and only if the original 3-partition problem has a solution. Therefore, Q_2 and Q_1 have the same answer. \square

Theorem 3.3.5. *Computing the worst social welfare (i.e., maximal total delay) among PNEs of a DNC is NP-hard.*

Proof. Again, we reduce from the 3-partition problem. Recall that in the 3-partition problem, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$. The question Q_1 is: Can S be partitioned into m sets S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$?




Figure 3-5: Illustration of the DNC instance corresponding to a 3-partition problem, for proving NP-hardness of computing the worst social welfare among PNEs. Double-line edges are slow edges, dashed edges are fast edges, and other edges except the ones that connect to the source s have no delay. Non-unit-length edges have labels to indicate their lengths.

We build on the construction of DNC used in the proof of Theorem 3.3.3. We create a new vertex s as the source and connect s to t_0 and s_i where $1 \leq i \leq 3m$:

$$\begin{aligned}
 w_{(s, t_0)} &= 1 & d_{(s, t_0)}(x) &= \mathbb{1}_{x \geq m+1} \cdot R \\
 w_{(s, s_i)} &= T + i - a_i + 1 & d_{(s, s_i)}(x) &= \mathbb{1}_{x \geq 2} \cdot R
 \end{aligned}
 \quad \text{where } R = 9m + 2$$


The delay functions on fast and slow edges are changed to $d_{\text{fast}}(x) = 2\mathbb{1}_{x \geq 2} + 2\mathbb{1}_{x \geq 3}$ and $d_{\text{slow}}(x) = 3$. Figure 3-5 presents the constructions.

There are $4m$ players in this game with a distance bound $b = T + 3m + 1$. Since R is greater than the delay on any path from s_i or t_0 to the sink, we can assume WLOG that player i chooses (s, s_i) where $1 \leq i \leq 3m$, and players $3m + 1, \dots, 4m$ all choose (s, t_0) . The first $3m$ players generate a total delay of $D_0 = d_{\text{slow}} \cdot 3m(3m - 1)/2 = 9m(3m - 1)/2$ where player i occupies one fast edge and $3m - i$ slow edges. Each of the last m players occupies 2 or 3 fast edges in their best response. Occupying one fast edge incurs 4 total delay among all players because one of the first $3m$ players also uses that edge. Therefore, each of the last m players contributes $9m + 2$ or $9m + 3$ to the total delay. We ask the question Q_2 : Is there a PNE where the total delay is at least $D_0 + m(9m + 3)$? From our analysis, we can see that Q_2 and Q_1 have the same answer. \square


3.3.2 Complexity Results for DNCDA

Theorem 3.3.6. *DNCDA is PLS-complete.*

Proof. The best response of a player in DNCDA can be computed in polynomial time similarly to DNC. Now we prove its PLS-completeness by presenting a reduction from the quadratic threshold game. We modify the network layout in the proof of Theorem 3.3.2 as follows. We assign zero length to the vertical edges $(v_{ij}, v_{i+1,j})$, edges for r_{ij} , and edges in the set $\cup_{1 \leq i \leq n} \{(s_i, v_{i1}), (v_{ni}, t_i), (t_i, t)\}$. We also redefine the lengths of some other edges: $w_{(s_i, t_i)} = i(i - 1) + \mathbb{1}_{i=1}$ and $w_{(s, s_i)} = b - w_{(s_i, t_i)}$. The distance bound is $b = n(n - 1) + 1$. Note that after replacing weighted edges with a chain of unit-length edges to build the DNCDA network, some vertices will only have one unit-length outgoing edge, which violates the requirements of default action. In this case, we add auxiliary vertices with zero-length edges and a resource r_∞ that has sufficiently large delay to disincentivize any player from taking the zero-length auxiliary edges. Figure 3-6(b) illustrates this construction. In this DNCDA instance, player i will choose edges (s, s_i) and (t_i, t) . They then choose between the right-down



(a) The graph structure



(b) A gadget to add a default action for u

Figure 3-6: Illustration of the DNCDA instance corresponding to a four-player quadratic threshold game. The distance bound is $b = 13$. Non-zero-length edges have labels to indicate their lengths.

path from s_i to t_i or the edge (s_i, t_i) which correspond to the two strategies in the quadratic threshold game respectively. \square

Theorem 3.3.7. *Computing the best social welfare (i.e., minimal total delay) among PNEs of a DNCDA is NP-hard.*

Proof. Again, we reduce from the 3-partition problem. Recall that in the 3-partition problem, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$. The question Q_1 is: Can S be partitioned into m sets S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$?

We adopt the reduction given in the proof of Theorem 3.3.3. We modify the edge weights as $w_{(t_i, s_{i+1})} = 1$, $w_{(s_i, t_i)} = a_i - 1$, and $w_{(t_i, t_{i+1})} = 0$, as illustrated in Figure 3-7. Same as in the proof of Theorem 3.3.3, the DNCDA instance has two types of edges with non-zero delay: fast edge and slow edge, with delay functions $d_{\text{fast}}(x) = \mathbb{1}_{x \geq 1} + 2\mathbb{1}_{x \geq 2}$ and $d_{\text{slow}}(x) = 2$.




Figure 3-7: Illustration of the DNCDA instance corresponding to a 3-partition problem. Double-line edges are slow edges, dashed edges are fast edges, and other edges have no delay. Non-zero-length edges have labels to indicate their lengths. Deciding whether the total delay can be bounded by $6m - 3$ is NP-complete.

Similar to the construction used in the proof of Theorem 3.3.6, for vertices that do not have a zero-length outgoing edge, we use the gadget presented in Figure 3-6(b) to add auxiliary vertices with zero-length edges and a resource r_∞ that has sufficiently large delay to disincentivize any player from taking the zero-length auxiliary edges. We add a source vertex s before t_0 , and let the edge from s to t_0 have unit-length and zero delay. The distance bound is $b = T + 1$.

We ask the question Q_2 : Is there a PNE in the above game where the total delay is no more than $m(6m - 3)$? Following the exact same argument as in the proof of Theorem 3.3.3, we know that Q_2 and Q_1 have the same answer.

□

Theorem 3.3.8. *Computing the worst social welfare (i.e., maximal total delay) among PNEs of a DNCDA is NP-hard.*

Proof. Again, we reduce from the 3-partition problem. Recall that in the 3-partition problem, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$. The question Q_1 is: Can S be partitioned into m sets S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$?

We adopt the reduction given in the proof of Theorem 3.3.5. Figure 3-8 presents the construction. For edges connecting with the source s , we have for $1 \leq i \leq 3m$:

$$\begin{aligned}
 w_{(s, t_0)} &= 1 & d_{(s, t_0)}(x) &= \mathbb{1}_{x \geq m+1} \cdot R \\
 w_{(s, s_i)} &= T + 1 - a_i + 1 & d_{(s, s_i)}(x) &= \mathbb{1}_{x \geq 2} \cdot R
 \end{aligned}
 \quad \text{where } R = 9m + 2$$




Figure 3-8: Illustration of the DNCDA instance corresponding to a 3-partition problem, for proving NP-hardness of computing the worst social welfare among PNEs. Double-line edges are slow edges, dashed edges are fast edges, and other edges except the ones that connect to the source s have no delay. Non-zero-length edges have labels to indicate their lengths.

The delay functions on fast and slow edges are $d_{\text{fast}}(x) = 2\mathbb{1}_{x \geq 2} + 2\mathbb{1}_{x \geq 3}$ and $d_{\text{slow}}(x) = 3$. Similar to the construction used in the proof of Theorem 3.3.6, for vertices that do not have a zero-length outgoing edge, we use the gadget presented in Figure 3-6(b) to add auxiliary vertices with zero-length edges and a resource r_∞ that has sufficiently large delay to disincentivize any player from taking the zero-length auxiliary edges. There are $4m$ players in this game with a distance bound $b = T + 1$.

We ask the question Q_2 : Is there a PNE where the total delay is at least $9m(3m - 1)/2 + m(9m + 3)$? Following the exact same argument as in the proof of Theorem 3.3.5, we know that Q_2 and Q_1 have the same answer. \square

Theorem 3.3.9. *Computing the optimal global welfare of pure strategies in DNCDA is NP-hard.*

Proof. We use the same reduction from the 3-partition problem as presented in the proof of Theorem 3.3.7. Recall that in the 3-partition problem, we are given a multiset of $3m$ positive integers $S = \{a_i \in \mathbb{Z}^+ \mid 1 \leq i \leq 3m\}$ and a number T such that $\sum a_i = mT$ and $T/4 < a_i < T/2$. The question Q_1 is: Can S be partitioned into m sets

S_1, \dots, S_m such that $\sum_{a_i \in S_j} a_i = T$ for all $1 \leq j \leq m$?

We use the same construction of the DNCDA game instance as in the proof of Theorem 3.3.7. We ask the question Q_2 : Is there a pure strategy profile in the constructed game where the total delay is no more than $m(6m - 3)$? Following the same argument as the proof of Theorem 3.3.7, we can see that the optimal global welfare of any “centralized” solution (where players cooperate to minimize total delay instead of selfishly minimizing their own delay) achieves $m(6m - 3)$ if and only if the original 3-partition problem has a solution. Therefore, Q_2 and Q_1 have the same answer. \square

3.4 Impact of Player Capability on Social Welfare

In this section, we study the impact of player capability on social welfare at pure Nash equilibria in the context of DNCDA and GMG. We first introduce four *capability preference* properties that characterize the impact of player capability on social welfare for general games (Section 3.4.1). Then, in Sections 3.4.2 and 3.4.3, we prove sufficient and necessary conditions on the delay functions such that each of the capability preference properties holds for any network topology in the context of DNCDA and GMG. Finally, in Section 3.4.4, we fully characterize how social welfare at equilibrium varies with player capability for a specific version of GMG called the alternating ordering game.

3.4.1 Capability Preference Properties

We first introduce four capability preference properties for general games. Given a game with a finite hierarchy of player capabilities, we use \mathcal{L}_b to denote the strategy space when player capability is bounded by b . Assuming the maximal capability is \bar{b} (e.g. in DNCDA, \bar{b} is the length of the longest $s-t$ simple path). We use $\text{Equil}(b) \subseteq \mathcal{L}_b^n$ to denote the set of all PNEs at capability level b . We define $W_b^+ \stackrel{\text{def}}{=} \max_{\mathbf{s} \in \text{Equil}(b)} W(\mathbf{s})$ to be the best social welfare at equilibrium and $W_b^- \stackrel{\text{def}}{=} \min_{\mathbf{s} \in \text{Equil}(b)} W(\mathbf{s})$ the worst social welfare.

Definition 3.4.1. *A game is capability-positive if social welfare at equilibrium cannot decrease as players become more capable, i.e., $\forall 1 \leq b < \bar{b}, W_b^+ \leq W_{b+1}^-$.*

Definition 3.4.2. *A game is max-capability-preferred if the worst social welfare at equilibrium under maximal player capability is at least as good as any social welfare at equilibrium under lower player capability, i.e., $\forall 1 \leq b < \bar{b}, W_b^+ \leq W_b^-$.*

Note that capability-positive implies max-capability-preferred. But max-capability-preferred does not imply capability-positive; we present an example Gold and Mines Game that is max-capability-preferred but not capability-positive in Example 3.4.3.

Example 3.4.3. *This game has $K = 3$ lines. The set of resources \mathcal{E} contains three gold and no mine, with one gold on each line. There are $n = 3$ players. The payoff function r_g is given by $r_g(1) = 11, r_g(2) = 5, r_g(3) = 4$. In this game, the maximal capability $\bar{b} = 3$.*

When $b = 1$, the PNEs in this game are each player chooses a different line (with each assignment of players to lines corresponding to a PNE), so $W_1^+ = W_1^- = 3 \cdot r_g(1) = 33$. When $b = 2$, in every PNE, each player covers two gold and each gold is covered by two players, so $W_2^+ = W_2^- = 3 \cdot 2 \cdot r_g(2) = 30$. When $b = 3$, the only PNE is each player covers all three gold, so $W_b^+ = W_b^- = 3 \cdot 3 \cdot r_g(3) = 36$. Since $W_1^+ > W_2^-$, the game is not capability-positive. Since $W_1^+ \leq W_b^-$ and $W_2^+ \leq W_b^-$, the game is max-capability-preferred.

We then define analogous properties for games where less capable players lead to better outcomes:

Definition 3.4.4. *A game is capability-negative if social welfare at equilibrium cannot increase as players become more capable, i.e., $\forall 1 \leq b < \bar{b}, W_{b+1}^+ \leq W_b^-$.*

Definition 3.4.5. *A game is min-capability-preferred if the worst social welfare at equilibrium under minimal player capability is at least as good as any social welfare at equilibrium under higher player capability, i.e., $\forall b \geq 2, W_b^+ \leq W_1^-$.*

	DNCDA (Section 3.4.2)	GMG (Section 3.4.3)
Resource layout	On a directed graph	On parallel horizontal lines
Strategy space	Paths from s to t	Piecewise-constant functions
Delay (payoff)	Non-negative non-decreasing	$r_g(\cdot)$ positive, $r_m(\cdot)$ negative
capability-positive	$d(\cdot)$ is a constant function	$r_g(\cdot), r_m(\cdot)$ are constant functions
max-capability-preferred	$d(\cdot)$ is a constant function	$w(x) = xr_g(x)$ attains maximum at $x = n$
capability-negative	$d(\cdot)$ is the zero function	Never
min-capability-preferred	$d(\cdot)$ is the zero function	Never

Table 3.1: Necessary and sufficient conditions on the delay or payoff functions such that the capability preference properties hold universally.

3.4.2 Sufficient and Necessary Conditions for Capability Preference Properties in DNCDA

We focus on a restricted version of DNCDA where all edges share the same delay function; formally, we consider the case $\forall e \in \mathcal{E} : d_e(\cdot) = d(\cdot)$ where $d(\cdot)$ is non-negative and non-decreasing. We call this game *distance-bounded network congestion game with default action and shared delay (DNCDA)*. We aim to find sufficient and necessary conditions on $d(\cdot)$ under which each of the capability preference properties hold *universally* (i.e., for all network configurations of DNCDA). This means that if $d(\cdot)$ satisfies the proven condition, then the target property (e.g. capability-positive) holds for all possible network configurations; if $d(\cdot)$ does not satisfy the proven condition, then for any such delay function, we can always find a network configuration where the target property (e.g. capability-positive) does not hold for the game. Table 3.1 summarizes the results.

Theorem 3.4.6. *DNCDA is universally capability-positive if and only if $d(\cdot)$ is a constant function.*

Proof. If $d(\cdot)$ is a constant function, the total delay achieved by a strategy is not affected by the load condition of each edge (thus not affected by other players' strategies). So each player's strategy in any PNE is the one in \mathcal{L}_b that minimizes the total delay under the game layout. Denote this minimum delay as $\delta(b)$. For any $b \geq 1$,

we have $\mathcal{L}_b \subseteq \mathcal{L}_{b+1}$, so $\delta(b) \geq \delta(b+1)$. And for any $\mathbf{s} \in \text{Equil}(b)$, $W(\mathbf{s}) = -n\delta(b)$. Hence $W_b^+ \leq W_{b+1}^-$.

If $d(\cdot)$ is not a constant function, we show that there exists an instance of DNCDAS with delay function $d(\cdot)$ that is not capability-positive. Define $v = \min \{x \mid d(x) \neq d(x+1)\}$. It follows that $d(v') = d(v)$ for all $v' \leq v$. We consider the cases $d(v) = 0$ and $d(v) > 0$ separately.

Case 1: $d(v) > 0$ Denote $\rho = \frac{d(v+1)}{d(v)}$. Since $d(\cdot)$ is non-decreasing, $\rho > 1$. We construct a game with the network layout in Figure 3-9a with $n = v + 1$ players. This game is a counterexample for the capability-positive property with $b = 1$ (i.e., $W_1^+ > W_2^-$).


First, it is easy to see that the PNEs when $b = 1$ are a players take the upper path and $v + 1 - a$ players take the lower path, where $1 \leq a \leq v$. All PNEs achieves a social welfare of $W_1 = -(v + 1)(N_1 + N_2 + 3)d(v)$.

We set the constants $N_1 = \lfloor \frac{1}{\rho-1} \rfloor$ and $N_2 = \lfloor (N_1 + 2)\rho \rfloor - 1$, which ensures $N_1 > \frac{1}{\rho-1} - 1$ and $(N_1 + 2)\rho - 2 < N_2 \leq (N_1 + 2)\rho - 1$. We claim that one PNE when $b = 2$ is that all players choose the path from upper left to lower right using the switching edge in the middle. Under this strategy profile, each player has a total delay of $\delta = (2N_1 + 3)d(v + 1)$. This is an equilibrium because if any player changes their strategy to the upper or the lower horizontal path (the only two alternative strategies) the new total delay is $\delta' = (N_1 + 1)d(v + 1) + (N_2 + 2)d(v) > \delta$ because $\frac{\delta' - \delta}{d(v)} = N_2 - ((N_1 + 2)\rho - 2) > 0$. The social welfare is $W_2 = -(v + 1)(2N_1 + 3)d(v + 1)$. Note that


$$\begin{aligned} \frac{W_1 - W_2}{(v + 1)d(v)} &= (N_1 + 1)(\rho - 1) - (N_2 + 2 - (N_1 + 2)\rho) \\ &> \left(\frac{1}{\rho - 1} - 1 + 1\right)(\rho - 1) - ((N_1 + 2)\rho - 1 + 2 - (N_1 + 2)\rho) = 0. \end{aligned}$$

Hence $W_1^+ \geq W_1 > W_2 \geq W_2^-$.


Case 2: $d(v) = 0$ We construct a game with the network layout in Figure 3-



(a) Counterexample for capability-positive in the case $d(v) > 0$.



(b) Counterexample for capability-positive in the case $d(v) = 0$. Edges between filled nodes have non-zero delay in an equilibrium when $b = 2$.



(c) Counterexample for capability-negative.

Figure 3-9: Counterexamples when $d(\cdot)$ does not meet the conditions. Dashed arrows denote unit-length edges and solid arrows denote zero-length edges (default action). Every edge shares the same delay function $d(\cdot)$.

9b where there are $2v$ players. With $b = 1$, half of the players choose the upper path and the others choose the lower path, which has a social welfare $W_1 = 0$.

With the bound $b = 2$, we consider this strategy: (i) v players take the path $(s, N_1$ edges, lower right N_2 edges, $t)$; and (ii) the other v players take the path $(s, \text{lower left } N_2$ edges, N_1 edges, $t)$. We choose N_1 and N_2 to be positive integers that satisfy $\frac{N_2+1}{N_1} > \frac{d(2v)}{d(v+1)}$. There are $2v$ players occupying the N_1 edges that incur a delay of $N_1 d(2v)$ on each player. The social welfare $W_2 = -2v N_1 d(2v) < W_1$. The above strategy is an equilibrium because for each player the alternative strategy to avoid the N_1 congestion edges is to take the lower path $(s, N_2$ edges, N_2 edges, $t)$ which has a delay of $(N_2 + 1)d(v + 1) > N_1 d(2v)$. Hence $W_1^+ \geq W_1 > W_2 \geq W_2^-$. \square

Theorem 3.4.7. *DNCDAS is universally max-capability-preferred if and only if $d(\cdot)$ is a constant function.*

Proof. The “if” part follows from Theorem 3.4.6 since a capability-positive game is also max-capability-preferred. The constructed games in the proof of Theorem 3.4.6

also serve as the counterexamples to prove the “only if” part. \square

Theorem 3.4.8. *DNCDAS is universally capability-negative if and only if $d(\cdot)$ is the zero function.*

Proof. If $d(\cdot) = 0$, then all PNEs have welfare 0, which implies capability-negative. If $d(\cdot)$ is not the zero function, denote $v = \min \{x \mid d(x) \neq 0\}$. We construct a game with the network layout shown in Figure 3-9c with $n = v$ players. When $b = 1$, all players use the only strategy with a social welfare $W_1 = -3vd(v)$. When $b = 2$: if $v = 1$, the player will choose both dashed paths and achieves $W_2 = -2d(1)$; if $v \geq 2$, the players will only experience delay on the first edge by splitting between the default path and the shortcut dashed path, which achieves a welfare $W_2 = -vd(v)$. In both cases, the game is not capability-negative since $W_1^- \leq W_1 < W_2 \leq W_2^+$. \square

The same argument can also be used to prove the following result:

Theorem 3.4.9. *DNCDAS is universally min-capability-preferred if and only if $d(\cdot)$ is the zero function.*

Proof. The “if” part follows from Theorem 3.4.8 since a capability-negative game is also min-capability-preferred. The constructed games in the proof of Theorem 3.4.8 also serve as the counterexamples to prove the “only if” part. \square

3.4.3 Sufficient and Necessary Conditions for Capability Preference Properties in GMG

Similar as in Section 3.4.2, here we aim to find sufficient and necessary conditions on the payoff functions $r_g(\cdot)$ and $r_m(\cdot)$ under which each of the capability preference properties hold *universally* (i.e., for all game layout) in the context of GMG. The results are summarized in Table 3.1.

Theorem 3.4.10. *GMG is universally capability-positive if and only if both $r_g(\cdot)$ and $r_m(\cdot)$ are constant functions.*

Proof. If both r_g and r_m are constant functions, the same proof as in Theorem 3.4.6 applies to show that capability-positive holds universally here.

If r_g and r_m are not both constant functions, we show that there is always some instance of GMG with payoff r_g and r_m that is not capability-positive.

If r_g is not a constant function, let $v = \min \{x \mid r_g(x) \neq r_g(x+1)\}$. Denote $\rho = \frac{r_g(v+1)}{r_g(v)}$, then $\rho \neq 1$, and $r_g(v') = r_g(v)$ for all $v' \leq v$. We show that we can always construct a game with $n = v + 1$ players and $K = 2$ lines that is a counterexample for the capability-positive property with $b = 1$ (i.e., $W_1^+ > W_2^-$).

Case 1. $\rho < 1$ The layout of the constructed game is shown in Figure 3-10. All the resources are gold. Block B is constructed as follows. Following the increasing order of x , $y_t = k$ means the t -th point is on line $y = k$. $N_0(t)$ ($N_1(t)$) denotes the number of points on line $y = 0$ ($y = 1$) within the first t points. Denote $D(t) = \rho N_0(t) - N_1(t)$. We use the following algorithm to position the gold:

1. while $N_0(t) < \frac{3}{1-\rho}$, do:
 - If $D(t) \leq 1$, put $y_{t+1} = 0$; else, put $y_{t+1} = 1$;
 - $t \leftarrow t + 1$;
2. while $D(t) \leq 1 + \rho$, do: $y_{t+1} = 0, t \leftarrow t + 1$.

Denote the total number of gold put down as N . Under the above construction, the following properties hold:

- For all $t = 1 \dots N$, $D(t) > 0$ (all prefixes are better)
- For all $t = 0 \dots N - 1$, $D(N) - D(t) > 0$ (all suffixes are better)
- $D(N) < 3$
- $N_0(N) > \frac{3}{1-\rho}$

Block B' is obtained by flipping B in both x and y direction.




Figure 3-10: Counterexample for the case $\rho < 1$ for r_g and $\rho > 1$ for r_m . Each dot is a resource. The upper (lower) line is $y = 1$ ($y = 0$). Block B and B' are centrosymmetric. The direction of increasing x is from left to right.

For $b = 2$, a PNE is all players choose $y = 0$ in B and $y = 1$ in B' . This is a PNE because of the prefix and suffix properties. Its social welfare $W_2 = (v + 1) \cdot 2N_0(N)r_g(v + 1)$. For $b = 1$, the PNEs are a players chooses $y = 0$ and $v + 1 - a$ players choose $y = 1$, where $1 \leq a \leq v$. All PNEs have the same social welfare $W_1 = (v + 1) \cdot (N_0(N) + N_1(N))r_g(v)$. Then

$$\frac{W_1 - W_2}{(v + 1)r_g(v)} = N_0(N) + N_1(N) - 2\rho N_0(N) = N_0(N)(1 - \rho) - D(N) > 3 - 3 = 0.$$

Therefore $W_1 > W_2$, which implies $W_1^+ > W_2^-$.

Case 2. $\rho > 1$ The layout of the constructed game is in Figure 3-11. We choose $N > \frac{\rho}{\rho - 1}$. A PNE for $b = 2$ is all players choose the first N gold on $y = 0$ and the N gold on $y = 1$, which has welfare $W_2 = (v + 1) \cdot 2Nr_g(v + 1)$. A PNE for $b = 1$ is all players choose $y = 0$, which has welfare $W_1 = (v + 1) \cdot (2N + 1)r_g(v + 1)$. Clearly $W_1 > W_2$, so $W_1^+ > W_2^-$.




Figure 3-11: Counterexample for the case $\rho > 1$ for r_g and $\rho < 1$ for r_m .

If r_m is not a constant function, let $v = \min \{i \mid r_m(i) \neq r_m(i + 1)\}$. Denote $\rho = \frac{r_m(v+1)}{r_m(v)}$.

Case 1. $\rho < 1$ We use the same layout as in Figure 3-11 with all resources being mines and $v + 1$ players. We choose $N > \frac{\rho}{1-\rho}$. A PNE for $b = 2$ is all players chooses the bottom-right $N + 1$ mines and avoiding all other mines, which has welfare $W_2 = (v + 1) \cdot (N + 1)r_m(v + 1)$. A PNE for $b = 1$ is all players choose $y = 1$, which has welfare $W_1 = (v + 1) \cdot Nr_m(v + 1)$. Clearly $W_1 > W_2$.

Case 2. $\rho > 1$ We use the same layout as in Figure 3-10 with all resources being mines and $v + 1$ players. Block B is constructed in a similar way with $D(t) = \frac{1}{\rho}N_0(t) - N_1(t)$, and the algorithm is:

1. while $N_1(t) < \frac{3}{1-1/\rho}$, do:
 - If $D(t) \leq 1$, put $y_{t+1} = 0$; else, put $y_{t+1} = 1$;
 - $t \leftarrow t + 1$;
2. while $D(t) \leq 1 + \frac{1}{\rho}$, do: $y_{t+1} = 0, t \leftarrow t + 1$.

The properties following the construction becomes:

- For all $t = 1 \dots N$, $D(t) > 0$ (all prefixes are better)
- For all $t = 0 \dots N - 1$, $D(N) - D(t) > 0$ (all suffixes are better)
- $D(N) < 3$
- $N_1(N) > \frac{3\rho}{\rho-1}$

For $b = 2$, a PNE is all players choose $y = 1$ in B and $y = 0$ in B' , whose social welfare $W_2 = (v + 1) \cdot 2N_1(N)r_m(v + 1)$. For $b = 1$, the PNEs are a players choose $y = 0$ and $v + 1 - a$ players choose $y = 1$, where $1 \leq a \leq v$. All PNEs have the same social welfare $W_1 = (v + 1) \cdot (N_0(N) + N_1(N))r_m(v)$. Then

$$\frac{W_1 - W_2}{-(v + 1)r_m(v)} = 2\rho N_1(N) - N_0(N) - N_1(N) = (\rho - 1)N_1(N) - \rho D(N) > 3\rho - 3\rho = 0.$$

Therefore $W_1 > W_2$, which implies $W_1^+ > W_2^-$.

□

Theorem 3.4.11. *Define welfare function for gold as $w_g(x) \stackrel{\text{def}}{=} x \cdot r_g(x)$. GMG is universally max-capability-preferred if and only if w_g attains its maximum at $x = n$ (i.e., $\max_{x \leq n} w_g(x) = w_g(n)$), where n is the number of players.*

Proof. We first notice that there is only one PNE when $b = \bar{b}$, which is all players cover all gold and no mines. This is because r_g is a positive function and r_m is a negative function, and since all x_e 's are distinct, each player can cover an arbitrary subset of the resources when $b = \bar{b}$. So $W_{\bar{b}} = nM_g r_g(n)$ where M_g is the number of gold in the game.

If $\max_{x \leq n} w_g(x) = w_g(n)$, we show that $W_{\bar{b}}$ is actually the maximum social welfare over all possible strategy profiles of the game. For any strategy profile \mathbf{s} of the game, the social welfare

$$W(\mathbf{s}) = \sum_{i \in \mathcal{N}} \sum_{e \in E_i} r_e(x_e) = \sum_{e \in \mathcal{E}} x_e \cdot r_e(x_e) \leq \sum_{e \in \mathcal{E}_g} n \cdot r_g(n) = nM_g r_g(n) = W_{\bar{b}}.$$

Therefore, the game is max-capability-preferred.

If $\max_{x \leq n} w_g(x) > w_g(n)$, denote $n' = \arg \max_{x \leq n} w_g(x)$, $n' < n$. We construct a game with the corresponding $r_g(\cdot)$ that is not max-capability-preferred. The game has n players, $K = n$ lines, and each line has one gold. The only PNE in $\text{Equil}(\bar{b})$ is all players cover all gold, which achieves a social welfare of $W_{\bar{b}} = n \cdot w_g(n)$. When $b = n'$, one PNE is each player covers n' gold, with each gold covered by exactly n' players. This can be achieved by letting player i cover the gold on lines $\{y = (j \bmod n)\}_{j=i}^{i+n'-1}$. To see why this is a PNE, notice that any player's alternative strategy only allows them to switch to gold with load larger than n' . For all $x > n'$, since $w_g(n') \geq w_g(x)$, $r_g(n') > r_g(x)$. So such change of strategy can only decrease the payoff of the player. The above PNE achieves a social welfare $W_{n'} = n \cdot w_g(n') > W_{\bar{b}}$, so the game is not max-capability-preferred. □

Theorem 3.4.12. *For any payoff functions $r_g(\cdot)$ and $r_m(\cdot)$, there exists an instance of GMG where min-capability-preferred does not hold.*

Proof. It is trivial to construct such a game with mines. Here we show that for arbitrary $r_g(\cdot)$, we can actually construct a game with only gold that is not min-capability-preferred.

Let $r_{\min} = \min_{x \leq n} r_g(x)$ and $r_{\max} = \max_{x \leq n} r_g(x)$. We construct a game with $K = 2$ lines and $N + 1$ gold where $N > \frac{r_{\max}}{r_{\min}}$. In the order of increasing x , the first N gold is on $y = 0$ and the final gold is on $y = 1$. When $b = 1$, for an arbitrary player, denoting the payoff of choosing $y = 0$ (resp. $y = 1$) as r_0 (resp. r_1). Then $r_0 = \sum_{e \in \mathcal{E}_0} r_g(x_e) \geq \sum_{e \in \mathcal{E}_0} r_{\min} = Nr_{\min} > r_{\max} \geq r_1$, where \mathcal{E}_0 is the set of resources on $y = 0$. So all the players will choose $y = 0$ in the PNE. The social welfare is $W_1 = nNr_g(n)$. When $b = 2$, all the players will choose to cover all the gold in the PNE. So the social welfare is $W_2 = n(N + 1)r_g(n) > W_1$. Therefore, the game is not min-capability-preferred. □

Corollary 3.4.13. *For any payoff functions $r_g(\cdot)$ and $r_m(\cdot)$, there exists an instance of GMG where capability-negative does not hold.*

Proof. The same construction used in the proof of Theorem 3.4.12 can be used here, since any game that is not min-capability-preferred is also not capability-negative. □

3.4.4 Complete Characterization for the Alternating Ordering Game

In this section, we present a special form of GMG called the alternating ordering game. We derive exact expressions of the social welfare at equilibrium with respect to the player capability. The analysis provides insights on the factors that affect the trend of social welfare at equilibrium over player capability.

Definition 3.4.14. *The alternating ordering game is a special form of the GMG, with $n = 2$ players and $K = 2$ lines. The layout of the resources follows an alternating ordering of gold and mines as shown in Figure 3-12. Each line has M mines and $M + 1$ gold. The payoff functions satisfy $0 < r_g(2) < \frac{r_g(1)}{2}$ (reflecting competition*

when both players occupy the same gold) and $r_m(1) = r_m(2) < 0$. WLOG, we consider normalized payoff where $r_g(1) = 1, r_g(2) = \rho, 0 < \rho < \frac{1}{2}, r_m(1) = r_m(2) = \mu < 0$.




Figure 3-12: Resource layout for the alternating ordering game. Each dot (resp. cross) is a gold (resp. mine). The dashed lines represent a PNE when $b = 2$ (with $-2 + \rho < \mu < -\rho$).

Let's consider the cases $b = 1$ and $b = 2$ to build some intuitive understanding. When $b = 1$, the PNE is that each player covers one line, which has social welfare $W_1 = 2M + 2M\mu + 2$. When $b = 2$ (and $-2 + \rho < \mu < -\rho$), one PNE is shown in Figure 3-12, where the players avoid one mine but cover one gold together, which has social welfare $W_2 = W_1 - 1 - \mu + 2\rho$. Whether the social welfare at $b = 2$ is better depends on the sign of $2\rho - \mu - 1$. In fact, we have the following general result:

Theorem 3.4.15. *If $-2 + \rho < \mu < -\rho$, then for any level b strategy space \mathcal{L}_b , all PNEs have the same social welfare*

$$W_{\text{Equil}}(b) = \begin{cases} (2M + 1)(1 + \mu) + 2(1 - \rho) + (2\rho - \mu - 1)b & \text{if } b \leq 2M + 1 \\ (4M + 4)\rho & \text{if } b \geq 2M + 2 \end{cases}.$$

The full proof is lengthy and involves analyses of many different cases. We present the main idea here. The full proof is deferred to Section 3.4.5 for readers who are interested.

Proof idea. We make three arguments for this proof: (i) Any function in a PNE must satisfy some specific form indicating where it can switch lines; (ii) Any PNE under \mathcal{L}_b must consist of only functions that use exactly b segments; and (iii) For any function with b segments that satisfies the specific form, the optimal strategy for the other player always achieves the same payoff. \square

Remark $-2 + \rho < \mu < -\rho$ is in fact a necessary and sufficient condition for all PNEs having the same social welfare for any b and M :

Theorem 3.4.16. *For the alternating ordering game, if*

- *for all $b \in \mathbb{Z}^+$ and $M \in \mathbb{Z}^+$, all PNEs in the game with M mines in each line and strategy space \mathcal{L}_b have the same social welfare,*

then $-2 + \rho < \mu < -\rho$.

The proof for Theorem 3.4.16 is deferred to Section 3.4.5 for readers who are interested.

Depending on the sign of $2\rho - \mu - 1$, $W_{\text{Equil}}(b)$ can increase, stay the same, or decrease as b increases until $b = 2M + 1$. $W_{\text{Equil}}(b)$ always decreases at $b = 2M + 2$ and stays the same afterwards. Figure 3-13 visualizes this trend. Figure 3-14 summarizes how the characteristics of the PNEs varies in the ρ - μ space.

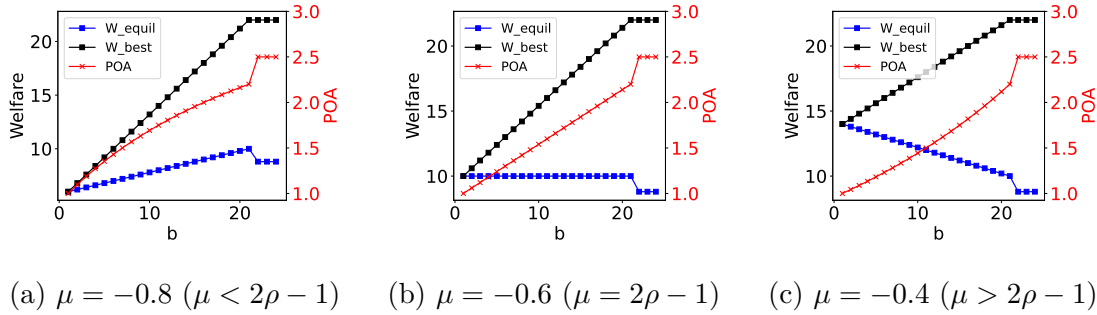


Figure 3-13: W_{Equil} , W_{best} , POA varying with b . $M = 10, \rho = 0.2$.

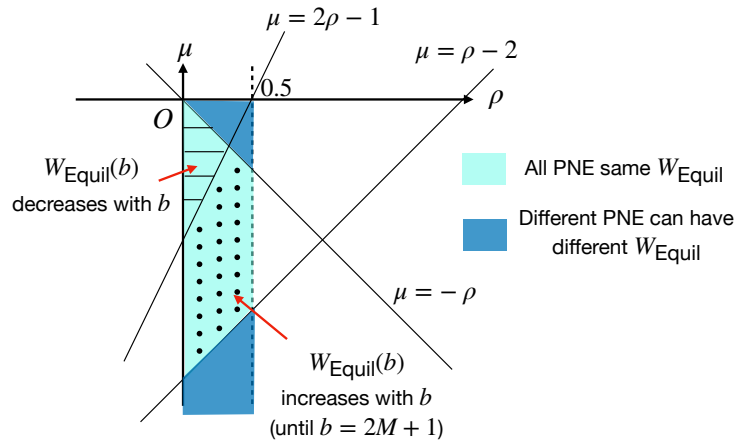


Figure 3-14: Characteristics of PNEs over the ρ - μ landscape.

Price of Anarchy The price of anarchy (POA) [56] is the ratio between the best social welfare achieved by any centralized solution and the worst welfare at equilibria: $\text{POA}(b) = \frac{W_{\text{best}}(b)}{W_{\text{Equil}}(b)}$. The following theorem shows the exact expression of the best social welfare achieved by any centralized solution.

Theorem 3.4.17. *For the alternating ordering game, the best social welfare achieved by any centralized solution under \mathcal{L}_b is*

$$W_{\text{best}}(b) = \begin{cases} 2M + 2 + (2M + 1)\mu - \mu b & \text{if } b \leq 2M + 1 \\ 2M + 2 & \text{if } b \geq 2M + 2 \end{cases}.$$

Proof. First, we notice that a function with c changes from $y = 0$ to $y = 1$ covers at least $M - c$ mines.

For $b = 2c, c \leq M$, we can construct two functions (\hat{f}_1, \hat{f}_2) achieving $W_{\text{best}}(b)$. \hat{f}_1 starts from $y = 1$ and ends at $y = 0$, covers all gold on $y = 1$ except the rightmost one, and the rightmost gold on $y = 0$; \hat{f}_2 starts from $y = 0$ and ends at $y = 1$ covers all gold on $y = 0$ except the rightmost one, and the rightmost gold on $y = 1$. Both use the line changes to avoid mines, with \hat{f}_1 covering $M - c + 1$ mines and \hat{f}_2 covering $M - c$ mines. No other functions can achieve a better welfare, since they cannot jointly cover more gold, and $w_g = xr_g(x)$ attains its maximum at $x = 1$. And to jointly cover fewer mines, both functions need to start from $y = 0$ and ends at $y = 1$, which can reduce the number of mines covered by at most 1. But then the two functions can only jointly cover at most $2M$ gold, which makes the social welfare lower than that of (\hat{f}_1, \hat{f}_2) .

For $b = 2c + 1, c \leq M$, we can construct two functions (\hat{f}_1, \hat{f}_2) achieving $W_{\text{best}}(b)$. They jointly cover all gold with no overlap, and $M - c$ mines each. The construction is simply let \hat{f}_1 covers all gold on $y = 1$, \hat{f}_2 covers all gold on $y = 0$. No other functions can achieve a better welfare, since they cannot jointly cover more gold or less mines.

For $b \geq 2M + 2$, let \hat{f}_1 covers all gold on $y = 1$ and no mines, and \hat{f}_2 covers all gold on $y = 0$ and no mines. This achieves $W_{\text{best}}(b)$, which is in fact the maximum possible social welfare of this game.

□

Combining Theorems 3.4.15 and 3.4.17, we obtain the following result regarding the price of anarchy:

Corollary 3.4.18. *For the alternating ordering game, if $-2 + \rho < \mu < -\rho$, the price of anarchy under \mathcal{L}_b is*

$$\text{POA}(b) = \begin{cases} 1 + \frac{(1-2\rho)(b-1)}{2M+2+2M\mu+(2\rho-\mu-1)(b-1)} & \text{if } b \leq 2M + 1 \\ \frac{1}{2\rho} & \text{if } b \geq 2M + 2 \end{cases},$$

Remark POA(b) increases with b up to $b = 2M + 2$, then stays the same. Figure 3-13 presents the relationships.

Interpretation The results in this part surface an interesting and somewhat counter-intuitive phenomenon that in some situations, increasing player capabilities may deliver a worse overall outcome. There are two opposing factors that affect whether increased capability is beneficial for social welfare or not. With increased capability, players can improve their payoff in a non-competitive way (e.g. avoiding mines), which is always beneficial for social welfare; they can also improve payoff in a competitive way (e.g. occupying gold together), which may reduce social welfare. The joint effect of the two factors determines the effect of increasing capability.

3.4.5 Full Proofs for the Alternating Ordering Game

Here we present the full proofs of Theorem 3.4.15 and Theorem 3.4.16. Readers who are not interested in the detailed proofs can skip this section.

Preliminaries

We first establish some notational convenience for the subsequent analyses:

- Following the increasing order of x , we use (x_t, y_t, α_t) to denote the location and type of the t -th resource.

- Each player's strategy can be represented as a function over the integer domain $[0, 4M + 1]$, specifying the function value at each x_t for $t \in [0, 4M + 1]$. We represent a function compactly as the set of intervals $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ over t with $f(t) = 1$. For example, $f = \{[0, 2], [5, 5]\}$ represents the function of $f(t) = 1$ if $t \in [0, 2]$ or $t = 5$, and $f(t) = 0$ otherwise.

- We use the **canonical** representation of f throughout this section, where

- a_ξ^-, a_ξ^+ are integers for all ξ , and $a_0^- \geq 0, a_{c-1}^+ \leq 4M + 1$;
- $a_\xi^- \leq a_\xi^+, a_{\xi+1}^- - a_\xi^+ \geq 2$ for all ξ , i.e., the representation uses the least number of intervals.

- Denote the set of functions with exactly k segments as \mathcal{F}_k . Then $\mathcal{L}_b = \bigcup_{k \leq b} \mathcal{F}_k$.

The following general form covers all functions within \mathcal{F}_k :

- For $k = 2c + 1$ (odd k), $\mathcal{F}_k = \left\{ \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c \mid a_0^- = 0, a_c^+ = 4M + 1 \right\} \cup \left\{ \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1} \mid a_0^- > 0, a_{c-1}^+ < 4M + 1 \right\}$
- For $k = 2c$ (even k), $\mathcal{F}_k = \left\{ \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1} \mid a_0^- = 0, a_{c-1}^+ < 4M + 1 \right\} \cup \left\{ \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1} \mid a_0^- > 0, a_{c-1}^+ = 4M + 1 \right\}$

Proof of Theorem 3.4.15

Lemma 3.4.19. *Given an arbitrary set of strategies used by the other players \mathbf{f}_{-i} and any $k_0 \geq 1$, denote $f_i^* = \arg \max_{f_i \in \mathcal{L}_{k_0}} u_i(f_i, \mathbf{f}_{-i})$ as the optimal strategy for player i , then $f_i^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ must satisfy the following condition:*

$$\forall \xi, a_\xi^- = 0 \text{ or } a_\xi^- = 4j_\xi^- + 3 \text{ for some } j_\xi^- \in \mathbb{Z}, a_\xi^+ = 4M + 1 \text{ or } a_\xi^+ = 4j_\xi^+ \text{ for some } j_\xi^+ \in \mathbb{Z}.$$

Proof. This lemma essentially states that the segments of an optimal strategy can only start and end at particular locations in the sequence. We prove this lemma by showing that if $f_i^* \in \mathcal{F}_k$ does not satisfy the above condition, then there exists f_i' which uses $k' \leq k$ segments and achieves a payoff $u_i' = u_i(f_i', \mathbf{f}_{-i})$ higher than $u_i^* = u_i(f_i^*, \mathbf{f}_{-i})$, therefore contradicting the fact that f_i^* is optimal.

If f_i^* does not satisfy the given condition, then either there exists a ξ_0 where $a_{\xi_0}^- \neq 0$ and $a_{\xi_0}^- \neq 4j + 3$ for all j , or there exists a ξ_0 where $a_{\xi_0}^+ \neq 4M + 1$ and $a_{\xi_0}^+ \neq 4j$ for all j . We consider each case separately here.

1. There exists a ξ_0 where $a_{\xi_0}^- \neq 0$ and $a_{\xi_0}^- \neq 4j + 3$ for all j . Consider the value of $a_{\xi_0}^-$:

- $a_{\xi_0}^- = 4j$ for some $j \neq 0$. Since f_i^* is in canonical form, $f_i^*(4j - 1) = 0$. We know $y_{4j-1} = 0, \alpha_{4j-1} = \text{mine}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^- = 4j - 1$, then the payoff achieved by f'_i is $u'_i = u_i^* - \mu > u_i^*$. And the number of segments of f'_i $k' \leq k$.
- $a_{\xi_0}^- = 4j + 1$ for some j . We know $y_{4j+1} = 0, \alpha_{4j+1} = \text{gold}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^- = 4j + 2$ (if this makes interval ξ_0 empty, then remove interval ξ_0), then the payoff achieved by f'_i is $u'_i = u_i^* + r_g(x'_{4j+1}) > u_i^*$, and $k' \leq k$.
- $a_{\xi_0}^- = 4j + 2$ for some j . We know $y_{4j+2} = 1, \alpha_{4j+2} = \text{mine}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^- = 4j + 3$ (if this makes interval ξ_0 empty, then remove interval ξ_0), then the payoff achieved by f'_i is $u'_i = u_i^* - \mu > u_i^*$, and $k' \leq k$.

2. There exists a ξ_0 where $a_{\xi_0}^+ \neq 4M + 1$ and $a_{\xi_0}^+ \neq 4j$ for all j . Consider the value of $a_{\xi_0}^+$:

- $a_{\xi_0}^+ = 4j + 1$ for some $j < M$. We know $y_{4j+1} = 0, \alpha_{4j+1} = \text{gold}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^+ = 4j$, then the payoff achieved by f'_i is $u'_i = u_i^* + r_g(x'_{4j+1}) > u_i^*$, and $k' \leq k$.
- $a_{\xi_0}^+ = 4j + 2$ for some j . We know $y_{4j+2} = 1, \alpha_{4j+2} = \text{mine}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^+ = 4j + 1$, then the payoff achieved by f'_i is $u'_i = u_i^* - \mu > u_i^*$, and $k' \leq k$.
- $a_{\xi_0}^+ = 4j + 3$ for some j . We know $y_{4j+4} = 1, \alpha_{4j+4} = \text{gold}$. Let f'_i be identical to f_i^* except changing $a_{\xi_0}^+ = 4j + 4$, then the payoff achieved by f'_i is $u'_i = u_i^* + r_g(x'_{4j+4}) > u_i^*$, and $k' \leq k$.

All the cases imply that f_i^* cannot be optimal within \mathcal{L}_{k_0} , which is a contradiction. Therefore, f_i^* must satisfy the given condition. \square

Lemma 3.4.20. *Following Lemma 3.4.19, if the optimal strategy f_i^* belongs to \mathcal{F}_k where $k \leq 2M + 1$, then it must satisfy the following properties:*

- *If $k = 2c + 1$ (odd k), then f_i^* satisfies either condition S_1 or condition S_2 , and f_i^* always covers $M + c + 1$ gold and $M - c$ mines.*
 - S_1 : $f_i^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0, a_c^+ = 4M + 1$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^c, \{j_\xi^- < j_\xi^+\}_{\xi=1}^{c-1}, \{j_\xi^+ \leq j_{\xi+1}^-\}_{\xi=0}^{c-1}, j_c^- < M, j_0^+ \geq 0$ for some $\{j_\xi^-\}_{\xi=1}^c$ and $\{j_\xi^+\}_{\xi=0}^{c-1}$.
 - S_2 : $f_i^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=0}^{c-1}, \{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{j_\xi^- < j_\xi^+\}_{\xi=0}^{c-1}, \{j_\xi^+ \leq j_{\xi+1}^-\}_{\xi=0}^{c-2}, j_{c-1}^+ \leq M, j_0^- \geq 0$ for some $\{j_\xi^-\}_{\xi=0}^{c-1}$ and $\{j_\xi^+\}_{\xi=0}^{c-1}$.
- *If $k = 2c$ (even k), then either f_i^* satisfies condition S_3 , in which case f_i^* always covers $M + c + 1$ gold and $M - c + 1$ mines, or f_i^* satisfies condition S_4 , in which case f_i^* always covers $M + c$ gold and $M - c$ mines.*
 - S_3 : $f_i^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^{c-1}, \{j_\xi^- < j_\xi^+\}_{\xi=1}^{c-1}, \{j_\xi^+ \leq j_{\xi+1}^-\}_{\xi=0}^{c-2}, j_{c-1}^+ \leq M, j_0^+ \geq 0$ for some $\{j_\xi^-\}_{\xi=1}^{c-1}$ and $\{j_\xi^+\}_{\xi=0}^{c-1}$.
 - S_4 : $f_i^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_{c-1}^+ = 4M + 1$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-2}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=0}^{c-1}, \{j_\xi^- < j_\xi^+\}_{\xi=0}^{c-2}, \{j_\xi^+ \leq j_{\xi+1}^-\}_{\xi=0}^{c-2}, j_0^- \geq 0, j_{c-1}^- < M$ for some $\{j_\xi^-\}_{\xi=0}^{c-1}$ and $\{j_\xi^+\}_{\xi=0}^{c-2}$.

Proof. We prove for each case.

1. $k = 2c + 1$ (odd k)

Here, f_i^* is either of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0$ and $a_c^+ = 4M + 1$, or of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- > 0$ and $a_{c-1}^+ < 4M + 1$.

1.1 If f_i^* is of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0$ and $a_c^+ = 4M + 1$, then according to Lemma 3.4.19, we have $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}$ and $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^c$ for some $\{j_\xi^-\}_{\xi=1}^c$ and $\{j_\xi^+\}_{\xi=0}^{c-1}$. And since f_i^* is canonical, we have $\{j_\xi^- < j_\xi^+\}_{\xi=1}^{c-1}$, $\{j_\xi^+ \leq j_{\xi+1}^-\}_{\xi=0}^{c-1}$, $j_c^- < M$, $j_0^+ \geq 0$. Therefore, in this case, f_i^* satisfies condition S_1 .

Now consider the number of gold and mines covered by f_i^* that satisfies condition S_1 . First consider $c \geq 1$. Segment $[a_\xi^-, a_\xi^+]$ (where f_i^* has value 1) covers $j_\xi^+ - j_\xi^-$ gold and $j_\xi^+ - j_\xi^- - 1$ mines, for $\xi = 1, \dots, c-1$. Segment $[a_{\xi-1}^+ + 1, a_\xi^- - 1]$ (where f_i^* has value 0) covers $j_\xi^- - j_{\xi-1}^+ + 1$ gold and $j_\xi^- - j_{\xi-1}^+$ mines, for $\xi = 1, \dots, c$. Segment $[a_0^-, a_0^+]$ covers $j_0^+ + 1$ gold and j_0^+ mines, and segment $[a_c^-, a_c^+]$ covers $M - j_c^-$ gold and $M - j_c^- - 1$ mines. Therefore, the number of gold covered by f_i^* is

$$j_0^+ + 1 + \sum_{\xi=1}^{c-1} (j_\xi^+ - j_\xi^-) + \sum_{\xi=1}^c (j_\xi^- - j_{\xi-1}^+ + 1) + M - j_c^- = M + c + 1,$$

and the number of mines covered by f_i^* is

$$j_0^+ + \sum_{\xi=1}^{c-1} (j_\xi^+ - j_\xi^- - 1) + \sum_{\xi=1}^c (j_\xi^- - j_{\xi-1}^+) + M - j_c^- - 1 = M - c.$$

It's straightforward to check that the above expressions also hold for $c = 0$.

1.2 If f_i^* is of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- > 0$ and $a_{c-1}^+ < 4M + 1$, then similar to case 1.1, Lemma 3.4.19 and the fact that f_i^* is canonical imply that f_i^* satisfies condition S_2 .

Consider the number of gold and mines covered by f_i^* that satisfies condition S_2 . Segment $[a_\xi^-, a_\xi^+]$ (where f_i^* has value 1) covers $j_\xi^+ - j_\xi^-$ gold and $j_\xi^+ - j_\xi^- - 1$ mines, for $\xi = 0, \dots, c-1$. Segment $[a_{\xi-1}^+ + 1, a_\xi^- - 1]$ (where f_i^* has value 0) covers $j_\xi^- - j_{\xi-1}^+ + 1$ gold and $j_\xi^- - j_{\xi-1}^+$ mines, for $\xi = 1, \dots, c-1$. Segment $[0, a_0^- - 1]$ (where f_i^* has value 0) covers $j_0^- + 1$ gold and j_0^- mines, and segment $[a_{c-1}^+ + 1, 4M + 1]$ (where f_i^* has value 0) covers $M - j_{c-1}^+ + 1$ gold and $M - j_{c-1}^+$ mines. Therefore, the number of

gold covered by f_i^* is

$$j_0^- + 1 + \sum_{\xi=0}^{c-1} (j_\xi^+ - j_\xi^-) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+ + 1) + M - j_{c-1}^+ + 1 = M + c + 1,$$

and the number of mines covered by f_i^* is

$$j_0^- + \sum_{\xi=0}^{c-1} (j_\xi^+ - j_\xi^- - 1) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+) + M - j_{c-1}^+ = M - c.$$

2. $k = 2c$ (even k)

Here, f_i^* is either of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$ and $a_{c-1}^+ < 4M + 1$, or of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- > 0$ and $a_{c-1}^+ = 4M + 1$.

2.1 If f_i^* is of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$ and $a_{c-1}^+ < 4M + 1$, then similar to case 1.1, Lemma 3.4.19 and the fact that f_i^* is canonical imply that f_i^* satisfies condition S_3 .

Consider the number of gold and mines covered by f_i^* that satisfies condition S_3 . Segment $[a_\xi^-, a_\xi^+]$ (where f_i^* has value 1) covers $j_\xi^+ - j_\xi^-$ gold and $j_\xi^+ - j_\xi^- - 1$ mines, for $\xi = 1, \dots, c-1$. Segment $[a_{\xi-1}^+ + 1, a_\xi^- - 1]$ (where f_i^* has value 0) covers $j_\xi^- - j_{\xi-1}^+ + 1$ gold and $j_\xi^- - j_{\xi-1}^+$ mines, for $\xi = 1, \dots, c-1$. Segment $[a_0^-, a_0^+]$ (where f_i^* has value 1) covers $j_0^+ + 1$ gold and j_0^+ mines, and segment $[a_{c-1}^+ + 1, 4M + 1]$ (where f_i^* has value 0) covers $M - j_{c-1}^+ + 1$ gold and $M - j_{c-1}^+$ mines. Therefore, the number of gold covered by f_i^* is

$$j_0^+ + 1 + \sum_{\xi=1}^{c-1} (j_\xi^+ - j_\xi^-) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+ + 1) + M - j_{c-1}^+ + 1 = M + c + 1,$$

and the number of mines covered by f_i^* is

$$j_0^+ + \sum_{\xi=1}^{c-1} (j_\xi^+ - j_\xi^- - 1) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+) + M - j_{c-1}^+ = M - c + 1.$$

2.2 If f_i^* is of form $\{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- > 0$ and $a_{c-1}^+ = 4M + 1$, then similar to case 1.1, Lemma 3.4.19 and the fact that f_i^* is canonical imply that f_i^* satisfies condition S_4 .

Consider the number of gold and mines covered by f_i^* that satisfies condition S_4 . Segment $[a_\xi^-, a_\xi^+]$ (where f_i^* has value 1) covers $j_\xi^+ - j_\xi^-$ gold and $j_\xi^+ - j_\xi^- - 1$ mines, for $\xi = 0, \dots, c-2$. Segment $[a_{\xi-1}^+ + 1, a_\xi^- - 1]$ (where f_i^* has value 0) covers $j_\xi^- - j_{\xi-1}^+ + 1$ gold and $j_\xi^- - j_{\xi-1}^+$ mines, for $\xi = 1, \dots, c-1$. Segment $[0, a_0^- - 1]$ (where f_i^* has value 0) covers $j_0^- + 1$ gold and j_0^- mines, and segment $[a_{c-1}^-, a_{c-1}^+]$ (where f_i^* has value 1) covers $M - j_{c-1}^-$ gold and $M - j_{c-1}^- - 1$ mines. Therefore, the number of gold covered by f_i^* is

$$j_0^- + 1 + \sum_{\xi=0}^{c-2} (j_\xi^+ - j_\xi^-) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+ + 1) + M - j_{c-1}^- = M + c,$$

and the number of mines covered by f_i^* is

$$j_0^- + \sum_{\xi=0}^{c-2} (j_\xi^+ - j_\xi^- - 1) + \sum_{\xi=1}^{c-1} (j_\xi^- - j_{\xi-1}^+) + M - j_{c-1}^- - 1 = M - c.$$

□

Lemma 3.4.21. *If one player (call it player A) covers g_A gold, player B covers g_B gold and m_B mines, and $g_A + g_B \geq 2M + 2$, then there is an upper bound on player B's payoff:*

$$u_B \leq (1 - \rho)(2M + 2 - g_A) + \rho g_B + \mu m_B$$

Proof. Among the gold covered by player B, denote the number of them also covered by player A as d . Since the total number of gold is $2M + 2$, we have $g_A + g_B - d \leq 2M + 2$, i.e., $d \geq g_A + g_B - 2M - 2$. Therefore,

$$\begin{aligned} u_B &= g_B - d + d\rho + m_B\mu \\ &\leq g_B - (1 - \rho)(g_A + g_B - 2M - 2) + m_B\mu \\ &= (1 - \rho)(2M + 2 - g_A) + \rho g_B + \mu m_B. \end{aligned}$$

□

We define that a strategy achieves *best coverage* if it covers all the gold that is not covered by the other player.

Lemma 3.4.22. *Given one player (call it player A) covers g_A gold, if a strategy f_B for player B covers g_B gold and m_B mines and achieves best coverage, then any strategy f'_B that covers g'_B gold and m'_B mines will achieve a lower payoff than f_B , if*

$$g'_B \leq g_B \wedge m'_B > m_B, \text{ or } g'_B < g_B \wedge m'_B \geq m_B$$

Proof. Since f_B achieves best coverage, it covers $2M + 2 - g_A$ gold that is not covered by player A, and $g_B + g_A - 2M - 2$ gold that is covered by A. So the payoff achieved by f_B is

$$\begin{aligned} u_B &= 2M + 2 - g_A + \rho(g_B + g_A - 2M - 2) + m_B\mu \\ &= (1 - \rho)(2M + 2 - g_A) + \rho g_B + \mu m_B. \end{aligned}$$

Consider the payoff of f'_B . By Lemma 3.4.21,

$$u'_B \leq (1 - \rho)(2M + 2 - g_A) + \rho g'_B + \mu m'_B.$$

Since $\rho > 0$ and $\mu < 0$, we can see that if $g'_B \leq g_B \wedge m'_B > m_B$, or $g'_B < g_B \wedge m'_B \geq m_B$,

$$u'_B < (1 - \rho)(2M + 2 - g_A) + \rho g_B + \mu m_B = u_B.$$

□

Lemma 3.4.23. *If both players' strategy space is \mathcal{L}_b ($b \leq 2M + 2$), then for all PNE (f_1^*, f_2^*) , $f_1^*, f_2^* \in \mathcal{F}_b$, i.e. both strategies in the equilibria must use exactly b segments.*

Proof. We prove by induction.

Base case For $b = 1$, there is only two possible strategies in \mathcal{L}_1 : $f^0 = \{\}$ and $f^1 = \{[0, 4M + 1]\}$. Both uses exactly 1 segment. So the statement holds.

Induction step Consider the case $b = k$. First we show that if one of the strategies in a PNE uses exactly k segments, then the other strategy must also use exactly k segments. Without loss of generality, let $f_1^* \in \mathcal{F}_k$.

1. $k = 2c + 1$ (odd k)

By Lemma 3.4.20, f_1^* must satisfy condition S_1 or S_2 .

1.1 If f_1^* satisfies condition S_1 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0, a_c^+ = 4M + 1, \{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^c$, and f_1^* covers $M + c + 1$ gold and $M - c$ mines. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by setting $\{\hat{a}_\xi^- = 4j_\xi^+ + 3, \hat{a}_\xi^+ = 4 \cdot \max(j_{\xi+1}^-, j_\xi^+ + 1)\}_{\xi=0}^{c-1}$ (note that here j_ξ^+ and j_ξ^- are the values used by f_1^*). It's easy to check that \hat{f}_2 satisfies condition S_2 and covers $g_2 = M + c + 1$ gold and $m_2 = M - c$ mines. In particular, among the gold covered by \hat{f}_2 , $2c$ of them are also covered by f_1^* , and $M - c + 1$ of them are covered by \hat{f}_2 only. Therefore, \hat{f}_2 achieves best coverage. \hat{f}_2 achieves a payoff of

$$\hat{u}_2 = M - c + 1 + 2c\rho + (M - c)\mu.$$

We show here that any $f'_2 \in \mathcal{F}_{k'}$ where $k' < k$ will achieve a payoff $u'_2 < \hat{u}_2$, therefore f_2^* must use exactly k segments. If $k' = 2c' + 1$, then $c' \leq c - 1$, and by Lemma 3.4.20, f'_2 covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c'$ mines. We have $g'_2 < g_2$ and $m'_2 > m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

If $k' = 2c'$, then $c' \leq c$. By Lemma 3.4.20, f'_2 either covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c' + 1$ mines, in which case $g'_2 \leq g_2$ and $m'_2 > m_2$, or $g'_2 = M + c'$ gold and $m'_2 = M - c'$ mines, in which case $g'_2 < g_2$ and $m'_2 \geq m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

1.2 By symmetry, the above proof also applies to the case where f_1^* satisfies condition S_2 (symmetry with respect to inverting the direction of x and y axis).

2. $k = 2c$ (**even** k), $c \leq M$

By Lemma 3.4.20, f_1^* must satisfy condition S_3 or S_4 .

2.1 If f_1^* satisfies condition S_3 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}$, $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^{c-1}$, and f_1^* covers $M+c+1$ gold and $M-c+1$ mines. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by sequentially setting $\hat{a}_0^-, \hat{a}_0^+, \hat{a}_1^-, \hat{a}_1^+, \dots, \hat{a}_{c-1}^-, \hat{a}_{c-1}^+$ with $\hat{a}_0^- = \max(4j_0^+ - 1, 3)$, $\{\hat{a}_\xi^- = \max(4j_\xi^+ - 1, \hat{a}_{\xi-1}^+ + 3)\}_{\xi=1}^{c-1}$, $\{\hat{a}_\xi^+ = \max(4 \cdot j_{\xi+1}^-, \hat{a}_\xi^- + 1)\}_{\xi=0}^{c-2}$, $\hat{a}_{c-1}^+ = 4M + 1$. This \hat{f}_2 satisfies condition S_4 and covers $g_2 = M + c$ gold and $m_2 = M - c$ mines. In particular, among the gold covered by \hat{f}_2 , $2c - 1$ of them are also covered by f_1^* , and $M - c + 1$ of them are covered by \hat{f}_2 only. Therefore, \hat{f}_2 achieves best coverage. \hat{f}_2 achieves a payoff of

$$\hat{u}_2 = M - c + 1 + (2c - 1)\rho + (M - c)\mu.$$

We show here that any $f'_2 \in \mathcal{F}_{k'}$ where $k' < k$ will achieve a payoff $u'_2 < \hat{u}_2$, therefore f_2^* must use exactly k segments. If $k' = 2c' + 1$, then $c' \leq c - 1$, and by Lemma 3.4.20, f'_2 covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c'$ mines. We have $g'_2 \leq g_2$ and $m'_2 > m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

If $k' = 2c'$, then $c' \leq c - 1$. By Lemma 3.4.20, f'_2 either covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c' + 1$ mines, in which case $g'_2 \leq g_2$ and $m'_2 > m_2$, or $g'_2 = M + c'$ gold and $m'_2 = M - c'$ mines, in which case $g'_2 < g_2$ and $m'_2 > m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

2.2 If f_1^* satisfies condition S_4 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_{c-1}^+ = 4M + 1$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-2}$, $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=0}^{c-1}$, and f_1^* covers $M + c$ gold and $M - c$ mines. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by setting $\hat{a}_0^- = 0$, $\hat{a}_0^+ = 4j_0^-$, $\{\hat{a}_\xi^- = 4j_{\xi-1}^+ + 3, \hat{a}_\xi^+ = 4 \cdot \max(j_\xi^-, j_{\xi-1}^+ + 1)\}_{\xi=1}^{c-1}$. This \hat{f}_2 satisfies condition S_3 and covers $g_2 = M + c + 1$ gold and $m_2 = M - c + 1$ mines. In particular, among the gold covered

by \hat{f}_2 , $2c - 1$ of them are also covered by f_1^* , and $M - c + 2$ of them are covered by \hat{f}_2 only. Therefore, \hat{f}_2 achieves best coverage. \hat{f}_2 achieves a payoff of

$$\hat{u}_2 = M - c + 2 + (2c - 1)\rho + (M - c + 1)\mu.$$

We show here that any $f'_2 \in \mathcal{F}_{k'}$ where $k' < k$ will achieve a payoff $u'_2 < \hat{u}_2$, therefore f_2^* must use exactly k segments. If $k' = 2c' + 1$, then $c' \leq c - 1$, and by Lemma 3.4.20, f'_2 covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c'$ mines. We have $g'_2 < g_2$ and $m'_2 \geq m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

If $k' = 2c'$, then $c' \leq c - 1$. By Lemma 3.4.20, f'_2 either covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c' + 1$ mines, in which case $g'_2 < g_2$ and $m'_2 > m_2$, or $g'_2 = M + c'$ gold and $m'_2 = M - c'$ mines, in which case $g'_2 < g_2$ and $m'_2 \geq m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

3. $k = 2c$ (even k), $c = M + 1$

By Lemma 3.4.20, f_1^* must satisfy condition S_3 or S_4 . In fact, in this case, no function satisfies S_4 , and there is only one function satisfies S_3 , which is the function that covers all $2M + 2$ gold and no mine. Construct \hat{f}_2 to be the same as f_1^* , which covers all $g_2 = 2M + 2$ gold and $m_2 = 0$ mine. Since f_1^* already covers all gold, \hat{f}_2 trivially achieves best coverage.

We show here that any $f'_2 \in \mathcal{F}_{k'}$ where $k' < k$ will achieve a payoff $u'_2 < \hat{u}_2$, therefore f_2^* must use exactly k segments. If $k' = 2c' + 1$, then $c' \leq c - 1$, and by Lemma 3.4.20, f'_2 covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c'$ mines. We have $g'_2 < g_2$ and $m'_2 \geq m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

If $k' = 2c'$, then $c' \leq c - 1$. By Lemma 3.4.20, f'_2 either covers $g'_2 = M + c' + 1$ gold and $m'_2 = M - c' + 1$ mines, in which case $g'_2 < g_2$ and $m'_2 > m_2$, or $g'_2 = M + c'$ gold and $m'_2 = M - c'$ mines, in which case $g'_2 < g_2$ and $m'_2 \geq m_2$. Therefore by Lemma 3.4.22, $u'_2 < \hat{u}_2$.

Now we have shown that if one of the strategies in a PNE uses exactly k segments,

then the other strategy must also use exactly k segments. What is left to show is that there is no PNE where both strategies use less than k segments.

We prove by contradiction. Assume there is a PNE (f_1^*, f_2^*) where both f_1^* and f_2^* use less than k segments. By the induction hypothesis, $f_1^*, f_2^* \in \mathcal{F}_{k-1}$.

1. $k - 1 = 2c + 1$ (even k), $c \leq M$

By Lemma 3.4.20, f_1^* must satisfy condition S_1 or S_2 . If f_1^* satisfies condition S_1 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0, a_c^+ = 4M + 1, \{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^c$. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^c$ according to f_1^* by setting $\hat{a}_0^- = 0, \hat{a}_0^+ = 0, \{\hat{a}_\xi^- = 4j_{\xi-1}^+ + 3, \hat{a}_\xi^+ = 4 \cdot \max(j_\xi^-, j_{\xi-1}^+ + 1)\}_{\xi=1}^c$. It is easy to check that \hat{f}_2 uses k segments, covers $g_2 = M + c + 2$ gold and $m_2 = M - c$ mines, and achieves best coverage. By Lemma 3.4.20, f_2^* covers $g_2^* = M + c + 1$ gold and $m_2^* = M - c$ mines. Thus $g_2^* < g_2$ and $m_2^* \geq m_2$. By Lemma 3.4.22, $u_2^* < \hat{u}_2$, therefore (f_1^*, f_2^*) cannot be a PNE, contradiction.

By symmetry, the above proof also applies to the case where f_1^* satisfies condition S_2 (symmetry with respect to inverting the direction of x and y axis).

2. $k - 1 = 2c$ (odd k), $c \leq M$

By Lemma 3.4.20, f_1^* must satisfy condition S_3 or S_4 . If f_1^* satisfies condition S_3 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^{c-1}$. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^c$ according to f_1^* by sequentially setting $\hat{a}_0^-, \hat{a}_0^+, \hat{a}_1^-, \hat{a}_1^+, \dots, \hat{a}_c^-, \hat{a}_c^+$ with $\hat{a}_0^- = 0, \hat{a}_0^+ = 0, \hat{a}_1^- = \max(4j_0^+ - 1, 3), \{\hat{a}_\xi^- = \max(4j_{\xi-1}^+ - 1, \hat{a}_{\xi-1}^+ + 3)\}_{\xi=2}^c, \{\hat{a}_\xi^+ = \max(4 \cdot j_\xi^-, \hat{a}_\xi^- + 1)\}_{\xi=1}^{c-1}, \hat{a}_c^+ = 4M + 1$. It is easy to check that \hat{f}_2 uses k segments, covers $g_2 = M + c + 1$ gold and $m_2 = M - c$ mines, and achieves best coverage. By Lemma 3.4.20, f_2^* either covers $g_2^* = M + c + 1$ gold and $m_2^* = M - c + 1$ mines, in which case $g_2^* \leq g_2$ and $m_2^* > m_2$, or $g_2^* = M + c$ gold and $m_2^* = M - c$ mines, in which case $g_2^* < g_2$ and $m_2^* \geq m_2$. Therefore by Lemma 3.4.22, $u_2^* < \hat{u}_2$, which means (f_1^*, f_2^*) cannot be a PNE, contradiction.

If f_1^* satisfies condition S_4 , let $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_{c-1}^+ = 4M + 1$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-2}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=0}^{c-1}$. f_1^* covers $g_1^* = M + c$ gold and $m_1^* = M - c$

mines. We construct $\hat{f}_1 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^c$ according to f_1^* by setting $\hat{a}_0^- = 0, \hat{a}_0^+ = 0, \{\hat{a}_\xi^- = a_{\xi-1}^-, \hat{a}_\xi^+ = a_{\xi-1}^+\}_{\xi=1}^c$, i.e. \hat{f}_1 is identical to f_1^* except $\hat{f}_1(0) = 1$. \hat{f}_1 uses k segments, and covers exactly the same set of gold and mines as f_1^* plus the gold at $t = 0$. Therefore, \hat{f}_1 's payoff is strictly higher than f_1^* 's payoff. This means (f_1^*, f_2^*) cannot be a PNE, contradiction.

This finishes the proof that there exists no PNE where both strategies use less than k segments. We have also shown that if one strategy in a PNE uses k segments, the other strategy must also use k segments. This together shows that for all PNE, both strategies in the equilibria must use exactly k segments. This finishes the proof by induction. □

Now we are ready to prove Theorem 3.4.15.

Theorem 3.4.15. *If $-2 + \rho < \mu < -\rho$, then for any level b strategy space \mathcal{L}_b , all PNEs have the same social welfare*

$$W_{\text{Equil}}(b) = \begin{cases} (2M+1)(1+\mu) + 2(1-\rho) + (2\rho - \mu - 1)b & \text{if } b \leq 2M+1 \\ (4M+4)\rho & \text{if } b \geq 2M+2 \end{cases}.$$

Proof. We consider different values of b .

1. $b = 2c + 1, 0 \leq c \leq M$

By Lemma 3.4.23, both f_1^* and f_2^* use exactly b segments, i.e. $f_1^*, f_2^* \in \mathcal{F}_b$. By Lemma 3.4.20, both f_1^* and f_2^* must satisfy condition S_1 or S_2 , and each covers $M + c + 1$ gold and $M - c$ mines. If f_1^* satisfies S_1 , denote $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^c$ where $a_0^- = 0, a_c^+ = 4M + 1, \{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}, \{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^c$. Same as in the proof of Lemma 3.4.23, we construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by setting $\{\hat{a}_\xi^- = 4j_\xi^+ + 3, \hat{a}_\xi^+ = 4 \cdot \max(j_{\xi+1}^-, j_\xi^+ + 1)\}_{\xi=0}^{c-1}$. $\hat{f}_2 \in \mathcal{F}_b$ achieves best coverage and a payoff of $\hat{u}_2 = M - c + 1 + 2c\rho + (M - c)\mu$. Since f_2^* always covers $M + c + 1$ gold and $M - c$ mines, by Lemma 3.4.21, f_2^* 's payoff $u_2^* \leq \hat{u}_2$. But by definition of Nash

equilibrium, $u_2^* \geq \hat{u}_2$. Therefore, $u_2^* = \hat{u}_2$, i.e. all f_2^* must achieve the same payoff of $M - c + 1 + 2c\rho + (M - c)\mu$.

By symmetry (with respect to inverting the direction of x and y axis), the above proof can also be applied to show that if f_1^* satisfies S_2 , then all f_2^* must achieve the same payoff of $M - c + 1 + 2c\rho + (M - c)\mu$.

Therefore, in all cases, $u_2^* = M - c + 1 + 2c\rho + (M - c)\mu$. Similarly, $u_1^* = M - c + 1 + 2c\rho + (M - c)\mu$. So

$$\begin{aligned} W_{\text{Equil}}(b) &= 2M(1 + \mu) + 2 + 2(2\rho - \mu - 1)c \\ &= 2M(1 + \mu) + 2 + (2\rho - \mu - 1)(b - 1) \\ &= (2M + 1)(1 + \mu) + 2(1 - \rho) + (2\rho - \mu - 1)b. \end{aligned}$$

2. $b = 2c, 1 \leq c \leq M$

By Lemma 3.4.23, $f_1^*, f_2^* \in \mathcal{F}_b$. By Lemma 3.4.20, both f_1^* and f_2^* must satisfy condition S_3 or S_4 . If f_1^* satisfies S_3 , denote $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_0^- = 0$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-1}$, $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=1}^{c-1}$, and f_1^* covers $M + c + 1$ gold and $M - c + 1$ mines. Same as in the proof of Lemma 3.4.23, we construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by sequentially setting $\hat{a}_0^-, \hat{a}_0^+, \hat{a}_1^-, \hat{a}_1^+, \dots, \hat{a}_{c-1}^-, \hat{a}_{c-1}^+$ with $\hat{a}_0^- = \max(4j_0^+ - 1, 3)$, $\{\hat{a}_\xi^- = \max(4j_\xi^+ - 1, \hat{a}_{\xi-1}^+ + 3)\}_{\xi=1}^{c-1}$, $\{\hat{a}_\xi^+ = \max(4 \cdot j_{\xi+1}^-, \hat{a}_\xi^- + 1)\}_{\xi=0}^{c-2}$, $\hat{a}_{c-1}^+ = 4M + 1$. This \hat{f}_2 satisfies condition S_4 and achieves a payoff of $\hat{u}_2 = M - c + 1 + (2c - 1)\rho + (M - c)\mu$. For any f_2' that satisfies S_3 , by Lemma 3.4.20, it covers $M + c + 1$ gold and $M - c + 1$ mines. By Lemma 3.4.21, such f_2' 's payoff

$$\begin{aligned} u_2' &\leq (1 - \rho)(M - c + 1) + \rho(M + c + 1) + \mu(M - c + 1) \\ &= M - c + 1 + 2c\rho + (M - c + 1)\mu \\ &= \hat{u}_2 + \rho + \mu < \hat{u}_2. \end{aligned}$$

By the definition of Nash equilibrium, $u_2^* \geq \hat{u}_2$, so f_2^* cannot satisfy S_3 . Therefore, f_2^* must satisfy S_4 , and by Lemma 3.4.20, f_2^* covers $M + c$ gold and $M - c$ mines. So by Lemma 3.4.21, $u_2^* \leq \hat{u}_2$. Therefore, $u_2^* = \hat{u}_2$, i.e. f_2^* always satisfies S_4 and achieves a

payoff of $u_2^* = M - c + 1 + (2c - 1)\rho + (M - c)\mu$.

If f_1^* satisfies S_4 , denote $f_1^* = \{[a_\xi^-, a_\xi^+]\}_{\xi=0}^{c-1}$ where $a_{c-1}^+ = 4M + 1$, and $\{a_\xi^+ = 4j_\xi^+\}_{\xi=0}^{c-2}$, $\{a_\xi^- = 4j_\xi^- + 3\}_{\xi=0}^{c-1}$, and f_1^* covers $M + c$ gold and $M - c$ mines. We construct $\hat{f}_2 = \{[\hat{a}_\xi^-, \hat{a}_\xi^+]\}_{\xi=0}^{c-1}$ according to f_1^* by setting $\hat{a}_0^- = 0, \hat{a}_0^+ = 4j_0^-, \{\hat{a}_\xi^- = 4j_{\xi-1}^+ + 3, \hat{a}_\xi^+ = 4 \cdot \max(j_\xi^-, j_{\xi-1}^+ + 1)\}_{\xi=1}^{c-1}$. This \hat{f}_2 satisfies condition S_3 and achieves a payoff of $\hat{u}_2 = M - c + 2 + (2c - 1)\rho + (M - c + 1)\mu$. For any f_2' that satisfies S_4 , by Lemma 3.4.20, it covers $M + c$ gold and $M - c$ mines. Denote d as the number of gold that is covered by both f_1^* and f_2' , noting that both f_1^* and f_2' cannot cover the gold at $t = 0$ and $t = 4M + 1$, we have $M + c + M + c - d \leq 2M$, so $d \geq 2c$. Therefore, such f_2' 's payoff

$$\begin{aligned} u_2' &= M + c - d + d\rho + (M - c)\mu \\ &\leq M + c - 2c(1 - \rho) + (M - c)\mu \\ &= \hat{u}_2 - 2 + \rho - \mu < \hat{u}_2. \end{aligned}$$

By the definition of Nash equilibrium, $u_2^* \geq \hat{u}_2$, so f_2^* cannot satisfy S_4 . Therefore, f_2^* must satisfy S_3 , and by Lemma 3.4.20, f_2^* covers $M + c + 1$ gold and $M - c + 1$ mines. So by Lemma 3.4.21, $u_2^* \leq \hat{u}_2$. Therefore, $u_2^* = \hat{u}_2$, i.e. f_2^* always satisfies S_3 and achieves a payoff of $u_2^* = M - c + 2 + (2c - 1)\rho + (M - c + 1)\mu$.

Combining the above results, we can show that for any PNE (f_1^*, f_2^*) , one of f_1^* and f_2^* must satisfy S_3 and achieves a payoff of $M - c + 2 + (2c - 1)\rho + (M - c + 1)\mu$, and the other must satisfy S_4 and achieves a payoff of $M - c + 1 + (2c - 1)\rho + (M - c)\mu$. Therefore,

$$\begin{aligned} W_{\text{Equil}}(b) &= (2M + 1)(1 + \mu) + 2(1 - \rho) + 2(2\rho - \mu - 1)c \\ &= (2M + 1)(1 + \mu) + 2(1 - \rho) + (2\rho - \mu - 1)b. \end{aligned}$$

3. $b \geq 2M + 2$

Since f_1^* and f_2^* can have at most $2M + 2$ segments, when $b \geq 2M + 2$, $f_1^*, f_2^* \in \mathcal{F}_{2M+2}$. There is only one function in \mathcal{F}_{2M+2} , which is the function that covers all

gold and no mines, therefore both f_1^* and f_2^* must be this particular function. So $u_1^* = u_2^* = (2M + 2)\rho$, $W_{\text{Equil}}(b) = (4M + 4)\rho$.

□

Proof of Theorem 3.4.16

Theorem 3.4.16. *For the alternating ordering game, if*

- *for all $b \in \mathbb{Z}^+$ and $M \in \mathbb{Z}^+$, all PNEs in the game with M mines in each line and strategy space \mathcal{L}_b have the same social welfare,*

then $-2 + \rho < \mu < -\rho$.

Proof. We provide constructions showing that if $-2 + \rho < \mu < -\rho$ is not satisfied, there is always some b and M where different PNEs have different social welfare. If $\mu \geq -\rho$, for $M > \frac{2(\rho+\mu)}{1-\rho} + 1$ and $b = 2$, ($\{[0, 0]\}, \{[0, 4M]\}$) is a PNE, ($\{[0, 4 \cdot \lfloor M/2 \rfloor]\}, \{[4 \cdot \lfloor M/2 \rfloor + 3, 4M + 1]\}$) is another PNE, and their social welfare is different. If $\mu \leq -2 + \rho$, for $M > \frac{2(-\mu-\rho)}{1-\rho}$ and $b = 2$, ($\{[4M - 1, 4M + 1]\}, \{[3, 4M + 1]\}$) is a PNE, ($\{[4 \cdot \lfloor M/2 \rfloor + 3, 4M + 1]\}, \{[0, 4 \cdot \lfloor M/2 \rfloor]\}$) is a PNE, and their social welfare is different. □

3.5 Related Work

There has been research exploring the results of representing player strategies using formal computational models. [86] proposes using programs to represent player strategies and analyzes program equilibrium in a finite two-player game. [32] extends the results, representing strategies as Turing machines. Another line of research uses various kinds of automata to model a player's strategy in non-congestion games [5], such as repeated prisoner's dilemma [72, 81]. Automata are typically used to model bounded rationality [67, 73] or certain learning behavior [20, 43]. [68] presents asymptotic results on equilibrium payoff in repeated normal-form games when automaton sizes meet certain conditions. There has also been research exploring structural strategy spaces in congestion games. [4] considers a player-specific network congestion game

where each player has a set of forbidden edges. [19] studies computing mixed Nash equilibria in a broad class of congestion games with strategy spaces compactly described by a set of linear constraints. Unlike our research, none of the above research defines a hierarchy of player capabilities or characterizes the effect of the hierarchy on game outcomes.

The results in this research are obtained in the context of network congestion games. Congestion games are originally introduced in [77], and subsequently have been applied in many areas including drug design [70], load balancing [94], and network design [59]. There is a rich literature on different aspects of congestion games, including their computational characteristics [3], efficiency of equilibria [24, 22], and variants such as weighted congestion games [21] or games with unreliable resources [69]. There is also a large body of research in the context of network congestion games. [33, 29] establish the computational complexity of finding Nash equilibria. [79] focuses on the efficiency of equilibria in terms of the price of anarchy. [2, 92, 6] analyze the effect of information in network congestion games. Many variants and extensions of network congestion games have also been studied in the literature, including network congestion games with tolls [85, 26, 31, 23], weighted network congestion games [63, 34, 46], Bayesian network congestion games [15, 17, 88], and network congestion games with malicious players [54, 11]. To the best of our knowledge, this research is the first to introduce and study network congestion games with a distance bound, i.e., a bound on the number of edges each player can use.

This distance bound can also be applied to many of the variants of network congestion games referenced above and generate interesting new research problems. For example, for network congestion games with tolls [85, 26, 31, 23], the effect of tolls may depend on the distance bounds of the players. Intuitively speaking, when players have larger distance bounds (more capable), more extensive/comprehensive tolls may be needed to achieve a better regulation effect. It would be interesting to study how the effects of tolls change with the distance bounds.

In Bayesian network congestion games [15, 17, 88], the delay function of each edge depends on the state of nature, which is randomly sampled from some prior

distribution and hidden from the players. A *principal* can observe the state of nature and *signal* the players, i.e., reveal information on the state of nature to the players. The principal can reveal different information to different players, with the objective of minimizing the overall delay (maximizing social welfare). When the distance bounds of the players change, the optimal signaling scheme and the corresponding effect on the social welfare may change correspondingly. For example, when players have larger distance bounds (more powerful), since they have more options on the paths to use, the principal may need a more complex signaling scheme to achieve the optimal social welfare, but at the same time the optimal social welfare may be better than what can be achieved with a smaller distance bound. It would be interesting to study how distance bounds affect the optimal signaling scheme and its effect.

In network congestion games with malicious players [54, 11], there are players whose objective is to maximize the overall delay (minimize social welfare) instead of minimizing their own delay. Intuitively speaking, if we increase the distance bounds of the malicious players, the malicious players can occupy more edges and cause more severe congestions, which may cause social welfare to decrease. Increasing the distance bounds of the normal players may have a similar effect to increasing the distance bounds of the players in a DNC, which, as we have shown in Section 3.4, may cause social welfare to increase, stays the same, or decrease, depending on the game setting. It would be interesting to study the effect of varying the distance bounds of the malicious players on the social welfare, as well as how the relative capabilities between normal players and malicious players affect game outcomes.

Chapter 4

Network Congestion Games with Incomplete Information on Player Capability and Multi-Round Play

In this chapter, we extend the original DNC model introduced in Chapter 3 with incomplete information on player capability and multi-round play to obtain a richer set of game models with hierarchies of player capabilities. These new game models open up the space of research on new phenomena involving player capabilities. Section 4.1 formally introduces the new game models. Section 4.2 presents the results on the existence of different types of equilibrium solutions for the new game models. Section 4.3 establishes the computational complexity of finding a pure Nash equilibrium in the new game models. Section 4.4 studies the emergence of locally suboptimal play in an example game with incomplete information on player capability and repeated play. Section 4.5 discusses the related work.

4.1 Models

We propose several new models of games that extend the original DNC model with incomplete information on player capability and multi-round play.

4.1.1 DNC with Mixed Capability

The first model is DNC with mixed capability, where different players can have different distance bounds within the same game. This game belongs to the general class of mixed capability games as introduced in [52].

Definition 4.1.1. *An instance of Distance-bounded Network Congestion game with mixed capability (DNC-mixed) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (b_i)_{i \in \mathcal{N}}, (d_e)_{e \in \mathcal{E}})$ where:*

- $b_i \in \mathbb{N}$ is the distance bound for player i .
- All other symbols have the same meaning as in Definition 3.2.1.

4.1.2 DNC with Private Capability

The second model incorporates incomplete information on player capability, where players are uncertain about the distance bounds of the other players:

Definition 4.1.2. *An instance of Distance-bounded Network Congestion game with private capability (DNC-private) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (\mathcal{B}_i)_{i \in \mathcal{N}}, \mu, (d_e)_{e \in \mathcal{E}})$ where:*

- $\mathcal{B}_i \subseteq \mathbb{N}$ is the set of possible distance bounds for player i .
- $\mu \in \Delta \mathcal{B}$ is the joint distribution of the capabilities of all players. μ is common knowledge to all players.
- All other symbols have the same meaning as in Definition 3.2.1.

A pure strategy for player i is a function $s_i : \mathcal{B}_i \rightarrow \mathcal{L}$ (\mathcal{L} denoting the set of all $s-t$ simple paths) specifying the path to pick for each of their possible distance bound, satisfying $s_i(b) \in \mathcal{L}_b$ for all $b \in \mathcal{B}_i$. A strategy profile \mathbf{s} is a (Bayesian) Nash equilibrium if no player can improve their *expected* delay by unilaterally changing their strategy, i.e., $\forall i \in \mathcal{N}, \forall s'_i \neq s_i, \mathbb{E}_{(b_1, \dots, b_n) \sim \mu} c_i(\mathbf{s}_{-i}(\mathbf{b}_{-i}), s'_i(b_i)) \geq \mathbb{E}_{(b_1, \dots, b_n) \sim \mu} c_i(\mathbf{s}_{-i}(\mathbf{b}_{-i}), s_i(b_i))$.

4.1.3 Sequential DNC

The third model is a sequential version of DNC, where in each round of the game, every player simultaneously chooses the next edge in their path:

Definition 4.1.3. *An instance of sequential Distance-bounded Network Congestion game (seq-DNC) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (b_i)_{i \in \mathcal{N}}, (d_e)_{e \in \mathcal{E}})$ where the symbols have the same meaning as in Definition 4.1.1.*

The game is played in rounds. Before the first round, all players start at s . In each round, every player simultaneously chooses the next edge in their path. Denote the (partial) path chosen by player i before the beginning of round τ as $h_i^\tau = (h_i^\tau[0], \dots, h_i^\tau[k])$ where $h_i^\tau[0] = s, k = \min(b_i, \tau - 1), (h_i^\tau[\kappa], h_i^\tau[\kappa + 1]) \in \mathcal{E}$ for $0 \leq \kappa \leq k - 1$. Denote $|h_i^\tau| = k$ as the length of h_i^τ . A history in the game at round τ is a partial path profile consisting of the partial paths chosen by each player before round τ : $h^\tau \stackrel{\text{def}}{=} (h_1^\tau, \dots, h_n^\tau)$. Denote \mathcal{H}^τ as the set of possible histories at round τ , and $\mathcal{H} = \bigcup_\tau \mathcal{H}^\tau$ as the set of all possible histories of the game. A pure strategy for player i specifies the next edge to take in their path for every possible history, $s_i : \mathcal{H} \rightarrow \mathcal{V} \cup \{\perp\}$. We require a valid strategy to satisfy the following requirements.

- *The choice of next edge follows the partial path so far: $(h_i^\tau[|h_i^\tau|], s_i(h^\tau)) \in \mathcal{E}$.*
- *It is possible to reach sink t with a path under the distance bound b_i following the chosen next edge: $\exists (p_1, \dots, p_k)$ s.t. $(s_i(h^\tau), p_1) \in \mathcal{E}, (p_j, p_{j+1}) \in \mathcal{E}$ for all $1 \leq j \leq k - 1, p_k = t$, and $k \leq b_i - |h_i^\tau| - 1$.*

A player is finished if they already reached sink t , i.e., $h_i^\tau[|h_i^\tau|] = t$, or there is no valid choice according to the above requirements. We use a placeholder to set $s_i(h^\tau) = \perp$ after player i is finished. The game ends after all players are finished. The delay experienced by each player is then determined in the same way as in the standard DNC, based on the paths chosen by each player. This game is finite since the game ends in at most $\max_i(b_i)$ rounds.

4.1.4 Sequential DNC with Private Capability

The fourth model is the sequential version of DNC where players are uncertain about the distance bounds of the other players:

Definition 4.1.4. *An instance of **sequential Distance-bounded Network Congestion** game with **private** capability (seq-DNC-private) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (\mathcal{B}_i)_{i \in \mathcal{N}}, \mu, (d_e)_{e \in \mathcal{E}})$ where the symbols have the same meaning as in Definition 4.1.2.*

At the beginning of the game, nature draws the distance bounds for each player (b_1, \dots, b_n) from the distribution μ . Then game is played in the same way as seq-DNC (see Definition 4.1.3). In seq-DNC-private, a pure strategy for player i specifies the next edge to take in their path for every possible history and every possible distance bound of theirs, $s_i : \mathcal{H} \times \mathcal{B}_i \rightarrow \mathcal{V} \cup \{\perp\}$. The goal of each player is to minimize their expected delay over capabilities (b_1, \dots, b_n) sampled from the distribution μ (see Definition 4.1.2).

4.1.5 Repeated DNC

The fifth model is repeated DNC where players repeatedly play a DNC for a finite number of rounds:

Definition 4.1.5. *An instance of **repeated Distance-bounded Network Congestion** game (rep-DNC) is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (b_i)_{i \in \mathcal{N}}, (d_e)_{e \in \mathcal{E}}, T)$ where:*

- $T \in \mathbb{N}^+$ is the number of rounds of the game.
- All other symbols have the same meaning as in Definition 4.1.1.

In each round, players simultaneously choose their path for that round, and the delay for each player in that round is determined in the same way as in the standard DNC. The goal of each player is to minimize the sum of their delay over all T rounds. In this game, the history at the beginning of round τ consists of the paths chosen by each player in each of the previous rounds. Similar to seq-DNC (Definition 4.1.3), we denote the paths chosen by player i in the history at the beginning of round τ as

$h_i^\tau = (h_i^\tau[1], \dots, h_i^\tau[\tau - 1])$, where $h_i^\tau[j]$ is the path chosen by i in round j . Then the history at round τ is denoted by $h^\tau \stackrel{\text{def}}{=} (h_1^\tau, \dots, h_n^\tau)$. The set of possible histories is denoted by \mathcal{H}^τ and $\mathcal{H} = \bigcup_\tau \mathcal{H}^\tau$. A pure strategy for player i specifies the path to choose in the next round for every possible history, $s_i : \mathcal{H} \rightarrow \mathcal{L}_{b_i}$.

4.1.6 Repeated DNC with Private Capability

The last model is repeated DNC where players are uncertain about the distance bounds of the other players:

Definition 4.1.6. *An instance of **repeated Distance-bounded Network Congestion game with private capability (rep-DNC-private)** is a tuple $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (\mathcal{B}_i)_{i \in \mathcal{N}}, \mu, (d_e)_{e \in \mathcal{E}}, T)$ where:*

- $T \in \mathbb{N}^+$ is the number of rounds of the game.
- All other symbols have the same meaning as in Definition 4.1.2.

At the beginning of the game, nature draws the distance bounds for each player (b_1, \dots, b_n) from the distribution μ . This set of distance bounds applies for every round of the game. Then the game is played in the same way as rep-DNC (Definition 4.1.5). Here, a pure strategy for player i specifies the path to choose in the next round for every possible history and every possible distance bound of theirs, $s_i : \mathcal{H} \times \mathcal{B}_i \rightarrow \mathcal{L}$, satisfying $s_i(\cdot, b) \in \mathcal{L}_b$ for all $b \in \mathcal{B}_i$. The goal of each player is to minimize their expected sum of delays of all rounds over capabilities (b_1, \dots, b_n) sampled from the distribution μ (see Definition 4.1.2).

4.2 Existence of Equilibrium Solutions

In this section, we study the existence of different types of equilibrium solutions for the game models introduced in Section 4.1. We consider 1) Nash equilibrium, 2) subgame-perfect equilibrium [82], and 3) sequential equilibrium [57]. We consider subgame-perfect equilibrium and sequential equilibrium only for game models with

	NE	PNE	SPE	PSPE	SE	PSE
DNC-mixed		Yes [77]	Reduce to NE and PNE for one-shot games			
DNC-private	Yes [66]	Yes (Theorem 4.2.1)				
seq-DNC		Yes (Theorem 4.2.2)	Yes [83]	Not always (Proposition 4.2.3)	Yes [83, 57]	Not always (Corollary 4.2.4)
seq-DNC-private		Not always (Proposition 4.2.5)		Not always (Corollary 4.2.6)		Not always (Corollary 4.2.7)
rep-DNC		Yes (Theorem 4.2.8)		Yes (Corollary 4.2.9)		Yes (Corollary 4.2.10)
rep-DNC-private		Not always (Theorem 4.2.11)		Not always (Corollary 4.2.12)		Not always (Corollary 4.2.13)

Table 4.1: Existence results for different equilibrium classes. NE: Nash equilibrium, PNE: pure Nash equilibrium, SPE: subgame-perfect equilibrium, PSPE: pure strategy subgame-perfect equilibrium, SE: sequential equilibrium, PSE: pure strategy sequential equilibrium.

multi-round play, since these concepts degenerate to standard Nash equilibrium in one-shot games. We also consider the pure-strategy version for each of the above types of equilibrium.

The results are summarized in Table 4.1. Since all the game models we consider here are finite games, Nash equilibrium (NE) exists [66]. Since all the multi-round games we consider here can be represented as finite extensive form games with perfect recall, subgame-perfect equilibrium (SPE) and sequential equilibrium (SE) exists [57, 83]. Since DNC-mixed is a congestion game, pure Nash equilibrium (PNE) exists [77]. We establish the results on the existence of pure strategy equilibria for the rest of the game models in the following sections.

4.2.1 DNC-private

Theorem 4.2.1. *Every DNC-private has a pure Nash equilibrium.*

Proof. For any instance of DNC-private $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (\mathcal{B}_i)_{i \in \mathcal{N}}, \mu, (d_e)_{e \in \mathcal{E}})$, we can construct a bigger congestion game, which we call the *super-game*, as follows. For every possible capability profile $(b_1, \dots, b_n) \in \mathcal{B}_1 \times \dots \times \mathcal{B}_n$, we construct a copy of the network $(\mathcal{V}^{(b_1, \dots, b_n)}, \mathcal{E}^{(b_1, \dots, b_n)})$ with the same structure as the original network but all delay functions scaled by a factor $d_e^{(b_1, \dots, b_n)}(\cdot) = \mu(b_1, \dots, b_n) \cdot d_e(\cdot)$. The super-game has the same set of players as the original DNC-private. The strategy for each player $i \in \mathcal{N}$ in the super-game is to choose a path for each of their possible capabilities $b_i \in \mathcal{B}_i$, which will be used in the set of network copies corresponding to that capability of him. The total delay of a player in the super-game is the sum of delays of their chosen paths in all network copies. It is easy to see that there is a one-to-one mapping between the strategy space in the super-game and the strategy space in the original game for each player, and that the payoffs in the super-game equal the expected payoffs in the original game under the corresponding strategy profiles. So any (pure) Nash equilibrium in the super-game corresponds to a (pure) Nash equilibrium in the original game. Since the super-game is a congestion game, PNE exists. Therefore, PNE exists for the original DNC-private game. \square

4.2.2 seq-DNC

Theorem 4.2.2. *Every seq-DNC has a pure Nash equilibrium.*

Proof. For any instance of seq-DNC $G = (\mathcal{V}, \mathcal{E}, \mathcal{N}, s, t, (b_i)_{i \in \mathcal{N}}, (d_e)_{e \in \mathcal{E}})$, we can take a PNE $\mathbf{s}^* = (s_1^*, \dots, s_n^*)$ of the corresponding one-shot DNC-mixed game and construct a PNE for the sequential version using the following idea: at each decision point for player i , the action is the next edge on the path s_i^* , regardless of the other players' action histories. Concretely, we construct the pure strategy for player i as $s_i(h^\tau) = s_i^*[\tau]$ for all $h^\tau \in \{(h_1^\tau, \dots, h_n^\tau) \mid h_i^\tau = (s_i^*[0], \dots, s_i^*[\tau - 1])\}$, where $s_i^*[k]$ denotes the k -th point in s_i^* . For all other h^τ , $s_i(h^\tau)$ can be set to an arbitrary choice. We can see that under this constructed strategy profile, no player can obtain a better

delay by unilaterally changing their strategy, since doing so will not affect the realized play of other players. Therefore, this constructed strategy profile is a PNE for the seq-DNC instance. \square

Proposition 4.2.3. *Not every seq-DNC has a pure strategy subgame-perfect equilibrium.*




Figure 4-1: Example subgame of seq-DNC that has no pure strategy subgame-perfect equilibrium. There are two players in the game. Player 1 starts at node a and has distance bound 5, player 2 starts at node b and has distance bound 3. Sink is t . Numbers on each edge specifies the delay function in the format of $(d_e(1), d_e(2))$.

Proof. Figure 4-1 shows an example subgame of seq-DNC that has no pure strategy subgame-perfect equilibrium. To prove by contradiction, assume there is a pure strategy subgame-perfect equilibrium (s_1, s_2) . Consider each player's choice at round 1, i.e. $(s_1(h^1), s_2(h^1))$. There are two possible choices for each player, $s_1(h^1) = c$ or $s_1(h^1) = d$, $s_2(h^1) = d$ or $s_2(h^1) = f$. For every pair of possible $(s_1(h^1), s_2(h^1))$, there is a unique pure strategy Nash equilibrium for the subsequent subgame following the corresponding h^2 . We can compute the equilibrium delay for each player under each of the possible pairs of $(s_1(h^1), s_2(h^1))$, and show that in each case, one of the players can unilaterally change their strategy at h^1 to obtain a lower delay:

- $(s_1(h^1) = c, s_2(h^1) = d)$: $(c_1 = 9, c_2 = 15)$, $s_2(h^1)$ will change to below
- $(s_1(h^1) = c, s_2(h^1) = f)$: $(c_1 = 15, c_2 = 12)$, $s_1(h^1)$ will change to below
- $(s_1(h^1) = d, s_2(h^1) = f)$: $(c_1 = 14, c_2 = 25)$, $s_2(h^1)$ will change to below
- $(s_1(h^1) = d, s_2(h^1) = d)$: $(c_1 = 10, c_2 = 24)$, $s_1(h^1)$ will change to the first.

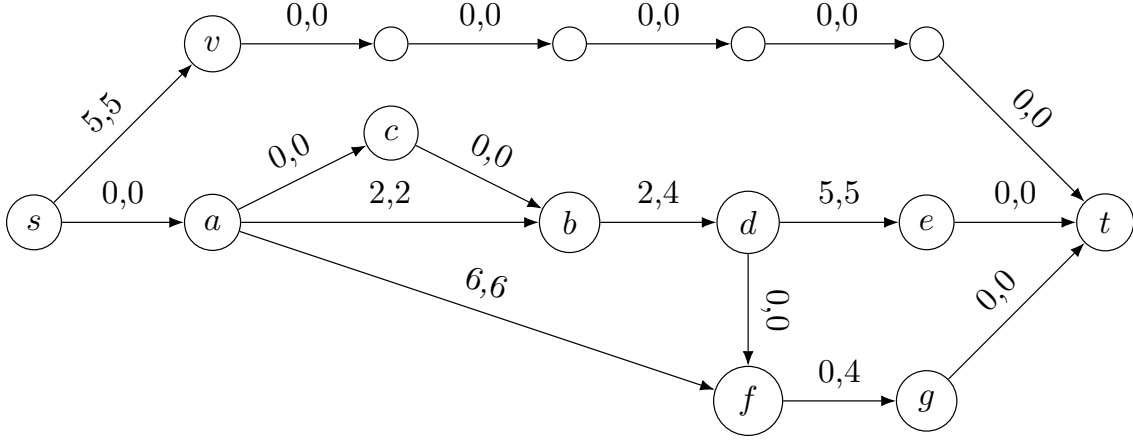


Figure 4-2: Example seq-DNC-private that has no pure Nash equilibrium. There are 2 players. Player 1 has two possible distance bounds $\mathcal{B}_1 = \{4, 7\}$ with $\mu(b_1 = 4) = \epsilon, \mu(b_1 = 7) = 1 - \epsilon$, ϵ is a small number; player 2 has one possible distance bound $\mathcal{B}_2 = \{5\}$. Numbers on each edge specifies the delay function in the format of $(d_e(1), d_e(2))$.

□

Since a pure strategy sequential equilibrium is always a pure strategy subgame-perfect equilibrium, we have:

Corollary 4.2.4. *Not every seq-DNC has a pure strategy sequential equilibrium.*

4.2.3 seq-DNC-private

Proposition 4.2.5. *Not every seq-DNC-private has a pure Nash equilibrium.*

Proof. Figure 4-2 shows an instance of seq-DNC-private that does not have a pure Nash equilibrium. Notice that with distance bound $b = 4$, only one path is available $s \rightarrow a \rightarrow f \rightarrow g \rightarrow t$; with $b = 5$, two paths are available $s \rightarrow a \rightarrow f \rightarrow g \rightarrow t$ and $s \rightarrow a \rightarrow b \rightarrow d \rightarrow e \rightarrow t$. With $b = 7$, all paths are available. Assume there is a PNE (s_1, s_2) . Player 1's choice at the beginning when $b_1 = 7$ has two possibilities: $s_1(h^1, b_1 = 7) \in \{v, a\}$; player 2's choice when both 1 and 2 have taken $s \rightarrow a$ has two possibilities: $s_2\left(h^2 = ((s, a), (s, a))\right) \in \{b, f\}$. So there are four possibilities regarding the values of $s_1(h^1, b_1 = 7)$ and $s_2\left(h^2 = ((s, a), (s, a))\right)$ in the PNE, which

we denote as $(v, b), (v, f), (a, b), (a, f)$. We show here that for each case, one of the players can change their strategy to obtain a lower expected delay:

- (v, b) . Since this is a PNE, player 2 will choose f when 1 took $s \rightarrow v$ and 2 took $s \rightarrow a$, i.e. $s_2\left(h^2 = \left((s, v), (s, a)\right)\right) = f$. So the expected delay is $u_1 = 5 + \epsilon, u_2 = 6 + 3\epsilon$. If player 1 changes its strategy to when $b_1 = 7$, choose $s \rightarrow a$ at the beginning, and then follows $s \rightarrow a \rightarrow c \rightarrow b \rightarrow d \rightarrow f \rightarrow g \rightarrow t$, then its expected delay is $u'_1 = 4 + 2\epsilon < u_1$. So such case cannot be a PNE.
- (a, b) . For this to be a PNE, player 1 will follow $s \rightarrow a \rightarrow c \rightarrow b \rightarrow d \rightarrow f \rightarrow g \rightarrow t$ when $b_1 = 7$. So the expected delay is $u_1 = 4 + 2\epsilon, u_2 = 11 - 2\epsilon$. If player 2 changes its strategy to $s_2\left(h^2 = \left((s, a), (s, a)\right)\right) = f$, then its expected delay is $u'_2 \leq 10 < u_2$. So such case cannot be a PNE.
- (a, f) . For this to be a PNE, player 1 will follow $s \rightarrow a \rightarrow c \rightarrow b \rightarrow d \rightarrow f \rightarrow g \rightarrow t$ when $b_1 = 7$. The expected delay is $u_1 = 6 + 4\epsilon, u_2 = 10$. If player 1 changes its strategy to $s_1(h^1, b_1 = 7) = v$, their expected delay $u'_1 \leq 5 + 5\epsilon < u_1$. So such case cannot be a PNE.
- (v, f) . For this to be a PNE, when player 1 chose v , player 2 will choose $a \rightarrow f$. This has expected delay $u_1 = 5 + 5\epsilon, u_2 = 6 + 4\epsilon$. If player 2 changes to choose b when 1 took $s \rightarrow a$, its expected delay is $u'_2 = 6 + 3\epsilon < u_2$. So such case cannot be a PNE.

□

Since a pure strategy sequential equilibrium/pure strategy subgame-perfect equilibrium is always a PNE, we have:

Corollary 4.2.6. *Not every seq-DNC-private has a pure strategy subgame-perfect equilibrium.*

Corollary 4.2.7. *Not every seq-DNC-private has a pure strategy sequential equilibrium.*




Figure 4-3: Example rep-DNC-private that does not have a PNE. This game has $T = 2$ rounds, $n = 2$ players. Player 1 has two possible distance bounds $\mathcal{B}_1 = \{2, 3\}$ with $\mu(b_1 = 2) = 0.5$ and $\mu(b_1 = 3) = 0.5$; player 2's distance bound is fixed $\mathcal{B}_2 = \{4\}$. Numbers on each edge specifies the delay function in the format of $(d_e(1), d_e(2))$.

4.2.4 rep-DNC

Theorem 4.2.8. *Every rep-DNC has a pure Nash equilibrium.*

Proof. We can construct a PNE of rep-DNC from any PNE of the corresponding stage game: every player always chooses the path in the PNE of the stage game in every round regardless of the history. Formally, given a PNE (s_1^*, \dots, s_n^*) of the corresponding stage game, (s_1, \dots, s_n) where $s_i(\cdot) = s_i^*$ for all $i \in \mathcal{N}$ is a PNE for the rep-DNC instance.

□

Notice that this constructed PNE is also a sequential equilibrium, so

Corollary 4.2.9. *Every rep-DNC has a pure strategy subgame-perfect equilibrium.*

Corollary 4.2.10. *Every rep-DNC has a pure strategy sequential equilibrium.*

4.2.5 rep-DNC-private

Theorem 4.2.11. *Not every rep-DNC-private has a pure Nash equilibrium.*

Proof. Figure 4-3 shows an instance of rep-DNC-private that does not have a PNE. To simplify notations, denote path $s \rightarrow a \rightarrow t$ as p_1 , $s \rightarrow a \rightarrow d \rightarrow t$ as p_2 ,

$s \rightarrow b \rightarrow c \rightarrow d \rightarrow t$ as p_3 , $s \rightarrow a \rightarrow e \rightarrow f \rightarrow t$ as p_4 . Assume there is a PNE (s_1, s_2) . Player 1's choice at round 1 when $b_1 = 3$ has two possibilities $s_1(h^1, b_1 = 3) \in \{p_1, p_2\}$; player 2's choice at round 2 when player 1 chose p_1 and player 2 chose $s_2(h^1)$ in round 1 has two possibilities $s_2\left(h^2 = \left((p_1), (s_2(h^1))\right)\right) \in \{p_3, p_4\}$, since p_3 and p_4 strictly dominates p_1 or p_2 at this information set. So there are four possibilities regarding the values of $s_1(h^1, b_1 = 3)$ and $s_2\left(h^2 = \left((p_1), (s_2(h^1))\right)\right)$ in the PNE, which we denote as $(p_1, p_3), (p_1, p_4), (p_2, p_4), (p_2, p_3)$. We show here that for each case, one of the players can change their strategy to obtain a lower expected delay:

- (p_1, p_3) . Since this is a PNE, $s_1\left(h^2 = \left((p_1), (s_2(h^1))\right), b_1 = 3\right) = p_2$, as p_2 strictly dominates p_1 at this information set. So the expected total delay of player 2 is $u_2 = c_2(p_1, s_2(h^1)) + 0.5 \cdot 20 + 0.5 \cdot 25 = c_2(p_1, s_2(h^1)) + 22.5$, where $c_1(\cdot), c_2(\cdot)$ is the utility function of the stage game. If we change player 2's strategy to $s_2\left(h^2 = \left((p_1), (s_2(h^1))\right)\right) = p_4$ and keeping other parts the same, we have $u'_2 = c_2(p_1, s_2(h^1)) + 21 < u_2$. So such case cannot be a PNE.
- (p_1, p_4) . Since this is a PNE, $s_1\left(h^2 = \left((p_1), (s_2(h^1))\right), b_1 = 3\right) = p_2$. So $u_1 = c_1(p_1, s_2(h^1)) + 0.5 \cdot 60 + 0.5 \cdot 50$. If we change player 1's strategy to $s_1(h^1, b_1 = 3) = p_2$, then $u'_1 \leq 0.5 \cdot c_1(p_1, s_2(h^1)) + 0.5 \cdot c_1(p_2, s_2(h^1)) + 0.5 \cdot 60 + 0.5 \cdot 50 < u_1$, since $c_1(p_2, s_2(h^1)) < c_1(p_1, s_2(h^1))$. So such case cannot be a PNE.
- (p_2, p_4) . Under this PNE, $u_2 = u_2^1 + 0.5 \cdot (21 + u_2^2)$, where u_2^1 is the expected delay of player 2 in round 1 and u_2^2 is the delay of player 2 in round 2 after player 1 chose p_2 in round 1 (i.e. $h_1^2 = (p_2)$), under the strategy profile of this PNE. If we change player 2's strategy to $s_2\left(h^2 = \left((p_1), (s_2(h^1))\right)\right) = p_3$ and keep the rest the same, $u'_2 = u_2^1 + 0.5 \cdot (20 + u_2^2) < u_2$. So such case cannot be a PNE.
- (p_2, p_3) . Since this is a PNE, $s_1\left(h^2 = \left((p_2), (s_2(h^1))\right), b_1 = 3\right) = p_2$. Thus, $s_2\left(h^2 = \left((p_2), (s_2(h^1))\right)\right) = p_4$. So $u_1 = 0.5 \cdot c_1(p_1, s_2(h^1)) + 0.5 \cdot c_1(p_2, s_2(h^1)) + 0.5 \cdot (40 + 50)$. If we change player 1's strategy to $s_1(h^1, b_1 = 3) = p_1$ and $s_1\left(h^2 = \left((p_1), (s_2(h^1))\right), b_1 = 3\right) = p_2$, then $u'_1 = c_1(p_1, s_2(h^1)) + 0.5 \cdot (40 + 35)$.

So $u'_1 - u_1 = 0.5 \cdot c_1(p_1, s_2(h^1)) - 0.5 \cdot c_1(p_2, s_2(h^1)) - 0.5 \cdot 15$. Since $c_1(p_1, s_2(h^1)) - c_1(p_2, s_2(h^1)) \leq 10$, $u'_1 < u_1$. So such case cannot be a PNE.

□

Corollary 4.2.12. *Not every rep-DNC-private has a pure strategy subgame-perfect equilibrium.*

Corollary 4.2.13. *Not every rep-DNC-private has a pure strategy sequential equilibrium.*

4.3 Complexity Results

This section presents results on the complexity of finding a pure Nash equilibrium for the game models introduced in Section 4.1. We focus on the game models where PNE is guaranteed to exist: DNC-mixed, DNC-private, seq-DNC, rep-DNC. We prove the complexity of finding a PNE for all four are PLS-complete.

4.3.1 DNC-mixed

Theorem 4.3.1. *Finding a PNE in DNC-mixed is PLS-complete.*

Proof. We can show that finding a PNE in DNC-mixed belongs to PLS by applying the same proof that shows finding a PNE in DNC belongs to PLS as in Lemma 3.3.1. For PLS-completeness, DNC is a special version of DNC-mixed, so the reduction is trivial. □

4.3.2 DNC-private

Theorem 4.3.2. *Finding a PNE in DNC-private is PLS-complete.*

Proof. First we show that finding a PNE in DNC-private is in PLS. Theorem 4.2.1 provides a construction of a super-game for any instance of DNC-private, where any PNE in the super-game corresponds to a PNE in the original DNC-private. The size

of the super-game is polynomial to the size of the original DNC-private instance, since representing a general μ in DNC-private requires size $\prod_{i=1}^n |\mathcal{B}_i|$, which equals the number of network copies in the super-game. Therefore, we only need to show that there exists a polynomial time algorithm for computing the best response in the super-game. To solve a best response for player i in the super-game, it suffices to solve a best response path for each of their possible capabilities $b_i \in \mathcal{B}_i$ in the set of network copies corresponding to b_i . Such best response paths can be solved in polynomial time using a dynamic programming algorithm (see Lemma 3.3.1).

To show PLS-completeness, we can trivially reduce any DNC-mixed instance to DNC-private by setting $|\mathcal{B}_i| = \{b_i\}$ for all player i . This completes the proof. \square

4.3.3 rep-DNC

Theorem 4.3.3. *Finding a PNE in rep-DNC is PLS-complete.*


Proof. The proof of Theorem 4.2.8 provides a construction of a PNE for any rep-DNC from a PNE of the stage game. Since finding a PNE in the stage game is in PLS, finding a PNE in rep-DNC is in PLS. We note that in any PNE of rep-DNC, the strategy profile at the final round under the realized play history must form a PNE of the stage game, since otherwise one of the player can change their strategy at the final round under the realized play history to improve their payoff. Therefore, we can obtain a PNE of the stage game from any PNE in the rep-DNC, which proves PLS-completeness. \square

4.3.4 seq-DNC


Theorem 4.3.4. *Finding a PNE in seq-DNC is PLS-complete.*

Proof. The proof of Theorem 4.2.2 provides a construction of a PNE for any seq-DNC from a PNE of the corresponding one-shot DNC-mixed game. Since finding a PNE in DNC-mixed is in PLS, finding a PNE in seq-DNC is in PLS.

To prove PLS-completeness, we reduce quadratic threshold games [3] to seq-DNC, similar to the proof of PLS-completeness of the original DNC model in Theorem 3.3.2.



(a) The graph structure



(b) Splitting the vertex containing resource r_{ij}

Figure 4-4: The seq-DNC instance corresponding to a four-player quadratic threshold game. All players share the same distance bound $b = 19$. Non-unit-length edges have labels to indicate their lengths. Dashed gray edges correspond to the S_i^{out} strategies.

Figure 4-4 presents the structure of the seq-DNC instance for the reduction. All constructions are the same as in the proof of Theorem 3.3.2, except we shift the lengths of the edges between s to s_i to both the gray edges between s_i to t_i and the horizontal edges between s_i to v_{1i} .

We aim to show that in any PNE of this seq-DNC instance, the paths taken by each player corresponds to a PNE of the original quadratic threshold game. The same argument used in the proof of Theorem 3.3.2 still works to show that in any PNE of this seq-DNC, 1) in the first round, each player chooses a different s_i , and 2) the player that goes through s_i eventually exits through t_i . The distance bound ensures that the player that goes through s_1 and t_i can only choose from two possible paths in between: the gray path and the right-down path. In this construction, the distances between s and each s_i are the same (distance 1). Therefore, all players reach s_i 's in the same round, so they need to simultaneously choose which of their

two alternative paths to take between s_i and t_i . This is then the same problem as the original DNC constructed in the proof of Theorem 3.3.2. Therefore, any PNE of this seq-DNC corresponds to a PNE of the original DNC in the proof of Theorem 3.3.2, which in turn corresponds to a PNE of the original quadratic threshold game.

□

4.4 Emergence of Locally Suboptimal Play in Repeated DNC with Private Capability

In Chapter 2, we studied the emergence of locally suboptimal play in finitely repeated games with complete information. In such complete information games, local suboptimality occurs due to ‘threats’ between players. In this chapter, we introduced games with incomplete information on player capabilities and multi-round play. For such games, there can be another type of motivation for locally suboptimal play: players may sacrifice some payoff in earlier rounds to hide their capability from other players, in order to get better payoff in the future and maximize their total payoff. In this section, we show an example rep-DNC-private game where local suboptimality emerges from rational play and provide a complete characterization of how it occurs. This result surfaces the social phenomena of concealment and deception.

The example rep-DNC-private game has T rounds, $n = 2$ players. The network structure and delay functions are shown in Figure 4-5. Player 1 has two possible distance bounds $\mathcal{B}_1 = \{2, 3\}$, with $\mu(b_1 = 2) = \phi$ and $\mu(b_1 = 3) = 1 - \phi$, $\phi > 0.5$; player 2’s distance bound is fixed $b_2 = 4$. Denote path $s \rightarrow a \rightarrow t$ as p_1 , $s \rightarrow a \rightarrow d \rightarrow t$ as p_2 , $s \rightarrow b \rightarrow c \rightarrow d \rightarrow t$ as p_3 , $s \rightarrow a \rightarrow e \rightarrow f \rightarrow t$ as p_4 . We refer to the K -th round of the game as the K -th to last round.

In this game, when $b_1 = 3$, player 1 has two available paths p_1 and p_2 , and p_2 strictly dominates p_1 in the stage game. However, as we will show in the following, in any sequential equilibrium of this game, player 1 will play p_1 in all but the last round even when $b_1 = 3$, which means local suboptimality always emerges from rational




Figure 4-5: Example rep-DNC-private where local suboptimality occurs. Numbers on each edge specifies the delay function in the format of $(d_e(1), d_e(2))$.

play in this game.

Lemma 4.4.1. *In any sequential equilibrium, if at the beginning of any round of the game it is common knowledge that $b_1 = 3$, then in the rest of the game (including this round), player 1 takes p_2 and player 2 takes p_4 in the realized play.*

Proof. We rephrase the statement as ‘if it is common knowledge that $b_1 = 3$ at K -th round of the game, then in the final K rounds, player 1 takes p_2 and player 2 takes p_4 in the realized play’. We prove by induction on K . When $K = 1$, p_2 is the dominant strategy for player 1 in the stage game, and p_4 is the best response to p_2 . For general K , the induction hypothesis tells us that regardless of the play in the K -th round, player 1 takes p_2 and player 2 takes p_4 in the realized play in the final $K - 1$ rounds. Therefore, player 1 and player 2 will play according to the only Nash equilibrium of the stage game in round K , which is player 1 takes p_2 and player 2 takes p_4 . This finishes the proof. \square

Theorem 4.4.2. *In any sequential equilibrium, the realized play is: player 2 takes p_3 for all the T rounds; player 1 takes p_1 in the first $T - 1$ rounds, and in the final round takes p_1 if $b_1 = 2$, p_2 if $b_1 = 3$.*

Proof. We prove by induction on T . For $T = 1$, when $b_1 = 3$, player 1’s dominant strategy is p_2 ; when $b_1 = 2$, player 1’s only strategy is p_1 . Since $\phi > 0.5$, player 2’s best response is p_3 . So the statement holds for $T = 1$.

For a general T , consider the T -th round. If $b_1 = 2$, player 1's only strategy is p_1 ; if $b_1 = 3$, denote player 1's strategy at T -th round as choosing p_2 with probability ϵ and p_1 with probability $1 - \epsilon$. Given player 1 chose p_2 in the T -th round, according to Lemma 4.4.1, in the rest $T - 1$ rounds, player 1 takes p_2 and player 2 takes p_4 in the realized play. Given player 1 chose p_1 in the T -th round, player 2's belief on b_1 at the start of $(T - 1)$ -th round according to Bayes rule is $\theta(b_1 = 2) = \frac{\phi}{\phi + (1 - \phi)(1 - \epsilon)} \geq \phi > 0.5$. So by the induction hypothesis, the realized play in the rest $(T - 1)$ rounds is player 2 takes p_3 and player 1 takes p_1 except the final round if $b_1 = 3$. Since the equilibrium strategies in the last $(T - 1)$ rounds does not depend on player 2's choice at T -th round, player 2 will play its best response in the stage game in the T -th round, which is p_3 . Therefore, player 1's total delay under the above strategy profile is:

$$u_1 = \phi \cdot 8T + (1 - \phi) \cdot [\epsilon \cdot (9(T - 1) + 7) + (1 - \epsilon) \cdot (8(T - 1) + 7)]$$

u_1 attains minimum at $\epsilon = 0$, which means any $\epsilon > 0$ cannot be a sequential equilibrium. This finishes the induction proof. □

Corollary 4.4.3. *In any sequential equilibrium, player 1's expected total delay is $u_1^* = 8T - (1 - \phi)$, player 2's expected total delay is $u_2^* = 4T + 2(1 - \phi)$.*

Consider a *frank* player who plays the dominant strategy in the stage game at every round. Then

Proposition 4.4.4. *If player 1 is a frank player, their expected total delay in any sequential equilibrium is $u_1' = 8T + (1 - \phi) \cdot (T - 2)$. It's worse than u_1^* of the rational player 1 for all $T \geq 2$.*

Proof. If player 1 is a frank player, they always play p_1 if $b_1 = 2$ and always play p_2 if $b_1 = 3$. In any sequential equilibrium, player 2 will play p_3 in the first round, then if player 1 plays p_1 , player 2 continues to always play p_3 ; if player 1 plays p_2 , player 2 plays p_4 in all the rest of the game. So the expected total delay of player 1

is $u'_1 = \phi \cdot 8T + (1 - \phi) \cdot (9T - 2) = 8T + (1 - \phi) \cdot (T - 2)$. When $T \geq 2$, $u'_1 > u_1^*$; when $T = 1$, $u'_1 = u_1^*$. \square

Remark Theorem 4.4.2 shows that local suboptimality occurs in *every* sequential equilibrium for the game we consider in this section. In contrast, for finitely repeated games with complete information (the game model we consider in Chapter 2), players repeatedly playing the same stage-game Nash equilibrium is always a subgame-perfect equilibrium of the repeated game, so there always exists a subgame-perfect equilibrium where local suboptimality does not occur. Therefore, incomplete information on player capabilities makes such universal occurrence of local suboptimality possible; with complete information, such universal occurrence of local suboptimality can never happen in finitely repeated games.

Remark For the case $\phi \leq 0.5$, we cannot apply the above proof to show that locally suboptimal play occurs with certainty in any sequential equilibrium. But we can still show that locally suboptimal play occurs with strictly positive probability in any sequential equilibria where $T \geq 4$. In another word, the frank strategy of player 1 which always plays p_2 (the dominant strategy of the stage game) when $b_1 = 3$ cannot be a part of a sequential equilibrium. Assuming this strategy is part of a sequential equilibrium. Then player 2's strategy must be playing p_4 in the first round, then: if player 1 played p_1 in the first round, always play p_3 ; if player 1 played p_2 in the first round, always play p_4 . Then the frank strategy achieves an expected total delay of $u_1 = \phi \cdot (12 + 8(T - 1)) + (1 - \phi) \cdot 9T$. Player 1 can change its strategy to play p_1 until the final round where it plays p_2 if $b_1 = 3$. This will achieve an expected total delay of $u'_1 = 12 + 8(T - 1) - (1 - \phi)$. Thus, $u'_1 - u_1 = (1 - \phi) \cdot (3 - T) < 0$, which means frank strategy of player 1 cannot be part of a sequential equilibrium. Therefore, when $T \geq 4$ and $b_1 = 3$, player 1 needs to mix the locally suboptimal play p_1 in their strategy in any sequential equilibrium.

4.5 Related Work

The research in this chapter presents several new models of games that extend the original DNC model introduced in [93] with incomplete information on player capability and multi-round play. Readers can refer to Section 3.5 for a discussion on the related literature on network congestion games and their variants. To the best of our knowledge, the research in this chapter is the first to consider distance-bounded network congestion games with incomplete information on player capability and multi-round play and establish results regarding the existence of equilibrium solutions and the complexity of finding a pure Nash equilibrium.

Chapter 5

Future Work

In this chapter, we discuss several directions for future work that follow from the research in this thesis. Sections 5.1 to 5.3 presents several directions that stem from our research on the emergence of locally suboptimal behaviors. Sections 5.4 and 5.5 presents directions that stem from our research on the impact of player capability on game outcome.

5.1 Minimum Number of Rounds for Local Suboptimality to Emerge

In Chapter 2, we prove sufficient and necessary conditions on the stage game G for there exists some T and some subgame-perfect equilibria (SPE) of the repeated game $G(T)$ where local suboptimality occurs for 2-player games (Theorems 2.3.1, 2.4.1 and 2.5.1). These results mean that if G satisfies the proven conditions, there exists some T where local suboptimality occurs in $G(T)$; if G does not satisfy the proven conditions, local suboptimality can never occur in any $G(T)$ with any T . For stage games G that satisfy the proven conditions, we further establish values of T above which there always exists some SPE of $G(T)$ where local suboptimality occurs (Corollaries 2.3.4, 2.4.4 and 2.5.3). Formally, we establish expressions for $\bar{T}(G)$ such that for all $T \geq \bar{T}(G)$, there exists some SPE of $G(T)$ where local suboptimality

occurs.

One question remains open: for stage games G that satisfy the proven conditions, what is the minimum T for local suboptimality to occur in $G(T)$? Denote this minimum T for local suboptimality to occur in $G(T)$ as $T_{\min}(G)$. The above $\bar{T}(G)$ established in Corollaries 2.3.4, 2.4.4 and 2.5.3 provides an upper bound for $T_{\min}(G)$. When $T = 1$, $G(1)$ reduces to the stage game G , so local suboptimality cannot happen. Therefore, a straightforward lower bound for $T_{\min}(G)$ is 2. An interesting direction for future research is to obtain the exact expression for $T_{\min}(G)$. In Corollaries 2.3.4, 2.4.4 and 2.5.3, we prove the upper bound $\bar{T}(G)$ by designing constructions of SPEs with $\bar{T}(G)$ rounds where local suboptimality occurs. Proving the exact expression for $T_{\min}(G)$ would additionally require proving that it is impossible to construct an SPE where local suboptimality occurs with less than $T_{\min}(G)$ rounds.

5.2 Sufficient and Necessary Condition for Local Suboptimality for n -Player Games

In Chapter 2, we prove a separate sufficient condition and a separate necessary condition on the stage game G for there exists some T and some subgame-perfect equilibria (SPE) of the repeated game $G(T)$ where local suboptimality occurs for n -player games when mixed strategies are allowed (Theorems 2.8.2 and 2.8.3). In the 2-player case, we are able to prove several properties that hold for 2-player games (see the remark in Section 2.8 for detailed discussions), which allows the proof of the sufficient and necessary condition for the general case where mixed strategies are allowed. It is not clear whether similar properties hold for n -player games. Therefore, what is a sufficient and necessary condition for n -player games when mixed strategies are allowed remains an open problem for future research. One angle of attack is to prove/disprove the generalized versions of the properties we proved for 2-player games that enable the proof of the sufficient and necessary condition. If we can prove the generalized versions of the properties, we can reuse the proof for the 2-player case to prove a

sufficient and necessary condition for the n -player case. Another angle of attack is to use some new proof structure and possibly new mechanisms for constructing SPEs with local suboptimality, different from what are used in the 2-player case.

5.3 Emergence of Local Suboptimality in Other Contexts

In Chapter 2, we thoroughly studied the emergence of local suboptimality in subgame-perfect equilibria of finitely repeated games with complete information. A direction for future research is to study similar research questions (Question 1.1.1, sufficient and necessary conditions for the emergence of local suboptimality, and Question 1.1.2, how to computationally decide if local suboptimality can occur) in other contexts/game models. We present here several examples of such contexts for future research.

The first is to study the emergence of local suboptimality in *Nash equilibria* (NEs) of finitely repeated games with complete information. Unlike SPEs, NEs do not put restrictions on the strategies off the equilibrium paths. The set of NEs is a superset of the set of SPEs. Therefore, the set of stage games where local suboptimality can occur in some NEs of some repeated game is a superset of the set of stage games where local suboptimality can occur in some SPEs of some repeated game. So we anticipate the sufficient and necessary conditions for the emergence of local suboptimality in NEs to be broader than the conditions for SPEs that we proved in Chapter 2.

The second is to study the emergence of local suboptimality in NEs/SPEs of *infinitely repeated games* with complete information, with/without discounting. The set of equilibrium solutions is different between finitely repeated games and infinitely repeated games. A result that is widely used by many textbooks and lecture notes to demonstrate this difference is on repeated prisoner's dilemma [40, 44, 71]: cooperation can arise in SPEs of infinitely repeated prisoner's dilemma but can never arise in SPEs of finitely repeated prisoner's dilemma. In general, it is easier to construct threats in infinitely repeated games than in finitely repeated games. Therefore, we anticipate

the sufficient and necessary conditions on the stage games for the emergence of local suboptimality in infinitely repeated games to be broader than the conditions for finitely repeated games.

The third is to study the emergence of local suboptimality in finitely repeated games with *incomplete information*. In repeated games with incomplete information [8], each player has a set of possible types. The types determine the stage-game payoff matrix. At the start of the game, nature samples and assigns a type for each player for the whole repeated game from a prior (joint) distribution. The prior distribution is known to all players, while the type assignment of each player is only known by themselves. To maximize the expected total payoff in the repeated game, players need to infer the types of the other players and choose strategies based on such beliefs throughout the game. The repeated DNC with private capabilities (rep-DNC-private) game model we introduced in Chapter 4 is an example of finitely repeated games with incomplete information. In Section 4.4, we characterized the sequential equilibria of an example rep-DNC-private game. From that analysis, we can see that analyzing equilibrium solutions for repeated games with incomplete information tends to be more complex than analyzing equilibrium solutions for repeated games with complete information due to the presence of types and beliefs. Therefore, we anticipate the problem of finding a sufficient and necessary condition for the emergence of local suboptimality in finitely repeated games with incomplete information to be technically more challenging than the complete information case as studied in Chapter 2. A viable first step may be to study a restricted case such as each player only has a small number of types (e.g., two types).

5.4 Impact of Player Capability on Learning Dynamics

In Chapter 3, we studied the impact of player capability on social welfare at Nash equilibria. An interesting direction for future research is to study the impact of player

capability on other aspects of the game besides equilibrium solutions. One important aspect is the learning dynamics. The learning dynamics of a game considers the behaviors of the players when they repeatedly play this game while adjusting their strategies over time as a result of their experience in the past play [38, 51, 18]. Compared with equilibrium analyses, the study of learning dynamics provides an alternative perspective on predicting game outcomes.

In the analysis of learning dynamics, we first assume a model for how players adapt their strategies over time. Then, given this model of adaptive behaviors, we can analyze: 1) do the strategies of the players converge over time, 2) if so, what solutions does the learning dynamics converge to, and 3) how quickly does the dynamics converge. Some examples of such models of adaptive behaviors include *fictitious play* and *no-regret learning*. In fictitious play [16], players assume that each of the other players is using a stationary (i.e., time independent) mixed strategy. Players count the empirical frequencies of plays of the other players in the past and use these empirical distributions as their beliefs of the mixed strategies used by the other players. Players then play their best responses to these beliefs in the next round. In no-regret learning [60, 12, 30], players aim to minimize *regret*, which is the difference between the actual total payoff obtained in the repeated play and the best possible total payoff that could have been obtained in hindsight by playing a single action throughout the repeated play.

An interesting direction for future research is to analyze the impact of player capability on learning dynamics in the context of distance-bounded network congestion games (DNC) as introduced in Chapter 3. For each type of learning dynamics, we can study 1) does varying player capability affect whether the learning dynamics converges or not, 2) does varying player capability affect the set of solutions the learning dynamics converges to, and 3) does varying player capability affect how quickly the learning dynamics converges. Intuitively speaking, as players become more capable, they have access to a larger strategy space, which is potentially more difficult to search/optimize over. So the convergence of the learning dynamics may be slower when players are more capable. It would be interesting to study under what condi-

tions of the game will increasing player capability lead to slower, the same, or faster convergence rates of the learning dynamics.

5.5 Impact of Player Capability in Other Game Models

In Chapter 3, we present a general framework for studying the impact of player capability on game outcome. We consider player capability as the size of the strategy spaces, with more capable players having access to a larger strategy space. We performed the study in the context of network congestion games. An interesting direction for future research is to study the impact of player capability in other game models. We present two example game models here to provide some starting ideas.

In auctions [64, 58, 55], we can measure player capability as the granularity of prices that players can bid on. For example, a more capable player can bid in price increments of 1 dollar, while a less capable player can only bid in price increments of 10 dollars. Intuitively speaking, if a player is more capable than their opponents, they have a better chance to win the bid without paying a large premium for outbidding the others. It would be interesting to study how varying player capabilities affects equilibrium outcomes and how the relative capabilities between players affect their payoffs at equilibrium.

In repeated games, we can model the capability of a player as the number of rounds into the past that they can remember. For a player who can only remember up to k rounds into the past, their behavior strategy profile in round T can only depend on the realized play in rounds $T - k$ to $T - 1$. Such ‘memory’ capability limits the space of strategies a player can implement in the repeated games. A player with lower capability (shorter memory) cannot implement strategies that involve long-term dependencies. It would be interesting to study how varying this memory capability of players affects the equilibrium solutions of the repeated games.

Chapter 6

Conclusion

Surfacing, modeling, and explaining different phenomena on how strategic agents interact in various social/economical/political settings has been a long term research goal in game theory. This thesis considers two topics in game theory. The first topic is the emergence of locally suboptimal behavior in finitely repeated games. The second topic is the impact of player capability on game outcome. The research in both topics surfaces and characterizes interesting phenomena of how strategic players interact in game theoretic settings.

6.1 Emergence of Locally Suboptimal Behavior in Finitely Repeated Games

Locally suboptimal behavior refers to players play suboptimally in some rounds of the repeated game (i.e., not maximizing their payoffs in those rounds) while maximizing their total payoffs in the whole repeated game. The emergence of locally suboptimal behavior reflects some fundamental psychological and social phenomena, such as delayed gratification, threats, and incentivized cooperation.

We focus on local suboptimality in subgame-perfect equilibria (SPE) of finitely repeated games with complete information. For 2-player games, we prove sufficient and necessary conditions on the stage game G that ensure that, for all T and all

subgame-perfect equilibria of the repeated game $G(T)$, the strategy profile at every round of $G(T)$ forms a Nash equilibrium of the stage game G . We prove the sufficient and necessary conditions for three cases: 1) only pure strategies are allowed, 2) the general case where mixed strategies are allowed, and 3) one player can only use pure strategies and the other player can use mixed strategies. This is the first sufficient and necessary condition for off-(stage-game)-Nash plays to occur in SPEs of 2-player finitely repeated games.

We further study the effect of changing from pure strategies to mixed strategies on the emergence of local suboptimality. We prove that if local suboptimality can occur before the change, then after changing any player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can still occur. On the other hand, we show that there are games where local suboptimality can never occur before the change, but after changing one player (or both players) from pure-strategies-only to mixed-strategies-allowed, local suboptimality can occur. In addition, we present complete characterizations on when changing any player (or both players) from pure-strategies-only to mixed-strategies-allowed would affect the emergence of local suboptimality, by proving sufficient and necessary conditions on the stage game G such that local suboptimality can never occur before the change but can occur after the change. Our characterizations are fine-grained based on the number of payoff values attainable at stage-game NEs for each player.

We also consider the computational aspect of the problem: Given an arbitrary stage game G , how to (algorithmically) decide if there exists some T and some SPE of $G(T)$ where local suboptimality occurs? Is this problem decidable? We propose an algorithm for the above problem for 2-player games for the general case where mixed strategies are allowed and analyze the computational complexity of this algorithm. This shows that the above problem is decidable for 2-player games where mixed strategies are allowed. The algorithm is based on the sufficient and necessary condition established earlier. We design several efficient methods for checking different parts of the condition by utilizing properties we prove for general games. Naive methods for checking these parts of the condition take exponential time in the worst

case, whereas our methods for checking these parts of the condition take polynomial time in the worst case.

Finally, for n -player games, we prove a sufficient and necessary condition for off-(stage-game)-Nash plays to occur in SPEs for the pure strategy case (i.e., only pure strategies are allowed for all players); for the general case where mixed strategies are allowed, we prove a separate sufficient condition and a separate necessary condition. These conditions are both tighter than what is previously known in the literature. The question of what is a sufficient and necessary condition for n -player games when mixed strategies are allowed remains open.

This part of the thesis provides a comprehensive answer to the question “what are the stage games where locally suboptimal behavior can arise in the corresponding finitely repeated games?” The results in this part may be helpful for analysts who need to predict the outcomes of certain activities that involve the structure of repeated games between strategic agents. For example, if the stage game belongs to the side where locally suboptimal behavior can never occur, the analyst can then focus their attention on stage-game Nash equilibria when predicting the outcome of the repeated game. The results in this part may also provide guidance for regulators and rule/game designers if they want to encourage/discourage locally suboptimal behaviors and threats.

6.2 Impact of Player Capability on Game Outcome

Varying player capabilities can significantly affect the outcomes of strategic games. Developing a comprehensive understanding of how different player capabilities affect the dynamics and overall outcomes of strategic games is therefore an important long-term research goal in the field.

The type of player capability we consider in this research is the size of the strategy space: higher capability means players have access to a larger strategy space.

We propose a framework of using programs in a domain-specific language (DSL) to compactly represent player strategies. Bounding the sizes of the programs available to the players creates a natural capability hierarchy, with more capable players able to deploy more diverse strategies defined by larger programs. Building on this foundation, we study the effect of increasing or decreasing player capabilities on game outcomes such as social welfare at equilibrium.

We propose four capability preference properties that characterize the impact of player capability on social welfare at equilibrium. We call a game capability-positive (resp. capability-negative) if social welfare at equilibrium does not decrease (resp. increase) when players become more capable. A game is max-capability-preferred (resp. min-capability-preferred) if the worst social welfare at equilibrium when players have maximal (resp. minimal) capability is at least as good as any social welfare at equilibrium when players have lower (resp. higher) capability. These are general properties applicable to any types of games.

We introduce a new game, the *Distance-bounded Network Congestion game (DNC)*, as the basis of our study. DNC is a symmetric network congestion game in which each player is subject to a distance bound — i.e., a bound on the number of edges that a player can use. We prove that finding a pure Nash equilibrium (PNE) in DNC is PLS-complete. This is interesting because finding a PNE in symmetric network congestion games is in P [29]. So with the addition of a distance bound, the problem becomes harder. We further prove that computing the best/worst social welfare among PNEs of a DNC is NP-hard and computing the best social welfare among all pure strategy profiles of a DNC is NP-hard.

We instantiate our framework on two variants of DNC, the *Distance-bounded Network Congestion game with Default Action (DNCDA)* and the *Gold and Mines Game (GMG)*, where we define simple DSLs that compactly represent the strategy spaces. For these two DNC variants, for each of the four capability preference properties, we prove sufficient and necessary conditions on the delay/payoff functions such that the property holds for any network topology/game layout. This means that if the delay/payoff function satisfies the proven condition, then the target property (e.g.

capability-positive) holds for all possible network configurations/game layouts; if the delay/payoff function does not satisfy the proven condition, then for any such delay/-payoff function, we can always find a network configuration/game layout where the target property (e.g. capability-positive) does not hold for the game.

For a specific version of GMG called the alternating ordering game, we fully characterize how social welfare at equilibrium varies with player capability by proving the functional form of $W_{\text{equil}}(b)$, where b is the player capability and W_{equil} is the social welfare at equilibrium, in terms of the game parameters. We identify situations where social welfare at equilibrium increases, stays the same, or decreases as players become more capable. This result surfaces an interesting phenomenon that in some situations, increasing player capabilities may deliver a worse overall outcome of the game. This phenomenon occurs since players engage in harmful/wasteful competitions due to their selfish nature. And the level of competition increases as players become more capable, which leads to a decrease in the overall social welfare.

We further extend the DNC model with incomplete information on player capability and multi-round play. For incomplete information on player capability, we adopt the Bayesian game theory framework where we let different player types correspond to different player capabilities. We consider two types of multi-round play: sequential play and repeated play. With these extensions, we introduce 6 new game models: DNC with mixed capability (DNC-mixed), DNC with private capability (DNC-private), sequential DNC (seq-DNC), sequential DNC with private capability (seq-DNC-private), repeated DNC (rep-DNC), and repeated DNC with private capability (rep-DNC-private).

We establish (algorithmic) game theoretic properties in these extensions, regarding the existence of different types of equilibrium solutions and the complexity of finding equilibrium solutions. We prove the existence (or non-existence) of 1) pure Nash equilibrium, 2) pure strategy subgame-perfect equilibrium, and 3) pure strategy sequential equilibrium, for each of the 6 new game models. We further prove the complexity of finding a PNE for DNC-mixed, DNC-private, rep-DNC, and seq-DNC: all four are PLS-complete. PNE does not in general exist for seq-DNC-private and

rep-DNC-private.

Finally, we present an example rep-DNC-private game where local suboptimality emerges from rational play and provide a complete characterization of how it occurs. In complete information games as studied in the first part of this thesis, local suboptimality occurs due to ‘threats’ between players. The analysis of this example rep-DNC-private game shows that, in games with incomplete information on player capabilities and multi-round play, there can be another type of motivation for locally suboptimal play: players may sacrifice some payoff in earlier rounds to hide their capability from other players, in order to get better payoff in the future and maximize their total payoff. This result reflects the phenomena of concealment and deception.

This part of the thesis presents the first systematic analysis of the effect of different player capabilities on the outcomes of strategic games. The results of this research may provide insights and guidance to regulators/rule designers for understanding/predicting how varying player capability would affect the overall social welfare. This may help regulators/rule designers to make better decisions on when to restrict the power of the players. For example, for situations where harmful/wasteful competitions prevail, regulators may consider imposing restrictions on the power/capabilities of the players. Game/rule designers may also improve the designs by reducing the opportunities for such harmful competitions.

Bibliography

- [1] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society*, pages 1041–1063, 1990.
- [2] Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Informational braess’ paradox: The effect of information on traffic congestion. *Operations Research*, 66(4):893–917, 2018.
- [3] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM (JACM)*, 55(6):1–22, 2008.
- [4] Heiner Ackermann and Alexander Skopalik. Complexity of pure Nash equilibria in player-specific network congestion games. *Internet Mathematics*, 5(4):323–342, 2008.
- [5] Sally Almanasra, Khaled Suwais, and Muhammad Rafie. The applications of automata in game theory. In *Intelligent Technologies and Techniques for Pervasive Computing*, pages 204–217. IGI Global, 2013.
- [6] Richard Arnott, Andre De Palma, and Robin Lindsey. Does providing information to drivers reduce traffic congestion? *Transportation Research Part A: General*, 25(5):309–318, 1991.
- [7] Charles Audet, Pierre Hansen, Brigitte Jaumard, and Gilles Savard. Enumeration of all extreme equilibria of bimatrix games. *SIAM Journal on Scientific Computing*, 23(1):323–338, 2001.
- [8] Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- [9] Robert J Aumann and Lloyd S Shapley. *Long-term competition—a game-theoretic analysis*. Springer, 1994.
- [10] David Avis, Gabriel D. Rosenberg, Rahul Savani, and Bernhard von Stengel. Enumeration of nash equilibria for two-player games. *Economic Theory*, 42(1):9–37, 2010.

- [11] Moshe Babaioff, Robert Kleinberg, and Christos H Papadimitriou. Congestion games with malicious players. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 103–112, 2007.
- [12] David E Bell. Regret in decision making under uncertainty. *Operations research*, 30(5):961–981, 1982.
- [13] J-P Benoit and Vijay Krishna. Nash equilibria of finitely repeated games. *International Journal of Game Theory*, 16(3):197–204, 1987.
- [14] Jean-Pierre Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53(4):905–922, 1985.
- [15] Umang Bhaskar, Yu Cheng, Young Kun Ko, and Chaitanya Swamy. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 479–496, 2016.
- [16] George W Brown. Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation*, 13(1):374, 1951.
- [17] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Signaling in bayesian network congestion games: the subtle power of symmetry. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5252–5259, 2021.
- [18] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [19] Hau Chan and Albert Xin Jiang. Congestion games with polytopal strategy spaces. In *IJCAI*, pages 165–171, 2016.
- [20] Raymond Chiong and Michael Kirley. Co-evolutionary learning in the n-player iterated prisoner’s dilemma with a structured environment. In *Australian Conference on Artificial Life*, pages 32–42. Springer, 2009.
- [21] George Christodoulou, Martin Gairing, Yiannis Giannakopoulos, Diogo Poças, and Clara Waldmann. Existence and complexity of approximate equilibria in weighted congestion games. In *ICALP*, 2020.
- [22] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73, 2005.
- [23] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 521–530, 2003.

- [24] José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. On the inefficiency of equilibria in congestion games. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 167–181. Springer, 2005.
- [25] Ghislain-Herman Demeze-Jouatsa. A complete folk theorem for finitely repeated games. *International Journal of Game Theory*, 49(4):1129–1142, 2020.
- [26] Robert B Dial. Network-optimized road pricing: Part i: A parable and a model. *Operations Research*, 47(1):54–64, 1999.
- [27] John Dickhaut and Todd Kaplan. A program for finding nash equilibria. In *Economic and Financial Modeling with Mathematica®*, pages 148–166. Springer, 1993.
- [28] Prajit K Dutta. A folk theorem for stochastic games. *Journal of Economic Theory*, 66(1):1–32, 1995.
- [29] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 604–612, 2004.
- [30] Peter C Fishburn. The foundations of expected utility. *Theory & Decision Library*, 1982.
- [31] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 277–285. IEEE, 2004.
- [32] Lance Fortnow. Program equilibria and discounted computation time. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '09, page 128–133, New York, NY, USA, 2009. Association for Computing Machinery.
- [33] Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The structure and complexity of nash equilibria for a selfish routing game. In *Automata, Languages and Programming: 29th International Colloquium, ICALP 2002 Málaga, Spain, July 8–13, 2002 Proceedings*, pages 123–134. Springer, 2002.
- [34] Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.
- [35] James Friedman. Oligopoly theory. *Handbook of mathematical economics*, 2:491–534, 1982.
- [36] James W Friedman. A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12, 1971.

- [37] Drew Fudenberg, David Levine, and Eric Maskin. The folk theorem with imperfect public information1. *Econometrica*, 62(5):997–1039, 1994.
- [38] Drew Fudenberg and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- [39] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.
- [40] Drew Fudenberg and Jean Tirole. *Game theory*. MIT press, 1991.
- [41] Drew Fudenberg and Yuichi Yamamoto. The folk theorem for irreducible stochastic games with imperfect public monitoring. *Journal of Economic Theory*, 146(4):1664–1683, 2011.
- [42] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. Mathematical Sciences Series. W. H. Freeman, 1979.
- [43] Rawan Ghnemat, Saleh Oqeili, Cyrille Bertelle, and Gérard HE Duchamp. Automata-based adaptive behavior for economic modelling using game theory. In *Emergent properties in natural and artificial dynamical systems*, pages 171–183. Springer, 2006.
- [44] Robert Gibbons. A primer in game theory. 1992.
- [45] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [46] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 142–151. IEEE, 2005.
- [47] Julio González-Díaz. Finitely repeated games: a generalized nash folk theorem. *Games and Economic Behavior*, 55(1):100–111, 2006.
- [48] Olivier Gossner. The folk theorem for finitely repeated games with mixed strategies. *International Journal of Game Theory*, 24(1):95–107, 1995.
- [49] JOHN C HARSANYI. Games with incomplete information played by " bayesian" players, i-iii. *MANAGEMENT SCIENCE*, 14(3), 1967.
- [50] J.E. Hopcroft, R. Motwani, and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Always learning. Pearson Education, 2014.
- [51] Luis R. Izquierdo, Segismundo S. Izquierdo, and Fernando Vega-Redondo. *Learning and Evolutionary Game Theory*, pages 1782–1788. Springer US, Boston, MA, 2012.
- [52] Kai Jia, Martin Rinard, and Yichen Yang. Mixed capability games, 2022.

- [53] David S Johnson, Christos H Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of computer and system sciences*, 37(1):79–100, 1988.
- [54] George Karakostas and Anastasios Viglas. Equilibria for networks with malicious users. *Mathematical Programming*, 110(3):591–613, 2007.
- [55] Paul Klemperer. Auction theory: A guide to the literature. *Journal of economic surveys*, 13(3):227–286, 1999.
- [56] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413, 1999.
- [57] David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, 1982.
- [58] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [59] Shuting Le, Yuhu Wu, and Mitsuru Toyoda. A congestion game framework for service chain composition in NFV with function benefit. *Information Sciences*, 514:512–522, 2020.
- [60] Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, 92(368):805–824, 1982.
- [61] Jean-François Mertens, Sylvain Sorin, and Shmuel Zamir. *Repeated games*, volume 55. Cambridge University Press, 2015.
- [62] Carol A Meyers and Andreas S Schulz. The complexity of welfare maximization in congestion games. *Networks*, 59(2):252–260, 2012.
- [63] Igal Milchtaich. Network topology and equilibrium existence in weighted network congestion games. *International Journal of Game Theory*, 44:515–541, 2015.
- [64] Paul R Milgrom et al. Auction theory. In *Advances in economic theory: Fifth world congress*, volume 1, page 32. Cambridge University Press Cambridge, 1987.
- [65] JF NASH. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
- [66] John F Nash Jr. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- [67] Abraham Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners’ dilemma. *Economics letters*, 19(3):227–229, 1985.
- [68] Abraham Neyman. Cooperation, repetition, and automata. In Sergiu Hart and Andreu Mas-Colell, editors, *Cooperation: Game-Theoretic Approaches*, pages 233–255, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.

- [69] Julian Nickerl and Jacobo Torán. Pure Nash equilibria in a generalization of congestion games allowing resource failures. In *International Symposium on Algorithmic Game Theory*, pages 186–201. Springer, 2021.
- [70] Natalia Nikitina, Evgeny Ivashko, and Andrei Tchernykh. Congestion game scheduling for virtual drug screening optimization. *Journal of computer-aided molecular design*, 32(2):363–374, 2018.
- [71] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- [72] Christos H Papadimitriou. On players with a bounded number of states. *Games and Economic Behavior*, 4(1):122–131, 1992.
- [73] Christos H Papadimitriou and Mihalis Yannakakis. On complexity as bounded rationality. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 726–733, 1994.
- [74] Thomas Quint and Martin Shubik. A theorem on the number of nash equilibria in a bimatrix game. *International Journal of Game Theory*, 26:353–359, 1997.
- [75] Roy Radner. Repeated principal-agent games with discounting. *Econometrica: Journal of the Econometric Society*, pages 1173–1198, 1985.
- [76] Roy Radner, Roger Myerson, and Eric Maskin. An example of a repeated partnership game with discounting and with uniformly inefficient equilibria. *The Review of Economic Studies*, 53(1):59–69, 1986.
- [77] Robert W Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [78] Tim Roughgarden. How unfair is optimal routing? In *Symposium on Discrete Algorithms: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, volume 6, pages 203–204, 2002.
- [79] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.
- [80] Ariel Rubinstein. Equilibrium in supergames with the overtaking criterion. *Journal of economic theory*, 21(1):1–9, 1979.
- [81] Ariel Rubinstein. Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 39(1):83–96, 1986.
- [82] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfragerträgeit: Teil i: Bestimmung des dynamischen preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft / Journal of Institutional and Theoretical Economics*, 121(2):301–324, 1965.

- [83] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory*, 4(1):25–55, mar 1975.
- [84] Lones Smith. Necessary and sufficient conditions for the perfect finite horizon folk theorem. *Econometrica*, 63:425–30, 02 1995.
- [85] Chaitanya Swamy. The effectiveness of stackelberg strategies and tolls for network congestion games. In *SODA*, pages 1133–1142. Citeseer, 2007.
- [86] Moshe Tennenholtz. Program equilibrium. *Games and Economic Behavior*, 49(2):363–373, 2004.
- [87] J v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.
- [88] Shoshana Vasserman, Michal Feldman, and Avinatan Hassidim. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [89] J Von Neumann and O Morgenstern. Theory of games and economic behavior. 1944.
- [90] Bernhard von Stengel. New maximal numbers of equilibria in bimatrix games. *Discrete & Computational Geometry*, 21(4):557–568, 1999.
- [91] Bernhard Von Stengel. Computing equilibria for two-person games. *Handbook of game theory with economic applications*, 3:1723–1759, 2002.
- [92] Manxi Wu, Saurabh Amin, and Asuman E Ozdaglar. Value of information in bayesian routing games. *Operations Research*, 69(1):148–163, 2021.
- [93] Yichen Yang, Kai Jia, and Martin Rinard. On the impact of player capability on congestion games. In *Algorithmic Game Theory: 15th International Symposium, SAGT 2022, Colchester, UK, September 12–15, 2022, Proceedings*, pages 311–328. Springer, 2022.
- [94] Fenghui Zhang and Michael Mao Wang. Stochastic congestion game for load balancing in mobile-edge computing. *IEEE Internet of Things Journal*, 8(2):778–790, 2020.